Annales Academiæ Scientiarum Fennicæ
Mathematica
Volumen 28, 2003, 3–34

ATTRACTING DYNAMICS OF EXPONENTIAL MAPS

Dierk Schleicher

International University Bremen, School of Engineering and Science
Postfach 750 561, D-28725 Bremen, Germany; dierk@iu-bremen.de

Abstract. We give a complete classification of hyperbolic components in the space of iterated exponential maps $z \mapsto \lambda \exp(z)$, and we describe a preferred parametrization of those components. More precisely, we associate to every hyperbolic component of period n a finite symbolic sequence of length $n-1$, we show that every such sequence is realized by a hyperbolic component, and the hyperbolic component specified by any such sequence is unique. This leads to a complete classification of all exponential maps with attracting dynamics, which is a fundamental step in the understanding of exponential parameter space.

1. Introduction

This paper is part of the program to describe the dynamics of exponential maps $\lambda \exp$ 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 [DH1] 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, much of it based on methods and results developed initially for the Mandelbrot set. Among transcendental entire maps, the exponential family $\lambda \exp$ stands out as the simplest family. It is expected that a good understanding of this space of maps will guide the way for further progress on a study of more classes of entire maps.

For holomorphic spaces of rational maps, such as the space $z \mapsto z^2 + c$ of quadratic polynomials, it is known from the work of Mañe, Sad and Sullivan [MSS] that the set of *structurally stable* maps is open and dense: a map is structurally stable if it has a neighborhood in which all maps are topologically (and even quasiconformally) conjugate. The complementary locus is called the *bifurcation locus* within this space of maps; it is closed and nowhere dense. An investigation of a space of rational maps thus starts with a description of connected components in the space of structurally stable maps, called *stable components*. In most spaces of rational maps (those in which the bifurcation locus is non-empty and not every map has an indifferent orbit), all known stable components consist of *hyperbolic* maps: there is a uniformly expanding metric in a neighborhood of the Julia set, which is equivalent to the fact that all critical points converge to attracting or

2000 Mathematics Subject Classification: Primary 30D05, 33B10, 37B10, 37C15, 37C20, 37F10, 37F20, 37F45.

superattracting cycles [M2]. It is conjectured that stable components are always hyperbolic; unfortunately, this has not yet been confirmed for any space.

One reason why quadratic polynomials are the simplest non-trivial rational maps is because they have only one critical point (except for the fixed point ∞) and of lowest multiplicity; critical points determine the dynamics of rational maps to a large extent. For the space of quadratic polynomials, the bifurcation locus is the boundary of the Mandelbrot set. Douady and Hubbard [DH1] have developed a complete conjectural description of the topology of the Mandelbrot set and its boundary, and they showed that it is a true description if and only if the Mandelbrot set is locally connected. This would imply that every stable component was hyperbolic. They also provided a complete classification of hyperbolic components as part of the complete description of the topology. The importance of the study of the space of quadratic polynomials stems not only from the fact that they are the simplest class of rational maps, but also because renormalization theory [DH2] shows that quadratic polynomials have universal properties; as a consequence, every non-trivial bifurcation locus in spaces of rational maps contains infinitely many homeomorphic copies of the Mandelbrot set [Mc3] (or of *Multibrot sets*, which are the analogues for maps $z^d + c$ with $d > 2$).

It will be a long way to establish analogous results for transcendental maps, or even to find out which results have analogues and in which sense. The fact that structurally stable maps are dense in many spaces of transcendental entire maps has been established by Eremenko and Lyubich [EL2]; this includes the space of exponential maps. The decisive role of critical points (or critical values) for rational maps is, for transcendental maps, assumed by either critical values or asymptotic values, which are jointly known as singular values. Exponential maps $\lambda \exp$ have only one singular value of the simplest kind: they have no critical values, 0 is the only asymptotic value, and every $\lambda \exp$ is a universal cover $\mathbf{C} \to \mathbf{C} \setminus \{0\}$. This makes them good candidates for prototypes of transcendental maps. Probably for this reason, the space of exponential maps has been studied more than any other space of transcendental maps. An exponential map will be called *hyperbolic* if it has an attracting orbit (which necessarily attracts the singular orbit, so there can be at most one attracting orbit). A structurally stable component is called hyperbolic if it consists of hyperbolic maps (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 transcendental maps).

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 more general entire maps), but a description of the global structure of parameter space was in terms of pictures and conjectures. Eremenko and Lyubich conjectured that every structurally stable exponential map is hyperbolic, so the union of the hyperbolic components would be dense in λ-space.

For a hyperbolic exponential map, the Julia set has measure zero [EL1], [EL3] but Hausdorff dimension two [Mc2]. There are some results on the topology, in particular if the attracting orbit has period one [AO]. It seems possible to give a more complete description of the topology for exponential maps with attracting orbits of arbitrary periods.

In this paper, we give a complete description of the space of hyperbolic exponential maps. 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 [DH1] about the Mandelbrot set. An earlier version of this paper was circulated as [S3].

The bifurcation locus of exponential maps is not locally connected. It would be interesting to have a topological criterion (analogous to local connectivity of the Mandelbrot set) which would imply the validity of the conjecture of Eremenko and Lyubich. It seems possible that this could be done in terms of fibers as discussed in [S2]. There are two more conjectures in [EL2] which have now been established [S1], [RS]; they are explained at the end of Section 7.

The set of parameters λ for which there is a (necessarily unique) attracting periodic orbit is clearly open; connected components where this happens are *hyperbolic components*. The period of the attracting orbit is constant throughout the component. Our object is to classify hyperbolic components in the λ parameter plane, where λ ranges over $\mathbf{C}\setminus\{0\}$. It is known [BR], [EL2], [DGH] that all exponential maps with attracting orbits of period 1 are contained in a single hyperbolic component which is bounded in $\mathbf{C}$; it contains a neighborhood of 0. For period 2, there is a unique hyperbolic component which is contained in a left half plane and unbounded to the left; see also Section 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 1 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 $\lambda \exp$. Each of them is characterized by a sequence*

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

(its "intermediate external address"), where s_1, $s_2, \ldots, s_{n-2} \in \mathbf{Z}$ with $s_1 = 0$, and $s_{n-1} \in (\mathbf{Z}+\frac{1}{2})$. 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 corresponding sequences $s_1, s_2, \ldots, s_{n-1}$.

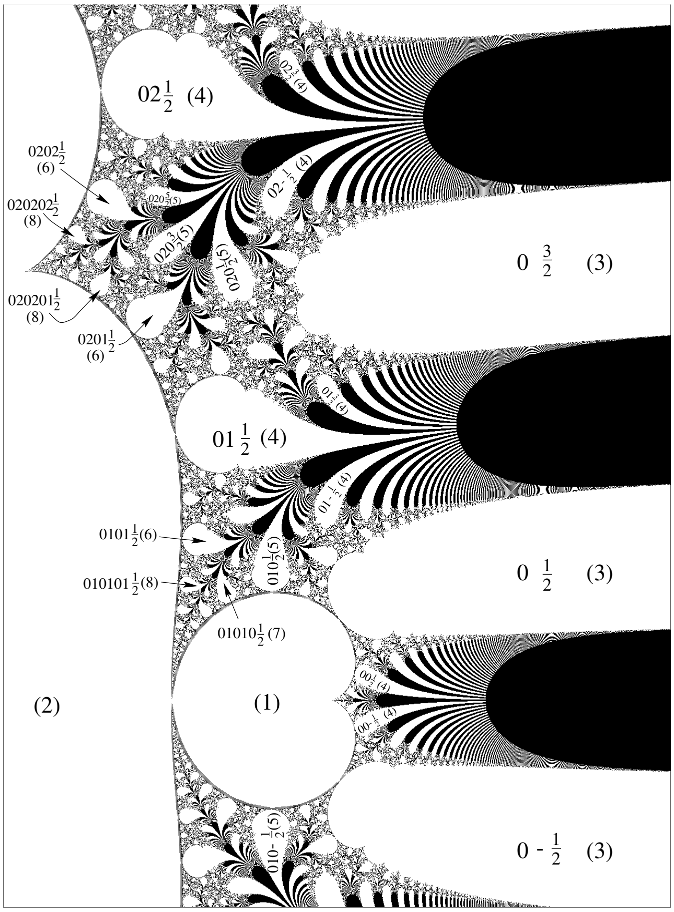

Figure 1. The space of parameters λ for exponential maps $\lambda\exp$, with hyperbolic components indicated in white. Various hyperbolic components are labeled by their intermediate external addresses, or briefly by their periods (in parentheses). The picture has kindly been contributed by Jack Milnor: for every pixel, an approximate test is performed whether or not the corresponding map $\lambda\exp$ 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.

In particular, between any pair of consecutive hyperbolic components of period n, there are infinitely many hyperbolic components of period $n+1$, ordered like $\mathbf{Z}$.

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

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

be the unique cycle of periodic Fatou components for $\lambda\exp$, where $0 \in U_1$, where $\lambda \in U_2$, and where U_n contains a left half plane. Here $\lambda\exp\colon U_n \to U_1 \setminus \{0\}$ is a universal cover, and all other $\lambda\exp\colon U_k \to U_{k+1}$ are conformal isomorphisms.

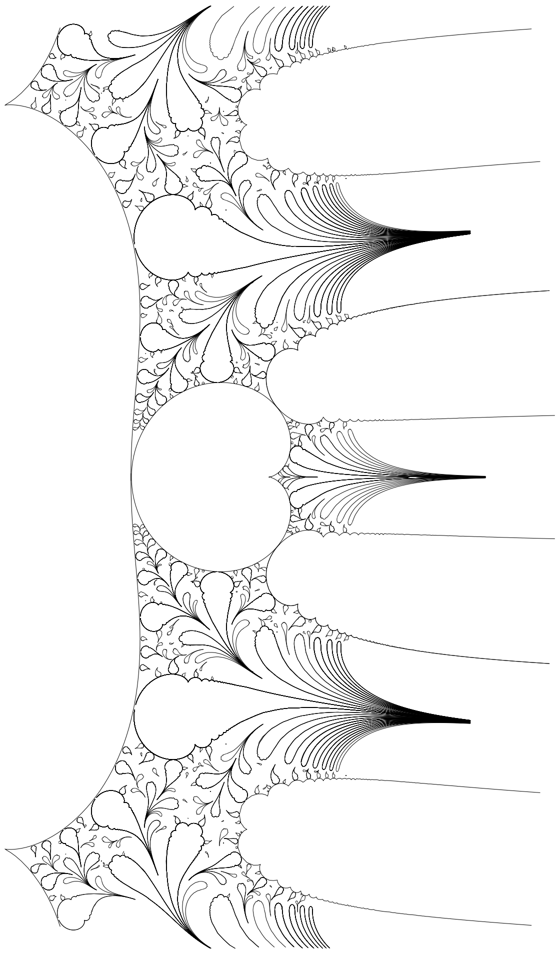

Figure 2. The same space as in Figure 1, drawn differently by a special purpose program of Günter Rottenfußer: this program traces out the boundaries of hyperbolic components, which is possible with arbitrary precision for any given hyperbolic component. Unlike for the Mandelbrot set, in the exponential case it is impossible 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.

Define the horizontal lines

$$L(s_k) := \big\{z \in \mathbf{C} : \mathrm{Im}(z) = 2\pi s_k - c\big\};$$

here $c = \mathrm{Im}\big(\log(\lambda)\big)$, choosing the branch with $|c| < \pi$. Then for $k = 1, \ldots, n-1$, the component U_k contains a curve which is asymptotic to $L(s_k)$ as $\mathrm{Re}(z) \to +\infty$ and which maps to $L(s_{n-1})$ under $(\lambda \exp)^{\circ(n-1-k)}$. Thus s_k specifies precisely which branch of $(\lambda \exp)^{-1}$ 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 $+\infty$.

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.

Our results are easily translated into $\kappa = \log(\lambda)$-space: then Theorem 1.1 holds for $n \geq 2$, and the condition $s_1 = 0$ is lifted. However, κ-space is not a true parameter space: every exponential map is represented countably often; adding a constant integer to every intermediate external address in κ-space yields the same map with a different branch of κ.

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 *at least one* is strengthened to *exactly one* in Section 4 using a variant of spider theory. This finishes the classification of hyperbolic components. The second main result (in Section 7) constructs, for every hyperbolic component, a preferred parametrization (which even extends to the boundary). The main difficulty is in breaking the symmetry and fixing an origin of the parametrization, which is accomplished using a "dynamic root" of every periodic Fatou component: this is a boundary point which is fixed under the first return map of the component and which is the landing point of at least two periodic dynamic rays. Existence and uniqueness of dynamic roots is shown in Section 6, while Section 5 provides the necessary combinatorial properties of periodic dynamic rays landing at a common point: whenever three or more periodic rays land at a common point, then the dynamics permutes all these rays transitively, and two rays in this orbit are singled out as " characteristic rays". We conclude this paper with a discussion of further results on hyperbolic components, in particular their boundary properties.

Some notation. We write our exponential maps as $z \mapsto E_\lambda(z) := \lambda e^z = \exp(z + \kappa)$ with $\lambda = \exp(\kappa)$, where $\lambda \in \mathbf{C} \setminus \{0\}$ and $\kappa \in \mathbf{C}$; usually we will use the branch $|\mathrm{Im}(\kappa)| \leq \pi$. We will often need $F(t) = e^t - 1$, in particular for $t \in \mathbf{R}$. Let $\mathbf{D} := \{z \in \mathbf{C} : |z| < 1\}$ and $\mathbf{D}^* := \mathbf{D} \setminus \{0\}$. We write that a curve or sequence in $\mathbf{C}$ 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). Helpful and constructive suggestions by Lasse Rempe and by an anonymous referee are also gratefully acknowledged.

2. Exponential dynamics

In this section, we will review known properties of exponential dynamics. The map $E_\lambda = \lambda \exp$ has no critical points or critical values, and a unique omitted value 0 also known as *singular value*. The singular value plays an equally decisive role for exponential dynamics as the critical values do for polynomial dynamics.

Points z with $E_\lambda^{\circ k}(z) \to \infty$ as $k \to \infty$ 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 $\mathbf{Z}$ (there is a well-understood exception if the singular value does escape; that case does not matter for our purposes). The dynamic ray at external address $\underline{s}$ is an injective curve $g_{\underline{s}}\colon]t_{\underline{s}}, \infty[\to \mathbf{C}$ with $\mathrm{Re}\big(g_{\underline{s}}(t)\big) \to +\infty$ as $t \to \infty$, while $\mathrm{Im}\big(g_{\underline{s}}(t)\big)$ is bounded. The quantity $t_{\underline{s}} \geq 0$ depends on $\underline{s}$ in a well-understood way; we only need bounded sequences $\underline{s}$, and those have $t_{\underline{s}} = 0$. We say that a ray $g_{\underline{s}}$ *lands* at a point $w \in \mathbf{C}$ if $\lim_{t \searrow t_{\underline{s}}} g_{\underline{s}}(t)$ exists and is equal to w. A *ray tail* is an unbounded subcurve of a ray: it is a curve $g_{\underline{s}}(\,[\tau, \infty[\,)$ for $\tau > t_{\underline{s}}$.

Every point on a dynamic ray is an escaping point, and every escaping point is on such a ray, or the unique limit point of such a ray. We have the dynamic relation

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

where σ is the shift map on external addresses, dropping the first entry. The meaning of the external address of a ray is the following: the set $E_\lambda^{-1}(\mathbf{R}^-)$ is a countable union of horizontal lines, spaced at distance $2\pi i\mathbf{Z}$, and $\mathbf{C} \setminus E_\lambda^{-1}(\mathbf{R}^-)$ are horizontal strips, labeled by $\mathbf{Z}$ 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 $s_1 s_2 s_3 \ldots$ 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: it consists of sequences satisfying a certain exponential growth condition [SZ2], and in particular it contains all bounded sequences.

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 orbits, except the unique attracting one, are repelling. We will need a construction and results from [SZ3, Section 4.3]: let $n \geq 2$ be the period of the attracting orbit, let $U_1, U_2, \ldots, U_n = U_0$ be the cycle of periodic Fatou components, labeled cyclically modulo n so that U_1 is the characteristic Fatou component, and let $a_1, a_2, \ldots, a_n$ 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. For $k = 0, 1, \ldots, n$, let

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

and $V := V_0 \cup V_1 \cup \cdots V_{n+1}$. Then V_n contains a left half plane, and for $k \in \{n-1, n-2, \ldots, 1\}$, V_k contains a neighborhood of a curve towards $+\infty$ with $V_1 \supset V_{n+1}$, while V_0 contains neighborhoods of infinitely many such curves towards $+\infty$, spaced equally at integer translates of $2\pi i$. The construction assures that $E_\lambda(V) \subset V$ and that all V_k are connected and simply connected.

Let $R := \mathbf{C} \setminus V_0$; it consists of countably many connected components which we will call "regions" $R_{\mathtt{u}}$: let R_0 be the region containing the singular value and $R_{\mathtt{u}} := R_0 + 2\pi i\mathtt{u}$, for $\mathtt{u} \in \mathbf{Z}$. Then $R = \bigcup_{\mathtt{u}} R_{\mathtt{u}}$. Any orbit (z_k) within the Julia set then has an associated *itinerary* $\mathtt{u}_1, \mathtt{u}_2, \mathtt{u}_3, \ldots$ such that $z_k \in R_{\mathtt{u}_k}$ for all k. 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. The itineraries as defined here are dynamically natural; compare [SZ3, Sections 4 and 5] for a discussion of the differences. (Note the different fonts for external addresses $s_1 s_2 \ldots$ and itineraries $\mathtt{u}_1\mathtt{u}_2\ldots$.)

For exponential dynamics with an attracting periodic orbit, every periodic dynamic ray lands at a repelling periodic point [SZ3, Theorem 3.1]; every repelling periodic point is the landing point of a finite positive number of periodic rays [SZ3, Theorem 5.4]; and a periodic ray lands at a periodic point if and only if ray and point have identical itineraries [SZ3, Proposition 4.5]. In particular, different periodic points have different itineraries, while different periodic rays have the same itinerary if and only if they land together.

The following results are known from [EL2], [EL3], [BR]: in λ-space, there is a unique hyperbolic component H_1 of period 1 which is a bounded neighborhood of the puncture $\lambda = 0$ of parameter space; it comes with a conformal isomorphism $\mathbf{D}^* \to H_1$, $\mu \mapsto \mu \exp^{-\mu}$ (so that $\lambda \exp$ has a fixed point with multiplier μ if and only if $\lambda = \mu \exp^{-\mu}$). All $\lambda \in \mathbf{C}^*$ with $|\lambda| < 1/e$ have $\lambda \in H_1$. There is a unique period 2 component H_2 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 $\lambda = r\exp(i\vartheta) \in H_2$). Every hyperbolic component H of period $n \geq 2$ is simply connected, and the multiplier map $\mu\colon H \to \mathbf{D}^*$ is a universal cover.

Lemma 2.1 (Strong attraction only at far parameters). *For every period $n \geq 2$ and $r < 1$, there is an $R > 0$ such that any parameter λ for which there is an attracting orbit of multiplier $|\mu| \leq r$ has $|\lambda| > R$.*

For every hyperbolic component H of period $n \geq 2$ and $\lambda \in H$, there is a unique homotopy class of curves $\gamma([0,\infty]) \to H \cup \{\infty\}$ with $\gamma(0) = \lambda$ and $\gamma(\infty) = \infty$ such that $\mu(\gamma(t)) \to 0$ as $t \to \infty$ (homotopy with fixed endpoints).

Proof. If $a_1, a_2, \ldots, a_n$ is any periodic orbit of period n under $\lambda\exp$, then $\big(E_\lambda^{\circ n}(a_m)\big)' = \prod_k a_k$. If the orbit is attracting, then there is some $|a_k| < 1$. Hence if $|\lambda| \le R$, then all $|a_k| < \xi$ for some ξ depending on R and n. Now $\mathrm{Re}(a_k) \ge -\xi$ implies that $|a_{k+1}| \ge |\lambda|\exp(-\xi) \ge \exp(-\xi)/e$ (if $|\lambda| < 1/e$, then $\lambda\exp$ has an attracting fixed point). Hence we have a lower bound for all $|a_k|$ and hence for $|\mu|$.

It follows that any curve $\gamma'\colon([0,\infty[) \to \mathbf{D}^*$ with $\lim_{t\to\infty}\gamma'(t) = 0$ lifts under any branch of the inverse multiplier map to a curve $\gamma\colon([0,\infty[) \to H$ with $\lim_{t\to\infty}\gamma(t) = \infty$. Conversely, any two curves $\gamma_1, \gamma_2\colon([0,\infty[) \to H$ which limit at ∞ and with $\mu\big(\gamma_{1,2}(t)\big) \to 0$ as $t \to \infty$ project under μ to two curves $\gamma_1', \gamma_2'\colon [0,\infty[\to \mathbf{D}^*$, and these are homotopic in $\mathbf{D}^*$; hence γ_1 and γ_2 are also homotopic in H (different branches of μ^{-1} can be compensated by loops around 0 in $\mathbf{D}^*$ and do not matter). □

Remark. This lemma does not show that any two curves $\gamma_{1,2}\colon([0,\infty])\to H \cup \{\infty\}$ with $\gamma_i(\infty) = \infty$ are homotopic if the condition $\mu\big(\gamma_{1,2}(t)\big) \to 0$ is dropped: it would even be conceivable that, for some $\vartheta \in \mathbf{R}$, a branch of $\mu^{-1}(\,]0, e^{i\vartheta}[\,)$ (an *internal parameter ray*, Definition 4.5) tends to ∞ as $|\mu| \to 1$. It was conjectured by Eremenko and Lyubich [EL2] that this does not happen. The proof of this is not easy: see [S1, Section V] and [RS].

3. Classification of hyperbolic components

Hyperbolic components of Multibrot sets have the helpful property that they have a unique "center" in which the dynamics is postcritically finite [Mc1], [M1], [ES]. If an exponential map has an attracting orbit, it can never be postsingularly finite (with multiplier 0); the center of hyperbolic components "is at ∞" (in the sense of Lemma 2.1). Fairly enough, it turns out that hyperbolic components of exponential maps have a different 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 half-integer.

Definition 3.1 (Intermediate external address). An *intermediate external address* of period $n \ge 2$ is a finite sequence $s_1 s_2 \ldots\ s_{n-2}, s_{n-1}$ with $s_k \in \mathbf{Z}$ for $k \le n-2$ and $s_{n-1} \in (\mathbf{Z} + \frac{1}{2})$.

The lexicographic order on external addresses (infinite sequences over $\mathbf{Z}$) extends naturally to intermediate external addresses such as $\underline{s} = s_1 s_2 \ldots s_{n-1}$. Intermediate external addresses of period n (which consist of $n-1$ numbers) label hyperbolic components of period n.

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 $n \ge 2$ and let $\underline{s} = s_1 s_2 \ldots s_{n-1}$ be an intermediate external address of period n. As always, let a_1 be the attracting periodic point in the characteristic Fatou component U_1. We say E_λ has an *attracting dynamic ray at external address* $\underline{s}$ if there is a curve $\gamma\colon [0,\infty[\to U_1$ such that the following hold:

- $\gamma(0)=a_1$;
- $\lim_{t\to\infty} E_\lambda^{\circ k}\gamma(t)=+\infty$ for $k=0,1,\ldots,n-2$;
- $\lim_{t\to\infty} E_\lambda^{\circ(n-1)}\gamma(t)=-\infty$;
- 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 every dynamic ray. Both γ and any given dynamic ray tend to $+\infty$ at bounded imaginary parts, so the ray must be *above* or *below* γ in the following sense: for real ξ sufficiently large, γ cuts the half plane $\{z\in\mathbf{C} : \mathrm{Re}(z)>\xi\}$ into two unbounded parts, one above and one below γ, and every dynamic ray must tend to $+\infty$ within one of these two parts.

Lemma 3.3 (Attracting dynamics has external addresses). *For every exponential map E_λ with an attracting orbit of period $n\geq 2$, there is a unique intermediate external address $\underline{s}=s_1s_2\ldots s_{n-1}$ of period n such that there is an attracting dynamic ray at external address $\underline{s}$. Every exponential map from the same hyperbolic component has an attracting dynamic ray at the same external address $\underline{s}$.*

Proof. The periodic Fatou component $U_0=U_n$ contains a left half plane. By simple connectivity, it contains a unique homotopy class of curves connecting the attracting periodic point a_n to $-\infty$ eventually within a left half plane, even eventually along $\mathbf{R}^-$. This homotopy class of curves can be pulled back $n-1$ steps to a preferred homotopy class of curves within U_1 connecting a_1 to $+\infty$. Choose one such curve $\gamma\subset U_1$. This curve avoids dynamic rays, and it is easy to check that the supremum of external addresses of dynamic rays below γ has well-defined $n-1$ initial entries in $\mathbf{Z}$: 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 $\mathbf{R}^-$ 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 the hyperbolic component of E_λ (the intermediate external address is a discrete object which depends continuously on the parameter and on γ). □

We thus have a combinatorial coding for every hyperbolic component (the unique component of period 1 is coded by the empty sequence), 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 4.4 (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 $\mathrm{Re}(\lambda) > h$ and the initial segment $z_1 = 0$, $z_2 = \lambda, \ldots, z_n$ of the singular orbit has the property that $|\mathrm{Im}(z_k)| < h$ for $1 \le k \le n$. Suppose moreover that z_n is the first point on the singular orbit with $\mathrm{Re}(z_n) < 0$. Then the map E_λ has an attracting periodic orbit of exact period n, and the attracting basin contains the left half plane $\mathrm{Re}(z) \le \mathrm{Re}(z_n) + 1$. As $\mathrm{Re}(\lambda) \to \infty$ 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, and its image is extremely close to 0. 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 things work 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 $S := \{z \in \mathbf{C} : |\mathrm{Im}(z)| < h\}$ 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 $|\lambda| > h > 1$: $|z_k| = |\lambda| \exp\big(\mathrm{Re}(z_{k-1})\big) > he^h$, so $\mathrm{Re}(z_k) > h$ for $k = 2, \ldots, n-1$ and $\mathrm{Re}(z_n) < -h$.

Now we show for $m \le n$

$$\prod_{k=2}^{m} (|z_k| + 1) < (|z_m| + 1)^2. \tag{1}$$

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

$$\begin{aligned}(|z_m| + 1)^2 &< \big(\sqrt{2}\ \mathrm{Re}(z_m) + 1\big)^2 \\ &< 3\exp\big(\mathrm{Re}(z_m)\big) < |\lambda|\,|\exp(z_m)| < |z_{m+1}| + 1.\end{aligned}$$

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

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

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

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

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 $z_{n-1}, \dots, D_1$ 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 $|z_{k+1}| + 1$ (using the inductive fact that $\mathrm{dist}(z_k, \partial D_k) \leq 1$). 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 $\varrho = \prod_{k=2}^{n}(|z_k| + 1)^{-1} > (|z_n| + 1)^{-2}$ by equation (1) above (since we have a bound on the contraction for every point on the disks, the claim follows without invoking Koebe's theorems).

On the other hand, all the points in D_n are contained in the left half plane $\mathrm{Re}(z) \leq \mathrm{Re}(z_n) + 1$. For every point z in this left half plane, $|E_\lambda(z)| \leq |\lambda| \exp\bigl(\mathrm{Re}(z_n) + 1\bigr) = e|\lambda| \exp\bigl(\mathrm{Re}(z_n)\bigr) < \varrho$ by equation (2) above. It follows that $E_\lambda(D_n) \subset D_1$ and hence $E_\lambda^{\circ n}(D_1) \subset D_1$. This is a proper inclusion, so 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, $\mathrm{Re}(\lambda)$ 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: the unique component of period 1 in λ-space is coded by the empty external address, and the component of period 2 is coded by the address $\underline{s} = 0$ of length 1.

Theorem 3.5 (Existence of hyperbolic components). *For every $n \geq 3$ and every intermediate external address $\underline{s} = s_1 s_2 \dots s_{n-1}$ of period n with $s_1 = 0$, there is a hyperbolic component in λ-space in which every exponential map has an attracting dynamic ray at external address $\underline{s}$.*

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.

Proof. Let $s^{\pm}_{n-1} := s_{n-1} \pm \frac{1}{2} \in \mathbf{Z}$ and define two periodic external addresses of period $n-1$ via $\underline{s}^- := \overline{s_1 s_2 \dots s^-_{n-1}}$ and $\underline{s}^+ := \overline{s_1 s_2 \cdots s^+_{n-1}}$. Let $A := 1 + \max_k\{|s_k|\}$.

In [SZ2, Proposition 3.4] (or [SZ1, Theorem 2.3]), the existence of dynamic rays $g_{\underline{s}^+}$, $g_{\underline{s}^-}$ was shown: these are curves $g_{\underline{s}^\pm}\colon\,]0,\infty[\to \mathbf{C}$ satisfying

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

for $t \geq 1 + 2\log(|\kappa| + 3)$, 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

$$\tau := 1 + 2\log(|\kappa| + 3 + 2\pi AC'),$$

then $\tau \geq \max\{1 + 2\log(|\kappa| + 3), \log 2 + \log(|\kappa| + 2 + 2\pi AC')\}$ so that (3) holds and

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

for $t \geq \tau$. After $n-2$ iterations, the ray tails $g_{\underline{s}^{\pm}}([\tau,\infty[)$ map to

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

with

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

with $|r^{\pm}(t)| < e^{-(t-\tau)} \leq 1$. Define a curve

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

it has the property that $E_\lambda(\gamma'_\kappa) \subset \mathbf{R}^-$. The construction assures that the two ray tails $g_{\sigma^{n-2}(\underline{s}^{\pm})}\big([F^{\circ(n-2)}(\tau),\infty[\big)$ are above respectively below $\gamma'_\kappa\big([F^{\circ(n-2)}(\tau),\infty[\big)$ (asymptotically by $i\pi$), and all three curves are disjoint. Moreover, every 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 $n-2$ times along equal branches of E_λ^{-1} on $\mathbf{C} \setminus \mathbf{R}_0^-$, it follows that there is a curve $\gamma_\kappa\colon [\tau,\infty[\to \mathbf{C}$ between the two rays $g_{\underline{s}^{\pm}}([\tau,\infty[)$ with

$$E_\lambda^{\circ(n-2)}\big(\gamma_\kappa(t)\big) = \gamma'_\kappa\big(F^{\circ(n-2)}(t)\big) = F^{\circ(n-2)}(t) - \kappa + 2\pi i s_{n-1}. \tag{4}$$

Just like dynamic rays, this curve inherits the bound for $t \geq \tau$

$$\gamma_\kappa(t) = t - \kappa + 2\pi i s_1 + r_\kappa(t) \quad \text{with } |r_\kappa(t)| < 2e^{-t}(|\kappa| + 2 + 2\pi AC') < e^{-(t-\tau)} \tag{5}$$

(it follows from the construction of γ_κ using $n-2$ pull-backs that γ_κ satisfies asymptotically the same bounds (3) as the two rays $g_{\underline{s}^{\pm}}$ which surround γ_k; this can also be verified in the proof of (3) in [SZ1], [SZ2]).

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 exponential map has.

For any $R \geq 0$, set $\kappa_R^{\pm} := R + 2\pi i s_1 \pm i\pi = R \pm i\pi$ and $I_R := [\kappa_R^-, \kappa_R^+]$ (a vertical interval of length 2π); then $|\kappa| \leq R + \pi$ for all $\kappa \in I_R$. Set $t_R := 1+2\log(R+\pi+3+2\pi AC') > \tau$. For $t \geq t_R$, the bound (5) implies $\operatorname{Im}\bigl(\gamma_{\kappa_R^-}(t)\bigr) < 0$ and $\operatorname{Im}\bigl(\gamma_{\kappa_R^+}(t)\bigr) > 0$. For all $\kappa \in I_R$, we have

$$\operatorname{Re}\bigl(\gamma_\kappa(t_R)\bigr) < t_R - R + 1.$$

Now fix R large enough so that $t_R - R + 1 < 0$; this is possible since $t_R = O(\log R)$. As κ moves from κ_R^- to κ_R^+, by continuity there must be an intermediate value κ^* where $\gamma_{\kappa^*}(t^*) = 0$ for some $t^* > t_R > \tau$. 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, $E_{\lambda^*}(0)$, $\ldots$, $E_{\lambda^*}^{\circ(n-1)}(0)$ are in the strip $|\operatorname{Im}(z)| \leq 2\pi A + \pi + 1$: the $2\pi A$ comes from the bound $|s_k| < A$; the π is the bound on $\operatorname{Im}(\kappa)$, 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 $\operatorname{Re}\bigl(E_{\lambda^*}^{\circ(n-1)}(0)\bigr) \ll 0$. Now if R is large enough, then Lemma 3.4 shows that there is an attracting orbit of exact period n for κ^*, and $E_{\lambda^*}^{\circ(n-1)}\bigl(\gamma_{\kappa^*}\bigl([t^*, \infty[\,\bigr)\bigr) \subset U_n$. Hence $\gamma_{\kappa^*}\bigl([t^*, \infty[\,\bigr)$ is in the attracting basin with $\gamma_{\kappa^*}(t^*) = 0$, so $\gamma_{\kappa^*}\bigl([t^*, \infty[\,\bigr) \subset U_1$.

Connect γ_{κ^*} to the attracting periodic point $a_1 \in U_1$ and call this resulting curve $\gamma\colon \bigl([0, \infty[\,\bigr) \to \mathbf{C}$. 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 $\underline{s} = s_1 s_2 \ldots s_{n-1}$.

By Lemma 3.3, every parameter λ in the hyperbolic component of λ^* has an attracting dynamic ray at external address $\underline{s}$. This finishes the existence part of the theorem.

To justify the asymptotics, start with large R, hence large $t^* > t_R = O(\log R)$, and observe that $\gamma_{\kappa^*}(t^*) = 0$ implies, using (5), that $\kappa^* = t^* + 2\pi i s_1 + r_{\kappa^*}(t^*)$ with $|r_{\kappa^*}(t^*)| \to 0$, hence $\operatorname{Im}(\kappa^*) \to 2\pi s_1 = 0$. For $n = 3$, we have $\lambda^* = E_{\lambda^*}\bigl(\gamma_{\kappa^*}(t^*)\bigr) = \gamma'_{\kappa^*}\bigl(F(t)\bigr) = F(t) - \kappa + 2\pi i s_2$, hence $\operatorname{Im}(\lambda^*) \to 2\pi s_2$. For $n > 3$, we argue similarly:

$$\begin{aligned}\lambda^* = E_{\lambda^*}\bigl(\gamma_{\kappa^*}(t^*)\bigr) &= g_{\sigma(\underline{s}^{\pm})}\bigl(F(t^*)\bigr) + o(1) = F(t^*) - \kappa^* + 2\pi i s_2 + o(1)\\ &= F(t^*) - t^* + 2\pi(s_2 - s_1) + o(1).\end{aligned}$$

This way, we have shown the existence of a map $R \mapsto (\kappa^*, t^*)$ (for sufficiently large R) with $\operatorname{Re}(\kappa^*) = R$ and $\gamma_{\kappa^*}(t^*) = 0$. Note that $E_{\lambda^*}^{\circ(n-1)}(t^*) \in \mathbf{R}^-$ for all $t^* = t^*(R)$. It is quite easy to check that $(\partial/\partial\kappa^*)\bigl(E_{\lambda^*}^{\circ(n-1)}(t^*)\bigr) \to \infty$ as $t^* \to \infty$, so the implicit function theorem shows that the graph of $R \mapsto \kappa^*(R)$ is analytic. It follows from Lemma 3.4 that the attracting multiplier tends to 0 as $R \to \infty$. □

We know from Lemma 2.1 that every hyperbolic component H has a preferred homotopy class of curves $\gamma\colon[0,\infty]\to H\cup\{\infty\}$ with $\gamma(0)\in H$ and $\gamma(\infty)=\infty$ such that multipliers tend to 0 along this curve. These preferred homotopy classes of curves give a natural vertical order to hyperbolic components, much as the order for dynamic rays: we say that *some hyperbolic component is above another hyperbolic component* if the corresponding preferred 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 curve γ_κ for $\underline{s}'$ is always above $g_{\underline{s}}$, while the corresponding curve for $\underline{s}''$ is always below $g_{\underline{s}}$. □

Remark. This vertical order can also be expressed in terms of parameter rays [S1], [F]: these are differentiable curves $G_{\underline{s}}\colon(t_{\underline{s}},\infty)\to\mathbf{C}$ with $G_{\underline{s}}(t)\to+\infty$ as $t\to\infty$ such that for $\lambda=G_{\underline{s}}(t)$, the singular value escapes with $0=g_{\underline{s}}(t)$; parameter rays are thus disjoint from hyperbolic components. With the notation of the proof of Corollary 3.6, there is a parameter ray $G_{\underline{s}}$, 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. Uniqueness of the classification

In the following lemma, we construct a curve $\gamma=\gamma_-\cup\{a_1\}\cup\gamma_+$ in the dynamical plane of every exponential map which has an attracting orbit with positive real multiplier. The curve γ_+ will be used in this section to construct fundamental domains for exponential dynamics, while γ_- will be used in Section 7 to construct a "dynamic root", which helps parametrizing hyperbolic components in a dynamically meaningful fashion.

Lemma 4.1 (Attracting dynamic ray to boundary fixed point). *Let E_λ be an exponential map which has an attracting periodic orbit with positive real multiplier and with period $n\geq 2$. Then there is a proper analytic curve $\gamma\colon\mathbf{R}\to U_1$ which contains the orbit of 0 under $E_\lambda^{\circ n}$. This curve is unique up to reparametrization and can be written $\gamma=\gamma_-\cup\{a_1\}\cup\gamma_+$ with two disjoint subcurves $\gamma_\pm$ which have the following properties:*

- *γ_+ is an attracting dynamic ray and contains the orbit of 0 under $E_\lambda^{\circ n}$; it connects a_1 to $+\infty$;*
- *γ_- starts at a_1 and lands at some $w\in\partial U_1$ with $E_\lambda^{\circ n}(w)=w$.*

Proof. There is a unique open neighborhood D of a_1 which corresponds to a round disk in linearizing coordinates around 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 $0 \in \overline{\gamma_D}$. Then the first return map $E_\lambda^{\circ n}$ of U_1 sends D into itself and γ_D into itself. Any proper analytic curve which contains the orbit of 0 under $E_\lambda^{\circ n}$ must be an extension of γ_D because the orbit of 0 accumulates at a_1, so γ is unique if it exists.

The point a_1 cuts γ_D into two radii; let γ'_+ be the one which ends at 0 and γ'_- the other one (then $\gamma_D = \gamma'_+ \cup \{a_1\} \cup \gamma'_-$). The set $E_\lambda^{\circ(-n)}(\gamma'_+) \cap U_1$ consists of countably many curves; let γ_+ be the one which extends γ'_+: it starts at a_1 and lands at ∞, running through 0. Since γ'_+ is differentiable at 0, every branch of $E_\lambda^{-1}(\gamma'_+)$ has bounded imaginary parts, and γ_+ approaches $+\infty$ along bounded imaginary parts. Then γ_+ is an attracting dynamic ray: it satisfies the first three conditions of Definition 3.2, and by Lemma 3.3, the intermediate external address of γ_+ is uniquely determined by the hyperbolic component containing E_λ.

There is a unique curve $\gamma_- \subset U_1 \setminus \{a_1\}$ 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 γ_- is analytic and $E_\lambda^{\circ n}\colon \gamma_- \to \gamma_-$ is a homeomorphism.

The curve γ_- can be parametrized (in many ways) as $\gamma_-\colon \mathbf{R} \to U_1$ so that $E_\lambda^{\circ n}\big(\gamma_-(t)\big) = \gamma_-(t+1)$. 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 the landing of external dynamic rays of polynomials. Let

$$U' := U_1 \setminus \overline{\bigcup_{k\geq 0} E_\lambda^{\circ kn}(0)} \quad \text{and} \quad U'' := E_\lambda^{-(\circ n)}(U') \cap U_1 .$$

Then both domains are open and $E_\lambda^{\circ n}\colon U'' \to U'$ is a holomorphic covering map, hence a local isometry with respect to the unique normalized hyperbolic metrics of U'' and U', and the inclusion $U'' \hookrightarrow U'$ is a contraction. Therefore, the hyperbolic distance in U' between any $\gamma_-(t-1)$ and $\gamma_-(t)$ is less than between $\gamma_-(t)$ and $\gamma_-(t+1)$. By continuity, there is an $s > 0$ such that the hyperbolic distance in U' between $\gamma_-(t)$ and $E_\lambda^{\circ n}\big(\gamma_-(t)\big)$ is at most s, for all $t < 0$. 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 $\big|\gamma_-(t) - E_\lambda^{\circ n}\big(\gamma_-(t)\big)\big| \to 0$ as $t \to -\infty$. Therefore, any limit point of γ_- is either ∞ or a fixed point of $E_\lambda^{\circ n}$. Since the limit set is non-empty and connected, while the set of fixed points is discrete, it follows that γ_- lands at a well-defined boundary point of U_1 which is either ∞ or fixed under $E_\lambda^{\circ n}$. If the landing point is a fixed point of $E_\lambda^{\circ n}$, then the curve of the claim is $\gamma := \gamma_- \cup \{a_1\} \cup \gamma_+$. All that remains is to show that the curve γ_- does not land at ∞.

Suppose that $\gamma_-(t) \to \infty$ as $t \to -\infty$. Then also $E_\lambda^{\circ n}\big(\gamma_-(t)\big) = \gamma_-(t-1) \to \infty$ as $t \to -\infty$. Since $\gamma_- \subset U_1$, it follows that $E_\lambda^{\circ(n-1)}(\gamma_-) \subset U_0 = U_n$ and

$E_\lambda^{\circ(n-1)}\big(\gamma_-(t)\big) \to \infty$ as $t \to -\infty$. If we had $\liminf \mathrm{Re}\big(E_\lambda^{\circ(n-1)}\big(\gamma_-(t)\big)\big) < +\infty$ as $t \to -\infty$, then $\liminf\big|E_\lambda^{\circ n}\big(\gamma_-(t)\big)\big| = \liminf |\gamma_-(t+1)| < \infty$, and the landing point of γ_- would be in $\mathbf{C}$. Therefore, $\mathrm{Re}\big(E_\lambda^{\circ(n-1)}\big(\gamma_-(t)\big)\big) \to +\infty$ as $t \to -\infty$, and this curve must have bounded imaginary parts because the translates of U_1 cut $\{z \in U_n : \mathrm{Re}(z) > 0\}$ into parts with bounded imaginary parts. But then, for $t \ll 0$ and $\mathrm{Re}\big(\gamma_-(t)\big) \gg 0$, the entire orbit of $\gamma_-(t)$ must have large real parts, and standard estimates similarly as in Lemma 3.4 show $\mathrm{Re}\big(E_\lambda^{\circ n}\big(\gamma_-(t)\big)\big) \gg \mathrm{Re}\big(\gamma_-(t)\big)$, which is a contradiction: for every t, the point $\gamma_-(t)$ converges to a_1 under iteration of $E_\lambda^{\circ n}$. □

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

Using the analytic curve γ from Lemma 4.1, we construct *fundamental domains* for E_λ as follows (provided the attracting multiplier is positive real): there is a subcurve $\gamma' \subset \gamma$ which connects the singular value 0 to $+\infty$ (in fact, $\gamma' = \gamma_+ \setminus (\gamma'_+ \cup \{0\})$); the full inverse image $E_\lambda^{-1}(\gamma')$ is a countable collection of analytic curves which connect $-\infty$ to $+\infty$ and which differ by translation by $2\pi i\mathbf{Z}$. The connected components of the complement in $\mathbf{C}$ of these curves are fundamental domains for E_λ; each of them is mapped by E_λ conformally onto $\mathbf{C} \setminus \gamma'$, and each has bounded imaginary parts because γ' is differentiable in its endpoint 0.

Theorem 4.2 (Conformal conjugation). *Suppose that two exponential maps have attracting orbits of equal period $n \geq 2$ with equal positive real multipliers, and both have attracting dynamic rays at the same intermediate external address $s_1 s_2 \dots s_{n-1}$. Suppose in addition that for both maps, the fundamental domains as constructed above are such that a single fundamental domain contains both the periodic point a_0 and the singular value. Then both maps are conformally conjugate.*

Proof. Step 1. Let E_λ and $E_{\lambda'}$ be two exponential maps satisfying the assumptions of the theorem. Let V_1 be an open round disk with respect to linearizing coordinates of a_1, large enough so as to contain 0 in its interior. For $k = 2, 3, \dots, n+1$, let $V_k \subset U_k$ be the domain $E_\lambda^{\circ(k-1)}(V_1)$, and let $V_0 := E_\lambda^{-1}(V_1)$. Then $V_1 \supset V_{n+1}$ and $V_0 \supset V_n$. Denote the corresponding sets for $E_{\lambda'}$ as U'_k and V'_k, where the size of V'_1 is chosen so that the dynamics on it is conformally conjugate to the dynamics on V_1, respecting the singular value.

Step 2. We will construct a quasiconformal homeomorphism $\varphi : \mathbf{C} \to \mathbf{C}$ 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 define $\varphi|_{V_k} : V_k \to V'_k$ as conformal isomorphisms for $k =$

$1, \ldots, n$ which respect the dynamics on V_k in the sense that for $z \in V_k$

$$\varphi\bigl(E_\lambda(z)\bigr) = E_{\lambda'}\bigl(\varphi(z)\bigr). \tag{6}$$

It becomes unique on V_1 by the requirement that $\varphi(0) = 0$, and in view of (6) it is unique on all V_k.

Since for $k = 1, \ldots, n$, the sets $\overline{V}_k$ are disjoint closed disks with analytic boundaries, the map φ can be extended from the conformal isomorphism $\varphi\colon \cup V_k \to \cup V'_k$ to a quasiconformal homeomorphism $\varphi\colon \mathbf{C} \to \mathbf{C}$. On the multiply connected domain $\mathbf{C} \setminus (\cup V_k)$, we need to specify the homotopy class of φ. We will do this using the attracting dynamic ray γ_+ from Lemma 4.1: let $\gamma_k := E_\lambda^{\circ(k-1)}(\gamma_+)$ for $k = 1, 2, \ldots, n$. Every $\gamma_k \subset U_k$, and it connects V_k to $+\infty$ (for $k = 1, 2, \ldots, n-1$) or to $-\infty$ (for $k = n$). Now $\varphi(\gamma_k)$ connects V'_k to ∞ for every k, and we require that $\varphi(\gamma_k)$ be homotopic to the analogous curve γ'_k relative to $V'_1 \cup \cdots \cup V'_n$. This fixes the homeomorphism $\varphi\colon \mathbf{C} \to \mathbf{C}$ uniquely up to homotopy, and it is well known that φ may be chosen so as to be quasiconformal. Note that we do not claim that the extension of φ away from the V_k respects the dynamics.

Step 3. Our goal is to promote φ to a conformal conjugation via a sequence of quasiconformal maps φ_j with $\varphi_0 := \varphi$ and

$$\varphi_j \circ E_\lambda = E_{\lambda'} \circ \varphi_{j+1}. \tag{7}$$

We will show that $\varphi_j \to \mathrm{id}$ uniformly on compact sets, which implies that $E_\lambda = E_{\lambda'}$. The construction is inspired by the theory of spiders [HS].

By induction over j, we will assure that $\varphi_j|_{V_k} = \varphi_0|_{V_k}$ for all j and all k, and in particular $\varphi_j(0) = 0$ and $\varphi_j(\lambda) = \lambda'$; moreover, $\varphi_j(\gamma_k)$ is homotopic to $\varphi_0(\gamma_k)$ and hence to γ'_k, relative to $V'_1 \cup \cdots \cup V'_n$, for all k and j.

To construct φ_{j+1} from φ_j, note that both $\varphi_j \circ E_\lambda$ and $E_{\lambda'}$ are universal covers from $\mathbf{C}$ to $\mathbf{C}^*$, so there is a homeomorphism $\varphi_{j+1}\colon \mathbf{C} \to \mathbf{C}$ which satisfies (7), and it becomes unique by fixing the value at a single point. Since $\varphi_j \circ E_\lambda(0) = \lambda' = E_{\lambda'}(0)$, we can and will require $\varphi_{j+1}(0) = 0$ to fix φ_{j+1} everywhere uniquely. Note that (7) implies $\varphi_{j+1}(z + 2\pi i) = \varphi_{j+1}(z) + 2\pi i$ for all z.

By (7) and (6), we have

$$(E_{\lambda'} \circ \varphi_{j+1})|_{V_k} = (\varphi_j \circ E_\lambda)|_{V_k} = (\varphi_0 \circ E_\lambda)|_{V_k} = (E_{\lambda'} \circ \varphi_0)|_{V_k} \subset V'_{k+1}.$$

To justify $\varphi_{j+1}|_{V_k} = \varphi_0|_{V_k}$, all that remains to show is that φ_{j+1} maps V_k onto V'_k and not onto another branch of $(E_{\lambda'})^{-1}(V'_{k+1})$. It is here that the assumption about the attracting dynamic rays comes in.

Step 3a: the case $k = 1, 2, \ldots, n-1$. Since $(E_{\lambda'} \circ \varphi_{j+1})(\gamma_k) = (\varphi_j \circ E_\lambda)(\gamma_k) = \varphi_j(\gamma_{k+1})$ is homotopic to $\gamma'_{k+1} = E_{\lambda'}(\gamma'_k)$ relative $V'_1 \cup \cdots \cup V'_n$, it follows that $\varphi_{j+1}(\gamma_k)$ is homotopic to (a $2\pi i\mathbf{Z}$-translate of) γ'_k relative $E_{\lambda'}^{-1}(V'_1 \cup \cdots \cup V'_n) \supset$

$V_1'\cup\cdots\cup V_n'$. Now $\varphi_{j+1}(0)=0$ implies $\varphi_{j+1}|_{V_1}=\varphi_0|_{V_1}$, and since γ_1 is attached to V_1, it follows that $\varphi_{j+1}(\gamma_1)$ is homotopic to γ_1'.

For $k=2,3,\ldots,n-1$, the number of $2\pi i\mathbf{Z}$-translates of γ_k between γ_k and γ_1 is coded in the intermediate external address, and it equals the number of $2\pi i\mathbf{Z}$-translates of γ_k' between γ_k' and γ_1'. This quantity must be respected by φ_{j+1} in the sense that the number of translates of γ_k' between γ_1' and $\varphi_{j+1}(\gamma_k)$ is the same as between γ_k and γ_1, so $\varphi_{j+1}(\gamma_k)$ is homotopic to γ_k' and $\varphi_{j+1}|_{V_k}=\varphi_0|_{V_k}$. For $k=2$, this implies that $\varphi_{j+1}(\lambda)=\lambda'$. (For later use in Corollary 4.4, we note that this reasoning is unchanged if a constant integer is added to all $s_{k'}$: only the differences $s_{k'}-s_1$ matter).

Step 3b: the case $k=0$. Since V_0 is invariant under translation by $2\pi i\mathbf{Z}$ and the curve γ_n runs towards $-\infty$, not $+\infty$, a different argument is needed to show that $\varphi_{j+1}|_{V_0}=\varphi_0|_{V_0}$. This is built into the construction: since $\varphi_{j+1}(0)=0$, it follows that φ_{j+1} maps the E_λ-fundamental domain containing 0 to the $E_{\lambda'}$-fundamental domain containing 0 (up to homotopies of the boundary curve). Since a_0 and a_0' are in the same fundamental domains of E_λ respectively of $E_{\lambda'}$ as 0, it follows $\varphi_{j+1}(a_0)=a_0'=\varphi_0(a_0)$; now $\varphi_{j+1}|_{V_0}=\varphi_0|_{V_0}$ follows from the covering property.

Finally, $\varphi_0(\gamma_n)$ connects a_n' to $-\infty$ in the complement of $V_1'\cup\cdots\cup V_{n-1}'$ and their associated curves $\gamma_{k'}'$; since there is only one homotopy class of such curves, it follows that $\varphi_{j+1}(\gamma_n)$ is homotopic to $\varphi_0(\gamma_n)$ rel $V_1'\cup\cdots\cup V_n'$.

Step 4. We have a sequence (φ_j) of homeomorphisms $\mathbf{C}\to\mathbf{C}$; it follows from (7) that all are quasiconformal with the same maximal dilatation as φ_0. All φ_j coincide on $\bigcup_k V_k$, and all fix the points 0, $2\pi i$, and ∞, and quasiconformal homeomorphisms with these properties form a compact space. On the domain $E_\lambda^{\circ(-(j-1))}\big(\bigcup_k V_k\big)$ with $j\geq 1$, the map φ_j is conformal and coincides with all $\varphi_{j'}$ for $j'\geq j$. Now $\bigcup_j E_\lambda^{\circ(-(j-1))}(\bigcup_k V_k)$ fills up the entire Fatou set, while the Julia set has measure zero by [EL1], [EL3]. Therefore, the support of the bounded quasiconformal dilatation of the φ_j converges to zero, so the φ_j converge uniformly on compact subsets of $\mathbf{C}$ to an automorphism of $\mathbf{C}$ fixing 0, $2\pi i$ and ∞, hence to the identity. Finally, $\varphi_j(\lambda)=\lambda'$ implies $E_\lambda=E_{\lambda'}$. This is what we claimed. □

In order to conclude that every hyperbolic component is uniquely described by its intermediate external address, we first state a lemma with routine proof.

Lemma 4.3 (Same fundamental domain). *Every hyperbolic component of exponential maps with period $n\geq 2$ contains a map for which the attracting orbit has positive real multiplier, and for which a single fundamental domain contains both the periodic point a_0 and the singular value.*

Proof. Let H be the hyperbolic component and let $\mu\colon H\to\mathbf{D}^*$ be the multiplier map. Since it is a universal covering, we may find a map $E_{\lambda_0}\in H$ with

$\mu(\lambda_0) > 0$. There is a curve in H starting at λ_0 so that the μ-image of this curve describes a circle in $\mathbf{D}^*$ around the origin in positive orientation; let $\lambda_1 \in H$ be the endpoint of the curve.

For E_{λ_0}, recall the analytic curve γ_+ from Lemma 4.1: it is an extension of a diameter in linearizing coordinates around a_1. Let γ_+^1 be the subcurve between 0 and $E_{\lambda_0}^{\circ n}(0)$. As the parameter λ is deformed from λ_0 to λ_1, the curve γ_+^1 extends to a homotopy $\gamma_+^1(\lambda)$ of curves within $U_1(\lambda)$ connecting 0 to $E_\lambda^{\circ n}(0)$ in the complement of the respective singular orbits. In particular, we obtain a curve $\gamma_+^1(\lambda_1)$ connecting 0 to $E_{\lambda_1}^{\circ n}(0)$ within $U_1(\lambda_1)$ in the complement of the singular orbit, and this curve is unique up to homotopy. This is not a closed curve, but the beginning and endpoint are on the same radius with respect to linearizing coordinates of a_1, so it makes sense to say that $\gamma_+^1(\lambda_1)$ winds once around a_1. For the map E_{λ_1}, Lemma 4.1 provides another curve which connects 0 to $E_{\lambda_1}^{\circ n}(0)$ (part of a linearizing radius), and this has winding number zero. Call this curve $\tilde{\gamma}_+^1(\lambda_1)$.

For $k = 1, 2, \ldots, n$, pull back both curves under $E_{\lambda_1}^{\circ k}$, choosing the branch ending at $E_{\lambda_1}^{\circ(n-k)}(0)$ (as in Lemma 4.1): this yields two curves $\gamma_+^0(\lambda_1)$ and $\tilde{\gamma}_+^0(\lambda_1)$ connecting 0 to $+\infty$; their winding numbers around a_1 differ by 1.

All the countably many E_{λ_1}-preimages of $\tilde{\gamma}_+^0(\lambda_1)$ bound the fundamental domains of E_{λ_1} as defined before Theorem 4.2, while the E_{λ_1}-preimages of $\gamma_+^0(\lambda_1)$ are deformations of the fundamental domain boundaries of E_{λ_0}: their difference shows that between E_{λ_0} and E_{λ_1}, the periodic point a_0 has jumped one fundamental domain up.

A finite repetition of this process (forward or backward) can bring a_0 into the fundamental domain which contains 0. □

Corollary 4.4 (Uniqueness of classification). *The intermediate external addresses associated to hyperbolic components associate a bijection between hyperbolic components and intermediate external addresses with $s_1 = 0$ of any given period.*

Proof. By Lemma 3.3, every hyperbolic component is associated to a unique intermediate external address $s_1, s_2, \ldots, s_{n-1}$; this defines a "classification map" from hyperbolic components to intermediate external addresses. Every component contains a parameter as described in Lemma 4.3, say with multiplier $\mu = \frac{1}{2}$. By Theorem 4.2 two such exponential maps are identical if their external addresses $s_1 s_2 \ldots s_{n-1}$ and $s'_1 s'_2 \ldots s'_{n-1}$ satisfy $s_k - s_1 = s'_k - s'_1$ for all k, so the classification map is injective. Since every address with $s_1 = 0$ is realized (Theorem 3.5), only addresses with $s_1 = 0$ are realized, and the classification map is a bijection. □

Remark. In the proof that the intermediate external address uniquely describes hyperbolic components, we singled out attracting exponential maps for which the same fundamental domain contains the periodic point a_0 and the singular value 0. This was convenient for the proof, but such maps are dynamically not special enough so that the corresponding locus in parameter space would stand out. We do not believe that these maps have any real significance other than that they are helpful in the proof. There is, however, a dynamically significant locus within every hyperbolic component, called "central internal ray": this ray has significance both dynamically and in parameter space and allows us to give a preferred parametrization of hyperbolic components. This will be discussed in Section 7, but some preparations are more conveniently done here.

Recall that the multiplier map $\mu\colon H \to \mathbf{D}^*$ is a universal covering map for every hyperbolic component H of period $n \geq 2$, hence there is a conformal isomorphism $\Phi\colon H \to \mathbf{H}$ with $\mu = \exp\circ\Phi$, and Φ is unique up to translation by $2\pi i\mathbf{Z}$ in the range (throughout this paper, $\mathbf{H}$ denotes the left half plane). For our preferred parametrization of the component, we have to specify a choice of this integer translation, which is a combinatorial problem.

Definition 4.5 (Internal rays of hyperbolic components). An *internal ray at angle* $\vartheta \in \mathbf{R}/\mathbf{Z}$ of a hyperbolic component H is any branch of $\mu^{-1}\big(\,]0, e^{2\pi i\vartheta}[\,\big)$, where $\mu\colon H \to \mathbf{D}^*$ is the multiplier map. In other words, an internal ray is a connected component of the locus in H where $\arg(\mu)$ is constant.

For every fixed angle $\vartheta \in \mathbf{R}/\mathbf{Z}$, there are countably many parameter rays with the same angle (except for the unique period 1 component in λ space), and on each of them the map $|\mu|$ induces a homeomorphism onto $]0,1[$.

The internal parameter rays at angle 0 are the loci where the multiplier is real and positive, and they can be distinguished dynamically by the following generalization of Lemma 4.3. Note that the conformal isomorphism $\Phi\colon H \to \mathbf{H}$ induces a vertical order of parameter rays induced by imaginary parts within $\mathbf{H}$; this order does not depend on the ambiguity in the definition of Φ (but since $\mathbf{H}$ is a left half plane, for hyperbolic components with unbounded positive real parts this order is "upside down" when parameter rays are ordered with respect to imaginary parts in their approach to $+\infty$).

Lemma 4.6 (Internal parameter rays at angle 0). *Every exponential map which has an attracting orbit of period $n \geq 2$ with real positive multiplier has an associated integer-valued index Δ' which specifies how many fundamental domains the periodic point a_0 is above the singular value (or below, if the index is negative). The index Δ' is constant for exponential maps on the same internal parameter ray. For every hyperbolic component, the index Δ' induces an order preserving bijection between $\mathbf{Z}$ and parameter rays at angle 0. In particular, for every conformal isomorphism $\Phi\colon H \to \mathbf{H}$ with $\mu = \exp\circ\Phi$, there is an integer $m = m(H,\Phi)$ depending only on H and Φ such that every $\lambda \in H$ with $\mu(\lambda) > 0$ has $\mathrm{Im}\big(\Phi(\lambda)\big) = 2\pi(\Delta' - m)$.*

Proof. Both a_0 and the singular value are in some fundamental domain, and the latter are ordered like $\mathbf{Z}$; therefore, there is a well-defined integer Δ' which specifies how many fundamental domains a_0 is above 0. By continuity, the index is constant along internal parameter rays at angle 0.

Moving one parameter ray up within the hyperbolic component (i.e. moving to a parameter ray on which $\mathrm{Im}(\Phi)$ is greater by 2π for any given Φ) can be achieved by a deformation as in the proof of Lemma 4.3, and this increases the index Δ' by one. □

5. 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 their 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 5.1 (Essential orbit, characteristic point and 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 [M3].

Lemma 5.2 (Permutation of dynamic rays). *Every essential periodic orbit has exactly one characteristic point and exactly two characteristic rays at every 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 given point z_k cut the complex plane into r open connected components which will be called the "sectors" at z_k. These dynamic rays connect z_k to $+\infty$, so exactly one of the sectors at z_k contains a left half plane.

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); this will be called the *singular index* of the sector. For the sector which does contain a left half plane, we let the singular index be $m:=0$. Clearly, any sector at z_k with index $m\geq 1$

maps homeomorphically onto a sector at z_{k+1} with index $m-1$ (for $m=1$, 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 $\mathbf{R}^-$). It follows that for each sector, the index equals the number of iterations it takes for this sector to map over a left half plane.

If the index of a sector S at some z_k equals 0 so that the sector contains a left half plane, then S will not map forward homeomorphically. If $g_{\underline{s}}$ and $g_{\underline{s}'}$ are the two dynamic rays bounding S, then $E_\lambda(g_{\underline{s}}) = g_{\sigma(\underline{s})}$ and $E_\lambda(g_{\underline{s}'}) = g_{\sigma(\underline{s}')}$ bound a sector S' at z_{k+1} containing the singular value; we call S' the image sector of S. If the image sector also contains a left half plane, then the index remains 0; otherwise, the rays at z_{k+1} separate the singular value from a left half plane, and the index is strictly greater than 0. Each z_k has exactly one sector with index 0 (the unique sector containing a left half plane).

Every sector at every point z_k is periodic (in the sense just defined, not as a subset of $\mathbf{C}$), and so is the sequence of the indices. The index sequence of each sector must of course contain at least one index 0, and it cannot be the constant sequence: if all the entries of one cycle of sectors were equal to 0, then all the other cycles of sectors could never have index 0, a contradiction.

Among the nr sectors based at the n points $z_1 \cdots z_n$, there is at least one with largest index. Let S be such a sector and let z_1 be the periodic point at which S is based (possibly after cyclic relabeling). If S contains a periodic point $z_k \neq z_1$, then at least one of the sectors at z_k has index at least as large as S; after replacing S with such a sector, we may assume that S contains no periodic point z_k. Let S_0 be the unique sector (at z_0) with image sector S. Then S_0 must contain a left half plane (or the index of S_0 would exceed that of S), hence S contains the singular value. This makes z_1 the characteristic point of its orbit, and the rays bounding S are the characteristic rays. This shows the first statement.

Let $\alpha_1 > \alpha_2 > \cdots > \alpha_r$ 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, $\alpha_r = 0$ is the sector containing a left half plane.

Consider any cycle C of sectors and let $\alpha > 0$ be the largest index within its period. Since indices are always decreasing unless they are equal to 0, 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 largest index 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 within the sector at z_k with index α, and so are all the sectors at z_1 with indices $\alpha_1 > \cdots > \alpha_{r-1}$. 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 (the maximum α is not assumed at z_1).

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 $r=2$. □

6. 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.

We will need the partition constructed in [SZ3, 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 nth entry 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 each 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 $U_2, \ldots, U_{n-1}$ (but not U_0). Since every U_i (for $i = 1, 2, \ldots, n-1$) is disjoint from its $2\pi i\mathbf{Z}$-translates, and U_{n-1} contains a preimage of an unbounded part of $\mathbf{R}^-$, each $U_i \neq U_0$ has bounded imaginary parts.

Lemma 6.1 (Periodic external address above U_1). *Let U_1 be the characteristic Fatou component of an exponential map with attracting orbit of period $n \geq 2$. Then*

$$\underline{s}^+ := \inf\{\underline{s} \in \mathcal{S} : g_{\underline{s}} \text{ is above } U_1\} \quad \text{and} \quad \underline{s}^- := \sup\{\underline{s} \in \mathcal{S} : g_{\underline{s}} \text{ is below } U_1\}$$

define two periodic external addresses of period n with equal itineraries.

Proof. It may not be clear a priori that $\underline{s}^+$ and $\underline{s}^-$ are well-defined sequences in $\mathbf{Z}^{\mathbf{N}}$, but their first entries are certainly integers $s_1^+ \geq s_1^-$. If $s_1^+ = s_1^-$, then iteration gives

$$\sigma(\underline{s}^+) = \inf\{\underline{s} \in \mathcal{S}\colon g_{\underline{s}} \text{ is above } U_2\} \quad \text{(and similarly for } \sigma(s^-))$$

so that $s_2^+ \geq s_2^-$ are well-defined. Repeating this argument shows that

$$\sigma^k(\underline{s}^+) = \inf\{\underline{s} \in \mathcal{S} : g_{\underline{s}} \text{ is above } U_{k+1}\}$$

where k is the first integer such that $s_{k+1}^+ > s_{k+1}^-$. Now U_{k+1} surrounds or contains a preimage of an unbounded part of $\mathbf{R}^-$, hence a preimage of a left half plane, and the next step is different:

$$\sigma^{k+1}(\underline{s}^+) = \inf\{\underline{s} \in \mathcal{S} : g_{\underline{s}} \text{ is below } U_1, \text{ and } U_{k+2} \text{ does not separate } g_{\underline{s}} \text{ from } U_1\}$$

(the problem is that the lower part of U_{k+1} maps above U_1, while the upper part maps below U_1, and we are interested only in the image of the upper part; another way of saying this is that $\sigma^{k+1}(\underline{s}^+)$ is the infimum of sequences which are above those parts of U_{k+2} that are below U_1). Note that the infimum still has finite first entry because some part of U_{k+2} is below U_1. We can continue to iterate this:

$$\begin{aligned}\sigma^{k+m}(\underline{s}^+) = \inf\{\underline{s} \in \mathcal{S} :& g_{\underline{s}} \text{ is below } U_m, \\ & \text{and } U_{k+m+1} \text{ does not separate } g_{\underline{s}} \text{ from } U_m\}\end{aligned}$$

where m is the first integer such that U_m surrounds or contains a preimage of an unbounded part of $\mathbf{R}^-$, or there is a preimage of an unbounded part of $\mathbf{R}^-$ below U_m which separates U_m from U_{k+m+1}. If that happens, we get

$$\sigma^{k+m+1}(\underline{s}^+) = \inf\{\underline{s} \in \mathcal{S} : g_{\underline{s}} \text{ is above } U_{k+m+2}\}$$

and we are back to the initial situation. Repeating these arguments for a total of $n-2$ iterations, we see that either

$$\sigma^{n-2}(\underline{s}^+) = \inf\{\underline{s} \in \mathcal{S} : g_{\underline{s}} \text{ is above } U_{n-1}\}, \tag{8}$$

or there is a $k' \leq n$ such that (8) is false, but

$$\begin{aligned}\sigma^{n-2}(\underline{s}^+) = \inf\{s \in \mathcal{S} :& g_{\underline{s}} \text{ is below } U_{k'}, \\ & \text{and } U_{n-1} \text{ does not separate } g_{\underline{s}} \text{ from } U_{k'}\}.\end{aligned} \tag{9}$$

In the case of (8), we get

$$\begin{aligned}\sigma^{n-1}(\underline{s}^+) = \inf\{s \in \mathcal{S} :& g_{\underline{s}} \text{ is in the same region } R_{\mathrm{u}_1} \text{ as } U_1, \\ & \text{and } U_n = U_0 \text{ does not separate } g_{\underline{s}} \text{ from } U_1\}\end{aligned}$$

(because U_{n-1} and the region surrounded by U_{n-1} together map homeomorphically over all of $\mathbf{C}$ except a subset of $R_{\mathtt{u}_1}$), and then

$$\sigma^n(\underline{s}^+) = \inf\{s \in \mathcal{S} : g_{\underline{s}} \text{ is above } U_1\}.$$

Similarly, in the case of (9), we get

$$\sigma^{n-1}(\underline{s}^+) = \inf\{s \in \mathcal{S} : g_{\underline{s}} \text{ is in the same region } R_{\mathtt{u}_{k'}+1} \text{ as } U_{k'+1}, \\ \text{and } U_n \text{ does not separate } g_{\underline{s}} \text{ from } U_{k'+1}\}$$

and again $\sigma^n(\underline{s}^+) = \inf\{s \in \mathcal{S} : g_{\underline{s}} \text{ is above } U_1\}$.

Thus in every case $\sigma^n(\underline{s}^+) = \underline{s}^+$, and this is a periodic sequence over $\mathbf{Z}$ with period n. The same applies to $\underline{s}^-$, and both have itineraries of period (dividing) n.

The first $n-1$ entries in the itinerary of g_{s^+} are $\mathtt{u}_1\mathtt{u}_2\cdots\mathtt{u}_{n-1}$, which are the same as in the itinerary of U_1 or of the singular value. The nth entry in the itinerary of U_1 is not defined (because the image component U_n extends through all fundamental domains), but we saw above that the nth entry in the itinerary of g_{s^+} is either $\mathtt{u}_1$ or $\mathtt{u}_{k'+1}$, and g_{s^-} has the same itinerary. □

Theorem 6.2 (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. The characteristic point of this periodic orbit is on the boundary of the characteristic Fatou component.*

Proof. Existence follows from Lemma 6.1: the periodic dynamic rays g_{s^+} and g_{s^-} have identical itineraries, so by the results of [SZ3] mentioned in Section 2, they land at a common 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 $\mathbf{C} \setminus \overline{\bigcup_{k\geq 0} E_\lambda^{\circ k}(0)}$.

Assume that $l > 0$. 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 has decreased. This contradiction shows that $l = 0$ and $z \in \partial U_1$.

Let z_1 be the characteristic point of the orbit of z (cf. Definition 5.1 and Lemma 5.2). The two characteristic rays separate the singular value from the orbit of z (except z_1 itself) and from a left half plane. If $z_1 \neq z$, then the characteristic rays would separate $z \in \partial U_1$ from a left half plane, and this is a contradiction. This proves the existence claim for the periodic Fatou component U_1, and for the others it follows easily.

Now we show uniqueness. The point z_1, together with the two characteristic rays landing at it, cut $\mathbf{C}$ into two open parts; let V be the one containing the singular value. 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 $z_1' \in V$ 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. □

Remark. The same statement holds also for parabolic dynamics; the proof requires only the same modifications as in [SZ3, Section 4.3].

Definition 6.3 (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 6.2) will be called the *dynamic root* of the characteristic Fatou component.

Lemma 6.4 (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. Every periodic Fatou component U_i has a unique point z_i on its boundary (Theorem 6.2). 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. 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 them 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 5.2 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.) □

7. Parametrization of hyperbolic components

Theorem 7.1 (Parametrization of hyperbolic components). *For every hyperbolic component H of period $n \geq 2$, there is a unique conformal isomorphism $\Phi_H\colon H \to \mathbf{H}$ with $\mu = \exp \circ \Phi_H$ such that if $\mu(\lambda) > 0$, then $\operatorname{Im}\bigl(\Phi_H(\lambda)\bigr) = 2\pi\Delta$, where the integer Δ specifies how many fundamental domains the periodic point $a_0 \in U_0$ is above the dynamic root z_0 of U_0 (or below if $\Delta < 0$).*

Proof. We already know from Lemma 4.6 the existence of a conformal isomorphism $\Phi\colon H \to \mathbf{H}$ with $\mu = \exp \circ \Phi_H$, which is unique up to addition of a constant in $2\pi i\mathbf{Z}$ in the range, so we only need to justify the combinatorial interpretation of imaginary parts if $\mu(\lambda) > 0$.

Choose any $\lambda \in H$. Let z_1 be the dynamic root of the characteristic Fatou component U_1 and let $\underline{s}_1 < \underline{s}_2$ be the external addresses of its two characteristic rays. Let z_0 be the dynamic root of U_0 (with $E_\lambda(z_0) = z_1$) and let $\underline{s}'_2 < \underline{s}'_1$ be the least and greatest external addresses (with respect to lexicographic ordering) of two dynamic rays landing at z_0. By continuity, these external addresses are the same for every $\lambda \in H$.

Since the rays $g_{\underline{s}_1}$ and $g_{\underline{s}_2}$ surround 0, they bound the characteristic sector at z_1, which is the image of the sector at z_0 containing a left half plane. It follows that $\sigma(\underline{s}'_1) = \underline{s}_1$ and $\sigma(\underline{s}'_2) = \underline{s}_2$, and there is an $m \in \mathbf{Z}$ with $\underline{s}'_2 = m\underline{s}_2$ and $\underline{s}'_1 = (m+1)\underline{s}_1$ (concatenation). By the translation symmetry, it follows that for every $m' \in \mathbf{Z}$, the rays at addresses $m'\underline{s}_2$ and $(m'+1)\underline{s}_1$ land together at $z_0 + 2\pi i(m'-m)$. Since $\underline{s}_1$ and $\underline{s}_2$ are characteristic external addresses, no $\sigma^k(\underline{s}_1)$ or $\sigma^k(\underline{s}_2)$ is contained in $]\underline{s}_1, \underline{s}_2[$; therefore, no $\sigma^k(\underline{s}_1)$ or $\sigma^k(\underline{s}_2)$ is contained in any $]m'\underline{s}_1, m'\underline{s}_2[$ for $m' \in \mathbf{Z}$. It follows that there is an $m_1 \in \mathbf{Z}$ with

$$m_1\underline{s}_2 < \underline{s}_1 < \underline{s}_2 < (m_1+1)\underline{s}_1;$$

since the ray pair $g_{\underline{s}_1}$ and $g_{\underline{s}_2}$ surround the singular value 0, it follows that the ray pair at addresses $m_1\underline{s}_2$ and $(m_1+1)\underline{s}_1$ surrounds 0 as well. We had seen above that the ray pair at addresses $m\underline{s}_2$ and $(m+1)\underline{s}_1$ lands at z_0.

Now suppose that $\mu(\lambda) > 0$. Then z_0 is exactly $m - m_1$ fundamental domains above 0. All we used for this reasoning are the external addresses $\underline{s}_1$, $\underline{s}_2$, $\underline{s}'_1$ and $\underline{s}'_2$ which are the same for any $\lambda \in H$, so the index $m - m_1$ is the same throughout H.

Recall the index Δ' from Lemma 4.6 which specifies, whenever $\mu(\lambda) > 0$, how many fundamental domains a_0 are above 0. Therefore, $\Delta := \Delta' - (m - m_1)$ specifies how many fundamental domains a_0 are above z_0. By Lemma 4.6, every conformal isomorphism $\Phi\colon H \to \mathbf{H}$ with $\mu = \exp\circ\Phi$ has an integer $m_\Phi \in \mathbf{Z}$ such that for all $\lambda \in H$ with $\mu(\lambda) > 0$, $\operatorname{Im}\bigl(\Phi(\lambda)\bigr) = 2\pi(\Delta' - m_\Phi)$. Setting $\Phi_H := \Phi + 2\pi i\bigl(m_\Phi - (m - m_1)\bigr)$ yields the desired conformal isomorphism. It is clearly unique. □

Having stated this theorem, we should outline why this parametrization is indeed a preferred one, for example over the one from Lemma 4.6 counting fundamental domains between a_0 and the singular value. The dynamic root clearly has dynamic significance, and the invariant curve γ^- from Lemma 4.1 is part of an analytic curve containing the singular orbit. Both might be linked for any particular exponential map: the following lemma states when this happens. It should come as no surprise that the locus of such maps stands out in parameter space; this is discussed at the end of this section.

Lemma 7.2 (Invariant curve lands at dynamic root). *For every exponential map which has an attracting orbit of period $n \geq 2$ with positive real multiplier,* *the invariant curve γ_- from Lemma 4.1 lands at the dynamic root $z_1 \in \partial U_1$ if and only if $\Delta = 0$.*

Proof. The curve γ_- starts at a_1 and lands at a point $w_1 \in \partial U_1$ with $E_\lambda^{\circ n}(w_1) = w_1$. Let $\gamma_0 \subset U_0$ be the branch of $E_\lambda^{-1}(\gamma_-)$ starting at a_0; it lands at a point $w_0 \in \partial U_0$ with $E_\lambda^{\circ n}(w_0) = w_0$. Clearly, $w_1 = E_\lambda(w_0)$, so w_0 is the dynamic root of U_0 if and only if w_1 is the dynamic root of U_1.

Recall the analytic curve $\gamma_+ \subset U_1$ from Lemma 4.1: it connects a_1 to $+\infty$ such that $\gamma_+' := E_\lambda^{\circ n}(\gamma_+) \subset \gamma_+$, and γ_+' contains the singular orbit within U_1. The fundamental domains for E_λ are bounded by $E_\lambda^{-1}(\gamma_+ \setminus \gamma_+')$.

Since γ_0 is disjoint from the boundary of the fundamental domains, it follows that w_0 is in the same fundamental domain as a_0, so w_0 can be the dynamic root $z_0 \in U_0$ only if a_0 and z_0 are in the same fundamental domain, i.e. only if $\Delta = 0$.

Conversely, we show that w_0 is the only boundary point of U_0 which is fixed by $E_\lambda^{\circ n}$ and which is in the same fundamental domain as a_0; this implies that whenever $\Delta = 0$, then γ_0 lands at z_0 and γ_- lands at z_1.

Suppose there is another point $w_0' \in \partial U_0$ with $E_\lambda^{\circ n}(w_0') = w_0'$ in the same fundamental domain as w_0. There is a curve $\gamma_0' \subset U_0$ with $E_\lambda^{\circ n}(\gamma_0') \supset \gamma_0'$ which lands at w_0' (take a point b_0' in a linearizing neighborhood of w_0' and pull it back repeatedly so as to obtain a sequence (b_k') converging to w_0'; these can be connected by a curve γ_0' as required). Let (b_k) be an analogous sequence of points converging to w_0.

Connect b_0 to b_0' by a differentiable curve $\Gamma_0 \subset U_0$ which avoids the fundamental domain boundaries and the curve $E_\lambda^{\circ(n-1)}(\gamma_+)$ which contains the singular orbit within U_0. Then $E_\lambda^{-n}(\Gamma_0)$ contains a curve Γ_1 connecting b_1 to b_1' and homotopic to Γ_0 in $U_0' := U_0 \setminus \big(E_\lambda^{-1}(\gamma_+) \cup \{a_0\}\big)$ (the fact that the same branch of E_λ^{-n} fixes both w_0 and w_0' uses the assumption that both points are within the same fundamental domain and would fail if w_0' was an arbitrary fixed point of $E_\lambda^{\circ n}$ in ∂U_0). This can be repeated, and the hyperbolic lengths of Γ_k within U_0' become shorter each time, while $b_k, b_k' \to \partial U_0$, and this implies $w_0 = w_0'$. □

For the following definition, recall the conformal isomorphism $\Phi_H: H \to \mathbf{H}$ which maps parameter rays of H to horizontal lines in $\mathbf{H}$.

Definition 7.3 (Height and central internal ray). Given a hyperbolic component H of period $n \geq 2$ with preferred conformal isomorphism $\Phi_H: H \to \mathbf{H}$, we define the *height* of a parameter ray as the number $h \in \mathbf{R}$ such that $\operatorname{Im}\big(\Phi_H(\lambda)\big) = 2\pi h$ for λ on this parameter ray.

The *central parameter ray* of H is the parameter ray at height $h = 0$.

Remark. The angle of a parameter ray is the projection of h to $\mathbf{R}/\mathbf{Z}$. If $\lambda \in H$ has $\mu(\lambda) > 0$, then clearly $h = \Delta$. Note that for components of period $n \geq 3$ in λ-space (and of period $n \geq 2$ in κ-space) a ray with larger height is *below* a ray with smaller height: see the discussion after Definition 4.5. This is unavoidable if among nearby rays one wants rays at larger heights to have larger angles of their multipliers.

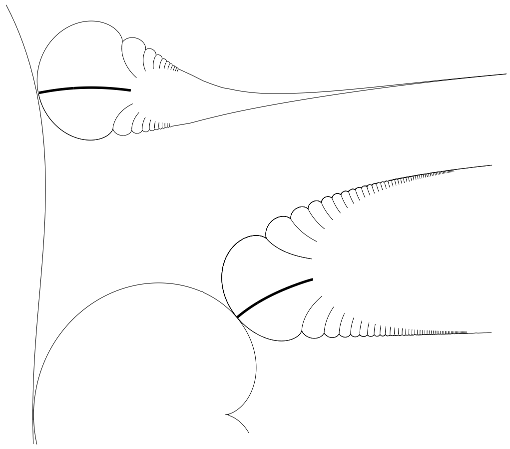

Figure 3. Hyperbolic components of periods 3 and 4, with internal parameter rays at integer heights drawn in. In both components, the central parameter rays are highlighted. Their landing points are the roots of the component (at which the components bifurcate from components of period 1 respectively 2).

Remark. We have classified hyperbolic components in terms of intermediate external addresses. A different coding would be in terms of the external addresses of the two characteristic dynamic rays landing at the root of the characteristic Fatou component: clearly, all exponential maps from the same hyperbolic component have the same external addresses at the dynamic rays landing at the root, so there is an algorithm to turn the intermediate external address of the component into the external addresses of the characteristic rays. The converse is much easier: knowing the external addresses of the two characteristic rays, it is not hard to write down the intermediate external address of the attracting dynamic ray and thus of the component. Both algorithms can be found in [RS]. While the coding in terms of intermediate external addresses leads to the easier classification, the characteristic external addresses are more easily related to the structure of the Multibrot sets $\mathcal{M}_d$ and their limiting configurations as $d \to \infty$.

The boundary of hyperbolic components and bifurcations. While this paper completely describes hyperbolic components of exponential maps, their boundary properties are discussed in [S1, Section V] and [RS]. To complete the picture, we mention some results here.

By Theorem 7.1, every hyperbolic component H (except the period 1 component in λ-space) comes with a preferred conformal isomorphism $\Phi_H\colon H \to \mathbf{H}$. It extends as a homeomorphism to the closures $\overline{\Phi}_H\colon \overline{H} \to \overline{\mathbf{H}}$. This result requires a much better understanding of the exponential parameter space. It implies that every hyperbolic component has connected boundary $\partial H = \overline{\Phi}_H^{-1}(i\mathbf{R})$, which was conjectured by Eremenko and Lyubich [EL2].

Different hyperbolic components may have common boundary points. This happens if and only if one hyperbolic component bifurcates from another, and then both components have a unique boundary point in common. The structure of bifurcations between hyperbolic components in exponential parameter space is cleared up in [RS]. In particular, this proves a third conjecture in [EL2] which states that there are infinitely many "trees" of hyperbolic components such that two components are in the same tree if and only if they can be connected via a finite chain of components so that adjacent components in the chain have common boundary.

Every internal parameter ray of H at height h *lands* at a well-defined parameter in ∂H with indifferent orbit so that the landing point depends continuously on h. The landing point of the central parameter ray ($h = 0$) is called the *root* of H, and it is significant in several ways: the root of H is the only boundary point which may simultaneously be a boundary point of another hyperbolic component H_0, and this happens if and only if H bifurcates from H_0. Moreover, every boundary point of H is the landing point of one or two *external parameter ray* (curves consisting of parameters for which the singular orbit converges to ∞ under iteration, see [S1, Chapter II] or [F]). Boundary points of H at positive heights are the landing points of parameter rays which come from $+\infty$ below H, while boundary points at negative heights are landing points of parameter rays above H. The root point of H is the landing point of two parameter rays, one above and one below H [S1, Chapter IV]. This shows once more the significance of our preferred parametrization from Theorem 7.1.

References

[AO] Aarts, J., and L. Overstegen: The geometry of Julia sets. - Trans. Amer. Math. Soc. 338, 1993, 897–918.

[BR] Baker, I.N., and P.J. Rippon: Iteration of exponential functions. - Ann. Acad. Sci. Fenn. Ser. A I Math. 9, 1984, 49–77.

[DFJ] Devaney, R., N. Fagella, and X. Jarque: Hyperbolic components of the complex exponential family. - Manuscript, 2000.

[DGH] Bodelón, C., R. Devaney, L. Goldberg, M. Hayes, J. Hubbard, and G. Roberts: Hairs for the complex exponential family. - Internat. J. Bifurc. Chaos Appl. Sci. Engrg. 9:8, 1999, 1517–1534.

[DH1] Douady, A., and J. Hubbard: Etude dynamique des polynômes complexes. - Prépublications mathématiques d'Orsay 2, 1984, and 4, 1985.

[DH2] Douady, A., and J. Hubbard: On the dynamics of polynomial-like mappings. - Ann. Sci. École Norm. Sup. 18, 1985, 287–343.

[EL1] Eremenko, A., and M. Lyubich: Iterates of entire functions. - Preprint, physico-technical institute of low-temperatures Kharkov 6, 1984 (Russian).

[EL2] Eremenko, A., and M. Lyubich: Structural stability in some families of entire functions. - Preprint, physico-technical institute of low-temperatures Kharkov 29, 1984 (Russian).

[EL3] Eremenko, A., and M. Lyubich: Dynamical properties of some classes of entire functions. - Ann. Inst. Fourier (Grenoble) 42:4, 1992, 989–1020.

[ES] Eberlein, D., and D. Schleicher: On the structure of Multibrot sets. - Manuscript, to be published.

[F] Förster, M.: Parameter Rays for Exponential Maps. - Diploma thesis, Technische Universität München, 2002.

[HS] Hubbard, J., and D. Schleicher: The spider algorithm. - In: Complex Dynamical Systems, edited by R. Devaney, Amer. Math. Soc., Providence, 1994, 155–180.

[M1] Milnor, J.: Hyperbolic components in spaces of polynomial maps. - Preprint, Institute for Mathematical Sciences at Stony Brook 3, 1992.

[M2] Milnor, J.: Dynamics in One Complex Variable: Introductory Lectures. - Vieweg, 1999.

[M3] Milnor, J.: Periodic orbits, external rays and the Mandelbrot set; an expository account. - Astérisque 261, 2000, 277–333.

[Mc1] McMullen, C.: Automorphisms of rational maps. - In: Holomorphic Functions and Moduli I, MSRI Publications 10, edited by D. Drasin, C. Earle, F. Gehring, I. Kra, and A. Marden, Springer-Verlag, 1988, 31–60.

[Mc2] McMullen, C.: Area and Hausdorff dimension of Julia sets of entire functions. - Trans. Amer. Math. Soc. 300:1, 1987, 329–342.

[Mc3] McMullen, C.: The Mandelbrot set is universal. - In: The Mandelbrot Set: Theme and Variations, edited by Tan Lei, London Math. Soc. Lecture Note Ser. 274, 2000.

[MSS] Mañe, R., P. Sad, and D. Sullivan: On the dynamics of rational maps. - Ann. Sci. École Normale Sup. 16, 1983, 193–217.

[RS] Rempe, L. and D. Schleicher: Bifurcations in the space of exponential maps. - Manuscript, 2001.

[S1] Schleicher, D.: On the Dynamics of Iterated Exponential Maps. - Habilitationsschrift, Technische Universität München, 1999.

[S2] Schleicher, D.: On fibers and local connectivity of compact sets in $\mathbf{C}$. - Preprint, Institute for Mathematical Sciences at Stony Brook 12, 1998.

[S3] Schleicher, D.: Attracting dynamics of exponential maps. - Preprint, Institute for Mathematical Sciences at Stony Brook 4, 2000.

[SZ1] Schleicher, D., and J. Zimmer: Dynamic rays for exponential maps. - Preprint, Institute for Mathematical Sciences at Stony Brook 9, 1999. Submitted for publication as [SZ2] and [SZ3].

[SZ2] Schleicher, D., and J. Zimmer: Escaping points of exponential maps. - J. London Math. Soc. (2) 67, 2003, 380–400.

[SZ3] Schleicher, D., and J. Zimmer: Periodic points and dynamic rays of exponential maps. - Ann. Acad. Sci. Fenn. Math. (to appear).

Received 30 May 2000