

A COMBINATORIAL CLASSIFICATION OF POSTSINGULARLY FINITE COMPLEX EXPONENTIAL MAPS

*Dedicated to the memory of Des Sheiham,
our inspiring instructor, valued colleague and great friend*

BASTIAN LAUBNER

Institut für Informatik, Humboldt-Universität zu Berlin
Unter den Linden 6, 10099 Berlin, Germany

DIERK SCHLEICHER

School of Engineering and Science, Jacobs University Bremen
Postfach 750 561, D-28725 Bremen, Germany

VLAD VICOL

Department of Mathematics, University of Southern California
Los Angeles, CA 90089, USA

ABSTRACT. We give a combinatorial classification of postsingularly finite exponential maps in terms of external addresses starting with the entry 0. This extends the classification results for critically preperiodic polynomials [2] to exponential maps. Our proof relies on the topological characterization of postsingularly finite exponential maps given recently in [14]. These results illustrate once again the fruitful interplay between combinatorics, topology and complex structure which has often been successful in complex dynamics.

Discrete and Continuous Dynamical Systems. Volume 22, Number 3 (2008), 663–682.

1. INTRODUCTION

We study the dynamical systems given by iteration of exponential maps $z \mapsto E_\lambda(z) := \lambda \exp(z)$ for non-zero complex parameters λ . The family of exponential maps is the simplest family of transcendental entire functions and has been investigated by many people (see for example [1, 3, 8]), often in analogy to quadratic polynomials as the simplest family of algebraic entire functions.

The dynamics of iterated holomorphic functions $f: \mathbb{C} \rightarrow \mathbb{C}$ is determined to a large extent by the dynamics of the singular values: these are values $a \in \mathbb{C}$ which have no neighborhood $U \subset \mathbb{C}$ so that f is an unramified covering over U . For polynomials, singular values are critical values. For transcendental functions, singular values can also be asymptotic values, or limit points of critical or asymptotic values. The exponential family is special because it has only one asymptotic value, just like unicritical polynomials (those conjugate to $z \mapsto z^d + c$) have only one critical value.

In any family of iterated holomorphic functions $f_\lambda: \mathbb{C} \rightarrow \mathbb{C}$, the easiest maps to understand are usually those for which all singular values have finite orbits, i.e. the singular orbits are periodic or preperiodic; such maps are called *postsingularly finite* (or, for polynomials, *postcritically finite*). Often they are also the maps which are most important for the structure of parameter space.

The main example is the Mandelbrot set: the *Branch Theorem* [5, 23] asserts that branch points (in a precise sense) within the Mandelbrot set are postcritically finite, and the entire topology of the Mandelbrot set is completely described by them (under the assumption of local connectivity). Moreover, if the unique critical orbit is periodic, the corresponding parameter is the center of a hyperbolic component, while if the critical orbit is preperiodic, the parameter is called a “Misiurewicz point”. For iterated rational functions,

Date: February, 2007.

1991 Mathematics Subject Classification. Primary: 30D05, 37F10, 37F20; Secondary: 37F45.

Key words and phrases. Exponential map, postsingularly finite, classification, kneading sequence, external address, spider.

there is a powerful theorem by Thurston [6] which helps to understand postcritically finite rational functions; a variant for polynomials is known as *spiders* [13]. While Thurston's theorem is deep and powerful, each time it is applied is usually a theorem in its own right. For instance, the classification of quadratic polynomials with periodic critical orbits in [13], the classification of general polynomials with preperiodic critical orbits in [2], and the classification of general postcritically finite polynomials in [16] are all derived from Thurston's theorem.

There are many analogies between the bifurcation locus of quadratic polynomials (the boundary of the Mandelbrot set) and the bifurcation locus of exponential functions. Again, one expects that much of the structure of the bifurcation diagram is determined by hyperbolic components and postsingularly finite exponential maps. For a recent survey about exponential parameter space, see [18, 19, 22].

Since for exponential maps, the singular value 0 is an omitted value, it can never be periodic, so hyperbolic components have no center. Hyperbolic components have been classified completely in [21]. Postsingularly finite exponential maps thus necessarily have preperiodic singular orbits and are sometimes equivalently called “postsingularly preperiodic”. Since Thurston's theorem applies only to rational maps, the investigation of postsingularly finite entire functions is much harder. Recently, [14] provided an extension of Thurston's theorem specifically to postsingularly finite exponential maps. We make essential use of that theorem. Our main result is a combinatorial classification of exponential functions $z \mapsto \lambda \exp(z)$ for which the singular value 0 is preperiodic. Our classification is in terms of preperiodic external addresses, i.e. preperiodic sequences over the integers. We should mention that Bergweiler (unpublished) used value distribution theory to estimate the density of postsingularly finite exponential maps.

Our result also contributes to answering (a generalization of) an old question of Euler [7]: for which values a does the limit $a^{a^{a^{\dots}}}$ exist? Euler asked this only for real $a > 0$, for which the answer is relatively simple; if a is allowed to be complex, the answer has a very rich structure. In order to be well-defined, the question needs to be rewritten: fixing a branch $\lambda = \log(a)$, then $a^{a^{a^{\dots}}} = e^{\lambda e^{\lambda e^{\dots}}}$, and we are asking for which λ the sequence $e^\lambda, e^{\lambda e^\lambda}, e^{\lambda e^{\lambda e^\lambda}}, \dots$ has a limit; except for the final exponentiation step, this is asking for which values of λ the exponential map $z \mapsto \lambda e^z$ has a converging singular orbit. The answer to this comes in three parts: (a) convergence in $\mathbb{C}$ without being eventually constant; (b) eventually constant convergence; (c) convergence to ∞ . Part (a) is easy to answer: this happens iff $\lambda = \mu e^{-\mu}$ with $|\mu| < 1$ or μ a root of unity (E_λ has an attracting or parabolic fixed point). Part (c) has been answered in [10]: the corresponding locus in parameter space consists of uncountably many curves in λ -space called parameter rays (see Proposition 3.3). Finally, part (b) are exactly the postsingularly finite exponential maps, and their classification is our main result.

This paper grew out of the Bachelor's theses of Bastian and Vlad at International University Bremen in Spring 2005. We would like to thank Nikita Selinger and an anonymous referee for many helpful comments.

2. DEFINITIONS AND CLASSIFICATION THEOREM

In this section, we introduce the necessary background from exponential and symbolic dynamics, we state our main theorems, and we present a global overview of the argument and thus of the entire paper.

Notation. We set $\mathbb{C}^* := \mathbb{C} \setminus \{0\}$, $\mathbb{C}' := \mathbb{C}^* \setminus \mathbb{R}^-$, $\overline{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$ and let $\mathbb{D}$ be the open unit disk in $\mathbb{C}$. We will denote by $f^{\circ n}$ the n^{th} iterate of the function f , and by $\mathbb{S}^2$ a 2-sphere with two distinguished points 0 and ∞ . A holomorphic exponential map will be written as $E_\lambda(z) := \lambda \exp(z)$.

We will use the following standard concepts on exponential dynamics; compare with [25, 26].

Definition 2.1. (Escaping point)

For an entire holomorphic function f , an escaping point is a point $z \in \mathbb{C}$ with $f^{\circ n}(z) \rightarrow \infty$ as $n \rightarrow \infty$; its orbit is an escaping orbit.

A holomorphic exponential map E_λ will be called *postsingularly finite* if its singular value 0 has a finite orbit, which means that the singular orbit is preperiodic (we use the word “preperiodic” in the strict sense, excluding the periodic case; the term “(pre)periodic” is used to mean either “periodic” or “preperiodic”).

The following discussion applies to exponential maps for which the singular orbit does not escape; only this case is of interest to us.

Definition 2.2. (Dynamic ray)

A dynamic ray of f is a maximal injective curve $\gamma: (0, \infty) \rightarrow \mathbb{C}$ with $\gamma(t) \rightarrow \infty$ as $t \rightarrow \infty$ so that $\gamma(t)$ is an escaping point for each $t \in (0, \infty)$. The dynamic ray lands at a point $a \in \mathbb{C}$ if $\lim_{t \rightarrow \infty} \gamma(t) = a$.

In [25, Theorem 4.2], dynamic rays were defined as curves consisting of escaping points and satisfying certain asymptotic properties (as in Theorem 2.4 below). It was shown [25, Corollary 6.9] that every escaping point is either on a unique dynamic ray, or it is the landing point of a unique dynamic ray. Every path component of the set of escaping points is a dynamic ray [10, Corollary 4.3]; as such it comes with a parametrization as an injective curve. Therefore, our definition of dynamic rays given above coincides with the original one in [25] (and is easier to state).

The E_λ -image of any dynamic ray is contained in a dynamic ray; if the singular value 0 does not escape, then the E_λ -image of every dynamic ray equals a dynamic ray. A dynamic ray γ is *periodic* if there is an $n \geq 1$ such that $\gamma(0, \infty) \supset E_\lambda^{\circ n}(\gamma(0, \infty))$ and it is *preperiodic* if $E_\lambda^{\circ k}(\gamma(0, \infty))$ is periodic for some $k > 0$. Note that no point on a ray can be periodic or preperiodic since it escapes, but the curve as a set can be.

Definition 2.3. (External address)

An external address $\underline{s}$ is a sequence $\underline{s} = s_1 s_2 s_3 \dots$ over the integers. Let S be the space of all external addresses endowed with the lexicographic order, and let $\sigma: S \rightarrow S$ be the (left) shift map.

The meaning of the external address in the dynamics of the exponential map is as follows. For an exponential map E_λ and for each $j \in \mathbb{Z}$, we let

$$R_j = \{z \in \mathbb{C} : -\operatorname{Im}(\log \lambda) - \pi + 2\pi j < \operatorname{Im}(z) < -\operatorname{Im}(\log \lambda) + \pi + 2\pi j\},$$

using the convention that $-\pi < \operatorname{Im}(\log \lambda) \leq \pi$. On each R_j , E_λ is a conformal isomorphism onto $\mathbb{C}'$. The boundaries of the strips are the set $E_\lambda^{-1}(\mathbb{R}^-)$. This partition of the complex plane into strips is called the *static partition*. The choice of labels for the strips is so that $0 \in \overline{R}_0$.

A dynamic ray γ has external address $\underline{s}$ if for all $n \in \mathbb{N}$ there is $r_n \in \mathbb{R}^+$ so that $E_\lambda^{\circ n}(\gamma(t)) \in R_{s_n}$ if $t > r_n$ (dynamic rays may well cross the static partition, but they do so only for bounded values of t). By construction in [25, Theorem 4.2], different dynamic rays have necessarily different external addresses: external addresses are the analog to external angles for polynomial rays. Different rays may land at the same point; these rays will then of course have different external addresses; see Section 3. A dynamic ray is (pre)periodic if and only if its external address is.

In [25, Theorem 4.2 and Corollary 6.9], a complete classification of escaping points and thus of dynamic rays was given. For our purposes, the following special case is sufficient.

Theorem 2.4. (Dynamic ray at external address)

If the singular orbit does not escape, then for every bounded external address $\underline{s}$ there is a unique injective continuous curve $g_\underline{s}: (0, \infty) \rightarrow \mathbb{C}$ (the dynamic ray at external address $\underline{s}$) consisting of escaping points such that:

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

satisfying

$$E_\lambda(g_\underline{s}(t)) = g_{\sigma(\underline{s})}(F(t)), \quad \forall t > 0$$

and

$$g_\underline{s}(t) = t - \log \lambda + 2\pi i s_1 + r_\underline{s}(t)$$

with $|r_\underline{s}| < 2e^{-t}(|\log \lambda| + C)$ and $F(t) = e^t - 1$; here C is a universal constant and $\log \lambda$ denotes a branch with $|\operatorname{Im}(\log \lambda)| \leq \pi$.

For our combinatorial classification of postsingularly finite exponential maps, we need a few concepts from symbolic dynamics. In what follows, terms like $t_1 \underline{t}$ (with $t_1 \in \mathbb{Z}$ and $\underline{t} \in S$) will denote concatenation.

Let $\underline{t} = t_1 t_2 \cdots \in \mathcal{S}$, and suppose $\underline{t}$ is not a constant sequence. Then either $\underline{t} \in (t_1 \underline{t}, (t_1 + 1) \underline{t}) \subset \mathcal{S}$ or $\underline{t} \in ((t_1 - 1) \underline{t}, t_1 \underline{t}) \subset \mathcal{S}$; denote the interval that contains $\underline{t}$ by I_0 . For $u \in \mathbb{Z}$, define the intervals

$$I_u := \{s_1 s_2 s_3 \cdots \in \mathcal{S} : (s_1 - u) s_2 s_3 \cdots \in I_0\}.$$

Then $\bigcup_{u \in \mathbb{Z}} I_u$ is a partition of $\mathcal{S} \setminus \{\sigma^{-1}(\underline{t})\}$. Using this, we can define combinatorial itineraries.

Definition 2.5. (Itinerary $\text{It}(\underline{s} \mid \underline{t})$)

Consider two sequences $\underline{s}, \underline{t} \in \mathcal{S}$. The itinerary of $\underline{s}$ with respect to $\underline{t}$, denoted $\text{It}(\underline{s} \mid \underline{t})$, is the sequence $\underline{u} = u_1 u_2 u_3 \dots$ over $\mathbb{Z}$ such that $\sigma^k(\underline{s}) \in I_{u_{k+1}}$ for $k \geq 0$, where the I_u are defined as above. If $\sigma^k(\underline{s}) \in \partial I_u$ for some k (hence $\sigma^{k+1}(\underline{s}) = \underline{t}$), we leave the itinerary undefined; this case will not be needed here.

In order to motivate this formal definition, consider the dynamic ray γ at external address $\underline{t}$ and suppose it does not contain the singular value. The countably many E_λ -preimages of γ are dynamic rays at external addresses $k\underline{t}$ for $k \in \mathbb{Z}$ (where again $k\underline{t}$ denotes concatenation). These preimage rays subdivide right half planes $\{z \in \mathbb{C} : \text{Re}(z) > x\}$ (for sufficiently large x) into countably many components, and every unbounded component contains unbounded parts of exactly the dynamic rays $g_{\underline{s}}$ at external addresses $\underline{s} \in I_u$ for one particular choice of $u \in \mathbb{Z}$, or equivalently those rays $g_{\underline{s}}$ whose itineraries $\text{It}(\underline{s} \mid \underline{t})$ have a given first entry u . The lexicographic order of addresses corresponds exactly to the vertical order of rays in their approach to ∞ . In Section 4, we will show that this concept makes particular sense for postsingularly finite exponential maps.

We are now ready to state the main theorems. Together, they give a complete combinatorial coding of postsingularly finite exponential maps: we construct a map from preperiodic external addresses to postsingularly finite exponential maps. The first theorem shows that the map is well-defined and surjective, the second one measures how injective this map is and thus defines an equivalence relation on preperiodic external addresses in terms whether or not they describe the same map.

Theorem 2.6. (Combinatorial coding of exponential maps)

For every preperiodic external address $\underline{s}$ starting with the entry 0, there is a unique postsingularly finite exponential map such that the dynamic ray at external address $\underline{s}$ lands at the singular value.

Every postsingularly finite exponential map is associated in this way to a positive finite number of preperiodic external addresses starting with 0.

Theorem 2.7. (Different codings)

For any two preperiodic external addresses $\underline{s}$ and $\underline{s}'$, the following are equivalent:

- (1) there is a postsingularly finite exponential map E_λ so that in its dynamic plane, the dynamic rays at external addresses $\underline{s}$ and $\underline{s}'$ land at the singular value;
- (2) the parameter rays at external addresses $\underline{s}$ and $\underline{s}'$ (see Section 3) land at the same parameter λ ;
- (3) $\text{It}(\underline{s}' \mid \underline{s}) = \text{It}(\underline{s} \mid \underline{s})$;
- (4) $\text{It}(\underline{s} \mid \underline{s}') = \text{It}(\underline{s}' \mid \underline{s}')$;
- (5) $\text{It}(\underline{s} \mid \underline{s}') = \text{It}(\underline{s} \mid \underline{s}) = \text{It}(\underline{s}' \mid \underline{s}) = \text{It}(\underline{s}' \mid \underline{s}')$;


In all these cases, $\underline{s}$ and $\underline{s}'$ have equal period and equal preperiod.

If $\underline{s}$ is a preperiodic external address with preperiod l and period k , then the itinerary $\text{It}(\underline{s} \mid \underline{s})$ (the kneading sequence of $\underline{s}$) has also preperiod l and period k' dividing k . The exact number of external addresses which yield the same postsingularly finite exponential map is equal to k/k' if $k > k'$, and it equals 1 or 2 if $k = k'$.

The above two theorems give a complete classification of postsingularly finite exponential maps in terms of external addresses. With some more combinatorial efforts, one can turn this into a classification by internal addresses as defined in [15, 24, 18]; in this setting, every postsingularly finite exponential map is described by a unique internal address, which is a strictly increasing sequences of positive integers for which the difference sequence is eventually periodic, and subject to a certain admissibility condition. We do not discuss this here (see the section on unicritical polynomials in [4]).

The proof of our classification result uses the main result of Hubbard, Schleicher and Shishikura [14], which is an extension of Thurston’s fundamental theorem on rational maps to the setting of exponential maps: their theorem is used in the existence part of our statement.

The global approach to our results is illustrated in the following commutative diagram. Our exponential maps are always assumed to be postsingularly finite and the external addresses to be preperiodic.



We start our classification in Section 3: for every postsingularly finite holomorphic exponential map, a finite positive number of preperiodic dynamic rays lands at the singular value; choose one such ray. Every dynamic ray has a unique associated external address; it turns out that dynamic rays landing at the singular value always have external addresses starting with 0. So far, this associates to every postsingularly finite exponential map a preperiodic external address (this involves a choice). In Section 3, we also discuss rays in parameter space.

In Section 4, we introduce some more concepts and algorithms from symbolic dynamics which we will need in the sequel. The main technical construction then comes in Section 5: for every preperiodic external address we first construct a graph with a continuous self-map and then extend it to a branched covering of $\mathbb{S}^2$ which we call a *topological exponential map*. Symbolic dynamics helps us to set things up so that there is no Thurston obstruction. Therefore, in Section 6 we can apply Thurston theory (applied to exponential maps) to find an equivalent holomorphic postsingularly finite exponential map, and again symbolic dynamics shows that the ray at external address $\underline{s}$ lands at the singular value. This finally shows that there is a well-defined and surjective map from preperiodic external addresses to postsingularly preperiodic exponential maps, so that the exponential map associated to an address $\underline{s}$ has the property that the dynamic ray $g_{\underline{s}}$ lands at the singular value. Finally, we investigate which external addresses give rise to the same holomorphic exponential map, thus describing exactly how far this map is from being injective.

3. FROM EXPONENTIAL MAP TO EXTERNAL ADDRESS

In this section, we start with a postsingularly finite exponential map E_λ . We show that a preperiodic dynamic ray lands at the singular value, and associate to E_λ the external address of the ray.

The hardest part of the work has conveniently been done in [26, Theorem 4.3], stated below.

Theorem 3.1. (Preperiodic ray at singular value)

For every postsingularly finite exponential map, at least one and at most finitely many preperiodic dynamic rays land at the singular value.

There can be several dynamic rays landing at the singular value. Our classification uses the fact that all of them start with the entry 0.

Proposition 3.2. (External address starts with 0)

If the dynamic ray $g_{\underline{s}}$ lands at the singular value for a postsingularly finite exponential map, then the external address $\underline{s}$ starts with 0.

This is not an obvious statement: the external address of a dynamic ray is defined using the asymptotics for large real parts; a priori, it seems quite possible that dynamic rays with non-zero first entries in their external addresses could make it to the singular value. We prove this result at the end of this section, but we will need to introduce parameter rays (and also for other purposes). Note that this happens in reversal of

Douady's famous principle "you first plough in the dynamical plane and then harvest in parameter space". We do this the other way around (like Rempe in [17]).

Similar to how dynamic rays give structure to dynamical planes, parameter space gets a lot of structure through parameter rays; the latter also help to understand bifurcations of exponential maps. Just as for quadratic polynomials and the Mandelbrot set [5], as well as for higher degree unicritical polynomials and Multibrot sets [9], there are deep relations between the structure in dynamical planes and in parameter space. We follow the arguments from [20, Section IV.6]. We will need the following special case of the main result in [11].

Proposition 3.3. (Parameter rays)

For every bounded sequence $\underline{s} \in \mathcal{S}$ starting with 0, there is an injective curve $G_{\underline{s}}: (0, \infty) \rightarrow \mathbb{C}^$ in parameter space, so that for every $t > 0$, the parameter $\lambda = G_{\underline{s}}(t)$ is the unique parameter λ so that for E_{λ} , the singular value $0 = g_{\underline{s}}(t)$. These parameter rays are disjoint for different external addresses $\underline{s}$.*

The general statement in [11] deals also with unbounded external addresses, but all we need here are preperiodic hence bounded addresses.

Theorem 3.4. (Landing of preperiodic parameter rays)

For every postsingularly finite exponential map E_{λ_0} and every preperiodic external address, the dynamic ray $g_{\underline{s}}$ lands at the singular value if and only if the parameter ray $G_{\underline{s}}$ lands at λ_0 .

Proof. Suppose E_{λ_0} has the property that the dynamic ray $g_{\underline{s}}$ at preperiodic external address $\underline{s}$ lands at the singular value; then the singular orbit for E_{λ_0} is preperiodic and terminates at a necessarily repelling periodic orbit. There is then a neighborhood $U \ni \lambda_0$ in parameter space and a unique holomorphic function $z: U \rightarrow \mathbb{C}$ so that for every $\lambda \in U$, the point $z(\lambda)$ is preperiodic with $z(\lambda_0) = 0$, and $z(\lambda)$ is still the landing point of the dynamic ray $g_{\underline{s}}$. This follows from the same arguments as in [12] for the polynomial case: it suffices to know that $z(\lambda_0)$ can be extended holomorphically as a repelling preperiodic point (this is the implicit function theorem) and that for fixed potentials $t > 0$, the point $g_{\underline{s}}(t)$ depends holomorphically on λ (and this follows from [25, Proposition 3.4]). If λ makes a small loop around λ_0 , there must be at least one parameter along this loop for which $g_{\underline{s}}$ contains the singular value 0: during one loop of λ around λ_0 , the landing point $z(\lambda)$ must loop some number $n \neq 0$ times around 0 (where n is the local degree of the holomorphic map $\lambda \mapsto z(\lambda)$); the same is thus true for points $g_{\underline{s}}(t)$ with very small potentials t . However, this is not so for large potentials t because of the asymptotics in Theorem 2.4, and this proves the claim. If $0 = g_{\underline{s}}(t)$ for E_{λ} , this means $\lambda = G_{\underline{s}}(t)$ by Proposition 3.3. Since this is true for arbitrarily small loops, λ_0 must be a limit point of $G_{\underline{s}}$.

Suppose that $\lambda_1 \in U$ was another limit point of $G_{\underline{s}}$ with $z(\lambda_1) \neq 0$. For this parameter, the dynamic ray $g_{\underline{s}}$ lands at $z(\lambda_1)$ by definition of U , and in particular the singular value is not on $g_{\underline{s}}$ or on one of the finitely many rays on the forward orbit of $g_{\underline{s}}$. Since $z(\lambda_1) \neq 0$, and the ray $g_{\underline{s}}$ together with its landing point form a compact set which changes continuously with λ (again in analogy to [12]), it follows that λ_1 has a neighborhood of parameters λ in which $0 \notin g_{\underline{s}}$. But this contradicts the assumption that $\lambda_1 \in U$ was a limit point of $G_{\underline{s}}$. Therefore, the only limit points of $G_{\underline{s}}$ within U can be λ_0 , plus possibly finitely many further parameters λ with $z(\lambda) = 0$. The set of limit points of any ray is always connected, so $G_{\underline{s}}$ lands at λ_0 .

Conversely, suppose λ_0 is the landing point of the parameter ray $G_{\underline{s}}$. Then by [26, Theorem 3.2], the dynamic ray $g_{\underline{s}}$ lands at a repelling preperiodic point z_0 . Similarly as above, ray and landing point are stable under perturbations. If $z_0 \neq 0$, then λ_0 could not even be a limit point of $G_{\underline{s}}$. $\square$

The following result is stated for convenient reference.

Corollary 3.5. (Landing properties of preperiodic parameter rays)

Every parameter ray $G_{\underline{s}}$ at preperiodic external address $\underline{s}$ lands at a postsingularly finite exponential map, and every preperiodic exponential map is the landing point of a finite positive number of parameter rays at preperiodic external addresses.

Proof. This follows immediately as soon as our classification theorems are proved (we will not use it before). $\square$

Proof of Proposition 3.2. It is shown in [11] (or [10, Corollary 3.2]) that if the singular value escapes on a dynamic ray $g_{\underline{s}}$, then the external address $\underline{s}$ starts with 0 (provided dynamic rays are parametrized so that $|\operatorname{Im} \log \lambda| < \pi$). If a parameter ray lands at a postsingularly finite exponential map, then rays and their parametrization change continuously. $\square$

It might seem that the statement of Proposition 3.2 makes sense only once a branch of $\log \lambda$ is chosen, which is not a dynamically well-defined quantity. However, this is not so: the proposition says that independently of any choice of branch, and any choice of labels of strips defining external addresses, any dynamic ray landing at 0 has asymptotic imaginary part in $(-\pi, \pi)$.

Definition 3.6. (External address of E_λ)

Let E_λ be a postsingularly finite holomorphic exponential map. Then we associate to E_λ the external address of a dynamic ray $g_{\underline{s}}$ which lands at 0 (this may involve a choice).

4. SYMBOLIC DYNAMICS AND KNEADING SEQUENCES

In Section 5 we aim to construct a topological exponential map f in which we encode all the combinatorial information of a given external address $\underline{s}$. In order to do this, we need a few more concepts from symbolic dynamics.

In Definition 2.5, we defined the space $\mathcal{S}$ of external addresses and, for every pair of sequences $\underline{s}, \underline{t} \in \mathcal{S}$, the itinerary $\operatorname{It}(\underline{s} \mid \underline{t})$ of $\underline{s}$ with respect to $\underline{t}$. Of special importance is the itinerary of a sequence with itself: the kneading sequence.

Definition 4.1. (Kneading sequence)

For a sequence $\underline{s} \in \mathcal{S}$ we call $K(\underline{s}) := \operatorname{It}(\underline{s} \mid \underline{s})$ the kneading sequence of $\underline{s}$.

The methods of symbolic dynamics and the concept of itineraries are especially useful for those exponential maps E_λ for which a dynamic ray $g_{\underline{s}}$ lands at the singular value: in particular, if the singular orbit is preperiodic (the main case of interest to us), then by Theorem 3.1 there are one or several dynamic rays at preperiodic external addresses landing at 0 (see also [26, Section 4] for a discussion of several other cases with similar properties). In this case, the countably many E_λ -preimages of $g_{\underline{s}}$ partition all of $\mathbb{C}$ and form what we call a *dynamic partition*. The components in this partition are translates of each other by $2\pi i\mathbb{Z}$; the imaginary parts of any component are in general unbounded (usually, the ray $g_{\underline{s}}$ spirals into its landing point 0).

There is always a unique component, called I_0 , which contains the singular value, and its vertical translate by $2\pi j$ is called I_j for $j \in \mathbb{Z}$. If $z \in \mathbb{C}$ is a point whose orbit is disjoint from $g_{\underline{s}}$, then we define the *itinerary* of z (with respect to the ray $g_{\underline{s}}$) as the sequence of component labels visited by the orbit of z .

We call this new partition the *dynamic partition* (as opposed to the static partition introduced in Section 2). The dynamic partition has the advantage that each dynamic ray is completely contained in one component (unless it is one of the rays forming the partition boundary), and all points on the ray and its possible landing point have the same itinerary. In fact, the itinerary of all points on the ray at external address $\underline{t}$ is $\operatorname{It}(\underline{t} \mid \underline{s})$, and the itinerary of the singular value (or of any point on any ray landing at the singular value) is the kneading sequence $K(\underline{s})$. The following result is shown in [26, Proposition 4.4].

Lemma 4.2. (Itinerary of landing points and rays)

For postsingularly finite exponential maps, no two (pre)periodic points have the same itinerary, and a (pre)periodic dynamic ray lands at a given periodic or preperiodic point if and only if ray and point have the same itinerary. In particular, two (pre)periodic dynamic rays land together if and only if they have the same itineraries.

Note that for the dynamic partition, unlike the static partition, several rays may have the same itinerary.

The following simple algorithm illustrates the close relation between external addresses and itineraries and shows in particular that their entries differ at most by 1 (up to simultaneous translation of all entries by the same integer).

Algorithm 4.3. (Construction of $\text{It}(\underline{s} \mid \underline{t})$)

Given external addresses $\underline{s}, \underline{t} \in \mathcal{S}$ so that $\underline{t}$ is non-constant and $\sigma^n(\underline{s}) \neq \underline{t}$ for all $n \geq 0$. Then $\underline{u} := \text{It}(\underline{s} \mid \underline{t})$ can be constructed as follows.

- (1) For $n \geq 1$ define $\delta_n := \begin{cases} -1 & \text{if } \sigma(\underline{t}) > \underline{t} \text{ and } \sigma^{on}(\underline{s}) < \underline{t} \\ 1 & \text{if } \sigma(\underline{t}) < \underline{t} \text{ and } \sigma^{on}(\underline{s}) > \underline{t} \\ 0 & \text{otherwise.} \end{cases}$
- (2) Construct $\underline{u} = u_1 u_2 \dots$ as $u_n = s_n - t_1 + \delta_n$.

Proof. It suffices to show that the first entry u_1 in $\text{It}(\underline{s} \mid \underline{t})$ is correct: the n -th entry equals by definition the first entry in $\text{It}(\sigma^{n-1}(\underline{s}) \mid \underline{t})$. Adding an integer k to the first entry s_1 of $\underline{s}$ will add k to u_1 , so we may assume that $s_1 = t_1$.

Suppose first that $\underline{s} > \underline{t}$. Then $u_1 = 0$ unless there is a preimage of $\underline{t}$ in $(\underline{t}, \underline{s})$; but since $s_1 = t_1$, this is equivalent to the condition $\underline{t} \in (\sigma(\underline{t}), \sigma(\underline{s}))$ or $\sigma(t) < t < \sigma(\underline{s})$; and exactly in this case, $u_1 = 1$. Similarly, if $\underline{s} < \underline{t}$, then $u_1 = 0$ unless $\sigma(t) > t > \sigma(\underline{s})$, and exactly in that case, $u_1 = -1$. $\square$

Notice that we have $\sigma^n(\underline{s}) = \underline{s}$ for some n if and only if $\underline{s}$ is periodic. Algorithm 4.3 therefore works for computing the kneading sequences of the preperiodic addresses that we are interested in. It will prove to be useful when we are trying to recover the external address from our constructed exponential map.

5. THE TOPOLOGICAL EXPONENTIAL MAP

In this section, we will start with a combinatorial object (external address) and turn this into a topological object (a postsingularly preperiodic topological exponential map). In the next section, we make the step from topology to the complex structure and find, whenever possible, a holomorphic exponential map which is equivalent, in a sense defined by Thurston, to the given topological exponential map.

Convention. All homeomorphisms and coverings in this paper will be orientation preserving.

Definition 5.1. (Topological exponential map)

A universal cover $f: (\mathbb{S}^2 \setminus \{\infty\}) \rightarrow (\mathbb{S}^2 \setminus \{\infty, 0\})$ is called a topological exponential map. It is called postsingularly finite if the orbit of 0 is finite, hence preperiodic. The postsingular set is $P_f := \bigcup_{n \geq 0} f^{\circ n}(0) \cup \{\infty\}$.

If a topological exponential map is holomorphic, then it is conformally conjugate to an exponential map $z \mapsto E_\lambda$.

Definition 5.2. (Thurston equivalence)

Two postsingularly finite exponential maps f and g with postsingular sets P_f and P_g are called Thurston equivalent if there are two homeomorphisms $\phi_1, \phi_2: \mathbb{S}^2 \rightarrow \mathbb{S}^2$ with $\phi_1|_{P_f} = \phi_2|_{P_f}$, $P_g = \phi_1(P_f) = \phi_2(P_f)$ and $\phi_1(\infty) = \phi_2(\infty) = \infty$ such that

$$\phi_1 \circ f = g \circ \phi_2 \quad \text{on } \mathbb{S}^2 \setminus \{\infty\}$$

and ϕ_1 is homotopic (or equivalently isotopic) to ϕ_2 on $\mathbb{S}^2$ relative to P_f .

Our goal will be to find, for every postsingularly finite topological exponential map, a postsingularly finite holomorphic exponential map which is Thurston equivalent. This is not always possible. In the case of rational mappings, Thurston [6] determined that this is impossible if and only if there is what is now called a Thurston obstruction; see also [2, 13]. The extension of this result to the case of exponential maps was done in [14]: in this case the possible obstructions have a much simpler form, called degenerate Levy cycles.

Definition 5.3. (Essential curves and levy cycle)

Let f be a topological exponential map with postsingular set P_f . A simple closed curve $\delta \subset \mathbb{S}^2 \setminus P_f$ is called essential if both components of $\mathbb{S}^2 \setminus \delta$ contain at least two points in P_f . Suppose there exist disjoint essential simple closed curves $\delta_0, \dots, \delta_k = \delta_0$ such that for each $i = 0, \dots, k-1$, δ_i is homotopic relative P_f to one component δ' of $f^{-1}(\delta_{i+1})$ and $f: \delta' \rightarrow \delta_{i+1}$ has degree 1. Then $\Lambda = \{\delta_0, \delta_1, \dots, \delta_k = \delta_0\}$ is called a Levy cycle.

Essential curves are important for the following reason: a simple closed curve $\delta \in \mathbb{S}^2 \setminus P_f$ is essential if and only if, for every homeomorphism $\phi: \mathbb{S}^2 \rightarrow \overline{\mathbb{C}}$, there is a lower bound of lengths (with respect to the hyperbolic metric of $\overline{\mathbb{C}} \setminus \phi(P_f)$) of simple closed curves homotopic to $\phi(\delta)$ relative to $\phi(P_f)$.

This section will be concerned with proving the following theorem.

Theorem 5.4. (External address yields topological exponential map)

Let $\underline{s}$ be a preperiodic external address. Then there exists a postsingularly finite topological exponential map f with the following properties:

- f has a preperiodic injective curve $\gamma: (0, \infty) \rightarrow \mathbb{S}^2$ connecting 0 to ∞ ;
- γ has itinerary $\text{It}(\underline{s} \mid \underline{s})$ with respect to the partition defined by $f^{-1}(\gamma)$,
- the vertical order of the rays $f^{on}(\gamma)$ coincides with the lexicographic order of the shifts of $\underline{s}$,
- f does not admit a Levy cycle.

Any two such postsingularly finite topological exponential maps for the same external address $\underline{s}$ are Thurston equivalent to each other.

Note that any injective curve connecting the singular value 0 to ∞ has countably many disjoint preimages under any topological exponential map, and this allows us to define a dynamic partition and thus dynamical itineraries just like for holomorphic exponential maps for which a dynamic ray lands at 0.

As always, the curve γ should be preperiodic as a set; its points need not be (except the endpoint). The preimage $f^{-1}(\gamma)$ is disjoint from all rays $f^{on}(\gamma)$. Let γ' be the unique component of $\gamma \cap \mathbb{D}$ starting at 0 and let p' be any component of $f^{-1}(\gamma')$. Then the rays $f^{on}(\gamma)$, as well as p' , are disjoint curves to ∞ and have a well-defined cyclic order. Removing p' induces a linear order among all rays, and this is the vertical order specified by the theorem; it does not depend on the choice of p' .

5.1. The graph map. Similarly as for polynomials in [2], we start by constructing an undirected graph Γ that encodes the combinatorial information given by $\underline{s}$. An important difference is that our graph is infinite.

We will construct an infinite topological graph Γ and later embed it into $\mathbb{S}^2$. Start with two vertices $\Gamma = \{e_\infty, e_{-\infty}\}$. For each $n \in \mathbb{Z}$, add disjoint edges p_n joining e_∞ to $e_{-\infty}$. Let k and l be the length of preperiod and period of $\underline{s}$ respectively. Add vertices $e_1, \dots, e_{k+l}$ to Γ , and for each e_n , add an edge γ_n connecting e_n and e_∞ , so that all edges are disjoint and all vertices are disjoint from each other and from all edges.

We will embed $\Gamma \setminus \{e_{-\infty}\}$ into $\mathbb{S}^2$ and define a graph map $\tilde{f}$ from the embedded graph to itself. By extending $\tilde{f}$ to a map $f: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{0, \infty\}$, we will obtain the desired topological exponential map.

It is straightforward to embed $\Gamma \setminus \{e_{-\infty}\}$ (as it has been constructed up to here) into $\mathbb{S}^2$ in a reasonable way, define a graph map, and extend it to a topological exponential map f on all of $\mathbb{S}^2$, so that f satisfies the first three properties of Theorem 5.4. The hard part is to make sure that f will not admit a Levy cycle. The following lemma tells us when to expect a Levy cycle.

Lemma 5.5. (Levy cycle and itineraries)

Consider a topological exponential map that satisfies the first two properties in Theorem 5.4. Then two or more postsingular points are surrounded by the same curve in a Levy cycle if and only if they are all periodic, and they have the same itinerary with respect to the partition consisting of preimages of the ray landing at the singular value.

Proof. Suppose that two or more postsingular points are surrounded by a simple closed curve δ in a degenerate Levy cycle. Note first that δ cannot surround the singular value 0: otherwise, the preimage of δ would

not contain any simple closed curve. After homotopy, we may thus assume that δ does not intersect the ray γ_1 connecting 0 to ∞ .

Taking preimages, no preimage curve δ' of δ can intersect the partition boundary, hence all postsingular points surrounded by δ' have the same first entries in their itineraries. Note that the number of postsingular points surrounded by δ' cannot be greater than that for δ ; this number could be smaller, depending on which branches of preimages are chosen. However, since δ is part of a degenerate Levy cycle and hence periodic (up to homotopy), the number of surrounded postsingular points must remain constant. It follows that all postsingular points surrounded by δ have the same periodic itinerary and are hence periodic points.

Conversely, assume that the periodic postsingular points $e_{i_1}, e_{i_2}, \dots, e_{i_k}$ have the same itinerary with respect to the dynamic partition of the plane. Surround these points (but no other postsingular points) by a simple closed curve δ ; this curve is automatically essential. Note that there are in general infinitely many homotopy classes of curves δ relative to the postsingular points. However, there is only a single homotopy class if we require that δ must not intersect the ray into any point which is not surrounded: the complement in $\mathbb{S}^2 \setminus \{\infty\}$ of all these rays is simply connected.

Since δ does not intersect the ray γ_1 , and all surrounded points are periodic and have the same itinerary, there is one preimage component of δ which surrounds all periodic preimages of the surrounded points, and again it does not intersect the rays into those points that it does not surround. Repeating this argument for one period of the itinerary, we obtain another curve which surrounds the same points as δ in the complement of the remaining rays, so this curve is homotopic to δ and we have a degenerate Levy cycle. $\square$

Corollary 5.6. (Levy cycle and itineraries)

A topological exponential map as in Lemma 5.5 admits a Levy cycle containing a curve surrounding the points e_i, e_j if and only if the kneading sequence $\underline{u} = \text{It}(\underline{s} \mid \underline{s})$ of the curve γ landing at the singular value has the property that $u_i u_{i+1} u_{i+2} \dots = u_j u_{j+1} u_{j+2} \dots$.

Proof. If $\underline{u} = u_1 u_2 \dots$ is the kneading sequence of the curve γ landing at the singular value, then the itineraries of the points $e_i = f^{o(i-1)}(0)$ and e_j are the appropriate shifts of the kneading sequence. Equality of itineraries can hold only if both points have equal periods and preperiods; since they are on the same preperiodic orbit, this implies that they can have identical itineraries only if they are periodic. $\square$

The Levy cycle obstruction warns us that rays γ_i and γ_j should really land together at a common point $e_i = e_j$. In order to solve this problem, we will simply “glue” points e_i and e_j in the graph: Define the equivalence relation

$$e_i \sim e_j \iff u_i u_{i+1} u_{i+2} \dots = u_j u_{j+1} u_{j+2} \dots$$

on points $e_k \in \Gamma$. Redefine Γ as $\Gamma / \sim$.

After embedding into $\mathbb{S}^2$, the new quotient graph Γ will have the property that no two different vertices have the same itinerary, so there can be no Levy cycle.

In order to check that the graph can be embedded, we have to verify the following unlinking property: it never happens that there are four external addresses $\underline{s}_1 < \underline{s}'_1 < \underline{s}_2 < \underline{s}'_2$ so that $\underline{s}_1$ and $\underline{s}_2$ have the same itinerary $\underline{u}$, and also $\underline{s}'_1$ and $\underline{s}'_2$ have the same itinerary $\underline{u}' \neq \underline{u}$. Suppose by contradiction that this problem does occur. Then (possibly after replacing all four addresses with the same shift), we may assume that $u_1 = u'_1$ but $u_2 \neq u'_2$, where $\underline{u} = u_1 u_2 u_3 \dots$ and $\underline{u}' = u'_1 u'_2 u'_3 \dots$. Without loss of generality, assume that $u_2 \neq 0$ (the argument is symmetric in $\underline{u}$ and $\underline{u}'$). Then $\sigma^{-1}(I_{u_2}) \cap I_{u_1}$ consists of a single interval which must contain $\underline{s}_1$ and $\underline{s}_2$, but it cannot contain $\underline{s}'_1$ or $\underline{s}'_2$. This proves the unlinking property.

After preparing our graph so that no Levy cycle can emerge, we embed Γ first into $\overline{\mathbb{C}}$ and then into $\mathbb{S}^2$. The complex structure on $\overline{\mathbb{C}}$ does not play any role, but it allows us to describe the construction more easily.

Define the embedding $\psi: \Gamma \rightarrow \overline{\mathbb{C}}$ as an injective continuous map as follows:

- (1) First let $\psi(e_1) = 0$ and $\psi(e_\infty) = \infty$.
- (2) Since the ψ -image of γ_1 must be a curve from 0 to ∞ , let $\psi(\gamma_1) = \mathbb{R}^+$.
- (3) For every $n \in \mathbb{Z}$ let $\psi(p_n)$ be the straight horizontal line with imaginary part $(2n - 1)\pi$.

- (4) The images of the edges p_n of Γ define a partition of $\mathbb{C}$ into strips Δ_n . We label the strips in vertical order so that Δ_n denotes the strip containing $2n\pi i$ (so $0 \in \Delta_0$).
- (5) Now we are ready to bring in the combinatorial information stemming from the kneading sequence $\underline{u}$. For each $n = 1, 2, \dots, l+k$, let $\psi(e_n)$ be a point in strip $\Delta_{\underline{u}_n}$ (taking into account that certain points e_n might be identified) and let $\psi(\gamma_n)$ be a curve in $\Delta_{\underline{u}_n}$ connecting $\psi(e_n)$ to $+\infty$. For simplicity, assume that imaginary parts of $\psi(\gamma_n)$ are eventually constant. Choose these curves so that they are disjoint from each other and from all endpoints except their own, and choose the eventually constant imaginary parts so that $\text{Im}(\psi(e_n)) < \text{Im}(\psi(e_m))$ if and only if $\sigma^{n-1}(\underline{s}) < \sigma^{m-1}(\underline{s})$ for each $n \neq m$; this ensures that the rays respect the order prescribed by the external address.

We have to justify that this can be done consistently if several endpoints are identified: it can never happen that two curves γ_n and γ_m have a common endpoint and separate $\mathbb{C}$ into two complementary components which both contain a curve $\gamma_{n'}$ and $\gamma_{m'}$ that should have a common endpoint different from $e_n = e_m$. This is exactly the unlinking property.

We now map $\overline{\mathbb{C}}$ homeomorphically to $\mathbb{S}^2$, mapping $0, \infty \in \overline{\mathbb{C}}$ to the two corresponding marked points $0, \infty \in \mathbb{S}^2$. From now on, we view ψ as a map from Γ to $\mathbb{S}^2$ and denote $\Gamma_{\mathbb{S}^2} := \psi(\Gamma) \subset \mathbb{S}^2$. For simplicity, we write e_n for $\psi(e_n)$, γ_n for $\psi(\gamma_n)$ and p_n for $\psi(p_n)$ (the vertices, the rays, and the partition boundaries). Since from now on, we only work with the embedded graphs, no confusion can arise.

Remark. In the construction of $\Gamma_{\mathbb{S}^2}$, the only combinatorial information coming from $\underline{s}$ are the kneading sequence $\underline{u}$ and the lexicographic (vertical) order of the set of external addresses: $\{\underline{s}, \sigma(\underline{s}), \dots, \sigma^{\circ(l+k-1)}(\underline{s})\}$. We will see in Algorithm 6.3 that this information gives us back a unique external address $\underline{s}$, normalized so that its first entry is 0.

We can now define a graph map $\tilde{f}: \Gamma_{\mathbb{S}^2} \rightarrow \Gamma_{\mathbb{S}^2}$ such that $\tilde{f}(\infty) = \infty$, as well as $\tilde{f}(e_j) = e_{j+1}$ and $\tilde{f}(\gamma_j) = \gamma_{j+1}$ for all j (counting indices modulo the period, so that $\tilde{f}(\gamma_{l+k}) = \gamma_{l+1}$). Furthermore, for all $k \in \mathbb{Z}$, define $\tilde{f}(p_k) = \gamma_1$.

Observe that $\tilde{f}$ is neither surjective nor injective, and that any two glued points have glued images, so the graph map respects the gluing. Under the map $\tilde{f}$, the orbit of 0 is necessarily preperiodic. The graph map $\tilde{f}$ is continuous everywhere except at ∞ ; this will not affect the extended map f , which is defined on $\mathbb{S}^2 \setminus \{\infty\}$.

5.2. Extension of the graph map. In order to prove Theorem 5.4, we need to do three things: we need to show that the graph map can be extended to a topological exponential map (Lemma 5.9), we need to show that it satisfies the conditions given in the theorem (Lemma 5.7), and we need to prove the uniqueness claims (Proposition 5.10). The second part is the easiest, so we do it first.

Lemma 5.7. (Every embedding satisfies conditions on rays)

If $f: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{0, \infty\}$ is a topological exponential map which extends $\tilde{f}: \Gamma_{\mathbb{S}^2} \rightarrow \Gamma_{\mathbb{S}^2}$, then f satisfies the following properties:

- the curve γ_1 connects 0 to ∞ and is preperiodic under f ;
- with respect to the dynamic partition induced by preimages of γ_1 , the curve γ_1 has itinerary $\underline{s}$;
- the vertical order of the rays $\gamma_n = f^{\circ(n-1)}(\gamma_1)$ coincides with the lexicographic order of the $\sigma^{n-1}(\underline{s})$;
- f does not admit a Levy cycle.

Proof. The first three properties follow directly from the construction. The last property follows from Corollary 5.6 because f can have a Levy cycle only if there are two different points e_i, e_j with identical itineraries, but such points have been glued together. $\square$

The next step is the extension of the graph map to a topological exponential map. In order to ensure that the graph map can be extended to a neighborhood of every vertex, we need the following lemma.

Lemma 5.8. ($\tilde{f}$ preserves cyclic order)

The cyclic order at ∞ of three or more rays with the same itinerary (landing at the same point) is preserved by the graph map $\tilde{f}$.

Proof. Suppose the endpoints $e_n, e_{n'}$ and $e_{n''}$ are identified, so that the three rays $\gamma_n, \gamma_{n'}, \gamma_{n''}$ land at a common point. The vertical order of these rays coincides with the lexicographic order of the addresses $\sigma^{n-1}(\underline{s}), \sigma^{n'-1}(\underline{s})$ and $\sigma^{n''-1}(\underline{s})$; suppose without loss of generality that $\sigma^{n-1}(\underline{s}) < \sigma^{n'-1}(\underline{s}) < \sigma^{n''-1}(\underline{s})$. Since the first entries in the itineraries of these three external addresses coincide, there is a $k \in \mathbb{Z}$ so that

$$k\underline{s} < \sigma^{n-1}(\underline{s}) < \sigma^{n'-1}(\underline{s}) < \sigma^{n''-1}(\underline{s}) < (k+1)\underline{s}$$

(where $k\underline{s}$ and $(k+1)\underline{s}$ denote adjacent preimages of $\underline{s}$ under the shift). But on every interval $(k\underline{s}, (k+1)\underline{s})$, the shift map is injective and preserves the cyclic order. $\square$

One key construction is Alexander's trick (compare e.g. [2]): if $f: \mathbb{S}^1 \rightarrow \mathbb{S}^1$ is an orientation-preserving homeomorphism, then there exists an orientation-preserving homeomorphism $\hat{f}: \overline{\mathbb{D}} \rightarrow \overline{\mathbb{D}}$ such that $\hat{f}|_{\partial\overline{\mathbb{D}}} = \hat{f}|_{\mathbb{S}^1} = f$: one such extension is given by $\hat{f}(re^{i\theta}) := rf(e^{i\theta})$.

Moreover if $f, g: \mathbb{S}^1 \rightarrow \mathbb{S}^1$ are isotopic rel some finite number of points in $\mathbb{S}^1$, then by extending the isotopy to the entire disk one gets an isotopy (rel the same points) between the extensions $\hat{f}$ and $\hat{g}: \overline{\mathbb{D}} \rightarrow \overline{\mathbb{D}}$.

Lemma 5.9. (Extension of graph map)

The graph map $\tilde{f}: \Gamma_{\mathbb{S}^2} \rightarrow \Gamma_{\mathbb{S}^2}$ can be extended to a topological exponential map $f: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{0, \infty\}$.

Proof. For convenience, let us adopt the convention that the edges p_n and γ_m contain their endpoints ∞ and e_m . Notice that for $n \in \mathbb{Z}$, the set $\mathbb{S}^2 \setminus (p_n \cup p_{n+1})$ is the union of three disjoint open topological disks; among them, Δ_n is the disk which does not intersect any p_k for $k \notin \{n, n+1\}$.

We shall first construct continuous maps $f_n: (\overline{\Delta_n} \setminus \{\infty\}) \rightarrow (\mathbb{S}^2 \setminus \{\infty\})$ for each n , such that $f_n(p_n) = f_n(p_{n+1}) = \gamma_1$ as above, and $f_n(\Gamma_{\mathbb{S}^2} \cap \Delta_n) = \tilde{f}(\Gamma_{\mathbb{S}^2} \cap \Delta_n)$. Moreover, the restriction of f_n to Δ_n is an orientation preserving homeomorphism onto $\mathbb{S}^2 \setminus \gamma_1$ which extends continuously to the boundary.

We distinguish four possible cases for Δ_n :

- (i) $\Delta_n \cap \Gamma_{\mathbb{S}^2} = \emptyset$;
- (ii) $\Delta_n \cap \Gamma_{\mathbb{S}^2} = \gamma_m$ for some $m \in \{1, 2, \dots, l+k\}$;
- (iii) $\Delta_n \cap \Gamma_{\mathbb{S}^2} = \gamma_{m_1} \cup \dots \cup \gamma_{m_i}$, for some $m_1, \dots, m_i \in \{1, 2, \dots, l+k\}$ such that all endpoints $e_{m_1}, \dots, e_{m_i}$ are glued;
- (iv) The general case: there may be combinations of several instances of case (ii) and (iii) on one domain Δ_n .

Case (i) is almost literally Alexander's trick: the graph map $\tilde{f}$ prescribes the boundary values on the topological disk Δ_n for the map f_n . The only problem is that the point ∞ appears twice on the boundary, mapping to 0 on the left and to ∞ on the right. This causes no problem (we do not define f_n or f on ∞).

In Case (ii) the idea of the construction is the same; we need to find a homeomorphism $f: \Delta_n \setminus \gamma_m \rightarrow \mathbb{S}^2 \setminus \gamma_{m+1}$ which coincides with the prescribed boundary values on p_n, p_{n+1}, γ_m and ∞ (at two sides). Note that $\Delta_n \setminus \gamma_m \approx \mathbb{S}^2 \setminus \gamma_{m+1} \approx \mathbb{D}$. In this case, the curve γ_m along with its endpoints occurs twice on the boundary, and f_n is already defined on γ_m via $\tilde{f}$.

Now we treat Case (iii). The set $\Delta_n \setminus (\gamma_{m_1} \cup \dots \cup \gamma_{m_i})$ consists of i domains, each of which is homeomorphic to $\mathbb{D}$; the same is true for $\mathbb{S}^2 \setminus (\gamma_1 \cup \gamma_{m_1+1} \cup \dots \cup \gamma_{m_i+1})$ (note that the endpoint $e_1 = 0$ is never glued with any other endpoint because its itinerary has longer preperiod than all others). As in Case (i), we want to extend the map f along its prescribed boundary values to the appropriate image domains. In order for this to be possible, we need to assure that for each of the i topological disks, the boundary rays map to the boundary rays of an appropriate image domain. Let a be the landing point of the rays $\gamma_{m_1}, \dots, \gamma_{m_i}$; then we need to make sure that the cyclic order of these rays at a coincides with the cyclic order of the image rays at $f(a)$. Our construction lets rays land together iff they have identical itineraries, so this fact is assured by Lemma 5.8. Therefore, the map f_n can be defined on Δ_n as well, and it is again a homeomorphism $f_n: \Delta_n \rightarrow \mathbb{S}^2 \setminus \gamma_1$.

Finally, the general case (iv) can incorporate several rays, and several groups of rays, within the same domain Δ_n . First observe that the vertical order of all rays $\gamma_{m_i} \subset \Delta_n$ and $p_n, p_{n+1} \subset \partial\Delta_n$ is compatible with the cyclic order of the image rays near infinity: the vertical order of the rays is determined by their external addresses, and all external addresses within one domain Δ_n belong to one interval on which the shift map is injective. Removing from Δ_n all rays, we obtain finitely many connected components which can be treated separately because they have compatible boundary values. For every connected component, the claim follows by combining the ideas from the previous cases.

Having constructed the maps f_n on all strips, we define a continuous map $f: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{0, \infty\}$, by $f|_{\Delta_n} := f_n$. The values on the boundaries of the Δ_n (except at ∞) match since here they coincide with the graph map $\tilde{f}$. We obtain a universal cover $f: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{0, \infty\}$. $\square$

5.3. Thurston-uniqueness. Of course, the construction of the extension is not unique; however, we have a uniqueness result in the following sense.

Proposition 5.10. (Thurston-uniqueness of topological exponential map) *Suppose that $f_1, f_2: (\mathbb{S}^2 \setminus \{\infty\}) \rightarrow (\mathbb{S}^2 \setminus \{0, \infty\})$ are topological exponential maps satisfying the four itemized properties of Theorem 5.4 for the preperiodic external address $\underline{s}$. Then f_1 and f_2 are postsingularly finite and Thurston equivalent.*

We start with the following lemma.

Lemma 5.11. (Uniqueness of graph map up to homotopy)

Suppose that $f: (\mathbb{S}^2 \setminus \{\infty\}) \rightarrow (\mathbb{S}^2 \setminus \{0, \infty\})$ is a postsingularly finite topological exponential map satisfying the four itemized properties of Theorem 5.4 for the preperiodic external address $\underline{s}$. Then the restriction of f to the collection of curves $\bigcup_{n \geq -1} f^{on}(\gamma)$ yields a graph map $\Gamma_{\mathbb{S}^2}$ which is homotopic relative to P_f to any graph map as constructed above.

With $\bigcup_{n \geq -1} f^{on}(\gamma)$, we mean the countably many preimage curves $f^{-1}(\gamma)$, together with the finitely many curves on the forward orbit of γ .

Proof. First we find a homotopy of $\mathbb{S}^2$ relative P_f sending the preimages $f^{-1}(\gamma)$ to the edges p_n of $\Gamma_{\mathbb{S}^2}$; this is possible because the points in P_f have prescribed itineraries. Then we can find homotopies within each complementary domain Δ_n of $\mathbb{S}^2 \setminus \bigcup_n p_n$ to match the graphs within each Δ_n ; this is possible because:

- the itinerary prescribes which legs γ_m are within which Δ_n ,
- the order at ∞ of the different legs within the same Δ_n is in both cases prescribed by the lexicographic order of appropriate shifts of $\underline{s}$,
- the same endpoints of legs are identified because of the non-existence of Levy cycles: Corollary 5.6 says which endpoints are identified in the topological exponential map, and the same endpoints are identified by construction in our graph map.

$\square$

Lemma 5.12. (Uniqueness of graph map extension up to homotopy)

Let $\tilde{f}: \Gamma_{\mathbb{S}^2} \rightarrow \Gamma_{\mathbb{S}^2}$ be a graph map and $f_1, f_2: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{0, \infty\}$ two different graph map extensions as in Lemma 5.9. Then the maps f_1 and f_2 are Thurston equivalent.

Proof. As in [2], we will define a map $\varphi: \mathbb{S}^2 \rightarrow \mathbb{S}^2$ such that $f_1 = f_2 \circ \varphi$ and φ is a homeomorphism homotopic to the identity rel the vertices of $\Gamma_{\mathbb{S}^2}$. The construction of φ is easy. Each component U of $\mathbb{S}^2 \setminus \Gamma_{\mathbb{S}^2}$ is homeomorphic to a disk, and by construction f_1 and f_2 coincide on their boundaries. Since both f_i restricted to these disks are homeomorphisms onto their images, they are homotopic to each other relative to the boundary. Since the postsingular set is contained in $\Gamma_{\mathbb{S}^2}$, f_1 and f_2 are Thurston equivalent in the sense of Definition 5.2. $\square$

Proof of Proposition 5.10. Since the curve γ is preperiodic as a set, and one of its endpoints is the singular value 0, it follows that f_1 and f_2 are postsingularly finite. By Lemma 5.11, the graph maps of f_1 and f_2

coincide up to homotopy relative to the postsingular set; the homotopy of graph maps extends to a homotopy of $\mathbb{S}^2$ because the graph maps are embedded into $\mathbb{S}^2$ in the same way, and by Lemma 5.12, we obtain a Thurston equivalence between f_1 and f_2 . $\square$

Proof of Theorem 5.4. We have first constructed a graph $\Gamma_{\mathbb{S}^2} \subset \mathbb{S}^2$ and a graph map $\tilde{f}: \Gamma_{\mathbb{S}^2} \rightarrow \Gamma_{\mathbb{S}^2}$, and we have then extended the graph map to a postsingularly finite topological exponential map (Lemma 5.9) which satisfies the four itemized properties in the theorem (Lemma 5.7). Finally, Proposition 5.10 shows that all topological exponential maps satisfying these conditions are Thurston equivalent. $\square$

6. HOLOMORPHIC EXPONENTIAL MAPS

In the previous section, we have constructed for every preperiodic external address a topological exponential map satisfying the properties of Theorem 5.4. In this section, we are going to find a holomorphic exponential map with the same properties. It is here that we use the main result of [14] stated below.

Theorem 6.1. (Characterization of exponential maps)

A postsingularly finite topological exponential map is Thurston equivalent to a (necessarily unique) postsingularly finite holomorphic exponential map if and only if it does not admit a Levy cycle.

Since the topological exponential map constructed in Theorem 5.4 has no Levy cycle, we have now associated to every preperiodic external address a unique postsingularly finite holomorphic exponential map. In order to close the loop of the argument (see the diagram at the end of Section 2), we need the following result.

Proposition 6.2. (Dynamic ray lands at singular value)

For every preperiodic external address $\underline{s}$, suppose that E_λ is a holomorphic exponential map which is Thurston equivalent to the topological exponential map constructed in Theorem 5.4 for external address $\underline{s}$. Then for E_λ , the dynamic ray at external address $\underline{s}$ lands at the singular value.

Proof. We prove this claim by translating Thurston equivalence into the language of spiders, and using results from [26].

For a postsingularly finite topological exponential map f , a *spider leg* is an injective curve $\gamma_i: [0, \infty] \rightarrow \mathbb{S}^2$ with $\gamma_i(0) = e_i$ and $\gamma_i(\infty) = \infty$, where $e_1 = 0, e_2, \dots, e_{l+k}$ is the postsingular orbit. A *spider* is a collection of spider legs $\Phi = \{\gamma_1, \dots, \gamma_{l+k}\}$ which are disjoint except possibly for their endpoints. Two spiders are *equivalent* if there is an isotopy of $\mathbb{S}^2$ relative to $P_f = \{e_1, \dots, e_{l+k}, \infty\}$ which moves one spider to the other.

Every spider Φ lifts under f to an image spider $\tilde{\Phi} = \{\tilde{\gamma}_1, \dots, \tilde{\gamma}_{l+k}\}$, where each $\tilde{\gamma}_i$ is the unique component of $f^{-1}(\gamma_{i+1})$ starting at e_i (counting indices modulo the period as always). It is easy to check that equivalent spiders have equivalent image spiders, so the spider map acts on equivalence classes of spiders.

If γ_1 is a curve connecting $e_1 = 0$ to ∞ and is preperiodic under f (as a curve) with preperiod l and period k , then the $l + k$ legs on the orbit of γ_1 obviously form an invariant spider.

Now suppose that f and E_λ are Thurston equivalent: $\phi_1 \circ f = E_\lambda \circ \phi_2$, where ϕ_1 and ϕ_2 are isotopic relative to P_f . Then $\Psi := \phi_1 \circ \Phi$ is a spider for E_λ . We can lift Ψ under E_λ to an image spider $\tilde{\Psi}$, and the condition $\phi_1 \circ f = E_\lambda \circ \phi_2$ implies that $\tilde{\Psi} = \phi_2 \circ \tilde{\Phi}$ (lifting spiders is compatible with Thurston equivalences). Since ϕ_1 and ϕ_2 are isotopic relative to P_f , it follows that the spiders Ψ and $\tilde{\Psi}$ are equivalent: the map E_λ has a fixed spider (up to equivalence) which is related by the Thurston equivalence to the preperiodic curves γ_i of f .

Now [26, Theorem 6.4] shows that this fixed spider of E_λ can be replaced by an equivalent spider consisting only of dynamic rays (this theorem is part of the proof that every postsingularly finite exponential map has a dynamic ray landing at the singular value; the statement reads a bit differently in that context, but what the theorem actually does is to take a periodic spider given by [26, Proposition 6.3] and turn it into a spider made of dynamic rays).

We now know that a dynamic ray landing at 0 has kneading sequence $K(\underline{s}) = \text{It}(\underline{s} \mid \underline{s})$, and the order of the image rays is as constructed in the topological case. At this point, Algorithm 6.3 below, comes in handy: it allows us to reconstruct $\underline{s}$ uniquely from its kneading sequence $\mathbf{u} = \text{It}(\underline{s} \mid \underline{s})$ and the order of the $\sigma^n(\underline{s})$. It proves that for the holomorphic map E_λ the dynamic ray at external address $\underline{s}$ lands at 0. $\square$

Algorithm 6.3. (Reversal of construction of kneading sequence)

Suppose we are given the kneading sequence $\mathbf{u}$ of a non-constant, non-periodic external address $\underline{s}$ and for all $n \geq 1$ we know the relative order of $\underline{s}$ and $\sigma^n(\underline{s})$. Then the following algorithm uniquely recovers $\underline{s}$:

- (1) For $n \geq 1$ define $\delta_n := \begin{cases} -1 & \text{if } \sigma(\underline{s}) > \underline{s} \text{ and } \sigma^{on}(\underline{s}) < \underline{s} \\ 1 & \text{if } \sigma(\underline{s}) < \underline{s} \text{ and } \sigma^{on}(\underline{s}) > \underline{s} \\ 0 & \text{otherwise.} \end{cases}$
- (2) Construct $\underline{s}$ as $s_n = u_n - \delta_n$.

Proof. Since the relative order of $\sigma^{on}(\underline{s})$ and $\underline{s}$ is known for all n , we can compute δ_n for all n . We can thus reverse the computation in Algorithm 4.3. $\square$

Now we can finish the proof of the first classification theorem.

Proof of Theorem 2.6. We have just finished the proof of the existence part of the theorem. For uniqueness, suppose there are two postsingularly finite exponential maps E_λ and $E_{\lambda'}$ which both have the property that their dynamic rays at external address $\underline{s}$ lands at the singular value. Then both maps have spiders consisting of this dynamic ray and its forward images, and there is a homeomorphism $\phi_1: \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$ which sends the spider of E_λ to the spider of $E_{\lambda'}$. Since $\phi_1(0) = 0$ and both maps are exponential maps with asymptotic value 0, it follows that ϕ_1 lifts to another homeomorphism $\phi_2: \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$ so that $\phi_1 \circ E_\lambda = E_{\lambda'} \circ \phi_2$. The spiders assure that ϕ_1 and ϕ_2 are homotopic to each other relative to the postsingular set, so E_λ and $E_{\lambda'}$ are Thurston equivalent. By Theorem 6.1, it follows that $\lambda = \lambda'$.

Finally, the fact that every postsingularly finite exponential map is actually associated to a finite positive number of preperiodic external addresses is Theorem 3.1. $\square$

A core ingredient in this proof was the combinatorial Algorithm 6.3 which allows to recover the external address $\underline{s}$ from its kneading sequence and the lexicographic order of the orbit of $\underline{s}$ under the shift. There are other ways to recover $\underline{s}$ which are perhaps more closely related to the dynamics of f ; see for example the proof of [26, Theorem 6.4].

We have now proved Theorem 2.6: for every preperiodic external address $\underline{s}$, there is a unique postsingularly finite exponential map for which the dynamic ray at external address $\underline{s}$ lands at the singular value, and this describes all postsingularly finite exponential maps.

It remains to determine combinatorially which external addresses yield the same exponential map; this is Theorem 2.7.

Proof of Theorem 2.7. Suppose $\underline{s}$ and $\underline{s}'$ give rise to the same exponential map E_λ : by definition, the dynamic rays at addresses $\underline{s}$ and $\underline{s}'$ for E_λ land together at 0, and hence they have the same itinerary with respect to either dynamic partition, hence $\text{It}(\underline{s} \mid \underline{s}') = \text{It}(\underline{s}' \mid \underline{s}')$ and $\text{It}(\underline{s}' \mid \underline{s}) = \text{It}(\underline{s} \mid \underline{s})$. It also follows that $\underline{s}$ and $\underline{s}'$ have the same period and the same preperiod.

Conversely, fix two preperiodic external addresses $\underline{s}, \underline{s}'$ of equal period and equal preperiod such that $\text{It}(\underline{s} \mid \underline{s}') = \text{It}(\underline{s}' \mid \underline{s}')$. Construct the holomorphic exponential map E_λ corresponding to $\underline{s}'$ using Theorem 2.6. E_λ has a dynamic ray at external address $\underline{s}$, and since $\text{It}(\underline{s} \mid \underline{s}') = \text{It}(\underline{s}' \mid \underline{s}')$, the rays at addresses $\underline{s}$ and $\underline{s}'$ land together at the same point 0. By uniqueness, Theorem 2.6 constructs the same map from $\underline{s}$. Hence, both addresses correspond to the same holomorphic exponential map.

This shows the equivalence of Conditions (1), (3), and (4) in the claim. The equivalence of Conditions (1) and (2) is Theorem 3.4. Before dealing with the remaining conditions, let us discuss the last statements of the theorem.

Dynamic rays landing at a common point always have equal period and equal preperiod. The construction of kneading sequences makes it clear that if $\underline{s}$ has preperiod l and period k , then $\text{It}(\underline{s} \mid \underline{s})$ has preperiod l and

period k' dividing k . Among the external addresses $\underline{s}, \sigma(\underline{s}), \dots, \sigma^{k+l}(\underline{s})$, only the periodic ones can have equal itineraries with respect to $\underline{s}$, and the number of those who do is obviously equal to k/k' . Therefore, the corresponding dynamic rays land in groups of k/k' ; by [21, Lemma 5.2], there can be no additional rays if $k > k'$, while if $k = k'$, then the number of rays can be 1 or 2, and both cases actually occur. Pulling back, this gives the number of preperiodic dynamic rays landing at the singular value.

Condition (5) implies the previous ones, so we now show that if $\underline{s}$ and $\underline{s}'$ generate the same postsingularly finite exponential map E_λ , then $\text{It}(\underline{s} \mid \underline{s}') = \text{It}(\underline{s} \mid \underline{s})$. For every $j \geq 0$, the dynamic rays $g_{\sigma^j(\underline{s})}$ and $g_{\sigma^j(\underline{s}')}$ land together and bound a component $U_j \subset \mathbb{C}$ with real parts bounded below. If $\text{It}(\underline{s} \mid \underline{s}')$ and $\text{It}(\underline{s} \mid \underline{s})$ differ in their k -th entries, this implies that $E_\lambda^{\circ k}(g_{\underline{s}}) \subset U_0$ and, since both rays land together, also $E_\lambda^{\circ k}(g'_{\underline{s}}) \subset U_0$ and $U_k \subset U_0$. We show that this is impossible.

There is a unique $k' \in \{0, 1, \dots, k-1\}$ so that the first entry in $\sigma^j(\underline{s})$ is the same as in $\sigma^j(\underline{s}')$, for every $j = k' + 1, \dots, k-2$, but not for $j = k'$ (if there was no such k' , then the rays $E_\lambda^{\circ k}(g_{\underline{s}})$ and $E_\lambda^{\circ k}(g'_{\underline{s}})$ would be further apart than the rays $g_{\underline{s}}$ and $g'_{\underline{s}}$, and they could not be contained in U_0).

Then for $j = k' + 1, \dots, k-2$, the restriction $E_\lambda: U_j \rightarrow U_{j+1}$ is a conformal isomorphism, and so is $E_\lambda: U_{k'} \rightarrow \mathbb{C} \setminus \overline{U}_{k'+1}$. It follows that $U_{k'+1}$ contains the singular value 0 and hence the rays $g_{\underline{s}}$ and $g'_{\underline{s}}$, hence $U_{k'+1} \supset U_0 \supset U_k$ (equality excluded). However, since the rays bounding $U_{k'}$ have more identical first entries in their external addresses than the rays bounding U_k , they must surround a smaller domain, and this is a contradiction. $\square$

REFERENCES

- [1] (0752391) Noel I. Baker and Philip Rippon, *Iteration of exponential functions*, Ann. Acad. Sci. Fenn. Ser. A I Math., **9** (1984), 49–77.
- [2] (1149891) Ben Bielefeld, Yuval Fisher and John Hubbard, *The classification of critically preperiodic polynomials as dynamical systems*, Journal of the AMS, **5** (1992), 721–762.
- [3] (1785056) Clara Bodelón, Robert Devaney, Michael Hayes, Gareth Roberts, Lisa Goldberg, John Hubbard, *Dynamical convergence of polynomials to the exponential*, J. Differ. Equations Appl., **6** (2000), 275–307.
- [4] Henk Bruin, Alexandra Kaffl and Dierk Schleicher, *Symbolic dynamics of quadratic polynomials*, monograph in preparation. Earlier version appeared as: Preprint No. 7, Institut Mittag-Leffler, 2001–2002, ISSN 1103–467X.
- [5] (0762431) Adrien Douady and John Hubbard, *Etude dynamique des polynômes complexes*, Prépublications Mathématiques d’Orsay, **2** (1984), **4** (1985).
- [6] (1251582) Adrien Douady and John Hubbard, *A proof of Thurston’s topological characterization of rational functions*, Acta Mathematica, **171** (1993), 263–297.
- [7] Leonhardt Euler, *De formulis exponentialibus replicatis*, Acta Acad. Petropolitanae, **1** (1777), 38–60.
- [8] (1196102) Alexandre Eremenko and Mikhail Lyubich, *Dynamical properties of some classes of entire functions*, Ann. Inst. Fourier (Grenoble), **42** (1992), 989–1020.
- [9] Dominik Eberlein and Dierk Schleicher, *Rational parameter rays of Multibrot sets*, manuscript, in preparation.
- [10] (2358507) Markus Förster, Lasse Rempe and Dierk Schleicher, *Classification of escaping exponential maps*, Proc. Amer. Math. Soc., **136** (2008), 651–663.
- [11] Markus Förster and Dierk Schleicher, *Parameter rays for the exponential family*, Ergod. Theory. Dyn. Sys., to appear. ArXiv:0505097
- [12] (1209913) Lisa Goldberg and John Milnor, *Fixed points of polynomial maps II*, Ann. Sci. École Norm. Sup., **26** (1993), 51–98.
- [13] (1315537) John Hubbard and Dierk Schleicher, *The spider algorithm*, in “Complex Dynamical Systems,” (ed. Robert Devaney), AMS, (2004), 155–180.
- [14] John Hubbard, Dierk Schleicher and Mitsuhiro Shishikura, *Exponential Thurston maps and limits of quadratic differentials*, Journal of the AMS, to appear.
- [15] Eike Lau and Dierk Schleicher, *Internal addresses in the Mandelbrot set and irreducibility of polynomials*, preprint **19** (1994), Institute for Mathematical Sciences, SUNY StonyBrook.
- [16] Alfred Poirier, *On postcritically finite polynomials, part I/II*, preprint **5/7** (2004), Institute for Mathematical Sciences, SUNY StonyBrook. ArXiv:9305207 / ArXiv:9307235
- [17] (2213743) Lasse Rempe, *A landing theorem for dynamic rays of exponential maps*, Proc. Amer. Math. Soc., **134** (2006), 2639–2648.
- [18] Lasse Rempe and Dierk Schleicher, *Combinatorics of bifurcations in exponential parameter space*, in “Transcendental Dynamics and Complex Analysis” (eds. Phil J. Rippon and Gwyneth M. Stallard), London Mathematical Society Lecture Note Series **348**, 317–370 (to appear). ArXiv:0408011

- [19] Lasse Rempe and Dierk Schleicher, *Bifurcations in the space of exponential maps*, preprint **3** (2004), Institute for Mathematical Sciences, SUNY StonyBrook. ArXiv:0311480
- [20] Dierk Schleicher, “On the Dynamics of Exponential Maps,” Habilitation thesis, Technische Universität München, (1999).
- [21] (1976827) Dierk Schleicher, *Attracting dynamics of exponential maps*, *Annales Academiæ Scientiarum Fennicæ*, **28** (2003), 3–34.
- [22] (2078079) Dierk Schleicher, *Hyperbolic components in exponential parameter space*, *C.R. Math. Acad. Sci. Paris*, **339** (2004), 223–228.
- [23] (2112117) Dierk Schleicher, *On fibers and local connectivity of Mandelbrot and Multibrot sets*, in “Fractal Geometry and Applications: A Jubilee of Benoît Mandelbrot, Part 1” (eds. Michel Lapidus and Machiel van Frankenhuysen), *Proceedings of Symposia in Pure Mathematics*, (2004), 477–517.
- [24] Dierk Schleicher, *Internal Addresses in the Mandelbrot set and Galois groups of polynomials*, submitted. ArXiv:9411238
- [25] (1956142) Dierk Schleicher and Johannes Zimmer, *Escaping points of exponential maps*, *Journal of the London Mathematical Society*, **67** (2003), 380–400.
- [26] (1996442) Dierk Schleicher and Johannes Zimmer, *Periodic points and dynamic rays of exponential maps*, *Annales Academiæ Scientiarum Fennicæ*, **28** (2003), 327–354.

E-mail address: dierk@jacobs-university.de

E-mail address: laubner@informatik.hu-berlin.de

E-mail address: vicol@usc.edu