

Periodic Points and Dynamic Rays of Exponential Maps

Dierk Schleicher and Johannes Zimmer

January 13, 2003

Abstract

We investigate the dynamics of exponential maps $z \mapsto \lambda e^z$; the goal is a description by means of dynamic rays. We discuss landing properties of dynamic rays and show that in many important cases, repelling and parabolic periodic points are landing points of periodic dynamic rays. For postsingularly finite exponential maps, we use spider theory to show that a dynamic ray lands at the singular value.

Contents

1	Introduction	2
2	Escaping Points and Dynamic Rays	3
3	Landing Properties of Dynamic Rays	5
4	Periodic Points and Dynamic Rays	9
4.1	Review of symbolic dynamics of polynomials	10
4.2	The Escape Case	12
4.3	The Postsingularly Finite Case	13
4.4	The Attracting and Parabolic Cases	15
5	(Pre-)Periodic Points are Landing Points	18
6	Spiders	27

Keywords: Iteration, exponential map, dynamic ray, landing point, symbolic dynamics.

MSC Classification (2000): 30D05, 33B10, 37B10, 37C25, 37F10, 37F20

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

Johannes Zimmer, Division of Engineering and Applied Science, Mail Stop 104-44, California Institute of Technology, Pasadena, CA 91125, USA, zimmer@caltech.edu.

1 Introduction

This paper is a contribution to the program to carry over results from the theory of iterated polynomials to iterated entire maps; we will focus on the family $z \mapsto \lambda \exp(z) = \exp(z + \kappa)$ with $\lambda \in \mathbb{C} \setminus \{0\}$ or $\kappa \in \mathbb{C}$. We believe that the methods will carry over to other entire maps as well.

In [SZ2] the existence of *dynamic rays* is proved. In this paper, we will show that these rays behave similar to those well known from the theory of iterated polynomials: for a large class of parameters (including the structurally important ones), periodic dynamic rays land at periodic points and periodic points are landing points of periodic dynamic rays; for those parameters, it is also possible to tell which dynamic rays land together.

After the pioneering work by Douady and Hubbard, various models for Julia sets and parameter spaces have been developed which all tried to describe the topology in terms of dynamic rays landing together: among them are Douady's pinched disks [Do], Thurston's laminations [T] and Milnor's orbit portraits [M2], all of which have been explored deeply only in the simplest context of quadratic polynomials. Our paper is intended as a first step of an investigation of exponential maps in a similar spirit, and many of the methods in these papers can now be applied to exponential maps.

In a sequel to this paper, it will be explained how the parameter space of exponential maps can profitably be studied in terms of parameter rays (external rays in parameter space) to bring out analogies and differences to the Mandelbrot set more clearly.

This paper is organized as follows: first, some definitions and results of [SZ2] are reviewed in Section 2. Landing properties of dynamic rays are then discussed in Section 3; we focus on periodic and preperiodic rays because those are known to describe the structure in the polynomial case. Section 4 discusses which rays land at which points in terms of itineraries with respect to dynamic partitions. These partitions must be constructed with respect to the kind of dynamics at hand; we cannot construct them in every case, but our methods cover those parameters that are most important for an understanding of parameter space. We want to prove that repelling and parabolic periodic points are landing points of periodic dynamic rays and reduce this to a combinatorial problem, which is then solved in Section 5. Technically the most difficult case is that of postsingularly finite exponential maps. In order to obtain a useful partition, we need to know that the singular value is the landing point of at least one preperiodic dynamic ray. We prove

this in Section 6 using spider theory.

An earlier version of this paper was first circulated as a preprint [SZ1]; it contains ideas from the Diploma thesis [Zi].

SOME NOTATION. Let $\mathbb{C}^* := \mathbb{C} \setminus \{0\}$, $\mathbb{C}' := \mathbb{C}^* \setminus \mathbb{R}^-$ and $\overline{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$. We use the notation $E_\lambda(z) := \lambda \exp(z) = \exp(z + \kappa)$, where κ fixes a particular choice of $\log \lambda$. While the exact choice of the logarithm is in principle inessential, we have written our estimates and combinatorics for $\text{Im}(\kappa) \in [-\pi, \pi]$; this is no loss of generality. Although many of our constructions depend on κ , we will usually suppress that from the notation.

The principal branch of the logarithm in \mathbb{C}' will be denoted Log . A frequently used abbreviation is $F(t) := \exp(t) - 1$.

We sometimes say that a sequence in \mathbb{C} converges to $+\infty$ to indicate that it converges to ∞ along bounded imaginary parts and with real parts diverging towards $+\infty$.

ACKNOWLEDGMENTS. We would like to thank John Milnor for the invitation to the Institute for Mathematical Sciences at Stony Brook where this work started, and to the hospitality there. We would also like to thank the Studienstiftung des deutschen Volkes for its support all along and in particular during the stay in Stony Brook. Moreover, we have enjoyed helpful discussions with Noel Baker, Bob Devaney, Adrien Douady, John Hubbard, Phil Rippon and Mitsuhiro Shishura.

2 Escaping Points and Dynamic Rays

In this section, we briefly recollect some facts from [SZ2].

Definition 2.1 (Escaping Point)

A point $z \in \mathbb{C}$ with $E_\lambda^{o_n}(z) \rightarrow \infty$ (in the Riemann sphere) as $n \rightarrow +\infty$ will be called an escaping point; its orbit will be called an escaping orbit.

REMARK. It follows from $|E_\lambda(z)| = |\lambda| \exp \text{Re}(z)$ that every escaping point z satisfies $\text{Re}(E_\lambda^{o_n}(z)) \rightarrow +\infty$ as $n \rightarrow +\infty$. Since every Fatou component is periodic or preperiodic and belongs to an attracting or parabolic orbit or to an orbit of Siegel disks [EL1, EL2], every escaping point belongs to the Julia set.

For $j \in \mathbb{Z}$ we define the strips

$$R_j := \{z \in \mathbb{C} : -\text{Im}(\kappa) - \pi + 2\pi j < \text{Im}(z) < -\text{Im}(\kappa) + \pi + 2\pi j\};$$

then $E_\lambda: R_j \rightarrow \mathbb{C}'$ is a conformal isomorphism for every j . The assumption $|\operatorname{Im}(\kappa)| \leq \pi$ implies that the singular value 0 is always in $\overline{R_0}$. The collection $\{R_j\}$ is a partition of the complex plane; the boundaries are the preimages of the negative real axis. The inverse of E_λ mapping \mathbb{C}' to R_j will be denoted by $L_{\kappa,j}$, so that $L_{\kappa,j}(z) = \operatorname{Log}z - \kappa + 2\pi ij$. As a consequence, $R_j \subset E_\lambda(R_k)$ for every $j \neq 0$ and every k . The strip R_0 is the only one which intersects the image of the boundary of an arbitrary strip.

The strips R_k are limits of sectors of angles $2\pi/d$ with vertices at $-d$, which are fundamental domains for the polynomials $z \mapsto \lambda(1 + z/d)^d$ which approximate $z \mapsto \lambda \exp(z)$.

The idea to define a partition by considering the preimages of the negative real axis can be found in [DGH]; we call such a partition a *static partition* (as opposed to various dynamic partitions introduced in Section 4).

Definition 2.2 (External Address)

Let $\mathcal{S} := \{s_1 s_2 s_3, \dots : \text{all } s_k \in \mathbb{Z}\}$ be the space of sequences over the integers, and let σ be the shift map on \mathcal{S} . We will often use the abbreviation $\underline{s} = s_1 s_2 s_3 \dots$. For any $z \in \mathbb{C}$ with $E_\lambda^{\circ k}(z) \notin \mathbb{R}^-$ for all $k \in \mathbb{N}$, the external address $S(z) \in \mathcal{S}$ is the sequence $s_1 s_2 s_3 \dots$ with $E_\lambda^{\circ k}(z) \in R_{s_{k+1}}$ for all $k \geq 0$.

Since we mainly consider periodic and preperiodic orbits in this paper, most external addresses will be bounded; we write $|\underline{s}| := \sup\{|s_k|\}$ for bounded $\underline{s} := s_1 s_2 s_3 \dots \in \mathcal{S}$.

The following simplified version of [SZ2, Theorem 4.2 and Theorem 6.5] will be sufficient for our purpose.

Theorem 2.3 (Dynamic Rays)

If the singular orbit does not escape, then for every bounded \underline{s} there is a unique injective and continuous curve $g_{\underline{s}}:]0, \infty[\rightarrow \mathbb{C}$ of external address \underline{s} satisfying

$$\lim_{t \rightarrow +\infty} \operatorname{Re}(g_{\underline{s}}(t)) = +\infty$$

which has the following properties: it consists of escaping points such that

$$E_\lambda(g_{\underline{s}}(t)) = g_{\sigma(\underline{s})}(F(t)) \quad \text{for every } t > 0, \quad (1)$$

and $g_{\underline{s}}(t) = t - \kappa + 2\pi i s_1 + r_{\underline{s}}(t)$ with $|r_{\underline{s}}(t)| < 2e^{-t}(|\kappa| + C)$, where $C \in \mathbb{R}$ depends only on a bound for \underline{s} .

If the singular orbit does escape, then the statement is still true for every bounded \underline{s} for which there is no $n \geq 1$ and $t_0 > 0$ such that $0 = g_{\sigma^n(\underline{s})}(t_0)$. For those exceptional \underline{s} , there is an injective curve $g_{\underline{s}}:]t_{\underline{s}}^*, \infty[\rightarrow \mathbb{C}$ with the same properties as before, where $t_{\underline{s}}^* \in \mathbb{R}^+$ is the largest number that has an $n \geq 1$ such that $F^{\circ n}(t_{\underline{s}}^*) = t_0$ and $0 = g_{\sigma^n(\underline{s})}(t_0)$.

REMARK. In fact, one can show more: dynamic rays yield a complete characterization of escaping points in the sense that every escaping point is either on a dynamic ray, or the landing point of a ray (or mapped to a point on a dynamic ray under iteration in the case of an escaping singular orbit). See [SZ2, Theorem 6.5].

The curve $g_{\underline{s}}:]0, \infty[\rightarrow \mathbb{C}$ (or $g_{\underline{s}}:]t_{\underline{s}}^*, \infty[\rightarrow \mathbb{C}$) will be called the *dynamic ray at external address \underline{s}* . We should note that in the exceptional case, it is no longer true that the $E_\lambda^{\circ k}$ -image of the \underline{s} -ray maps entirely onto the $\sigma^k(\underline{s})$ -ray; it only covers an unbounded part of the ray which terminates at some postsingular point. The dynamic ray $g_{\underline{s}}$ is called periodic or preperiodic whenever its external address \underline{s} is periodic or preperiodic; while no points $g_{\underline{s}}(t)$ are periodic, the ray as a set is periodic (unless the forward orbit of some $g_{\underline{s}}(t)$ contains the singular value).

3 Landing Properties of Dynamic Rays

For $\lambda \in \mathbb{C}^*$, we define the postsingular set as $P := P(E_\lambda) := \overline{\cup_{n \geq 0} E_\lambda^{\circ n}(0)}$. In this section, we will discuss landing properties of dynamic rays. In order to see that E_λ^{-1} is a contraction with respect to a hyperbolic metric on an appropriate domain, we need to make certain assumptions for the postsingular set; included are the important cases of maps with attracting and parabolic orbits, as well as postsingularly finite maps (for which the singular orbit is necessarily strictly preperiodic), and maps for which the singular orbit escapes to ∞ .

If a domain U admits a hyperbolic metric, we assume that this metric is normalized with constant curvature -1 . The hyperbolic distance of two points $z_1, z_2 \in U$ will be denoted by $d_U(z_1, z_2)$. We write $\rho_U(z)$ for the density of d_U with respect to the Euclidean metric $|dz|$.

Let us start with an auxiliary statement.

Lemma 3.1 (Bounded Ratio of Hyperbolic Densities)

Let $V \subset U \subset \mathbb{C}$ be two hyperbolic domains and $Z \subset V$ be an arbitrary subset. Suppose there is an $s > 0$ such that for every $z \in Z$ there exists $u \in U \setminus V$ with $d_U(u, z) < s$. Then there exists $\eta < 1$ such that $\frac{\rho_U(z)}{\rho_V(z)} < \eta$ for every $z \in Z$.

PROOF. By passing to a universal cover, we may replace U by \mathbb{D} and suppose that $z = 0$; the ratio $\rho_U(z)/\rho_V(z)$ remains unchanged. Within $U = \mathbb{D}$, we have $\rho_U(0) = 2$, and there is an upper bound (less than 1) for $|u|$ for some $u \in U \setminus V$ and hence a lower bound (greater than 2) for $\rho_V(0)$; these bounds depend only on s . This implies that $\rho_U(0)/\rho_V(0) < \eta < 1$, where η depends only on s . \square

The main result of this section is the following.

Theorem 3.2 (Landing of (Pre-)periodic Rays)

If the singular orbit is bounded in \mathbb{C} , then every periodic dynamic ray lands at a periodic point; moreover, every preperiodic dynamic ray lands at a preperiodic point unless some forward iterate of the preperiodic ray lands at the singular value, which is then necessarily preperiodic.

If the singular value escapes, then the periodic or preperiodic dynamic ray at external address \underline{s} lands at a periodic respectively preperiodic point, unless $g_{\sigma^n(\underline{s})}(t) = 0$ for some $n \geq 1$ and $t > 0$.

REMARK. This theorem does not cover all cases of exponential maps: the singular orbit might be unbounded without escaping. In particular if the singular orbit is dense in \mathbb{C} , we do not know whether an analogous claim holds.

PROOF. Let us consider periodic rays first. We will frequently pull back rays by the dynamics: for every periodic ray $g_{\underline{s}}$ which does not contain the singular value, $E_{\lambda}^{-1}(g_{\underline{s}})$ is a countable collection of disjoint rays, exactly one of which is periodic. This periodic ray can be pulled back again unless it contains the singular value, etc.. It follows that every periodic ray can be pulled back infinitely often, yielding a sequence of periodic rays, unless some forward image of the periodic ray contains the singular value; this exception is exactly the excluded case in the statement. With this exception, the cases of bounded and escaping singular orbits can be treated simultaneously.

First note that the postsingular set P does not intersect $g_{\underline{s}}([0, \infty[)$: by hypothesis, $g_{\underline{s}}([0, \infty[)$ cannot contain a point on the singular orbit, and if it contained a limit point of the singular orbit then P could neither be bounded nor discrete (but if the singular orbit escapes, then P is necessarily discrete).

By assumption, P cannot dissect the complex plane into two unbounded domains. Hence, the complement $\mathbb{C} \setminus P$ contains exactly one unbounded connected component; call it U . Since P contains at least the two points 0 and $\lambda \neq 0$, U carries a unique normalized hyperbolic metric. Let $V := E_{\lambda}^{-1}(U)$. The map $E_{\lambda}: V \rightarrow U$ is a holomorphic covering and thus a local hyperbolic isometry. Since P is forward invariant, we have $V \subset U$, and this is a proper inclusion because clearly $E_{\lambda}^{-1}(P) \neq P$. Therefore, the inclusion $V \hookrightarrow U$ is a strict contraction for the respective hyperbolic metrics. It follows that a curve with finite length l_1 with respect to the hyperbolic metric in U will be mapped by any branch of E_{λ}^{-1} to a curve with length $l_2 < l_1$, again with respect to the hyperbolic metric in U .

We extend the density $\rho_V(z)$ of the hyperbolic metric on V to a continuous map $\rho_U: U \rightarrow \mathbb{R}^+ \cup \{+\infty\}$ by setting $\rho_U(z) := +\infty$ on $U \setminus V$. Let $\eta: U \rightarrow \mathbb{R}$, $z \mapsto \rho_U(z)/\rho_V(z)$ be the quotient of the densities. We have $0 \leq \eta(z) < 1$ everywhere, and since η is continuous, it is bounded away from 1 on any compact subset of U .

Now consider a periodic ray $g_{\underline{s}}$ which on its forward orbit never hits the singular value. Denote its period by n and fix a point $w_0 \in g_{\underline{s}}([0, \infty[)$. Construct a sequence w_k of points on $g_{\underline{s}}([0, \infty[)$ such that $E_{\lambda}^{\circ n}(w_{k+1}) = w_k$; this defines the points w_k uniquely. Let $d_k := d_U(w_k, w_{k+1})$ for all $k \geq 0$. These distances are finite because P is closed and does not intersect $g_{\underline{s}}([0, \infty[)$. The preceding discussion of the hyperbolic metric in U shows that $d_k > d_{k+1} > 0$ for all k .

CLAIM. Suppose there is a set $K \subset U$ with bounded real parts which is closed in \mathbb{C} and which contains infinitely many w_k . Then the sequence (w_k) converges to a limit point in K .

PROOF OF CLAIM. Let $s > 0$ be arbitrary, let $K' := \{z \in U: d_U(z, K) \leq s\}$ and $\eta_0 := \sup_{z \in K'} \{\eta(z)\}$. We first show $\eta_0 < 1$.

There exist constants $M, M' > 0$ such that $|\operatorname{Re}(z)| < M$ for $z \in K$, and every $z \in P$ with $|\operatorname{Re}(z)| \leq 2M$ has $|\operatorname{Im}(z)| \leq M'$. Since $E_{\lambda}^{-1}(P)$ is $2\pi i$ -periodic and contains $2\pi i\mathbb{Z}$, there exists a constant $L > 0$ so that for every $z \in K$, there is a $w \in 2\pi i\mathbb{Z} \cap E_{\lambda}^{-1}(P) \setminus P$ with $|z - w| < L$. Since $\{z \in \mathbb{C}: |\operatorname{Re}(z)| < 2M, |\operatorname{Im}(z)| > M'\} \subset U$, we have $\rho_U(z) < 2/M$ for every

$z \in \mathbb{C}$ with $|\operatorname{Re}(z)| < M$ and $|\operatorname{Im}(z)| > M' + M$. Since $E_\lambda^{-1}(P) \setminus P \subset U \setminus V$, every $z \in K$ with $|\operatorname{Im}(z)| > M' + M + L$ has $d_U(z, \partial V) < 2L/M$. But $\{z \in K: |\operatorname{Im}(z)| \leq M' + M + L\}$ is compact, so there is an $s_1 > 0$ so that every $z \in K$ has $d_U(z, \partial V) < s_1$. By construction of K' , every $z \in K'$ has $d_U(z, \partial V) < s_1 + s$. Now the existence of $\eta_0 < 1$ follows from Lemma 3.1.

This implies that the contraction in K' is uniform in the following sense: there is a $\mu < 1$ such that, whenever some $w_k \in K'$, then $d_{k+1} \leq \mu d_k$. We will show that once the d_k are short enough, the sequence (w_k) can escape from K only so slowly that the rate of escape will be overcome by the uniform contraction in the neighborhood K' of K .

Let $\varepsilon := s(1 - \mu)$. There is an index m such that $w_m \in K$ and, due to the uniform contraction in $K' \supset K$, $d_m < \varepsilon$. If there is an index $l > m$ with $w_l \notin K'$, let l be the least such index. But then

$$d_U(w_m, w_l) \leq \sum_{k=m}^{l-1} d_k \leq d_m \sum_{k=0}^{l-m-1} \mu^k < \frac{d_m}{1 - \mu} < \frac{\varepsilon}{1 - \mu} = s$$

and $w_l \in K'$ contrary to our assumption. Hence all w_k are in K' for all $k \geq m$. But in this region, the contraction in every step is uniform, and the sequence (d_k) converges geometrically to zero. Therefore, the sequence (w_k) converges to a limit in K' . Since K contains infinitely many points w_k , it contains the limit point. This proves the CLAIM.

Now if (w_k) converges to some $w \in U$, then the ray $g_{\underline{s}}$ lands at w : this follows because the hyperbolic diameter of the ray segment between w_0 and w_1 is finite and an upper bound for the diameters of the segments between any w_k and w_{k+1} , and contraction is uniform in a neighborhood of w .

As a consequence of the CLAIM, we may concentrate on those dynamic rays for which the sequence (w_k) enters any closed set in U with bounded real parts just finitely often. This sequence must have a limit point in $\overline{\mathbb{C}}$, hence in $\overline{\mathbb{C}} \setminus U = P \cup \{\infty\}$.

In the case of an escaping singular value, the set P is discrete in \mathbb{C} and (w_k) cannot accumulate at more than one point in P (sufficiently small neighborhoods around points in P have arbitrarily large hyperbolic distances, greater than $d_U(w_k, w_{k+1})$ for any k), or simultaneously at a point in P and at ∞ (for a similar reason). Therefore, the entire sequence w_k converges either to some $w \in P$ or to ∞ . The latter is impossible because $|w_k| \rightarrow \infty$ implies $\operatorname{Re}(w_k) \rightarrow +\infty$, but points with large positive real parts will be mapped to points with smaller positive real parts under any branch of E_λ^{-1} .

Hence $w_k \rightarrow w$ for some $w \in P$. This implies again that g_s lands at w : the segments on g_s between w_k and w_{k+1} have hyperbolic diameters bounded above by the segment between w_0 and w_1 , and since $w_k \rightarrow w \in \partial U$, this implies that the Euclidean diameters of these ray segments tend to 0 and the ray lands at w .

If P is bounded, then (w_k) cannot simultaneously accumulate at P and at ∞ because then there would a closed set $K \subset U$ with bounded real parts containing infinitely many w_k . The sequence cannot converge to ∞ for the same reason as above, so the entire limit set of (w_k) is in P . Now the bounded hyperbolic distances $d_U(w_k, w_{k+1})$ imply Euclidean distances $|w_k - w_{k+1}| \rightarrow 0$ because the points w_k are near the boundary of U . Therefore, every limit point of w_k in P must be a fixed point of E_λ^{cn} . The set of such fixed points is discrete, while the set of limit points of (w_k) must be connected, so $w_k \rightarrow w \in P$. But this implies again that the ray g_s lands at w : the bounded hyperbolic diameters of the ray segments from w_k to w_{k+1} must have Euclidean diameters tending to 0.

For preperiodic rays, the statement follows by pulling back periodic rays; this is possible if the corresponding periodic rays land, and if the pull-back never runs through the singular value. \square

REMARK. The periodic landing point must be repelling or parabolic for the same reason as in the polynomial case (the Snail Lemma; compare again [M1, § 13 and § 18] or [St, § 6.1]).

4 Periodic Points and Dynamic Rays

For the symbolic description of the dynamics, in our case of exponential maps, it is of great importance to find an appropriate partition of the dynamic plane. We started with the static partition which is bounded by horizontal lines that are mapped to the negative real axis. Since the negative real axis itself is usually not distinguished by the dynamics, this partition is useful only in a far right half plane where the partition boundary is remote from its forward image, so that we have the Markov property that sector boundaries never map into sectors (but only for orbits which stay in this right half plane). The nice feature is that every periodic ray is encoded by a unique periodic symbolic sequence, different rays have different codings, and every periodic coding sequence is actually realized.

In some sense, this partition is too good to be useful: all rays are encoded differently. When considering landing properties of dynamic rays, it does happen that different rays land at a common point, and we would like to have a partition which describes this in such a way that rays have the same symbolic sequence if and only if they land together. We will propose such a partition which also labels periodic and preperiodic points, and the labels of these points are the same as the labels of the rays landing at them.

To explain the ideas, it may be helpful to compare this situation to the case of a monic quadratic polynomial with connected Julia set K . The following subsection briefly recollects the relevant parts of that theory. In particular, the polynomial analogues of the following cases will be discussed.

Definition 4.1 (Various Types of Parameters)

A parameter $\lambda \in \mathbb{C}^$ will be called attracting if the map E_λ has an attracting periodic orbit; it is called parabolic if there is a parabolic orbit. The parameter is called postsingularly preperiodic if the singular orbit is finite (and thus necessarily strictly preperiodic). Finally, it will be called an escape parameter if $E_\lambda^{on}(0) \rightarrow \infty$ as $n \rightarrow \infty$.*

In the attracting or parabolic case, the singular orbit converges to a unique non-repelling (attracting or parabolic) orbit. If the singular orbit is finite, it ends in a repelling cycle. In the escape case, the singular orbit is associated to a dynamic ray (see below). In all cases, every periodic orbit is repelling, with the obvious and only exception of the unique attracting or parabolic orbit if there is such an orbit.

4.1 Review of symbolic dynamics of polynomials

Let us consider a monic quadratic polynomial with connected Julia set K . The exterior of K is canonically foliated by dynamic rays which are labeled by their external angles in $\mathbb{S}^1 = \mathbb{R}/\mathbb{Z}$, and these angles can be described by the sequence of their binary digits. Using the unique fixed ray at angle 0 and its unique preimage (other than itself) at angle 1/2, we cut $\mathbb{C} \setminus K$ into two sectors which we label by 0 and 1. A dynamic ray is in sector 0 if its external angle is in $(0, 1/2)$, and it is in sector 1 if the external angle is in $(1/2, 1)$. The sequence of labels of the sectors some ray visits under iteration is exactly the sequence of the binary digits of its external angle. A crude approximation to the two dynamic rays at angles 0 and 1/2 could use straight radial lines at

angles 0 and $1/2$ (which are part of the real axis). This is the analog to our static partition and dynamically meaningful only near ∞ .

In any case, no two polynomial dynamic rays have the same external angle or binary sequence. This partition does not encode which rays land together. Therefore, we propose another partition. For simplicity, we suppose that the critical value is strictly preperiodic, so it is the landing point of a dynamic ray at some preperiodic angle ϑ . The two inverse image rays land together at the critical point; their angles are $\vartheta/2$ and $(\vartheta + 1)/2$. Now we use the partition formed by these two rays and label the sector containing the zero ray by 0 ; the other sector will be labeled 1 . We have to exclude points in the Julia set that will eventually hit the critical point under iteration, as well the rays landing at such points. Every remaining orbit will again have a symbolic sequence which is called the *itinerary* of the orbit. Since pairs of rays landing at a common point can never cross the partition boundary, rays landing at a common point will have the same itinerary, and the landing point also has the same itinerary. In fact, it is well known and not hard to see that different points have different itineraries and that a ray lands at some point in the Julia set if and only if ray and point have identical itineraries. Therefore, this partition reflects more of the dynamic properties. To stress the difference, we will use the font $0, 1, u, \dots$ for itineraries and the usual font $0, 1, s$ for external addresses.

There are some problems with itineraries, however: if there are several rays landing at the singular value, we obtain several inequivalent partitions, but all have the same properties. In many cases, however, there is no dynamic ray at the critical value at all, and it is not clear how to construct such a partition. In the hyperbolic or parabolic cases, where the critical value is in the interior of the filled-in Julia set, one can take dynamic rays landing on the boundary of the Fatou component, together with a curve within this Fatou component (and this choice is again far from unique). One more case where this construction works is when the critical orbit escapes: the critical value is then on some dynamic ray, and the two inverse images of this ray will hit the critical point, forming a useful partition. This case is technically the simplest: the Julia set is a Cantor set with uniform expansion and every bounded orbit has a well-defined itinerary. We will discuss analogues of all these cases for our exponential maps.

4.2 The Escape Case

Suppose that the singular orbit escapes to $+\infty$. It is shown in [SZ2, Theorem 6.5] that then the singular value 0 is either on a dynamic ray, or the escaping endpoint of a dynamic ray, as follows: there is a unique external address $\underline{s} \in \mathbb{Z}^{\mathbb{N}}$, a unique $\tau \geq 0$, and a unique curve $g_{\underline{s}}: [\tau, \infty[\rightarrow \mathbb{C}$ consisting of escaping points with $g_{\underline{s}}(\tau) = 0$ and $\lim_{t \rightarrow \infty} g_{\underline{s}}(t) = \infty$. Unlike our use of dynamic rays elsewhere in this paper, the external address \underline{s} need not be bounded (but it must satisfy certain growth conditions). Set $R_1 := g_{\underline{s}}([\tau, \infty[)$. It is shown in [SZ2, Theorem 6.5] that the escape of 0 and its ray is uniform in the following sense: for every real part $\xi > 0$, there is an $N_{\xi} \geq 0$ such that $\operatorname{Re}(E_{\lambda}^{\circ n}(z)) > \xi$ for every $z \in R_1$ and every $n \geq N_{\xi}$.

Each inverse image of the ray R_1 under E_{λ} is a curve starting at positive infinite real parts and bounded imaginary parts and stretches to negative infinite real parts. Together, all inverse images $E_{\lambda}^{-1}(R_1)$ form a partition of the complex plane. If the singular value 0 is not on the partition boundary, then it defines a unique sector S_0 . All the other sectors will be labeled S_j with $j \in \mathbb{Z}$ so that adjacent sectors are labeled by adjacent integers, and labels increase by 1 when the imaginary part of the sector is increased by 2π . (If the singular value is on the partition boundary, then it must be on a fixed ray, and we choose one of the two adjacent sectors to have label 0; it can be shown that this happens only for $\lambda \in \mathbb{R}$.) Any point $z \in \mathbb{C}$ and any dynamic ray will then have a well-defined itinerary (provided point or ray never land on the partition boundary); this includes all points with bounded orbits and in particular periodic and preperiodic points. A periodic ray can land at a periodic point, or two periodic rays can land together, only if their itineraries coincide. The converse is true as well:

Proposition 4.2 (Itineraries of Rays and Landing Points)

If the singular value escapes to $+\infty$, then no two periodic or preperiodic points have identical itineraries, and a periodic or preperiodic dynamic ray lands at a given periodic or preperiodic point if and only if ray and point have identical itineraries.

REMARK. Rays which have no itineraries cannot land because they will hit the singular value on their forward orbits.

PROOF. First we consider periodic orbits. Let (z_k) and (w_k) be two periodic orbits such that the itineraries of z_1 and w_1 coincide. We do not assume

that the periods of the two orbits are equal. Let n be an integer such that $z_1 = z_{n+1}$ and $w_1 = w_{n+1}$.

Let W be the complex plane from which the closure of the ray R_1 and all its (finitely or infinitely many) forward images are removed. By uniform escape of the points on R_1 , only finitely many forward images of R_1 may intersect any compact subset of \mathbb{C} , so the set W is open and connected. For an index $j \in \mathbb{Z}$, let $W_j := W \cap S_j$ be the connected and open subset within sector j . It carries a unique normalized hyperbolic metric. For every j , there is a branch of E_λ^{-1} mapping W into W_j . Restricting this branch to any $W_{j'}$, the resulting map $E_\lambda^{-1}: W_{j'} \rightarrow W_j$ is a contraction with respect to the hyperbolic metrics in $W_{j'}$ and W_j .

There is a common branch of the pull-back which maps z_k to z_{k-1} and w_k to w_{k-1} , and it shrinks hyperbolic distances. Repeating this n times, both periodic points are restored, but their hyperbolic distance has decreased. This is a contradiction unless $w_1 = z_1$. We also see that the period of the orbit equals the period of the itinerary: obviously, the period of the itinerary must divide the period of the orbit, and if it strictly divides it, then there are at least two periodic points with the same itinerary.

Therefore, any two different periodic orbits have different itineraries, and the period of the orbit equals the period of its itinerary. Any periodic ray lands at a periodic point by Theorem 3.2 (unless it hits the singular value on its forward orbit, but such rays have no itineraries), and the itineraries of the ray and its landing point must obviously be equal. It follows that any periodic point is the landing point of any ray that has the same itinerary.

Similar statements for preperiodic points are now immediate. □

We will show in Section 5 that (almost) every periodic respectively preperiodic point is the landing point of at least one periodic respectively preperiodic dynamic ray. The only thing we need for this proof is a combinatorial lemma to the effect that for every periodic itinerary (as given by the prospective landing point), there is a dynamic ray with that itinerary, and it suffices to describe the external address of this ray. We will provide this combinatorial result in Lemma 5.2, simultaneously for all cases.

4.3 The Postsingularly Finite Case

Suppose that the singular orbit is finite and thus preperiodic. This case has been investigated in [DJ] in a special case of external addresses called

“regular itineraries” there and “unreal” by Milnor (see the discussion at the end of Section 5). The following result will be of crucial importance:

Theorem 4.3 (Preperiodic Ray at Singular Value)

For every postsingularly finite exponential map, at least one preperiodic dynamic ray lands at the singular value.

In order to keep the arguments flowing and to maintain the parallel treatment of the different cases, we defer the proof of this theorem to Section 6: see Theorem 6.4.

Let R_1 be a dynamic ray which lands at the singular value. Then we get a similar partition as above, except that the imaginary parts of the partition boundary will usually become unbounded as the real parts approach $-\infty$ because the ray R_1 will usually spiral into its landing point 0. The singular value 0 will never be on the partition boundary and defines a unique sector S_0 , and the other sectors are labeled as above. Every periodic and preperiodic point has a well-defined itinerary, and also every dynamic ray which never iterates onto the ray landing at the singular value. We obtain the same statement as for the bounded escape case, but with a complication in the proof: there will be a periodic orbit on the forward orbit of the partition boundary, and the regions W_j need not be connected.

Proposition 4.4 (Itineraries of Rays and Landing Points)

If the singular orbit is finite, then no two periodic or preperiodic points have identical itineraries, and a periodic or preperiodic dynamic ray lands at a given periodic or preperiodic point if and only if ray and point have identical itineraries.

PROOF. Let (z_k) , (w_k) and n as in the proof for the escape case (Proposition 4.2) and let (u_k) be the common itinerary of (z_k) and (w_k) . Let W be the complex plane with the closure of a ray landing at the singular value and all its finitely many forward images removed. The set W is still open but it may fail to be connected if several different periodic forward images of the ray landing at the singular value land at the same point. For an index $j \in \mathbb{Z}$, let $W_j := W \cap S_j$ be the open subset within sector j . Every connected component of every W_j has a unique normalized hyperbolic metric. For every j , there is a branch of E_λ^{-1} mapping W into W_j . Restricting this branch to any connected component of any $W_{j'}$, it must contract the hyperbolic metrics.

If there is an index k such that z_k and w_k are in the same connected component of W_{u_k} , then the proof for the escape case goes through. However, if the orbits (w_k) and (z_k) are never in the same connected component of W_{u_k} , then their pull-backs must be synchronized with the pull-backs of the periodic orbit the singular orbit lands in: if z_k and w_k are in different connected components of W_{u_k} , then they are separated by a pair of rays landing at the same postsingular point, and both rays and their landing point are in the same strip S_{u_k} as z_k and w_k . The inverse image points z_{k-1} and w_{k-1} are in the same strip $S_{u_{k-1}}$ because their itineraries are equal; the inverse image of the ray pair in W_{u_k} can separate z_{k-1} and w_{k-1} only if it is in the same strip as well. Separating ray pairs may get fewer (if their inverse images are in different sectors), but there can never be new separations. This shows that z_k and w_k can be forever in different connected components of their W_{u_k} only if their itinerary is the same as that of a periodic postsingular point.

All we need to prove is the following: if some periodic point z_k has the same itinerary as a periodic postsingular point, then it is equal to this postsingular point. To prove this, we can connect z_k to a linearizable neighborhood of the periodic postsingular point in the same sector by a curve with finite hyperbolic length, and subsequent pull-backs will shrink this neighborhood to a point, while the hyperbolic length of the curve will remain at most the same, so its Euclidean length must shrink to zero because it is near the boundary of the domain.

The remaining steps are the same as in the escape case. □

4.4 The Attracting and Parabolic Cases

We will now discuss the case that there is an attracting or parabolic periodic orbit of some period n . Such an orbit will attract a neighborhood of the singular value; let $U_1, U_2, \dots, U_n = U_0$ be the periodic Fatou components such that $0 \in U_1$; then U_0 contains a left half plane. Let $a_1, a_2, \dots, a_n = a_0$ be the points on the attracting orbit such that $a_k \in U_k$.

Unlike in the previous two cases, we cannot construct a partition using dynamic rays landing at or crashing into the singular value. We will use closed subsets of periodic Fatou components in their place. Since we want all periodic points (except those on the attracting or parabolic orbit) to be in the complement, we cannot use closures of entire Fatou components. The construction of a partition will be specified in the proof.

Proposition 4.5 (Itineraries of Rays and Landing Points)

For attracting or parabolic parameters, no two non-attracting periodic or preperiodic points have identical itineraries, and a periodic or preperiodic dynamic ray lands at a given non-attracting periodic or preperiodic point if and only if ray and point have identical itineraries.

PROOF. First we consider the case of an attracting orbit and restrict ourselves to the case $n \geq 2$, leaving the easy modifications for $n = 1$ to the reader. Let V_{n+1} be a closed neighborhood of the point a_1 that corresponds to a disk in linearizing coordinates and that contains the singular value. For $k = 0, 1, \dots, n$, let

$$V_k := \{ z \in U_k : E_\lambda^{\circ(n+1-k)}(z) \in V_{n+1} \} \quad (2)$$

and $V := V_0 \cup \dots \cup V_{n+1}$, $W := \mathbb{C} \setminus V$ (compare Figure 1). Then $V_1 \supset V_{n+1}$, hence V is closed and forward invariant, so W is open and backward invariant: $E_\lambda^{-1}(W) \subset W$, and this is a proper inclusion. Since V_{n+1} is a neighborhood of the singular value, its pull-back V_n contains a left half plane; the next $n - 1$ pull-backs are all univalent and connect a_{n-1}, \dots, a_1 to $+\infty$ within their Fatou components, and the last of these pull-backs yields $V_1 \supset V_{n+1}$, which is unbounded and contains the singular value 0. The final pull-back yields V_0 , which is the only connected component of V disconnecting \mathbb{C} : in fact, $\mathbb{C} \setminus V_0$ consists of countably many connected components which are all translates of each other by integer multiples of $2\pi i$. Denote them by S_j such that $0 \in S_0$ and S_{j+1} is the $2\pi i$ -translate of S_j for every j (if $n = 1$, then $0 \in V_0$ and an arbitrary connected component of $\mathbb{C} \setminus V_0$ is labeled S_0).

Set $W_j := S_j \cap W$ for all j . Every W_j is open and connected and carries a unique normalized hyperbolic metric, and the same proof as above will go through once more.

If there is a parabolic orbit, rather than an attracting one, we have to define the set V_{n+1} somewhat differently: it should be a connected closed subset of the union of the Fatou set with the parabolic orbit, it should contain the singular value, and it should be forward invariant: $E_\lambda^{\circ n}(V_{n+1}) \subset V_{n+1}$. Such sets are easily constructed using Fatou coordinates near the parabolic orbit (compare [M1, § 8]): for example, in coordinates in which the parabolic dynamics corresponds to translation by $+1$, we can take a forward invariant horizontal strip that extends infinitely to the right and that contains the singular orbit. The set V is then defined similarly as in (2), except that z may be on the closure of a periodic Fatou component.




Figure 1: The construction of the partition in the attracting case, for period $n = 3$. Indicated are the circular region V_{n+1} , its immediate pull-back in dark gray, and the next n pull-back steps within periodic Fatou components in a lighter shade of gray. Together, they constitute the set V . Some Fatou components are also shown in light gray. The attracting orbit a_k is indicated, as well as the singular value.

With this modification, the given proof for the attractive case applies to all the repelling periodic and preperiodic points provided the quotient $W/2\pi i\mathbb{Z}$ is connected. It may happen, however, that the parabolic orbit disconnects this quotient (this happens if the multiplier of the parabolic orbit is a root of unity other than $+1$). As in the postsingularly finite case, one can show that two different periodic points with identical itineraries must be in different connected components of $W/2\pi i\mathbb{Z}$ during an entire period of the pull-back, and this is possible only if their pull-back is synchronized with that of the parabolic periodic orbit; but both periodic points must then be on the parabolic orbit: instead of linearizing neighborhoods in the repelling case, we use Fatou coordinates in repelling petals of the parabolic orbit. The details are routine. \square

5 (Pre-)Periodic Points are Landing Points

Now we want to show that every repelling periodic respectively preperiodic point is the landing point of at least one periodic respectively preperiodic dynamic ray. We prove this in the postsingularly finite, the attracting and parabolic, as well as in the escaping cases (in the latter case, there is a well-understood exception). The statement itself is true in much greater generality, as can be shown by a perturbation argument together with some knowledge about parameter space: if some repelling periodic point is the landing point of some periodic dynamic ray, then this will still be true for sufficiently small perturbations of the parameter (compare [GM] or [Sch1] for the proof of the analogous statement for quadratic polynomials).

The corresponding statement for polynomials is due to Douady and Yoccoz; see [H]. Their proof does not generalize to transcendental entire maps because it uses finiteness of the degree and the existence of the attracting basin of ∞ in an essential way.

Theorem 5.1 (The Douady-Yoccoz Landing Theorem)

If the Julia set of a polynomial is connected, then every repelling periodic point is the landing point of at least one and at most finitely many periodic dynamic rays, and analogously for preperiodic points. \square

If a polynomial Julia set is disconnected, then not every repelling periodic point is the landing point of a periodic dynamic ray; it may be the

landing point of no ray at all or of infinitely many non-periodic rays [GM, Appendix C]; however, the number of affected orbits is bounded by the number of critical orbits [LP]. Therefore, one would expect an analogous statement for exponential maps to have some exceptions at least in the case where the singular orbit escapes; see below.

We need a combinatorial lemma about symbolic dynamics on the symbol space $\mathcal{S} := \{s_1 s_2 s_3 \dots : \text{all } s_i \in \mathbb{Z}\}$ with the lexicographic order and the order topology. For $\underline{s}, \underline{t} \in \mathcal{S}$, let $(\underline{s}, \underline{t})$ be the open interval of elements of \mathcal{S} between \underline{s} and \underline{t} in this lexicographic order. The shift operator $\sigma: \mathcal{S} \rightarrow \mathcal{S}$ acts on this space in the usual way. It is a covering map, but it does not preserve the order. For any sequence $\underline{t} \in \mathcal{S}$ and $t_1 \in \mathbb{Z}$, terms like $t_1 \underline{t}$ or $(t_1 + 1) \underline{t}$ will denote the sequence starting with t_1 or $t_1 + 1$ and continuing with \underline{t} (concatenation of the first symbol with the remaining sequence).

We have to define itineraries of sequences $\underline{s} \in \mathcal{S}$ with respect to fixed symbolic sequences. Let \underline{t} be any sequence over \mathbb{Z} starting with $t_1 \in \mathbb{Z}$, and suppose that \underline{t} is different from the constant sequence. Then exactly one of the two intervals $(t_1 \underline{t}, (t_1 + 1) \underline{t})$ and $((t_1 - 1) \underline{t}, t_1 \underline{t})$ contains the sequence \underline{t} ; denote this interval I_0 (this interval is the first or second of the given two examples iff the first entry in \underline{t} different from t_1 is greater or smaller than t_1 , respectively). For $\mathbf{u} \in \mathbb{Z}$, let $I_{\mathbf{u}}$ be this same interval, except that in every sequence the first entry is increased by \mathbf{u} (then $\cup_{\mathbf{u} \in \mathbb{Z}} I_{\mathbf{u}}$ is a partition of \mathcal{S}). For a sequence $\underline{s} \in \mathcal{S}$ define its *itinerary with respect to \underline{t}* as the sequence $\underline{\mathbf{u}} = \mathbf{u}_1 \mathbf{u}_2 \mathbf{u}_3 \dots$ such that $\sigma^k(\underline{s}) \in I_{\mathbf{u}_{k+1}}$ for $k \geq 0$. This works unless \underline{s} eventually maps onto the boundary of the partition, which happens iff $\underline{t} = \sigma^k(\underline{s})$ for some $k \geq 1$. In this case, the corresponding entry in the itinerary will not be an integer, but a “boundary symbol” $\mathbf{u}_{\mathbf{u}}^{+1}$ indicating that the corresponding forward image of \underline{s} is on the boundary between the intervals $I_{\mathbf{u}}$ and $I_{\mathbf{u}+1}$.

If \underline{t} is a constant sequence, we can use one of the two intervals $(t_1 \underline{t}, (t_1 + 1) \underline{t})$ or $((t_1 - 1) \underline{t}, t_1 \underline{t})$ as I_0 and proceed as above; there is no preferred choice. We will continue to use the font $0, 1, \mathbf{u}, \dots$ for itineraries and the usual font $0, 1, s$ for external addresses in order to stress the difference.

The itinerary of any sequence \underline{t} with respect to itself will be called the *kneading sequence* of \underline{t} . By construction, the first symbol in any kneading sequence is 0 (except for sequences of period 1). If \underline{t} is periodic, then its orbit runs through the boundary of the partition, so the kneading sequence contains a boundary symbol (and only then). We will say that a kneading sequence $\underline{\mathbf{k}}$ with a boundary symbol $\mathbf{u}_{\mathbf{u}}^{+1}$ is *adjacent* to an itinerary $\underline{\mathbf{u}}$ if all

non-boundary entries in \underline{k} and \underline{u} coincide at corresponding positions, and if every boundary symbol \underline{u}^{+1} in \underline{k} is consistently replaced either by \underline{u} or by $\underline{u} + 1$ in \underline{u} .

The meaning of this construction is as follows: \underline{t} will be the external address of some dynamic ray landing at the singular value or hitting it. The partition boundaries above will then be dynamic rays (possibly truncated at the singular value) with external addresses $j\underline{t}$ (concatenation!) for integers j , and the itinerary of any dynamic ray having external address \underline{s} with respect to this partition is combinatorially determined as the itinerary of the sequence \underline{s} with respect to the sequence \underline{t} .

Lemma 5.2 (Combinatorics of Itineraries)

Given a periodic itinerary $\underline{u} \in \mathcal{S}$ and an arbitrary external address $\underline{t} \in \mathcal{S}$, there is a periodic external address $\underline{s} \in \mathcal{S}$ such that the itinerary of \underline{s} with respect to \underline{t} is \underline{u} , unless the kneading sequence of \underline{t} is periodic and is equal or adjacent to $\sigma^m(\underline{u})$ for some $m \geq 0$. The number of such external addresses \underline{s} is always finite, and it is 1 if \underline{u} contains no entry 0.

PROOF. Let n be the period of \underline{u} , write $\underline{u} = \overline{u_1 u_2 \dots u_n}$ and label indices modulo n . Moreover, write $\underline{t} = t_1 t_2 t_3 \dots$ and let $\underline{k} = k_1 k_2 k_3 \dots$ be the kneading sequence of \underline{t} . There is a least number $N \in \mathbb{N}$ with the following property:

$$\text{if for some } m \in \mathbb{N}, k_1 k_2 \dots k_m = u_{2-m} u_{3-m} \dots u_1, \text{ then } N \geq m \quad (3)$$

(clearly, if there are infinitely many m with this property, then \underline{k} is periodic and equal to a finite shift of \underline{u} , violating the hypothesis of the theorem; therefore, there is a finite N as claimed). Set $T := \{\underline{t}, \sigma(\underline{t}), \sigma^2(\underline{t}), \dots, \sigma^N(\underline{t})\}$ (if $N = 0$, then T is the empty set; this happens in particular if all $u_m \neq 0 = k_1$).

Since T is finite, the set $I_{\underline{u}_{n+1}} \setminus T$ consists of a finite number of open intervals, say $J_{n+1,1}, \dots, J_{n+1,M}$. Then there are M disjoint intervals $J_{n,1}, \dots, J_{n,M} \subset I_{\underline{u}_n}$ such that for every i , the shift $\sigma: J_{n,i} \rightarrow J_{n+1,i}$ is an order preserving homeomorphism: for this, all we need is that $\underline{t} \notin (I_{\underline{u}_{n+1}} \setminus T)$ because every interval $J \subset I_{\underline{u}_{n+1}}$ with $\underline{t} \notin J$ has a unique pull-back within $I_{\underline{u}_n}$; our condition (3) (with $m = 1$) assures that there is no problem in this step.

In a similar fashion, we construct intervals $J_{n-1,i}, \dots, J_{1,i}$, for $i = 1, \dots, M$ with order preserving homeomorphisms $\sigma^n: J_{1,i} \rightarrow J_{n+1,i}$ for every

i . This works unless there is an interval $J_{n+1,i}$ so that after $m-1$ pull-backs (for $m \in \{1, 2, \dots, n\}$), the interval $J_{n+1-(m-1),i} = J_{n+2-m}$ contains \underline{t} ; but again this is excluded by (3).

Note that $J_{1,i} \subset I_{u_1} = I_{u_{n+1}} \supset J_{n+1,i}$. The easiest case is if $J_{1,i} \subset J_{n+1,\rho(i)}$ for some $\rho(i) \in \{1, 2, \dots, M\}$. This does not always hold, but since the general case needs more notation, we discuss the special case first and indicate the necessary changes for the general case at the end.

This way, we have defined a map $\rho: \{1, 2, \dots, M\} \rightarrow \{1, 2, \dots, M\}$ such that $J_{1,i} \subset J_{n+1,\rho(i)}$ for all i . By finiteness of M , there is a number $i_0 \in \{1, 2, \dots, M\}$ and $p \leq M$ such that $\rho^{op}(i_0) = i_0$.

There are intervals $J'_0, \dots, J'_p \subset I_{u_1}$ such that $J'_0 = J_{n+1,i_0}$ and for $k = 1, 2, \dots, p$, we have $J'_k \subset J_{n+1,\rho^{o(k)}(i_0)}$ such that $\sigma^n: J'_k \rightarrow J'_{k-1}$ is a homeomorphism. Then $J'_p \subset J'_0$, and $\sigma^{pn}: J'_p \rightarrow J'_0$ is a homeomorphism.

All external addresses $\underline{s}' \in J'_p$ have identical first pn entries, say $s_1 s_2 \dots s_{pn}$. Similarly, there is a subinterval $J'_{2p} \subset J'_p$ such that $\sigma^{pn}: J'_{2p} \rightarrow J'_p$ is a homeomorphism, and all $\underline{s}' \in J'_{2p}$ have identical first $2pn$ entries $s_1 s_2 \dots s_{pn} s_1 s_2 \dots s_{pn}$, etc.. We set $\underline{s} := \overline{s_1 s_2 \dots s_{pn}}$; then clearly either (1) $\underline{s} \in (J'_0 \cap J'_p \cap J'_{2p} \cap \dots)$, or (2) \underline{s} is a common boundary point of all these intervals. The first pn entries in the itineraries of every $\underline{s}' \in J'_0$ are all equal to $u_1 u_2 \dots u_{pn}$; also, \underline{s} is periodic of period pn . Therefore, in case (1), \underline{s} has itinerary \underline{u} as claimed and we are done. In case (2), if the periodic orbit of \underline{s} does not visit \underline{t} and hence one of $\sigma^{-1}(\underline{t})$, then \underline{s} has itinerary \underline{u} as well; however, if $\underline{t} = \sigma^{om}(\underline{s})$ for some $m \geq 1$, then \underline{t} is periodic and its kneading sequence is adjacent to \underline{u} .

So far, this proof has been based on the assumption that all $J_{1,i}$ were entirely contained in some $J_{n+1,\rho(i)}$. Recall that $J_{1,i}$ is a pull-back of $J_{n+1,i}$ for n steps along branches prescribed by the itinerary $u_n u_{n-1} \dots u_1$. We claim that if $J_{1,i}$ is not contained in some $J_{n+1,\rho(i)}$, then one can pull back another n times, and after finitely many repetitions one obtains a subinterval of some $J_{n+1,\rho(i)}$. Indeed, $J_{n+1,i}$ can be pulled back along the branches prescribed by the itinerary \underline{u} until the external address \underline{t} is contained in an interval to be pulled back; if this happens after m pull-backs, then the kneading sequence \underline{k} must satisfy $\underline{k}_1 \underline{k}_2 \dots \underline{k}_m \underline{k}_{m+1} = u_{n-m+1} u_{n-m+2} \dots u_{n+1}$, and by Condition (3) this means that $N \geq m+1$, so $\sigma^{om}(\underline{t}) \in J_{n+1,i}$ in contradiction to the construction of T and $J_{n+1,i}$.

We thus see that $J_{n+1,i}$ can be pulled back arbitrarily many times for n steps; after at most N pull-backs, it will no longer contain an element of T , so eventually a pull-back of every $J_{n+1,i}$ will be contained in $J_{n+1,\rho(i)}$ for some $\rho(i)$. The argument above goes through as before, except that the number

of pull-back steps and hence the period of \underline{s} might be greater.

To see that there are only finitely many sequences \underline{s} , note that $\bigcup_i^M J_{1,i}$ is the set of sequences $\underline{s}' \in \mathcal{S}$ for which the itinerary starts with $u_1 u_2 \dots u_{n+1}$. Each $J_{1,i}$ can contain no more than one sequence \underline{s} with itinerary \mathbf{u} (any two of them would yield a whole interval of such sequences). Finally, if $k_j \neq 0$ for all j , then $M = 1$ in the construction. \square

This lemma can also be described in terms of a transition matrix between the intervals $J_{n+1,i}$. This has been elaborated in [DJ] for the special case of exponential maps which are postsingularly finite and for which the external address of the singular orbit is “regular” (see below).

Now we have all ingredients together to prove in many cases that every repelling or parabolic periodic point is the landing point of at least one dynamic ray.

Theorem 5.3 (Periodic Points are Landing Points, Cases 1&2)

For every postsingularly finite parameter, every repelling or parabolic periodic or preperiodic point is the landing point of at least one periodic resp. preperiodic dynamic ray.

For every parameter for which the singular orbit escapes on the end of a dynamic ray $g_{\underline{t}}$, the same is true unless the itinerary \underline{u} of the (pre-)periodic point is such that $\sigma^m(\underline{u})$ (for $m \geq 1$) is equal or adjacent to the kneading sequence of \underline{t} .

In any case, only finitely many periodic or preperiodic rays land at the same point.

PROOF. Similarly as for polynomials, all periodic rays landing at the same repelling or parabolic periodic point must have equal periods (which may be a multiple of the period of the periodic orbit): this is because the first return map near any repelling periodic point is a local homeomorphism. However, unlike for polynomials, the number of periodic rays of any fixed period is infinite, so finiteness of the number of periodic dynamic rays landing at a repelling or parabolic periodic point z requires an extra argument.

In all the cases covered by the theorem, a repelling periodic point z has a well-defined periodic itinerary \underline{u} with respect to a partition bounded by dynamic rays, so all rays landing at the same periodic point have identical periodic itineraries. But the periodic itinerary \underline{u} and the partition defining it give a bound on all entries of the external address \underline{s} of any dynamic ray

$g_{\underline{s}}$ landing at z . There are only finitely many external addresses with given period and given bound, hence the number of periodic dynamic rays landing at the repelling periodic point z is finite.

First we consider the postsingularly finite case. By Theorem 4.3, there is a preperiodic external address \underline{t} such that the dynamic ray $g_{\underline{t}}$ lands at the singular value; the kneading sequence of \underline{t} is strictly preperiodic. The partition constructed in Proposition 4.4 consists of the E_λ -preimages of $g_{\underline{t}}$ and defines itineraries of periodic points and of rays. Consider any periodic point z (necessarily repelling) and let \underline{u} be its itinerary. By Lemma 5.2, there is a finite non-empty collection of periodic sequences in \mathcal{S} which all generate the itinerary \underline{u} with respect to \underline{t} . (Note that the exception of the Lemma is inconsequential in this case.) For every such sequence, there is a unique dynamic ray with this external address, and each of these rays has itinerary \underline{u} . According to Proposition 4.4, these rays must land at z .

If w is preperiodic and $z := E_\lambda^l(w)$ is periodic, then we find a periodic dynamic ray $g_{\underline{s}}$ landing at z and pull it back along the orbit of w , yielding a preperiodic ray landing at w .

Now we consider the case where the singular value escapes on a periodic or preperiodic ray. Let \underline{t} be the external address so that the ray $g_{\underline{t}}$ contains the singular value; note that we make no restrictions on \underline{t} such as periodicity or boundedness. The partition constructed in Proposition 4.2 consists of preimages of $g_{\underline{t}}$ and defines itineraries once more. We argue as before: for a (necessarily repelling) periodic point z with itinerary \underline{u} , we need a periodic external address \underline{s} with itinerary \underline{u} with respect to \underline{t} , so that $g_{\underline{s}}$ lands at z ; Lemma 5.2 shows the existence of \underline{s} unless the kneading sequence of \underline{t} is equal or adjacent to a finite shift of \underline{u} .

If w is preperiodic and $z := E_\lambda^{ol}(w)$ is periodic, we pull back as above. However, we have to deal with the exceptions from Lemma 5.2; let \underline{k} be the kneading sequence of \underline{t} . The periodic point z may fail to be the landing point of a periodic dynamic ray if its itinerary (which is $\sigma^m(\underline{u})$ for some m) is equal or adjacent to \underline{k} ; even if z is the landing point of a periodic dynamic rays $g_{\underline{s}}$, then w may fail to be the landing point of a preperiodic dynamic ray if the pull-back of dynamic rays involves a preperiodic dynamic ray containing the singular value: but this implies $\underline{k} = \sigma^m(\underline{u})$ for some $m \geq 1$. \square

REMARK. The exceptions stated in this theorem for the escaping case correspond to cases for polynomials where periodic points in disconnected Julia sets fail indeed to be landing points of periodic dynamic rays. Therefore, we

believe that (at least for periodic points) the condition is sharp for arbitrary exponential maps.

In the case of attracting or parabolic exponential dynamics, the singular orbit is not associated to a dynamic ray with external address $\underline{s} \in \mathbb{Z}^{\mathbb{N}}$, but it is within the periodic Fatou components. There are two possibilities: as in [Sch2] or [SZ1], one can construct an “attracting dynamic ray” with an “intermediate external address” (which is a finite string of integers followed by a half-integer) and use this to construct appropriate symbolic dynamics as in Lemma 5.2. We will use a different possibility instead which is related to “spiders” (Section 6).

Theorem 5.4 (Periodic Points are Landing Points, Cases 3&4)

For every attracting or parabolic parameter, every repelling or parabolic periodic respectively preperiodic point is the landing point of at least one periodic respectively preperiodic dynamic ray; the number of (pre-)periodic rays landing at such a point is at most the period of attracting orbit, or the period of the Fatou components containing the parabolic orbit.

PROOF. As in Proposition 4.5 and Figure 1, let $U_1, U_2, \dots, U_n = U_0$ be the periodic cycle of Fatou components such that U_1 contains the singular value 0 and U_0 contains a left half plane, and let a_k be the attracting periodic point within U_k . Let V_{n+1} be a closed round disk around a_1 with respect to linearizing coordinates and, for $k = 0, 1, 2, \dots, n$, let $V_k := E_\lambda^{-(n+1-k)}(V_{n+1}) \cap U_k$; then $V_1 \supset V_{n+1}$ and $V_0 \supset V_n$. Let $V := \cup_{k=0}^{n+1} V_k$, $W := \mathbb{C} \setminus V$ and $W' := W \cup \{\infty\}$.

Let z be a repelling periodic point and let m be its period. The main steps of the proof are: (1) there is a finite positive number of homotopy classes of curves connecting z to ∞ within W' ; (2) at least one of these homotopy classes is periodic with respect to pull-backs, and is associated to a periodic external address \underline{s} ; (3) the dynamic ray $g_{\underline{s}}$ lands at z .

(1) Since every V_k (for $k = 0, 1, \dots, n - 1$) is unbounded, the number of homotopy classes of injective curves in W' connecting z to ∞ is at most n . It is easy to see that there is at least one homotopy class because all V_k have smooth boundaries and disjoint closures.

(2) Given a curve $\gamma: [0, \infty] \rightarrow W'$ with $\gamma(0) = z$ and $\gamma(\infty) = \infty$, there is a unique pull-back curve $\tilde{\gamma}: [0, \infty] \rightarrow W'$ with $E_\lambda(\tilde{\gamma}(t)) = \gamma(t)$ for all t so that $\tilde{\gamma}(0)$ is the unique periodic point in $E_\lambda^{-1}(z)$ (this uses the fact that

$E_\lambda^{-1}(V) \supset V$). We will consider homotopies of curves within W' , fixing the endpoints on the orbit of z and the condition $\gamma(\infty) = \infty$.

Any curve in W' homotopic to γ (with endpoints fixed) pulls back to a curve homotopic to $\tilde{\gamma}$, so the m -th iterate of the pull-back descends to a well-defined map on homotopy classes of injective curves in W' from z to ∞ . Since there are at most n such homotopy classes, there is a number $p \leq n$, a curve γ_{mp+1} with $\gamma_{mp+1}(0) = z$, a unique sequence of pull-backs γ_k (for $k = mp, mp-1, \dots, 1$) with $E_\lambda(\gamma_k(t)) = \gamma_{k+1}(t)$ for all t and $\gamma_1(0) = z$, such that γ_1 is homotopic to γ_{mp+1} .

Fix $\xi < 0$ such that $R :=]-\infty, \xi] \subset U_0$. Then $E_\lambda^{-1}(R)$ is a countable collection of horizontal lines which bound the static partition from Section 2 (at least for large real parts) with respect to which dynamic rays were defined. Since all γ_k are disjoint from R , they are all disjoint from $E_\lambda^{-1}(R)$, so each γ_k tends to ∞ within a well-defined strip R_{s_k} in the static partition (more precisely, for every γ_k there is a $\tau_k \geq 0$ such that $\gamma_k(t) \in R_{s_k}$ for $t \geq \tau_k$). The numbers s_k are respected by homotopies of the γ_k . Define a periodic external address $\underline{s} := \overline{s_1 s_2 \dots s_{mp}}$.

(3) We claim that the periodic dynamic ray $g_{\underline{s}}$ lands at z . In view of Proposition 4.5, it suffices to show that $g_{\underline{s}}$ has the same itinerary as z with respect to the partition $\mathbb{C} \setminus V_0$; both have period mp . The arguments are routine, so we only sketch the main idea. If both itineraries are not equal, then let $k \leq mp$ be the position of the first difference in the itineraries of $g_{\underline{s}}$ and γ_1 , so that $E_\lambda^{\circ(k-1)}(g_{\underline{s}})$ and $E_\lambda^{\circ(k-1)}(\gamma_1)$ are in different connected components of $\mathbb{C} \setminus V_0$. Since the external addresses of $E_\lambda^{\circ(k-1)}(g_{\underline{s}})$ and $E_\lambda^{\circ(k-1)}(\gamma_1)$ are identical, one can map forward under E_λ once and see that V_1 stretches to $+\infty$ “between” $E_\lambda^{\circ k}(g_{\underline{s}})$ and $E_\lambda^{\circ k}(\gamma_1)$ (in the sense that V_1 disconnects any sufficiently far right half plane into exactly two unbounded parts, one “above” and one “below” V_1 , and each part contains exactly one of $E_\lambda^{\circ k}(g_{\underline{s}})$ and $E_\lambda^{\circ k}(\gamma_1)$). After $n-2$ further forward iterations (again using equality of external addresses), it follows that V_{n-1} is between $E_\lambda^{\circ(k+n-2)}(g_{\underline{s}})$ and $E_\lambda^{\circ(k+n-2)}(\gamma_1)$; but this implies that the external addresses of $E_\lambda^{\circ k}(g_{\underline{s}})$ and $E_\lambda^{\circ k}(\gamma_1)$ are different after all and is a contradiction.

Preperiodic dynamic rays landing at repelling preperiodic points are obtained simply by pulling back.

Finally, the proof for the parabolic case is analogous to the attracting case; one only needs to require that the curves γ_k do not run through parabolic periodic points (except possibly at the endpoint $\gamma_k(0)$). \square

We conclude this section with a brief discussion of external addresses without entries 0. Such sequences and the corresponding rays have been investigated in a number of papers under the name of “regular itineraries”. The following two lemmas are known from [DGH] and are intended to show that the dynamical possibilities of “regular” itineraries are rather restricted. Milnor has suggested to call such itineraries “unreal” because by definition, the corresponding orbits must avoid a neighborhood of \mathbb{R} (and thus of the image of the partition boundary \mathbb{R}^-).

Lemma 5.5 (Dynamic Rays Intersecting Partition Boundary)

If a dynamic ray intersects the boundary of the static partition, then its external address must contain an entry 0.

PROOF. Suppose that a dynamic ray intersects a boundary of the static partition. This implies that the images of ray and sector boundary under E_λ intersect each other as well. Since the image of a sector boundary is the negative real axis which lies in sector zero of the static partition, the external address has to contain the entry zero. \square

Lemma 5.6 (Dynamic Rays Landing at Common Point)

If several dynamic rays with bounded external addresses land at a common point, then the external address of at least one of them contains the entry 0.

PROOF. Suppose that several dynamic rays with bounded external addresses land at a given point. If none of them contains the entry 0, then all these rays lie completely in one strip of the static partition by Lemma 5.5. This means they have the same first entry in their external address. The same argument applies to all forward images as well, so all these rays have the same external addresses. But there is a unique dynamic ray for any given external address by Theorem 2.3, a contradiction. \square

While rays with “regular” external addresses always land alone, the structure of Julia sets and of parameter space is largely determined by groups of rays landing together: for example, Douady’s pinched disk models [Do], Thurston’s laminations [T], Milnor’s orbit portraits [M2] are all based on pairs of rays landing together. Therefore, it is important to describe the landing properties of all rays, “regular” or not.

6 Spiders

In this section, we prove that for every postsingularly finite exponential map, the singular value is the landing point of at least one preperiodic dynamic ray. We will prove this result by a variant of the spider theory from [HS], which is an offspring of Thurston’s classification theory of rational maps [DH2]. Since we have used this theorem in Sections 4 and 5, we can use here only the results of Section 3. We use essential ideas from a common project with Mitsuhiro Shishikura [SS] which contains a systematic investigation of postsingularly finite exponential maps; however, this section is written so as to be self-contained.

The global plan of this section is the following: for a given postsingularly finite exponential map E_λ , we introduce “spiders” and an iteration procedure on the space of spiders (a “spider map”), and we show that there is a “periodic spider”, i.e., a spider that is “equivalent” to its N -th iterate for some $N \in \mathbb{N}$. Every spider consists of curves called “legs”, and we show that in a periodic spider we may replace the legs by dynamic rays: these rays will land at the postsingular orbit.

For this entire section, fix a postsingularly finite exponential map E_λ and let $p_1 = 0$, $p_2 = E_\lambda(0) = \lambda$, $p_3 = E_\lambda(E_\lambda(0))$, \dots be the singular orbit with preperiod l and period k , so that $p_{l+k+1} = p_{l+1}$. Let $P := \{p_1, \dots, p_{l+k}\}$ be the postsingular set.

The main characters of this section will be *spiders*. A *spider leg* is an injective piecewise continuously differentiable curve $\gamma: [0, \infty] \rightarrow \overline{\mathbb{C}}$ with $\gamma(0) \in P$, $\gamma(\infty) = \infty$ and $\gamma(t) \notin P \cup \{\infty\}$ for $t \in (0, \infty)$, subject to the condition that $\operatorname{Re}(\gamma(t)) \rightarrow +\infty$ as $t \rightarrow \infty$ and that $\operatorname{Im}(\gamma(t))$ be bounded. A *spider* is a collection $\{\gamma_i\}_{i=1,2,\dots,l+k}$ of spider legs with $\gamma_i(0) = p_i$ for every i , subject to the condition that all spider legs be disjoint (except for their common endpoint ∞).

Two spiders will be called *equivalent* if their legs are homotopic relative to $P \cup \{\infty\}$, where the homotopy should be through the space of spiders (in particular, homotopies must preserve the conditions on spider legs; it is immaterial if during a homotopy the disjointness of the legs is violated on bounded subsets of \mathbb{C}).

Note: Our definition of spiders is custom tailored for our purposes and a restricted variant of the more general theory from [HS]: for us, the endpoints of the legs are fixed throughout, and the legs themselves are restricted in their approach to ∞ . In general, the space of (equivalence classes of) spiders

is a Teichmüller space; here, the fixed location of the endpoints restricts the spiders to a discrete subset which corresponds to a single point in moduli space. (We may assure the reader that Teichmüller theory will not be used explicitly in this paper.)

Definition 6.1 (The Spider Map)

To a spider $\{\gamma_i\}_{i=1,2,\dots,l+k}$, associate its image spider $\{\tilde{\gamma}_i\}_{i=1,2,\dots,l+k}$, for which the leg $\tilde{\gamma}_i$ from p_i to ∞ is the unique inverse image of γ_{i+1} under E_λ which connects p_i to ∞ (by periodicity, we set $\gamma_{l+k+1} = \gamma_{l+1}$ here).

REMARK. Any leg γ_{i+1} connects p_{i+1} to ∞ ; since $E_\lambda(p_i) = p_{i+1}$, there is indeed a unique inverse image of γ_{i+1} which ends at p_i . The two new spider legs at p_l and p_{l+k} will both be different inverse images of $\gamma_{l+1} = \gamma_{l+k+1}$; on the other hand, the leg γ_1 landing at the singular value $p_1 = 0$ will not be used in the construction of the new spider (in fact, all the legs at preperiodic points will be thrown away eventually, but without them the point in Teichmüller space would not be specified completely).

Note that the image spider is indeed a spider: all new legs $\tilde{\gamma}_i$ are disjoint except for the endpoint ∞ , and they approach ∞ along increasing real parts with imaginary parts tending to a constant (which is in $\text{Im}(-\kappa) + 2\pi\mathbb{Z}$).

One important aspect in which our spider map differs from that in [HS] is that we have fixed the map E_λ and thus the postsingular orbit, and we choose the branch of $E_\lambda^{-1}(\gamma_{i+1})$ so that it ends at p_i ; we have no direct control over the order of the legs as they approach ∞ . In [HS], the branch of $E_\lambda^{-1}(\gamma_{i+1})$ is chosen with respect to the order at ∞ , and accordingly the positions of the endpoints change.

Lemma 6.2 (Spider Map on Equivalence Classes)

Under the spider map, equivalent spiders have equivalent images, so the map descends to a map on the set of equivalence classes.

PROOF. A homotopy between equivalent spiders yields a homotopy between the image spiders. □

Proposition 6.3 (Periodic Spider)

The iteration of the spider map, starting with any spider, will lead to a periodic spider after finitely many steps (that is, to a spider that is equivalent to its image spider after finitely many iterations of the spider map).

PROOF. Fix a real number $\xi > 0$ such that $|z| < \xi/2$ for every $z \in P$ and $|\lambda| \exp(\xi) > \xi$. Let $T := \{z \in \mathbb{C} : \operatorname{Re}(z) = \xi\}$ be a vertical line and $C := E_\lambda(T) = \{z \in \mathbb{C} : |z| = |\lambda| \exp(\xi)\}$ its image circle. Every spider leg must intersect C at least once; indeed, after replacing it with an equivalent spider, we may suppose that every leg intersects C exactly once. Fix numbers $\varepsilon_i > 0$ such that for all $p_i \in P$, the open disk neighborhoods $D_{2\varepsilon_i}(p_i)$ are disjoint and do not intersect C or T , and such that $E_\lambda(D_{\varepsilon_i}(p_i)) \supset D_{\varepsilon_{i+1}}(p_{i+1})$ (again identifying p_{l+k+1} with p_{l+1} etc.); this is possible because the singular orbit lands on a repelling periodic point.

Let $U := \mathbb{C} \setminus P$; this is an open domain which carries a unique normalized hyperbolic metric d_U . Define a compact set

$$K := \{z \in \mathbb{C} : |z| \leq |\lambda| \exp(\xi)\} \setminus \cup_i D_{\varepsilon_i}(p_i) .$$

For every leg γ_i , let ℓ_i be the d_U -length of γ_i restricted to K . Then all ℓ_i are finite positive numbers; we say that a spider $\Gamma := \{\gamma_i\}$ is an L -spider if all $\ell_i \leq L$.

CLAIM. *There is a length $L_0 > 0$ with the following property: if Γ is an L -spider and $L \geq L_0$, then the image spider $\tilde{\Gamma}$ is equivalent to an L -spider.*

We will prove this claim below; first we continue the proof of the proposition. Start with any spider Γ^0 and pick some $L \geq L_0$ such that Γ^0 is an L -spider. Starting with Γ^0 , apply the spider map and replace the image spider by an equivalent L -spider, say Γ^1 . Iterate this procedure so as to obtain a sequence $(\Gamma^n)_{n \in \mathbb{N}}$ of L -spiders.

Clearly, there are only finitely many homotopy classes of legs γ_i ending at p_i with $\ell_i \leq L$, and thus there are only finitely many homotopy classes of L -spiders. Therefore, there must be two indices $n' > n$ such that $\Gamma^{n'}$ is equivalent to Γ^n , and the iteration has reached a periodic spider. This proves the proposition modulo the CLAIM.

PROOF OF THE CLAIM. The intersection $C \cap T$ are two points which bound a closed vertical interval $I \subset T$. There is an $s > 0$ so that every $z \in I$ can be connected to a point $w \in C$ to the right of T (i.e., $\operatorname{Re}(w) > \xi$) by a differentiable curve in U of d_U -length at most s .

Let $V := E_\lambda^{-1}(U)$; as in Section 3, we have a holomorphic covering $E_\lambda: V \rightarrow U$ which is a local hyperbolic isometry, $V \subset U$, and this is a strict inclusion. Therefore, the densities of the hyperbolic metrics in U and V satisfy $\rho_U(z)/\rho_V(z) < 1$ for $z \in U$ (with $\rho_V(z) := +\infty$ for $z \in U \setminus V$). If

$\gamma \subset U$ is a curve with finite d_U -length ℓ and $\tilde{\gamma}$ is any branch of the pull-back $E_\lambda^{-1}(\gamma)$, then $\tilde{\gamma}$ has d_V -length ℓ and d_U -length less than ℓ . Since $K \subset U$ is compact, there is an $\eta < 1$ such that any continuously differentiable curve $\gamma \subset K$ with d_U -length ℓ pulls back to a curve $E_\lambda^{-1}(\gamma)$ of d_U -length at most $\eta\ell$, for any branch of E_λ^{-1} .

Set $K' := E_\lambda^{-1}(K)$; neither of K and K' is a subset of the other, but $K \setminus K'$ is the disk segment bounded by T and C .

Consider any spider leg γ_i with d_U -length ℓ_i within K . Then the restriction of the leg $\tilde{\gamma}_{i-1}$ to K' has d_U -length at most $\eta\ell_i$. Since $K \setminus K'$ is the disk segment as described above and any point $z \in I$ can be connected to C by a curve of length at most s , one can homotope $\tilde{\gamma}_{i-1}$ so that its restriction to K has d_U -length at most $\eta\ell_i + s$. Such homotopies can be applied to all spider legs of the image spider $\{\tilde{\gamma}_i\}$. Therefore, if $\{\gamma_i\}$ was an L -spider, then $\{\tilde{\gamma}_i\}$ is equivalent to an $\eta L + s$ -spider.

Now if $L_0 := s/(1 - \eta)$ and $L \geq L_0$, then $\eta L + s \leq L$. Therefore, the image spider $\{\tilde{\gamma}_i\}$ is equivalent to an L -spider. \square

In every spider, the legs have a vertical order: given a spider leg γ_i and a sufficiently large $\xi > 0$, then γ_i disconnects the half plane $\{z \in \mathbb{C}: \operatorname{Re}(z) > \xi\}$ into exactly two unbounded parts, exactly one of which contains an unbounded part of $\gamma_{i'}$ for any $i' \neq i$; depending on which part this is, we say that $\gamma_{i'}$ is *below* or *above* γ_i . This vertical order is respected by equivalence classes of spiders. Similarly, the dynamic rays g_s have a vertical order which coincides with the lexicographic order of their external addresses s .

Theorem 6.4 (Fixed Spider with Dynamic Rays)

Every postsingularly finite exponential map has a periodic spider for which all the legs consist of dynamic rays landing at the postsingular points.

PROOF. The proof consists of three steps: (1) given a periodic spider, we first associate an external address to every spider leg, then (2) we show that we can replace an unbounded “tail” of every leg by a tail of a dynamic ray, and finally (3) we show that the entire leg can be replaced by a dynamic ray.

(1) Let Γ^0 be a periodic spider with period N and let Γ^{n+1} be the image spider of Γ^n for every $n \geq 0$. Denote the i th leg of Γ^n by γ_i^n , and recall that every leg is parametrized as a curve $\gamma_i^n: [0, \infty] \rightarrow \overline{\mathbb{C}}$ with $\gamma_i^n(t) \rightarrow +\infty$ as $t \rightarrow \infty$.

Recall the static partition from Section 2: we introduced open horizontal strips R_j of width 2π bounded by E_λ -preimages of \mathbb{R}^- with $j \in \mathbb{Z}$. Since

every spider leg γ_{i+1}^n intersects \mathbb{R}^- only for bounded potentials t , every leg γ_i^{n+1} must “eventually” be contained in a single strip R_j (where “eventually” means that there is a $\tau \geq 0$ such that $\gamma_i^{n+1}(t) \in R_j$ for $t > \tau$). This property is not respected by homotopies of $\{\gamma_i^{n+1}\}$, i.e., by spiders equivalent to Γ^{n+1} , but it is respected by homotopies of Γ^n , followed by the spider map. Therefore, we have well-defined integers s_i^n (for $n > 0$) such that γ_i^n is eventually in the strip $R_{s_i^n}$. For every i , the sequence s_i^n is periodic in n with period N .

Define external addresses $\underline{s}(n, i) := s_i^n s_{i+1}^{n-1} s_{i+2}^{n-2} \dots s_{i+k}^{n-k} \dots$ (where the upper indices are evaluated modulo N). These are the external addresses of the dynamic rays from which we will assemble spider legs. They are all (pre-)periodic because the sequence of spiders is periodic with period N (upper index) and the sequence of legs is preperiodic with preperiod l and period k (lower index). More precisely, the external address $\underline{s}(n, i)$ is preperiodic for $i \leq l$ and periodic for $i > l$. Thus, there is an integer $s > 0$ with $|s_{i+k}^{n-k}| \leq s$ for all n, i, k .

Note that for every fixed n , the vertical order of the legs γ_i^n is the same as the lexicographic order of the sequences $\underline{s}(n, i)$: if $s_i^n > s_{i'}^n$, then γ_i^n is in a higher strip as $\gamma_{i'}^n$, and if $s_i^n = s_{i'}^n$, then both legs are in the same strip and E_λ maps them onto γ_{i+1}^{n-1} and $\gamma_{i'+1}^{n-1}$ respecting the vertical order, so the claim follows by induction.

(2) By Theorem 2.3 there are dynamic rays $g_{\underline{s}(n, i)}:]0, \infty[\rightarrow \mathbb{C}$ for every n, i , and there is a $\tau > 0$ so that $g_{\underline{s}(n, i)}(t) \in R_{s_i^n}$ and $\operatorname{Re}(g_{\underline{s}(n, i)}(t)) > \xi$ for every $t \geq \tau$ and every n, i , where $\xi > 0$ is such that $|z| < \xi/2$ for all $z \in P$.

After homotopies of our spiders, we may suppose that $\gamma_i^n(t) \in R_{s_i^n}$ for all $t \geq \tau$ and that $\gamma_i^n(\tau) = g_{\underline{s}(n, i)}(\tau)$, both for all n, i . We may therefore replace the unbounded tail $\gamma_i^n([\tau, \infty[)$ of the leg by the tail $g_{\underline{s}(n, i)}([\tau, \infty[)$ of the dynamic ray: both are entirely within the same strip $R_{s_i^n}$, and the vertical order of the legs within this strip is the same as the vertical order of the corresponding rays (because both coincide with the lexicographic order). Therefore, replacing the tails of legs by tails of the rays yields an equivalent spider.

The catch of this is that our construction is respected by the spider map (without any extra homotopies on the legs): the leg γ_i^n lands at p_i , and the image leg γ_{i-1}^{n+1} is the branch of $E_\lambda^{-1}(\gamma_i^n)$ which lands at p_{i-1} ; by construction of s_{i-1}^{n+1} , this is the same branch for which γ_{i-1}^{n+1} approaches $+\infty$ through the strip $R_{s_{i-1}^{n+1}}$, so this is also the branch that sends $g_{\underline{s}(n, i)}$ to $g_{\underline{s}(n+1, i-1)}$.

Therefore, once the spider Γ^n is endowed with tails of legs made of dynamic rays, this property is preserved automatically by the spider map. However, the tail of the leg consisting of a ray becomes longer: if the rays in the legs $\{\gamma_i^n\}$ start at potential $\tau > 0$, then in $\{\gamma_i^{n+1}\}$, the rays start at potential $\tau' := F^{-1}(\tau) < \tau$ because of the functional equation (1) in Theorem 2.3. Hence after sufficiently many iterations of the spider map, almost the entire legs consist of dynamic rays.

(3) By Theorem 3.2, we know that the periodic rays $g_{\underline{s}(n,i)}$ land at periodic points (those rays with $i > l$). We want to show that the landing points are the points p_i . Construct again linearizable disk neighborhoods $D_{\varepsilon_i}(p_i)$ around the points p_i such that $E_\lambda(D_{\varepsilon_i}(p_i)) \supset D_{\varepsilon_{i+1}}(p_{i+1})$, fix an index n and choose points $x_i \in \gamma_i^n \cap D_{\varepsilon_i}(p_i)$. Let $y_i \in \gamma_i^n$ be the points at which the legs turn into tails of rays. Then the hyperbolic d_U -distances between x_i and y_i are bounded above, and they decrease under pull-backs. Since the x_i converge to the postsingular set P under iteration of the spider map and $P \cap U = \emptyset$, it follows that the y_i must also converge to P . This means that after finitely many iterations, the legs $\gamma_i^{n'}$ outside of the disks $D_{\varepsilon_i}(p_i)$ consist entirely of dynamic rays, so the periodic rays among the $g_{\underline{s}(n,i)}$ land at the periodic points among the p_i . Now it follows that for every n , the rays $\{g_{\underline{s}(n,i)}\}_{i=1,\dots,l+k}$ form a spider equivalent to Γ^n , and in particular that all preperiodic rays $g_{\underline{s}(n,i)}$ with $i = 1$ land at the singular value. \square

This also concludes the proof of Theorem 4.3.

REMARK. The period of the postsingular orbit is k , and the period of the spider is N . Hence the dynamic rays landing at the postsingular periodic orbit have periods up to kN , and up to N rays land at the singular value.

One can extend this proof to show, without reference to the results in Sections 4 and 5, that in the postsingularly preperiodic case every periodic point z is the landing point of a periodic dynamic ray: one endows the spiders in this section with additional legs at the periodic orbit of z and applies the same arguments to the extended spider.

References

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

- [DGH] Clara Bodelón, Robert Devaney, Michael Hayes, Gareth Roberts, Lisa Goldberg and John Hubbard: *Hairs for the complex exponential family*. International Journal of Bifurcation and Chaos **9** 8 (1999), 1517–1534.
- [Do] Adrien Douady: *Compacts in \mathbb{C}* . In: Topological methods in modern mathematics. Publish or Perish, Houston (1992), 429–465.
- [DH1] Adrien Douady and John Hubbard: *Etude dynamique des polynômes complexes*. Publications mathématiques d’Orsay 84-02 (1984) (première partie) and 85-04 (1985) (deuxième partie).
- [DH2] Adrien Douady and John Hubbard: *A topological classification of postcritically finite rational maps*. Acta Math. **171** (1993), 263–297.
- [DJ] Robert Devaney and Xavier Jarque: *Misiurewicz points for complex exponential maps*. International Journal of Bifurcation and Chaos **7** 7 (1997) (1599–1615).
- [EL1] Alexandre Eremenko and Mikhail Lyubich: *Iterates of Entire Functions*. Soviet Mathematics Doklady **30** (1984), 592–594.
- [EL2] Alexandre Eremenko and Mikhail Lyubich: *Dynamical properties of some classes of entire functions*. Annales de l’Institut Fourier, Grenoble **42** 4 (1992), 989–1020.
- [GM] Lisa Goldberg and John Milnor: *Fixed points of polynomial maps II*. Ann. Scient. Ec. Norm. Sup., 4^e série, t. 26 (1993), 51–98.
- [H] John Hubbard: *Three theorems of J.-C. Yoccoz*. In: *Topological methods in modern mathematics*. Publish or Perish, Houston (1992), 467–511.
- [HS] John Hubbard and Dierk Schleicher: *The spider algorithm*. In: R. Devaney (ed.), *Complex Dynamical Systems*. American Mathematical Society (1994), 155–180.
- [LP] Gennadi Levin and Feliks Przytycki: *External rays to periodic points*. Israel Journal of Mathematics **94** (1995), 29–57.
- [M1] John Milnor: *Dynamics in one complex variable: introductory lectures*. Vieweg, Wiesbaden (1999).
- [M2] John Milnor: *Periodic orbits, external rays and the Mandelbrot set; an expository account*. Astérisque **261** (2000), 277–333.
- [Sch1] Dierk Schleicher: *Rational external rays of the Mandelbrot set*. Astérisque **261** (2000), 405–443.
- [Sch2] Dierk Schleicher: *Attracting dynamics of exponential maps*. Ann. Acad. Sci. Fenn., to appear.
- [SS] Dierk Schleicher and Mitsuhiro Shishikura: *Topology and combinatorics of postsingularly finite exponential maps*. In preparation.
- [SZ1] Dierk Schleicher and Johannes Zimmer: *Dynamic Rays of Exponential Maps*. Institute for Mathematical Sciences at Stony Brook, Preprint **9** (1999).

- [SZ2] Dierk Schleicher and Johannes Zimmer: *Escaping Points of Exponential Maps*. Journal of the London Mathematical Society **67** (2003), 1–21 (to appear).
- [St] Norbert Steinmetz: *Rational Iteration*. De Gruyter Studies in Mathematics **16** (1993).
- [T] William Thurston: *On the geometry and dynamics of iterated rational maps*. Preprint, Princeton University (1985).
- [Zi] Johannes Zimmer: *Über das Iterationsverhalten von postsingulär endlichen Exponentialabbildungen*. Diploma thesis, TU München (1997).