Dynamic Rays for Exponential Maps

Dierk Schleicher and Johannes Zimmer

Abstract

We discuss the dynamics of exponential maps $z \mapsto \lambda e^z$ from the point of view of *dynamic rays*, which have been an important tool for the study of polynomial maps. We prove existence of dynamic rays with bounded combinatorics and show that they contain all points which "escape to ∞" in a certain way. We then 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 the case of postsingularly finite exponential maps, this needs the use of spider theory.

1991 Mathematics Subject Classification: Primary 58F23

Contents

Stony Brook IMS Preprint #1999/9
June, 1999

1 Introduction

The study of the dynamics of iterated polynomials has been a success story for the last two decades, pioneered by Douady and Hubbard's study of quadratic polynomials [DH1]. As a main tool, they introduced dynamic rays which help to describe the topology of Julia sets. Our goal is to make dynamic rays available to iterated entire maps. As a prototype, we concentrate on the family $z \mapsto \lambda e^z$: this should help to bring out the ideas more clearly without obscuring them by technicalities; however, our methods apply to much larger classes of entire maps.

One main feature which helps to investigate polynomial dynamics is the existence of the attracting basin of the superattracting fixed point ∞ in the complement of the Julia set. In the case when the latter is connected, one gets dynamic rays for free. For many entire maps, the Julia set is the entire complex plane, so it is not clear how to identify dynamic rays there. In particular, this is the case for some of the maps which are understood best in this paper (the postsingularly finite and bounded escape cases).

There have been earlier attempts to study the dynamics of exponential maps, notably by Eremenko and Lyubich [EL2], Baker and Rippon [BR], Devaney, Goldberg and Hubbard [DGH], Devaney and Krych [DK], Devaney and Jarque [DJ], and Viana [Vi]. The latter four papers have also introduced variants of dynamic rays (under the name of "hairs"). However, most of the results in [DGH], [DK] and [DJ] have required the parameter λ to be real, or the combinatorics has been restricted to "regular" sequences (those containing no entry 0); this implies that no two rays can land together. On the other hand, it is our conviction that the most important rays for the structure of Julia sets are those which do land together. The most general discussion of rays is in [Vi]; here, the focus is on differentiability of the rays, not on the way they structure the Julia set.

We propose an approach which overcomes the mentioned restrictions for the parameter λ or for dynamic rays (at bounded combinatorics). The early part of this paper is closely related to earlier papers, particularly to that of Viana, and some known results are reproved. This approach allows to obtain a more complete theory: every orbit which escapes to ∞ with bounded imaginary parts is on a dynamic ray; 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. These results bring the dynamical study of exponential maps much closer to the study of polynomials.

In this paper, we restrict attention to bounded combinatorics. In Section 2, this is mostly convenience: the topological results work also for unbounded combinatorics, as is already known from Viana; some of our precise estimates become void. For the structure of Julia sets and parameter spaces, the important rays are for periodic or preperiodic combinatorics; they are discussed in later sections, and their combinatorics is bounded anyway.

The study of dynamic rays for polynomials has been pioneered by Douady and Hubbard [DH1]; later, 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 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: the fundamental construction of dynamic rays takes place in Section 2, first by constructing them sufficiently far to the right ("ray ends") and then pulling them back. We obtain a parametrization of the rays in terms of "potentials" and show that points on the rays at fixed potentials depend analytically on the parameter (some of the results here serve already as a preparation for an investigation of parameter space). The combinatorics is described in terms of a "static partition" which is useful only for orbits sufficiently far to the right. 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 which 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 and a result from [SS] which contains a systematic investigation of postsingularly finite exponential maps.

Some notation. Let $\mathbb{C}^* := \mathbb{C} - \{0\}$ and $\mathbb{C}' := \mathbb{C}^* - \mathbb{R}_-$. The principal branch of the logarithm in $\mathbb{C}'$ will be denoted Log. We will choose our parameters $\lambda \in \mathbb{C}^*$ and define our maps $E_\lambda\colon \mathbb{C} \to \mathbb{C}$ by $E_\lambda(z) = \lambda e^z$. For $E_1 = \exp$, we will simply write E. We will often need $\log(\lambda)$; we will suppose that together with λ, a branch of $\log(\lambda)$ has been chosen. While the exact choice is in principle inessential, we have written our estimates for $\mathrm{Arg}(\lambda) \in [-\pi, \pi]$. Although many of our constructions will depend on λ, we will usually suppress that from the notation. 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 $+\infty$.

Acknowledgements. 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 and Phil Rippon.

2 Ends of Dynamic Rays

In this section, we want to define dynamic rays with sufficiently large real parts for a parameter $\lambda \in \mathbb{C}^*$. For $k \in \mathbb{Z}$, define the horizontal strip R_k by

$$R_k := \{z \in \mathbb{C} \mid (2k-1)\pi - \operatorname{Im}(\log \lambda) < \operatorname{Im}(z) < (2k+1)\pi - \operatorname{Im}(\log \lambda)\} .$$

The union of all R_k is a partition of the complex plane; the boundaries are the preimages of the negative real axis since E_λ maps each strip R_k to the slit complex plane $\mathbb{C}'$. We will equivalently call the strips R_k "sectors" in analogy to the polynomial case: the strips are limits of polynomial sectors of angles $2\pi/d$ with vertices at $-d$, so we imagine them as sectors of angles 0 with vertex at $-\infty$.

The inverse of E_λ mapping $\mathbb{C}'$ to R_k will be denoted by $L_{\lambda,k}$, so that $L_{\lambda,k}(z) = \operatorname{Log} z - \log\lambda + 2\pi i k$. As a consequence, $R_k \subseteq E_\lambda(R_k)$ for every $k \neq 0$. Note that R_0 is the only sector having nonempty intersection with the image of the boundary of an arbitrary sector. 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). For any $z \in \mathbb{C}$ with $E_\lambda^{\circ n}(z) \notin \mathbb{R}^-$ for all $n \in \mathbb{N}$, the *external address* $S(z)$ is the sequence of numbers of the sectors containing $z, E_\lambda(z), E_\lambda^{\circ 2}(z), \ldots$ (this corresponds to the "binary expansion" of external angles of monic quadratic polynomials; compare the discussion at the beginning of Section 4). Let $\mathcal{S} = \{s_1 s_2 s_3, \ldots : \text{ all } s_k \in \mathbb{Z}\}$ be the space of sequences over the integers. If such a sequence $s := s_1 s_2 s_3 \ldots$ is bounded, we will write $|s| := \sup\{|s_k|\}$. The shift map on $\mathcal{S}$ will be denoted σ.

Definition 2.1 (Escaping Point)
A point $z \in \mathbb{C}$ with $\operatorname{Re}(E_\lambda^{\circ n}(z)) \to +\infty$ as $n \to +\infty$ will be called an escaping point.

It is one of our goals to describe the set of escaping points with given external address s. We will show in Theorem 2.3 that there exists a curve of escaping points sharing the given external address; in a suitably chosen right half plane, all escaping points with the given external address lie on that curve. Before stating the precise results, let us introduce a name for these curves.

Definition 2.2 (Ray End)
A ray end *with external address $s \in \mathcal{S}$ is a curve*

$$g_s \colon [1, \infty) \to \mathbb{C},$$

satisfying the following conditions: each point on the curve has external address s, $\operatorname{Re}\Big(E_\lambda^{\circ n}(g_s(t))\Big) \to +\infty$ for $n \to \infty$, and $\lim_{t \to +\infty} \operatorname{Re}(g_s(t)) = +\infty$.

Roughly speaking, a ray end is a curve of points sharing the same external address s. Ray ends stretch to $+\infty$ and lie entirely in the Julia set of E_λ (this follows from the fact that

$\lim_{n\to+\infty} \mathrm{Re}\Big(E_\lambda^{\circ n}(g_s(t))\Big) = +\infty$ since there is strong expansion for all orbits which start near a ray end, so there will always be nearby orbits which map far into a left half plane after many iterations and then get close to the origin; this is incompatible with locally uniform convergence in the Fatou set). Of course, ray ends depend on the parameter λ, but we suppress this in order not to overload notation. Every ray end satisfies a bound in the vertical direction depending only on the first entry in s. In [DGH], similar objects have been examined under the name (tails of) hairs.

Now it is time to state the main result of this section.

Theorem 2.3 (Existence of Ray Ends)
For every $\lambda \in \mathbb{C}^*$ and every bounded sequence $s \in \mathcal{S}$, there is a ray end having external address s. There is a positive $R = R(\lambda, s)$ such that every $z \in \mathbb{C}$ for which $S(z) = s$, $\mathrm{Re}(E_\lambda^{\circ n}(z)) > R$ for all n and $\mathrm{Re}(E_\lambda^{\circ n}(z)) \to +\infty$ as $n \to +\infty$ is a point on that ray.

Remark. This theorem generalizes a result of [DGH]: in that paper, the existence has been proved for so-called *regular sequences* — that is, a bounded $s \in \mathcal{S}$ with all $s_k \neq 0$. It is one of the goals of this paper to get rid of the restriction to regular sequences. We will see at the end of Section 5 that non-regular external addresses are the most interesting ones (for unbounded sequences satisfying certain growth conditions, see [DK]).

The rest of this section is devoted to the proof of Theorem 2.3. For E_λ, we will first prove the existence of a map $g_s(t)$ conjugating the dynamics of E_λ on a curve to the dynamics of $E\colon t \mapsto \exp(t)$ on some right end of $\mathbb{R}$. Recall that in the polynomial case, the conjugation to $z \mapsto z^d$ (the Riemann map) is used to define dynamic rays. We will use a similar approach here, but we cannot define it on any open set. Instead, we define inductively maps g_s^n on right ends of $\mathbb{R}$ as follows for $n \in \mathbb{N}$:

$$g_s^n(t) := L_{\lambda,s_1} \circ \ldots \circ L_{\lambda,s_n} \circ E^{\circ n}(t) \ .$$

We will show below that there is a $t_0 \in \mathbb{R}$ such that these maps are defined for all $t \geq t_0$ independently of n.

Lemma 2.4 (Bound on Real Parts)
For every $K > 0$ and for every $\lambda \in \mathbb{C}^*$ with $|\log\lambda| \leq K$, every $n \in \mathbb{N}$ and every external address s, the function g_s^n is defined for all $t \geq 2\log(K+2)$ and satisfies $\mathrm{Re}(g_s^n(t)) \geq t - (K+2)$ and $g_s^n(t) = \mathrm{Log}(g_{\sigma(s)}^{n-1}(e^t)) - \log(\lambda) + 2\pi i s_1$.

Proof. We start an induction with $g_s^0(t) = t$. It is defined for all real t, and hence $g_s^1(t) = t - \log\lambda + 2\pi i s_1$ is also defined for all real t. For $n \geq 1$, we have $g_s^n(t) = L_{\lambda,s_1}(g_{\sigma(s)}^{n-1}(e^t))$ and thus

$$\mathrm{Re}(g_s^n(t)) \quad = \quad \mathrm{Re}\left(\mathrm{Log}(g_{\sigma(s)}^{n-1}(e^t)) - \log\lambda + 2\pi i s_1\right)$$

$$
\begin{aligned}
&= \operatorname{Re}(\operatorname{Log}(g^{n-1}_{\sigma(s)}(e^t))) - \operatorname{Re}(\log\lambda) = \log\left|g^{n-1}_{\sigma(s)}(e^t)\right| - \operatorname{Re}(\log\lambda) \\
&\geq \log(e^t-(K+2)) - |\log\lambda| \\
&= t - \log\left(\frac{1}{1-(K+2)/e^t}\right) - |\log\lambda| \\
&> t - \frac{(K+2)/e^t}{1-(K+2)/e^t} - K = t - \frac{1}{e^t/(K+2)-1} - K \\
&\geq t - \frac{1}{K+1} - K > t-(K+2)\ .
\end{aligned}
$$

Therefore, $g^{n+1}_s(t) = L_{\lambda,s_1}(g^n_{\sigma(s)}(e^t))$ is defined for all $t \geq 2\log(K+2)$ as well, which concludes the inductive proof. The recursive relation for the g^n_s is built into the definition. □

Proposition 2.5 (Parametrization of Ray Ends)
Fix a constant $K>0$ and a bounded sequence $s := (s_1, s_2, \ldots) \in \mathcal{S}$, let $\lambda \in \mathbb{C}^*$ be a parameter with $|\log\lambda| \leq K$, and let $t^* := 2\log(K+2)$. Then the sequence of functions $g^n_s(t)$ is well defined for $t \geq t^*$ and converges uniformly in t to a limit function $g_s(t)$. This function is injective and continuous in t and depends for fixed $t \geq t^*$ analytically on λ. It satisfies the functional equation $g_{\sigma(s)}(e^t) = E_\lambda g_s(t)$. Moreover, $g_s(t) = t - \log\lambda + 2\pi i s_1 + r(t)$ with $|r(t)| < C(K+2\pi|s|)e^{-t}$ for some universal constant $C \in \mathbb{R}$.

Proof. For some fixed $t \geq t^*$, let $t_k := E^{\circ k}(t)$. It is then easy to verify by induction that $t_{k+1} - (K+2) > e^k$ for all k. Let $M := K + 2\pi|s|$. We have

$$g^{n+1}_s(t) - g^n_s(t) = L_{\lambda,s_1} \circ \ldots L_{\lambda,s_n} \circ L_{\lambda,s_{n+1}}(e^{t_n}) - L_{\lambda,s_1} \circ \ldots L_{\lambda,s_n}(t_n)\ .$$

Since $|L'_{\lambda,s_k}(z)| = 1/|z| \leq 1/\operatorname{Re}(z)$ for any k, the lemma gives

$$\operatorname{Re}(L_{\lambda,s_{k+1}} \circ \ldots L_{\lambda,s_n} \circ L_{\lambda,s_{n+1}}(e^{t_n})) = \operatorname{Re}(g^{n+1-k}_{\sigma^k(s)}(t_k)) \geq t_k - (K+2)\ .$$

Therefore, $L_{\lambda,s_{k+1}}$ is applied to arguments z_k with $\operatorname{Re}(z_k) \geq t_{k+1} - (K+2)$ and $\left|L'_{\lambda,s_{k+1}}\right| \leq 1/(t_{k+1}-(K+2))$. Since $|L_{\lambda,s_{n+1}}(e^{t_n}) - t_n| = |t_n - \log\lambda + 2\pi i s_{n+1} - t_n| \leq M$, we get for $n \geq 1$

$$|g^{n+1}_s(t) - g^n_s(t)| \leq M\left(\prod_{k=0}^{n-1}(t_{k+1}-(K+2))\right)^{-1} < M(e^t-(K+2))^{-1}\prod_{k=1}^{n-1} e^{-k}\ .$$

(Note that there is no problem with different branches of the logarithm because all logarithms are applied only to arguments within the right half plane.) This proves convergence of $g^n_s(t)$ to a limit $g_s(t)$ as $n \to \infty$. The convergence is uniform in t and λ provided $t > t^*$ and $|\log\lambda| \leq K$. This proves continuity of the limit g_s. Moreover, since all g^n_s are holomorphic in λ for fixed t, so is the limit g_s by Weierstraß' theorem. The functional equation $E_\lambda g_s(t) = g_{\sigma(s)}(e^t)$ follows from $E_\lambda g^n_s(t) = L_{\lambda,s_2} \circ \ldots \circ L_{\lambda,s_n}(E^{\circ n}(t)) = L_{\lambda,s_2} \circ \ldots \circ L_{\lambda,s_n}(E^{\circ(n-1)}(e^t)) = g^{n-1}_{\sigma(s)}(e^t)$.

Finally, we have

$$\begin{aligned} |r(t)| &= |g_s(t) - (t - \log\lambda + 2\pi i s_1)| = |g_s(t) - g_s^1(t)| \\ &< Me^{-t}(1-(K+2)e^{-t})^{-1}\sum_{n=1}^{\infty}\prod_{k=1}^{n-1} e^{-k} < MCe^{-t} , \end{aligned}$$

where

$$C = 2\sum_{n=1}^{\infty}\prod_{k=1}^{n-1} e^{-k} \approx 2.84$$

is a universal constant: since $t \geq 2\log(K+2)$ and $K > 0$, the term $(1-(K+2)e^{-t})^{-1}$ is bounded above by 2.

Injectivity of g_s follows like this: if $g_s(t_1) = g_s(t_2)$ for $t_1 \neq t_2$, then the functional equation implies $g_{\sigma(s)}(e^{t_1}) = g_{\sigma(s)}(e^{t_2})$; iterating this functional equation sufficiently often, we get a contradiction to the asymptotic bound for $g_{\sigma^n(s)}$. □

Proof of Theorem 2.3. The function $g_s(t)$ defined in Proposition 2.5 is constructed so that it parametrizes a ray end with the given external address: the estimate $g_s(t) = t - \log\lambda + 2\pi i s_1 + O(e^{-t})$ gives us $\mathrm{Re}(g_s(t)) \to +\infty$ for $t \to +\infty$. For t sufficiently large, it follows that

$$\mathrm{Re}\Big(E_\lambda^{\circ n}(g_s(t))\Big) = \mathrm{Re}\Big(g_{\sigma^n(s)}(E^{\circ n}(t))\Big) \to \infty \qquad \text{as} \qquad n \to \infty$$

since the bounds on $g_{\sigma^n(s)}$ need only the bounds on $|s|$. This shows that $g_s(t)$ is indeed a ray end with external address s.

It remains to show that every z_0 escaping to $+\infty$ within the half plane $\mathrm{Re}(z) > R$ and having external address s is on the ray end g_s. To see this, let $z_n := E_\lambda^{\circ n}(z_0)$ and $t_n := \mathrm{Re}(z_n)$. Then $|\mathrm{Im}(z_n)| < M$ for all n and we have

$$t_n = \mathrm{Re}(z_n) \geq \frac{|z_n|}{\sqrt{1+(M/R)^2}} = \frac{|\lambda|}{\sqrt{1+(M/R)^2}} e^{t_{n-1}} .$$

Since $t_{n-1} > R$, for sufficiently large R there is a $\delta > 0$ depending only on M and R such that $t_n - \delta > e^{t_{n-1}-\delta}$ and thus $t_n > E^{\circ n}(t_0 - \delta) + \delta$ by induction. Therefore, there is a sequence $(\tau_n) \in \mathbb{R}$ with $\tau_n > t_0 - \delta$ such that $E^{\circ n}(\tau_n) = t_n$. It is not hard to find an upper bound for the sequence (τ_n): we have $t_n \leq |\lambda| e^{t_{n-1}}$, so for R large enough there is a $\delta' > 0$ with $t_n + \delta' < e^{t_{n-1}+\delta'}$. Hence $t_n < E^{\circ n}(t_0 + \delta') - \delta'$ and $\tau_n < t_0 + \delta'$.

For all n, $|z_n - E^{\circ n}(\tau_n)|$ is bounded from above by $2\pi|s|$. Pulling the two points z_n and $E^{\circ n}(\tau_n)$ back under $L_{\lambda,s_1} \circ \ldots \circ L_{\lambda,s_n}$, we obtain z_0 respectively $g_s^n(\tau_n)$, and by the contraction of the logarithm as calculated in the proof of Proposition 2.5, it follows that $|z_0 - g_s^n(\tau_n)|$ converges exponentially fast to 0 as $n \to \infty$.

Let $\tau^* \in \mathbb{R}$ be a limit point of the sequence (τ_n) and fix an $\varepsilon > 0$. Then for n sufficiently large, we have $|z_0 - g_s^n(\tau_n)| < \varepsilon$. Moreover, by uniform convergence of g_s^n to g_s, we have $|g_s^n(\tau_n) - g_s(\tau_n)| < \varepsilon$ for sufficiently large n. Finally, for τ_n close enough to τ^*, we have $|g_s(\tau_n) - g_s(\tau^*)| < \varepsilon$. Combining this, it follows

$$|z_0 - g_s(\tau^*)| \leq |z_0 - g_s^n(\tau_n)| + |g_s^n(\tau_n) - g_s(\tau_n)| + |g_s(\tau_n) - g_s(\tau^*)| < 3\varepsilon \ .$$

Hence $g_s(\tau^*) = z_0$, which proves the theorem (and shows uniqueness of the limit point τ^* by injectivity of g_s). □

Any ray end can be pulled back by the dynamics unless it contains the singular value. The resulting object after as many pull-backs as possible will be called a *dynamic ray*: in the periodic case, any ray end will either hit the singular value, or it can be pulled back infinitely often. Preperiodic rays are simply iterated pre-images of periodic ones. In general, the ray for a bounded sequence s (whether or not it eventually becomes periodic) is the n-th inverse image of the ray end with external address $\sigma^{\circ n}(s)$, in the limit of $n \to \infty$. Since the construction of these rays involves only the bound on s, all the needed rays can be constructed starting at the same potential t^*. Therefore, any bounded external address s gives rise to an entire dynamic ray until the singular value interferes (and many statements can be carried over even to unbounded external address which do not diverge too fast).

However, we will lose the parametrization of the ray via the conjugation to the exponential map: the functional equation $g_{\sigma(s)}(e^t) = E_\lambda g_s(t)$ implied that, for a given external address s, the dynamics on the ray end with external address s was conjugated to $E(t) = e^t$. The positive real axis is a dynamic ray for the exponential map, but it does contain the singular value. This eventually destroys the conjugation because any real t will eventually run out of images under iterated logarithms. Therefore, we have to choose a different conjugation.

In any case, this conjugation was arbitrary: instead of $E(t) = e^t$, any other strictly monotonically increasing function could have been used. The choice $E(t)$ has the advantage that for large t, the conjugation approaches the identity (plus a constant) exponentially fast. In order to investigate landing properties of dynamic rays, a different conjugation will be more useful: we will use the map $F\colon u \mapsto e^u - 1$. It is strictly monotonically increasing and has a unique fixed point at 0, which is indifferent.

To see that we can use this conjugation just as well, observe first that $u \mapsto u + 1$ conjugates $F(u) = e^u - 1$ to $u \mapsto e^{u-1} = (1/e)e^u$. This latter map is in our family $\{E_\lambda\}$, and the positive real axis is a ray end for the periodic external address consisting only of entries 0, so its dynamics is conjugated to the dynamics of the exponential map on the positive real axis.

Near $+\infty$, the maps E and F do not differ much, but any real $t > 0$ can be pulled back infinitely often under F; it will converge to the fixed point 0. Any other exponential map $t \mapsto e^t - a$ could be used just as well provided $a \geq 1$, and the topology would not change much (except that the indifferent fixed point would become repelling for $a > 1$, which might yield a nicer conjugation near the fixed point).

We will from now on call the variable t the *potential* of a point on a ray, and it will always be with respect to the conjugation to $e^t - 1$. The conjugation to e^t was a preliminary construction and will no longer be used. Although there is no direct relation to potential theory on open domains, these names are intended to indicate the relation to dynamic rays of polynomials (which are usually parametrized by potentials).

3 Landing Properties of Dynamic Rays

For $\lambda \in \mathbb{C}^*$, we define the postsingular set as $P(E_\lambda) := \overline{\cup_{n\geq 0} E_\lambda^{\circ n}(0)}$. In this section, we will discuss landing properties of dynamic rays. In order to use contraction properties of hyperbolic metrics, 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 is on a dynamic ray with bounded external address.

The main result of this section is the following.

Theorem 3.1 (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 orbit is on a dynamic ray at bounded external address t, then the periodic or preperiodic dynamic ray at external address s lands at a periodic respectively preperiodic point unless $t = \sigma^n(s)$ for some $n \geq 1$.

Proof. First we discuss periodic rays. Since rays are simply connected, any inverse image of any point on some ray gives rise to a unique pull-back of the entire ray via analytic continuation, provided this ray does not contain the singular value. If the singular orbit is bounded, then it cannot be on any ray, and every ray can be pulled back infinitely often; if the singular orbit escapes on a ray, then any periodic ray can still be pulled back infinitely often unless some forward image of the periodic ray contains the singular value.

The postsingular set contains at least the two points 0 and $\lambda \neq 0$. The complement in $\mathbb{C}$ contains exactly one unbounded component U, say. It carries a unique normalized hyperbolic metric. Denote the hyperbolic distance of two points $z_1, z_2 \in U$ by $d_U(z_1, z_2)$. The map E_λ is a holomorphic covering from $E_\lambda^{-1}(U)$ onto U and thus a local hyperbolic isometry. Since $P(E_\lambda)$ is forward invariant, we have $E_\lambda^{-1}(U) \subset U$. This is a proper inclusion, for otherwise U would be an invariant open set and thus a Fatou component; since it contains dynamic rays and thus points escaping to $+\infty$, this Fatou component would have to be a domain at infinity, but such are known not to exist for our maps [EL1]. Therefore, the inclusion $E_\lambda^{-1}(U) \to U$ is a strict contraction for the respective hyperbolic metrics. Let $\rho(z)$ be the contraction factor of the densities of these two hyperbolic metrics. We have $0 < \rho(z) < 1$ everywhere, and since ρ is continuous, it is bounded away from 1 on any compact subset of U. It follows that any branch of E_λ^{-1} taking any ray to another ray is a strict contraction for the hyperbolic metric of U.

Now consider any periodic ray which on its forward orbit never hits the ray containing the singular value. Denote its period by n and let w_0 be a point on this ray. Construct a sequence

w_k of points on the same ray 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 the postsingular set is closed and does not hit the periodic ray under consideration. Obviously $d_k > d_{k+1} > 0$ 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 $\eta_0 := \sup_{z\in K'}\{\rho(z)\}$. We claim that $\eta_0 < 1$ (this would be obvious if K' was compact, but it may well be unbounded in the imaginary direction). To see this, consider a sequence of points $z_k \in K'$ such that $\rho(z_k) \to 1$. This sequence must leave every compact subset of $\mathbb{C}$, so its imaginary parts must diverge. However, since the postsingular set satisfies vertical bounds by assumption, the imaginary parts of $E_\lambda^{-1}(P) - P$ are distributed evenly (except within a compact subset of $\mathbb{C}$), and $\rho(z_n)$ cannot tend to 1.

There is an $\eta < 1$ such that, whenever some $w_k \in K'$, then $d_{k+1} \leq \eta d_k$. It follows that $d_k \searrow 0$. Although this alone does not imply convergence of the sequence (w_k), the idea is 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 $\varepsilon := s(1 - \eta)$. Then there is an index m such that $d_m < \varepsilon$ and $w_m \in K$. If there is an index $l > m$ with $w_l \notin K'$, let l be the least such index. But then

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

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, there is a uniform contraction in every step, and the sequence (d_k) converges geometrically to zero. Therefore, the sequence (w_k) converges to a limit in K'. Since s was arbitrary, the limit is in K. This proves the claim.

As a consequence, we may concentrate on those dynamic rays which enter any closed set in U with bounded real parts just finitely often. Such a ray must accumulate at (at least) one point in $P(E_\lambda) \cup \{\infty\}$. It cannot simultaneously accumulate at $P(E_\lambda)$ and at ∞ because that would require the ray to traverse regions with bounded real parts infinitely often. If it accumulates at $P(E_\lambda)$, then the bounded hyperbolic distances between w_k and w_{k+1} translate into Euclidean distances converging to zero because the points w_k are near the boundary of U. This implies that the ray must land at a periodic point of period dividing n (compare Milnor [M1, § 18]).

The final possibility is that $\mathrm{Re}(w_k) \to \pm\infty$ as $k \to \infty$. The limit cannot be $-\infty$ because any point with large negative real part has also large absolute value, and its pull-back then has large positive real part. But points with large positive real parts will land at smaller real parts under any branch of E_λ^{-1}. Therefore, the sequence of real parts of w_k has a bounded subsequence, and this finishes the proof of the theorem for periodic rays.

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. □

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 which 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 coded 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.

It may be helpful to compare this situation to the (possibly more familiar) case of 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 $\mathbb{C} - K$ 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. We can also partition the dynamic plane using straight radial lines at angles 0 and $1/2$ (which are part of the real axis). This is the analog to our static partition. Near infinity, there is not much of a difference to the ray partition, but it is usually not dynamically invariant and has no real dynamic significance. (However, the analogy is not complete: in the exponential case, there is interesting dynamics which takes place in a right half plane far from the negative real line, and here it is not very important whether this partition boundary is somewhat perturbed. The static partition can thus be used for rays.)

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 $\mathtt{0}$; the other sector will be labeled $\mathtt{1}$. We have to exclude points which ever hit the critical point under iteration, as well the rays landing at such points. Then every 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 $\mathtt{0}, \mathtt{1}, \mathtt{u}, \ldots$ 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.

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 a* bounded escape *parameter if the singular value is on a dynamic ray with bounded external address.*

In the attracting or parabolic case, the singular orbit will converge to a unique non-repelling (attracting or parabolic) orbit. If the singular orbit is finite, it ends in a repelling cycle. In the bounded escape case, the singular orbit diverges to $+\infty$ with bounded imaginary parts. 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 The Bounded Escape Case

Suppose that the singular orbit escapes to $+\infty$ with bounded imaginary parts. Then the singular value is on some dynamic ray with bounded external address (Theorem 2.3). Let R_1 be the dynamic ray to the singular value, restricted to potentials greater than and including the potential of the singular value.

Each inverse image of the ray R_1 is a curve starting at positive infinite real parts and bounded imaginary parts and stretches to negative infinite real parts (since typically the ray R_1 will be differentiable [Vi], the imaginary parts of every inverse image will be bounded as well). Together, all inverse images form a partition of the complex plane. If the singular value 0 is not on the partition boundary, then it defines a unique sector $S_{\mathtt{0}}$. All the other sectors will be labeled $S_{\mathtt{j}}$ with $\mathtt{j} \in \mathbb{Z}$ so that adjacent sectors are labeled by adjacent integers, and labels increase by $\mathtt{1}$ 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 $\mathtt{0}$; it can be shown that this happens only for $\lambda \in \mathbb{R}$.) Any point and any dynamic ray which never lands on the partition boundary will then have a well-defined itinerary; 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 $+\infty$ along bounded imaginary parts, 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. 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 $w_1 = w_{n+1}$ (then n is a common multiple of the periods of both orbits).

Let U be the complex plane from which the closure of the ray R_1 and all its (finitely or infinitely many) forward images are removed. Since only finitely many forward images of R_1 may intersect any compact subset of $\mathbb{C}$, the set U is open and connected. For an index $\mathtt{j} \in \mathbb{Z}$, let $U_{\mathtt{j}} := U \cap S_{\mathtt{j}}$ be the connected and open subset within sector $\mathtt{j}$. It carries a unique normalized hyperbolic metric. For every $\mathtt{j}$, there is a branch of E_λ^{-1} mapping U into $U_{\mathtt{j}}$. Restricting this branch to any $U_{\mathtt{j}'}$, it must contract the hyperbolic metric.

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 and shows that $w_1 = z_1$. 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 two periodic points with the same itinerary.

Therefore, any two 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 3.1 (unless it hits the singular value on its forward orbit, but such rays have no itineraries), and the itineraries of ray and its landing point must obviously be equal. It follows that any periodic point is the landing point of any ray which has the same itinerary.

Similar statements for preperiodic points are now immediate. □

We will show in Section 5 that (almost) every periodic or 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 three cases.

4.2 The Postsingularly Finite Case

Suppose that the singular orbit is finite and thus preperiodic. This case has been investigated in [DJ] in the special case of regular external addresses (called "itineraries" there). The following result will be of crucial importance:

Theorem 4.3 (Rays Land at Singular Orbit)
For every postsingularly finite exponential map, at least one dynamic ray lands at every point of the singular orbit.

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.

Let R_1 be the 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 $S_{\mathtt{0}}$, 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 $U_{\mathtt{j}}$ 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 bounded escape case and let $(\mathtt{u}_k)$ be the common itinerary of (z_k) and (w_k). Let U be the complex plane with the closure of the ray landing at the singular value and all its finitely many forward images removed. The set U 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 $\mathtt{j} \in \mathbb{Z}$, let $U_{\mathtt{j}} := U \cap S_{\mathtt{j}}$ be the open subset within sector $\mathtt{j}$. Every connected component of every $U_{\mathtt{j}}$ has a unique normalized hyperbolic metric. For every $\mathtt{j}$, there is a branch of E_λ^{-1} mapping U into $U_{\mathtt{j}}$. Restricting this branch to any connected component of any $U_{\mathtt{j}'}$, it must contract the hyperbolic metric.

If there is an index k such that z_k and w_k are in the same connected component of $U_{\mathtt{u}_k}$, then the proof for the bounded escape case goes through. However, if the orbits (w_k) and (z_k) are never in the same connected component of $U_{\mathtt{u}_k}$, 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 $U_{\mathtt{u}_k}$, 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 $S_{\mathtt{u}_k}$ as z_k and w_k. The inverse image points z_{k-1} and w_{k-1} are in the same strip $S_{\mathtt{u}_{k-1}}$ because their itineraries are equal; the inverse image of the ray pair in $U_{\mathtt{u}_k}$ can separate z_{k-1} and w_{k-1} only if it is in the same strip as well. Separating ray pairs can only get fewer if their inverse images are in different sectors, but there can never be new separations. This justifies the claim that z_k and w_k can forever be in different connected components of their $U_{\mathtt{u}_k}$ 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 bounded escape case. □

4.3 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, so there is a periodic Fatou component containing an entire left half plane. Let $a_1, a_2, \ldots, a_n = a_0$ be the points on the attracting orbit such that a_0 is in the Fatou component containing a left half plane. Then a_1 is in the Fatou component containing the singular value.

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 the closure of the 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. Let V_1 be a closed neighborhood of the point a_1 which corresponds to a disk in linearizing coordinates and which contains the singular value. Let

$$V := \{\, z \in \mathbb{C} : E_\lambda^{\circ k}(z) \in V_1 \text{ for some } k \le n+1 \\ \text{and } z \text{ is in a periodic Fatou component } \}$$

and $U := \mathbb{C} - V$ (compare Figure 1). Then V is closed and forward invariant, so U is open and backward invariant: $E_\lambda^{-1}(U) \subset U$, and this is a proper inclusion. Since V_1 is a neighborhood of the singular value, its pull-back contains a left half plane; the next $n-1$ pull-backs are all univalent and connect $a_{n-1}, \ldots, a_1$ to $+\infty$ within their Fatou components, and the last of these pull-backs yields a neighborhood of V_1. The final pull-back will connect the left half plane around a_0 to $+\infty$ along countably many open domains within the same periodic Fatou component, spaced integer multiples of $2\pi i$ apart. This is the only connected component of V disconnecting $\mathbb{C}$, so U consists of countably many connected components which are all translates of each other by integer multiples of $2\pi i$. Although none of these connected components contains the singular value, there is exactly one which surrounds it. Let $U_{\mathtt{0}}$ be this connected component and denote the others by $U_{\mathtt{j}}$ for integers $\mathtt{j}$ in the same way as before. Every $U_{\mathtt{j}}$ 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 somewhat differently: it should be a connected open subset of the Fatou set, it should contain the singular value, and it should be forward invariant such that $E_\lambda^{\circ n}$ maps V into itself; moreover, its closure should intersect the Julia set only in a parabolic periodic point. 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 which extends infinitely to the right and which contains the singular orbit.

With this modification, the given proof for the attractive case applies to all the repelling periodic and preperiodic points provided the quotient $U/2\pi i\mathbb{Z}$ is connected. It may happen, however, that the parabolic orbit disconnects this quotient. As in the postsingularly finite case, one can show that two different periodic points with identical itineraries must be in different connected components of $U/2\pi i\mathbb{Z}$ during an entire period of the pull-back, and this is possible only if their itinerary is the same as that of a parabolic periodic point. But such orbits must then be equal to 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 straightforward. □

5 (Pre-)periodic Points are Landing Points

Now we want to show that every repelling periodic or 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 cases, as well as in the bounded escape case if the singular

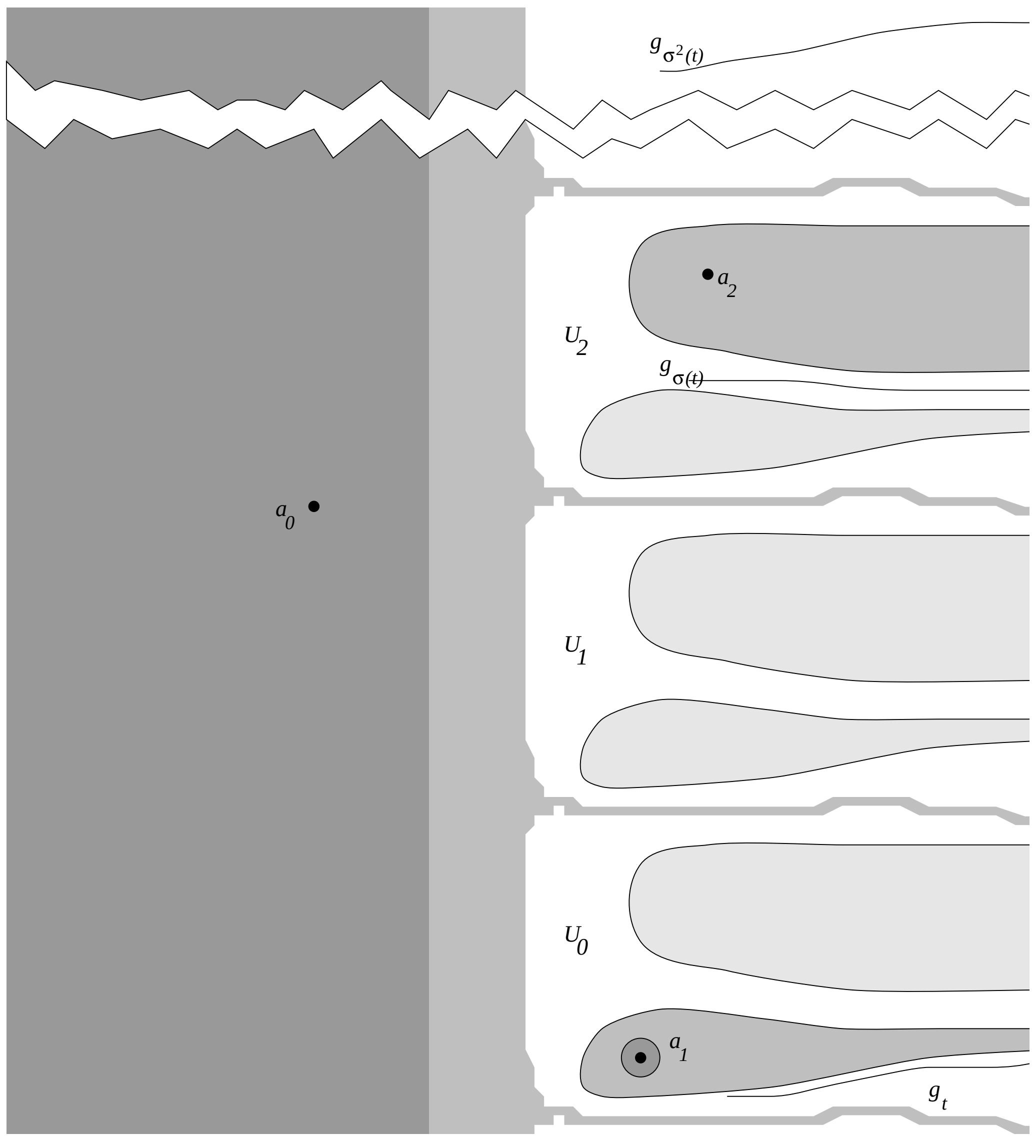

Figure 1: *The construction of the partition in the attracting case, for period $n = 3$. Indicated are the region V_1 and its immediate pull-back in dark grey, and the next n pull-back steps within periodic Fatou components in a lighter shade of gray. Together, they constitute the set V. Some preperiodic pull-backs are also shown in light grey. The attracting orbit a_k is indicated, as well as several dynamic rays used in the proof of Theorem 5.3.*

orbit is on a periodic or preperiodic ray (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 [Sch] for the proof of the analogous statement for quadratic polynomials).

The corresponding statement for polynomials is due to Douady and Yoccoz; see [H]. The proof does not generalize to 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. □

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 that the singular orbit escapes.

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

We need to define itineraries of sequences $s \in \mathcal{S}$ with respect to fixed symbolic sequences. In order to define them, let t be any sequence over $\mathbb{Z}$ starting with $t_1 \in \mathbb{Z}$, and suppose that t is different from the constant sequence. Then exactly one of the two intervals $(t_1 t, (t_1+1)t)$ and $((t_1-1)t, t_1 t)$ contains the sequence t; denote this interval $I_{\mathtt{0}}$ (this interval is the first or second of the given two examples iff the first entry in t different from t_1 is greater or smaller than t_1, respectively). For $\mathtt{j} \in \mathbb{Z}$, let $I_{\mathtt{j}}$ be this same interval, except that in every sequence the first entry is increased by $\mathtt{j}$. For a sequence $s \in \mathcal{S}$ we want to define its *itinerary with respect to* t: that should be the sequence of indices of intervals containing $\sigma^{\circ k}(s)$ for $k = 0, 1, 2, \ldots$. This works unless s ever maps onto the boundary of the partition, which happens iff it maps onto t in the next step. In this case, the corresponding entry in the itinerary will not be an integer, but a "boundary symbol" $\genfrac{}{}{0pt}{}{\mathtt{j+1}}{\mathtt{j}}$ indicating that the corresponding forward image of s is on the boundary between the intervals $I_{\mathtt{j}}$ and $I_{\mathtt{j+1}}$.

If t is a constant sequence, we can use one of the two intervals $(t_1 t, (t_1+1)t)$ or $((t_1-1)t, t_1 t)$ as $I_{\mathtt{0}}$ and proceed as above; there is no preferred choice.

The itinerary of any sequence t with respect to itself will be called the *kneading sequence* of t. By construction, the first symbol in any itinerary is $\mathtt{0}$. If t is periodic, then its orbit will run through the boundary of the partition, so the kneading sequence will contain a boundary symbol. We will say that a kneading sequence with a boundary symbol $\genfrac{}{}{0pt}{}{\mathtt{j+1}}{\mathtt{j}}$ is *adjacent* to an itinerary $\mathtt{u}$ if all non-boundary entries are equal, and if any boundary symbol $\genfrac{}{}{0pt}{}{\mathtt{j+1}}{\mathtt{j}}$ in one sequence corresponds to the same symbol or to an adjacent entry $\mathtt{j}$ or $\mathtt{j}+\mathtt{1}$ in the other sequence.

The meaning of this construction is as follows: 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 (or ends of dynamic rays) with external addresses jt (concatenation!) for integers j, and the itinerary of any dynamic ray having external address s with respect to this partition is combinatorially determined as the itinerary of the sequence s with respect to the sequence t.

Lemma 5.2 (Combinatorics of Itineraries)
For any periodic or preperiodic external address $t \in \mathcal{S}$ and any periodic or preperiodic itinerary $\mathtt{u} \in \mathcal{S}$, there is a periodic resp. preperiodic external address $s \in \mathcal{S}$ such that the itinerary of s with respect to t is $\mathtt{u}$, unless the kneading sequence of t is equal to or adjacent to a finite shift of $\mathtt{u}$. The number of sequences s with itinerary $\mathtt{u}$ always finite, and if $\mathtt{u}$ does not contain an entry $\mathtt{0}$, then there is a unique such sequence s.

Proof. First suppose that $\mathtt{u}$ is periodic with period $n \geq 1$ and let $\mathtt{u}_k$ be its k-th entry for $k \geq 1$. In order for the first entry in the itinerary of s to be $\mathtt{u}_1$, we need $s \in I_{\mathtt{u}_1}$. Since the forward orbit of t is finite by assumption, it cuts $I_{\mathtt{u}_1}$ into finitely many open sub-intervals; call them $J'_1, \dots, J'_m$ for some $m \geq 1$. Each of these sub-intervals can be pulled back n times so as to yield sub-intervals $J_1, \dots, J_m$ of $I_{\mathtt{u}_1}$ which generate n initial entries in their itineraries equal to $\mathtt{u}_1, \dots, \mathtt{u}_n$; compare Figure 2. The construction assures that the n-shift from any J_k to its image J'_k is monotone and that each J_k is completely contained in a unique $J'_{\rho(k)}$, thus defining a map ρ from $\{1, \dots, m\}$ to itself (it need not be injective).

Now we will construct sequences (m_k) with elements in $\{1, \dots, m\}$ as follows: start with any m_1 and let $m_{k+1} = \rho(m_k)$ for all k: then m_{k+1} is the index of the unique interval $J'_{m_{k+1}}$ which contains J_{m_k}. By finiteness of m, some index must repeat, and there is thus an interval J' among the J'_k such that J' contains a subinterval J which maps monotonically onto J' after qn shifts for some integer $q \geq 1$. If the intervals J and J' do not have a common boundary point, then the interior of J contains a sequence s which is fixed under the qn-th shift, so it is periodic. Obviously, this sequence s has itinerary $\mathtt{u}$. However, if the two intervals do have a common boundary point, then this boundary sequence (which is on the forward orbit of t) is fixed under the qn-th shift. Sequences within J sufficiently close to the boundary will have itineraries which coincide with $\mathtt{u}$ for arbitrarily many entries (at least n). If the boundary sequence itself has an itinerary, then it is equal to $\mathtt{u}$. This itinerary exists unless the boundary maps on its forward orbit through an immediate inverse image of t, which happens if and only

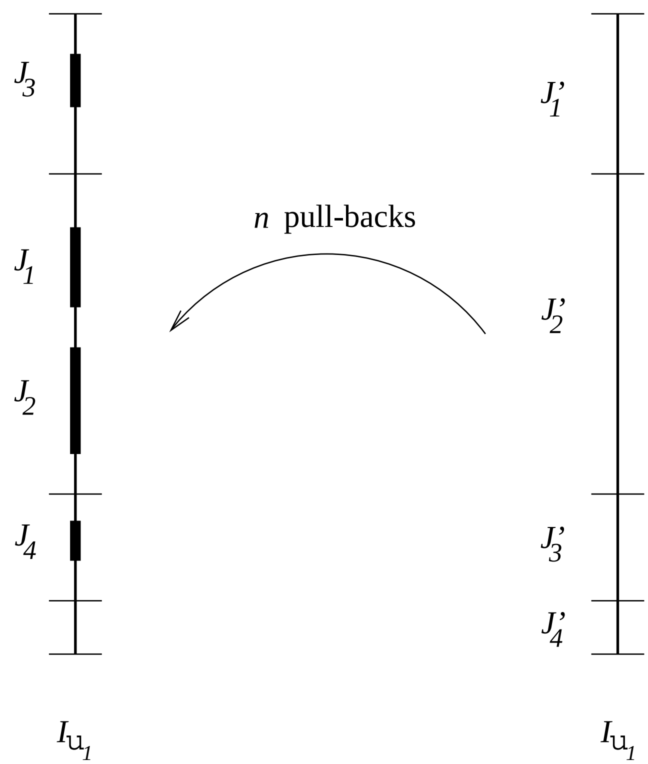

Figure 2: *The construction of the intervals $J_1, \ldots, J_m$ in the proof of Lemma 5.2.*

if it maps through t itself, and this happens only if t is periodic so that the itinerary of t is adjacent to a finite shift of $\mathtt{u}$. Everything is fine if t is preperiodic (its kneading sequence does not contain a boundary symbol). This finishes the proof for periodic $\mathtt{u}$.

For preperiodic sequences $\mathtt{u}$, we first construct a periodic sequence s' which generates a periodic itinerary $\mathtt{u}'$ on the forward orbit of $\mathtt{u}$; this is possible unless t is periodic and its itinerary is adjacent to a finite shift of $\mathtt{u}$. The preperiodic pull-back is possible if t is periodic; if it is preperiodic, it is impossible if the pull-back runs through t: this happens only if the itinerary of t is equal to a finite shift of $\mathtt{u}$. Combining everything, we see that there is a sequence s generating the itinerary $\mathtt{u}$ except if the itinerary of t is equal or adjacent to a finite shift of $\mathtt{u}$.

Finally, we need to prove uniqueness if $\mathtt{u}$ does not contain an entry $\mathtt{0}$. In this case, we can simply start with the interval $I_{\mathtt{u}_1}$ and keep pulling back. By assumption, the sequence of pull-backs will never be in sector $I_{\mathtt{0}}$ which contains t. Hence every pull-back step yields a contiguous interval, and it shrinks to a unique sequence in $\mathcal{S}$ when the pull-back is continued. (Comparing with the proof given above, the difference is this: the orbit of t may cut $I_{\mathtt{u}_1}$ into several intervals, but during the pull-backs the disconnecting forward images of t will eventually be on different branches than the various intervals themselves because the pull-back of $I_{\mathtt{u}_1}$ cannot map through $I_{\mathtt{0}} \ni t$; therefore, all the intervals J'_k will eventually pull back into a single interval, and all orbits under ρ land on a unique orbit.) This proves the lemma. $\square$

This lemma can also be described in terms of a transition matrix between the intervals J'_k. This has been elaborated in [DJ] for the case exponential maps which are postsingularly finite and for which the external address of the singular orbit is "regular".

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 (Every (Pre-)periodic Point is Landing Point)
For every postsingularly finite, attracting or parabolic 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 periodic or preperiodic dynamic ray, the same is true with the exception of periodic sequences t and periodic or preperiodic points for which a finite shift of the itinerary is equal to or adjacent to the kneading sequence of t.

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

Proof. The fact that only finitely many periodic (and thus preperiodic) rays may land together is the same as for polynomials and will thus be omitted; compare [M1, § 18].

First we consider the postsingularly finite case. We assume that some dynamic ray lands at the singular value; this will be proved in Theorem 6.8 independently of this section. Let t be the external address of such a ray. In Proposition 4.4, we have constructed partitions with respect to which every periodic point has an associated itinerary. Consider any periodic point z and let $\mathtt{u}$ 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 $\mathtt{u}$ with respect to t. For every such sequence, there is a unique dynamic ray with this external address, and each of these rays will have itinerary $\mathtt{u}$. According to Proposition 4.4, these rays must land at the periodic point z with the same itinerary. Preperiodic points are handled by pull-backs: the only problem may occur if a preperiodic point has an itinerary which runs through the kneading sequence of t under iterated shifts; but then the periodic forward orbit must be part of the singular orbit and the preperiodic point itself must map through the singular value. This is obviously impossible, and there is no problem.

Now we deal with the case that the singular value escapes on a periodic or preperiodic ray. Let t be the (periodic or preperiodic) external address of the ray containing the singular value. In Proposition 4.2, we have constructed partitions with respect to which every periodic point has an associated itinerary. Again, we consider a periodic point z with itinerary $\mathtt{u}$ and use Lemma 5.2 as before. The proof proceeds as in the postsingularly finite case, except that we have to deal with the exceptions from the lemma: if t is preperiodic, then every periodic point is the landing point of finitely many periodic dynamic rays (no exceptions so far); for preperiodic points, we can pull back unless the pull-back runs through the ray containing the singular value; in this bad case, a shift of the itinerary of the preperiodic point will be equal to the itinerary of the singular value, which is the kneading sequence of t. This case was excluded in the hypothesis.

Now we discuss the attracting case. We have constructed a partition in Proposition 4.5 (compare Figure 1), so every periodic and preperiodic point which is not on the attracting orbit is repelling and has a well-defined itinerary. The partition was not constructed using a dynamic ray through the singular value; instead we have used a subset of the Fatou component around the singular value which stretches out to $+\infty$. If n denotes again the period of the attracting orbit and thus of the periodic Fatou components, then the first $n-1$ entries in the itinerary of the singular value and its Fatou component are well defined. The n-th forward image is contained within the Fatou component defining the sector boundary, so this entry is not defined. The idea is to replace the left half plane by a ray which is very far up or down so that the partition and its combinatorics is the same; we can then apply our combinatorial lemma once more.

More specifically, let z be a repelling periodic point with itinerary $\mathtt{u}$ and let t be a periodic sequence of period n where the first $n-1$ entries are as in the itinerary of the singular orbit and the n-th entry is a very large positive integer which exceeds the absolute value of the largest entry in $\mathtt{u}$. Then the kneading sequence of t will contain at the n-th position an entry which is greater than all entries in $\mathtt{u}$, so Lemma 5.2 gives us a periodic sequence s such that its itinerary with respect to t is $\mathtt{u}$. Let g_s be the periodic ray for external address s; it exists by Theorem 2.3. We will argue that the itinerary of g_s is $\mathtt{u}$ also for the partition consisting of Fatou components. Indeed, if the itineraries are different, then some forward image of g_s must be "between" the dynamic t-ray and the singular value Fatou component (in the order near $+\infty$); since the first $n-1$ entries in the itineraries of the t-ray and the singular value Fatou component are equal, we may iterate forward $n-2$ times and preserve the order near $+\infty$ of the two rays and the Fatou component. But the $n-2$-th image of the singular value Fatou component is the Fatou component containing a left half plane, so one of the forward images of g_s must be "above" the $n-1$-th image of the t-ray; but this contradicts the bounds on the itinerary $\mathtt{u}$ of g_s. This contradiction shows that the itinerary of the ray g_s is indeed $\mathtt{u}$ with respect to the partition using Fatou components, and the ray lands at z by Proposition 4.5 (it also follows that the construction of s is independent of the choice for the n-th entry in the sequence t). For preperiodic points, we can simply pull back; there is no problem in this case.

Finally, the proof for the parabolic case is the same as for the attracting case; it works for all repelling orbits, as well as for the parabolic one. □

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.

Lemma 5.4 (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 had to contain the entry zero. □

Lemma 5.5 (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 5.4. 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. □

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.

6 Spiders

In this section, we will give a proof that for every postsingularly finite exponential map, the singular value is the landing point of at least one 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 earlier sections.

For this entire section, fix a postsingularly finite exponential map E_λ and let $s_1 = 0$, $s_2 = E_\lambda(0) = \lambda$, $s_3 = E_\lambda(E_\lambda(0))$, ... be the singular orbit with preperiod l and period k, so that $s_{l+k+1} = s_{l+1}$. Then there is an integer $N \neq 0$ such that $s_{l+k} - s_l = 2\pi i N$. Let $S := \{s_1, \ldots, s_{l+k}\}$ be the postsingular set.

A *spider* of the map E_λ will be a set of $l+k$ disjoint curves, one for each postsingular point s_j, which connects this point to ∞. These curves are not allowed to meet each other or any of the endpoints, other than those they connect to ∞. We will call these curves *spider legs* with *endpoints* s_i. Two spiders will be called *equivalent* if their endpoints are identical, and if their legs are homotopic relative to the endpoints. At ∞, the legs have a cyclic order which is fixed under homotopies. There are finitely many cyclic orders of the finitely many legs, and each of them defines a *spider space* as a set of equivalence classes of spiders. (Compare the more systematic discussion in [HS]. In general, the space of (equivalence classes of) spiders is a Teichmüller space; we are interested only in the discrete subset with fixed endpoints, which corresponds to a single point in moduli space. Usually, in Teichmüller space or spider space we have to divide by the action of Möbius transformations; however, we have already normalized so that 0 and ∞ are special and $s_{l+k} - s_l = 2\pi i N$.)

Definition 6.1 (The Spider Map with Fixed Endpoints)
Consider any spider which connects all the postsingular points of E_λ to ∞ and denote the leg which connects s_i to ∞ by γ_i, for $i = 1, 2, \dots, l+k$. Associate to this a new spider, the image spider*, which has the same endpoints s_i, and for which the leg $\tilde\gamma_i$ from s_i to ∞ is the unique inverse image of γ_{i+1} under E_λ which connects s_i to ∞ (by periodicity, we set $\gamma_{l+k+1} = \gamma_{l+1}$ here).*

Remark. Any leg γ_{i+1} connects s_{i+1} to ∞; since $E_\lambda(s_i) = s_{i+1}$, there is indeed a unique inverse image of γ_{i+1} which ends at s_i. The two new spider legs at s_l and s_{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 $s_1 = 0$ 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).

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. Any homotopy between equivalent spiders yields a homotopy between the image spiders. □

In general, the cyclic order of the spider legs near ∞ may be different for any spider and its image spider. The total number of possible cyclic orders is finite, so eventually one cyclic order must repeat. It is not obvious that then the image spiders will also have the same cyclic orders (the cyclic order alone does not determine the image spider or even its cyclic order). We want to prove more: that after some finite number of iterations, a periodic spider is reached.

Proposition 6.3 (Periodic Spider)
The iteration of any spider with endpoints at the singular orbit will lead to a periodic spider after finitely many steps (that is, to a spider which is equivalent to its image spider after finitely many iterations of the spider map).

This fact has been proved in [SS]. At this place, will not repeat the proof, which requires some background in hyperbolic geometry and Teichmüller theory. Instead, we will prove the analogous result for unicritical polynomials to illustrate the concepts in a much simpler context where a Hubbard tree is available (for postcritically preperiodic polynomials, the Julia set is a dendrite, while it is all of $\mathbb{C}$ for postsingularly finite exponential maps).

Suppose that we have a unicritical polynomial of degree d, written in the form $z \mapsto \lambda(1+z/d)^d$, and that the critical value 0 is strictly preperiodic. Then we can define spiders and a spider map in complete analogy to the exponential case (and in fact, this is possible much more generally).

Lemma 6.4 (Periodic Spider for Unicritical Polynomials)
The spider map for every postcritically preperiodic unicritical polynomial, started with an arbitrary spider, will lead to a periodic spider after finitely many steps (that is, to a spider which is equivalent to its image spider after finitely many iterations of the spider map).

Proof. In the dynamic plane of the polynomial, there is a unique Hubbard tree: that is the unique tree within the (arcwise connected) Julia set connecting the finite critical orbit. It is easy to modify the spider to a homotopic spider so that every leg intersects this tree in a finite number of points (this may be easier if the Hubbard tree is also considered only up to homotopy). It then follows easily that the image spider can have no more intersections with the Hubbard tree than the first spider, so we have a bound for the intersection numbers of all spiders on the iterated orbit of the first spider. But the number of spiders (up to homotopy) with any given number of intersections is finite, so the sequence of spiders must eventually become periodic as claimed (recall from Lemma 6.2 that we regard the spider map at this point as a map on equivalence classes). □

Remark. In fact, one can easily draw an initial spider which does not intersect the Hubbard tree at all, and it follows that the periodic spider will have the same property. After all, we are interested in dynamic rays landing at the singular orbit, and these do not intersect the Hubbard tree.

In order to make our spiders legs look more like dynamic rays, we will require from now on that they approach ∞ along bounded imaginary parts with real parts diverging to $+\infty$. We will call a spider with this property a *tamed spider*; tamed spiders will be called equivalent if they are equivalent as spiders, and if the homotopy between the spiders can be chosen so that it runs only through tamed spiders. Every spider has an equivalent tamed spider, and every equivalence class of spiders splits into finitely many equivalence classes of tamed spiders: the cyclic order of the legs at ∞ becomes refined as the order of the legs as they approach $+\infty$ (ordered by imaginary parts), and two tamed spiders are equivalent if they are equivalent as spiders and if the order of their legs at $+\infty$ is the same.

From now on, we will view the spider map on the level of spiders, rather than on equivalence classes of spiders, and we only need to consider spiders which represent the periodic equivalence class. Any spider is given by its legs, and its image spider needs one leg for each of the postsingular points: for the point s_n, we take an inverse image under E_λ of the leg at s_{n+1} and choose the branch of the inverse image (the logarithm) so that it maps s_{n+1} to s_n. If any spider is tamed, then the image spider will automatically be tamed.

Lemma 6.5 (Images of Tamed Spiders)
Two tamed spiders which are equivalent as spiders have tamed images which are equivalent as tamed spiders.

Proof. Take any tamed spider, take its lowest leg (with respect to the order at $+\infty$) and move it around the complex plane within the space of equivalent (non-tamed) spiders so that it becomes the highest leg: within the equivalence class of non-tamed spiders, we can thus change which leg is the lowest in the approach to $+\infty$, and we can arrange for any leg to be the lowest one after finitely many such moves. This way, we can turn any tamed spider into any tamed spider which is equivalent to it as a (non-tamed) spider. We need to show that the image spiders are equivalent even as tamed spiders.

There is a leg to the singular value, and all the inverse images of this leg under E_λ cut the complex plane into fundamental domains for the maps E_λ. When a leg turns around the complex plane so as to turn a tamed spider into a non-equivalent tamed spider, the images of both tamed spiders under the spider map will have their legs in the same sectors of the partition, and the claim follows easily. □

We can thus strengthen Proposition 6.3 as follows.

Corollary 6.6 (Periodic Tamed Spider)
For every postsingularly finite exponential map, there is a tamed spider which is mapped to an equivalent tamed spider after a finite number of iterations of the spider map.

Proof. Indeed, Proposition 6.3 supplies a spider which is mapped to an equivalent spider after finitely many spider iterations, and except possibly for the first time, the spider will be tamed and equivalent to its image tamed spider. □

As we see, spiders are tamed under iteration, and now they are tame enough so that we can replace the ends of spider legs by ends of dynamic ray. Let M be the number of spider iterations it takes to map the tamed spider to itself.

Proposition 6.7 (Periodic Spider with Ray Ends)
Every periodic tamed spider has an equivalent tamed spider for which the legs terminate in rays ends near $+\infty$ such that after M iterations of the spider map, the legs to the same postsingular points terminate in the same ray ends.

Proof. First we show that we can associate to each leg a canonical external address. Let R be a part of the negative real axis which does not meet the given periodic tamed spider. All its inverse images under E_λ form a static partition (within a right half plane) which does not meet any inverse image of a spider leg and in particular no leg of the image spider under the spider map (image spiders have legs which are constructed using inverse images of E_λ). Every spider leg will thus disappear to $+\infty$ within a well-defined sector of this partition, and this sector is the same for an entire equivalence class of tamed spiders. Label the sectors by the integers in the order of increasing imaginary parts so that the singular value 0 is in sector 0. This way, every leg of the image spider gets a well-defined label in $\mathbb{Z}$. Continuing for M iterations of the spider map, the tamed spider returns to an equivalent tamed spider, and the labeling becomes periodic. Since the period of the singular orbit is k, every periodic leg will have a kM-periodic sequence of labels.

For every leg, consider the periodic respectively preperiodic ray end with the same external address; such exists by Theorem 2.3. Since legs as well as ray ends are ordered near $+\infty$ according to the lexicographic order of their external addresses, we can replace the ends of all the legs by the ends of the ray ends with the same external addresses, thus obtaining an equivalent tamed spider. The construction assures that the spider map respects these ray ends so that after M iterations of the spider map, the leg into any given postsingular point always terminates in a ray end with the same external address. □

The proof of the following theorem is now routine.

Theorem 6.8 (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. Given any postsingularly finite exponential map E_λ, we can construct a tamed spider for which all the legs terminate in ends of dynamic rays as in Proposition 6.7. Iterating the spider map means pulling the ray ends back along their periodic respectively preperiodic inverse images. It suffices to consider the periodic ray ends.

We know from Theorem 3.1 that every periodic ray end extends to a periodic ray which lands at a repelling periodic point; all we want to prove is that the landing point is on the singular orbit. We will again use contraction of a hyperbolic metric.

Let U be the complex plane minus the singular orbit and $U' := E_\lambda^{-1}(U)$; then $E_\lambda\colon U' \to U$ is a holomorphic universal covering map, thus a local hyperbolic isometry. Since U' is strictly contained in U, every pull-back is also a local isometry from U to U' and the inclusion into U is a contraction. For every periodic leg, let z be a point on its ray end and let w be a point within the domain of linearization of the endpoint. Under iteration of the spider map, w will converge to the endpoint, while the hyperbolic length of the (continuously differentiable) leg between the orbits of z and w will decrease, and corresponding points on both orbits will always be on the same spider leg. Since w converges to the boundary of the domain U, hyperbolic neighborhoods of fixed radii will become Euclideanly small, so the orbit of z converges to the landing point of its leg. Since z is on a dynamic ray which lands, the ray must land at the same point as its leg, and we have a fixed spider made of dynamic rays. □

Remark. The period of the postsingular orbit is k, and the period of the tamed spider is M. Therefore, spider legs may have periods up to kM, and up to M legs (and thus rays) may land at the same postsingular point.

In Section 4, we promised a proof that every postsingular point is the landing point of a periodic or preperiodic dynamic ray. This is now trivial.

PROOF OF THEOREM 4.3. We have a spider consisting of dynamic rays which land at each point of the singular orbit. □

References

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

[DGH] Robert Devaney, Lisa Goldberg and John Hubbard: *A dynamical approximation to the exponential map by polynomials.* Preprint, Mathematical Sciences Research Institute, Berkeley (1986).

[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).

[DK] Robert Devaney and Michael Krych: *Dynamics of* $\exp(z)$. Ergodic Theory and Dynamical Systems **4** (1984), 35–52.

[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$^{\text{e}}$ 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.* Isreal Journal of Mathematics **94** (1995), 29–57.

[M1] John Milnor: *Dynamics in one complex variable: introductory lectures.* Preprint #1990/5, Institute for Mathematical Sciences, Stony Brook (1990).

[M2] John Milnor: *Periodic orbits, external rays and the Mandelbrot set; an expository account.* Preprint #1999/3, Institute for Mathematical Sciences, Stony Brook (1999), to appear in *Astérisque.*

[Sch] Dierk Schleicher: *Rational external rays of the Mandelbrot set.* Preprint #1997/13, Institute for Mathematical Sciences at Stony Brook (1997). To appear in Astérisque.

[SS] Dierk Schleicher and Mitsuhiro Shishikura: *Topology and combinatorics of postsingularly finite exponential maps.* In preparation.

[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).

[Vi] Marcelo Viana da Silva: *The differentiability of the hairs of* $\exp(z)$. Proceedings of the American Mathematical Society, 4 **103** (1988), 1179–1184

[Zi] Johannes Zimmer: *Über das Iterationsverhalten von postsingulär endlichen Exponentialabbildungen.* Diploma thesis, TU München (1997).

Dierk Schleicher and Johannes Zimmer
Fakultät für Mathematik
Technische Universität München
D-80290 München, Germany
dierk@mathematik.tu-muenchen.de
zimmer@mathematik.tu-muenchen.de