Dynamics of Exponential Maps

Dissertation

zur Erlangung des Doktorgrades
der Mathematisch-Naturwissenschaftlichen Fakultät
der Christian-Albrechts-Universität
zu Kiel

vorgelegt von

Lasse Rempe

Kiel

Mai 2003

Zusammenfassung

Diese Arbeit enthält verschiedene neue Resultate in der dynamischen Untersuchung der Exponentialfamilie $z \mapsto \exp(z) + \kappa$. Wir geben eine neue Beschreibung der Menge der entkommenden Punkte einer beliebigen Exponentialabbildung, die es erlaubt, Aussagen über deren Topologie zu beweisen. Unter anderem geben wir eine Antwort auf die Frage, welche externen Strahlen mit entkommenden Endpunkten in diesen Endpunkten differenzierbar sind, und zeigen, daß entkommende Punkte beliebig langsam entkommen können. Desweiteren beweisen wir mehrere Resultate über die Existenz von nichtlandenden Strahlen und Starrheit der Dynamik von Exponentialfunktionen auf den Mengen ihrer entkommenden Punkte.

Ferner zeigen wir, daß je zwei hyperbolische oder parabolische Exponentialfunktionen auf den Mengen ihrer entkommenden Punkte konjugiert sind. Genauer geben wir eine Beschreibung der Juliamenge einer solchen Funktion als Quotient der Juliamenge einer Exponentialfunktion mit einem anziehenden Fixpunkt. Dies ist ein Analogon zum "Pinched Disk"-Modell für Polynomfunktionen und erlaubt die Beschreibung der Dynamik dieser Funktionen anhand ihrer Kombinatorik.

Wir geben desweiteren eine vollständige Beschreibung der Bifurkationsstruktur von hyperbolischen Komponenten in der Exponentialfamilie und geben einen vereinfachten Beweis des Satzes von Schleicher, daß der Rand dieser Komponenten in $\mathbb{C}$ zusammenhängend ist, wie von Baker und Rippon [5] sowie Eremenko und Lyubich [30] vermutet wurde. (Dieser Teil der Dissertation ist eine gemeinsame Arbeit mit Dierk Schleicher.)

Ferner beweisen wir, daß periodische externe Strahlen von Exponentialabbildungen mit nichtentkommendem singulärem Wert stets landen. Dies ist ein Analogon eines Satzes von Douady und Hubbard für Polynome. Wir beweisen desweiteren, daß die Punkte höchstens eines periodischer Zykels einer solchen Exponentialfunktion nicht Landepunkte periodischer externen Strahlen sind.

Wir zeigen außerdem, daß der Rand von unbeschränkten Siegelscheiben dieser Familie immer den singulären Wert enthält; dies beantwortet eine Frage von Herman, Baker und Rippon in [13].

Die Arbeit verusucht desweiteren, eine Übersicht über den aktuellen Stand des Wissens auf dem Gebiet der Iteration von Exponentialfunktionen zu geben.

Abstract

This thesis contains several new results about the dynamics of exponential maps $z \mapsto \exp(z) + \kappa$. We give a new description of the set of escaping points of an arbitrary exponential map which allows us to prove statements about the topology of this set. Using this description, we give an answer to the question which external rays with escaping endpoints are differentiable in these endpoints and show that orbits may escape arbitrarily slowly. We also prove several results about the existence of nonlanding external rays and rigidity of exponential maps on their sets of escaping points.

Furthermore we show that any two hyperbolic or parabolic exponential maps are topologically conjugate on their sets of escaping points. More precisely, we give an analog of the "pinched disk model" for polynomials by describing the Julia set of any attracting or parabolic exponential map as the quotient of that of an exponential map with an attracting fixed point. This allows the description of the dynamics of such parameters purely in terms of their combinatorics.

We also give a complete description of the bifurcation structure of hyperbolic components in the space of exponential maps. This yields a simplified proof of Schleicher's theorem that the boundary of such a component in $\mathbb{C}$ is a curve, as was conjectured by Baker and Rippon [5] as well as Eremenko and Lyubich [30]. (This part of the thesis is joint work with Dierk Schleicher.)

Furthermore, we prove that periodic external rays of exponential maps with nonescaping singular value always land. This is an analog of a well-known theorem of Douady and Hubbard for polynomials. We also show that there is at most one periodic cycle of such a map whose points are not landing points of periodic external rays.

Finally, we show that the boundary of unbounded Siegel disks always contains the singular value. This answers a question of Herman, Baker and Rippon stated in [13]

In addition to the presentation of these new results, the thesis also aims to give an overview of the current state of knowledge on the dynamics of exponential maps.

Contents

List of Figures

Chapter 1

Introduction

This thesis studies the dynamics of exponential maps $E_\kappa : z \mapsto \exp(z) + \kappa$ and the structure of their parameter space. This simplest family among transcendental entire functions has received special attention over the years, much like the quadratic family has among polynomials, and it is hoped that an understanding of exponential dynamics will be useful in the study of more general classes. We wish to emphasize that such an understanding is important not only in its own right, but also because the iteration of transcendental function has links to many other areas in dynamical systems and function theory; we content ourselves here with giving three examples. Features of exponential dynamics appear in the study of parabolic implosion [79], which is one of the most prominent current topics of polynomial dynamics. The family $\lambda t e^{-t}$, a close relative of the exponential, is the second simplest model in population dynamics (the first being the logistic family). Finally, anyone interested in finding the roots of entire functions should consider studying their Newton's method, i.e. iteration of a transcendental meromorphic function.

The reason that the exponential family is a natural candidate to begin with is the same that has made the quadratic family a favorite object of study: in both cases the maps possess only one *singular value*. The singular values (i.e., the critical and asymptotic values) of a function play an important role in the study of its dynamics: a restriction on the number of singular values generally limits the amount of different dynamical features that can appear for the same map. Therefore, the simplest non-trivial parameter space of holomorphic functions is given by the quadratic family, in which each function has only a single simple critical point. Exponential maps are the only transcendental functions which have only one singular value (see e.g. [57, Appendix D] or Theorem 2.3.5), and thus form the simplest family consisting of transcendental maps. Furthermore, the exponential family is the limit — not only analytically, but dynamically — of the families of unicritical polynomials, parametrized as $z \mapsto (1 + \frac{z}{d})^d + c$. This makes exponential maps an excellent candidate for applying methods that have proved useful in the study of these polynomials, as first developed for the Mandelbrot Set in Douady and Hubbard's famous Orsay Notes [27].

Note that parameter space of exponential maps has often been studied under the parametrization $z \mapsto \lambda \exp(z)$, rather than our form $z \mapsto \exp(z) + \kappa$; compare the re-

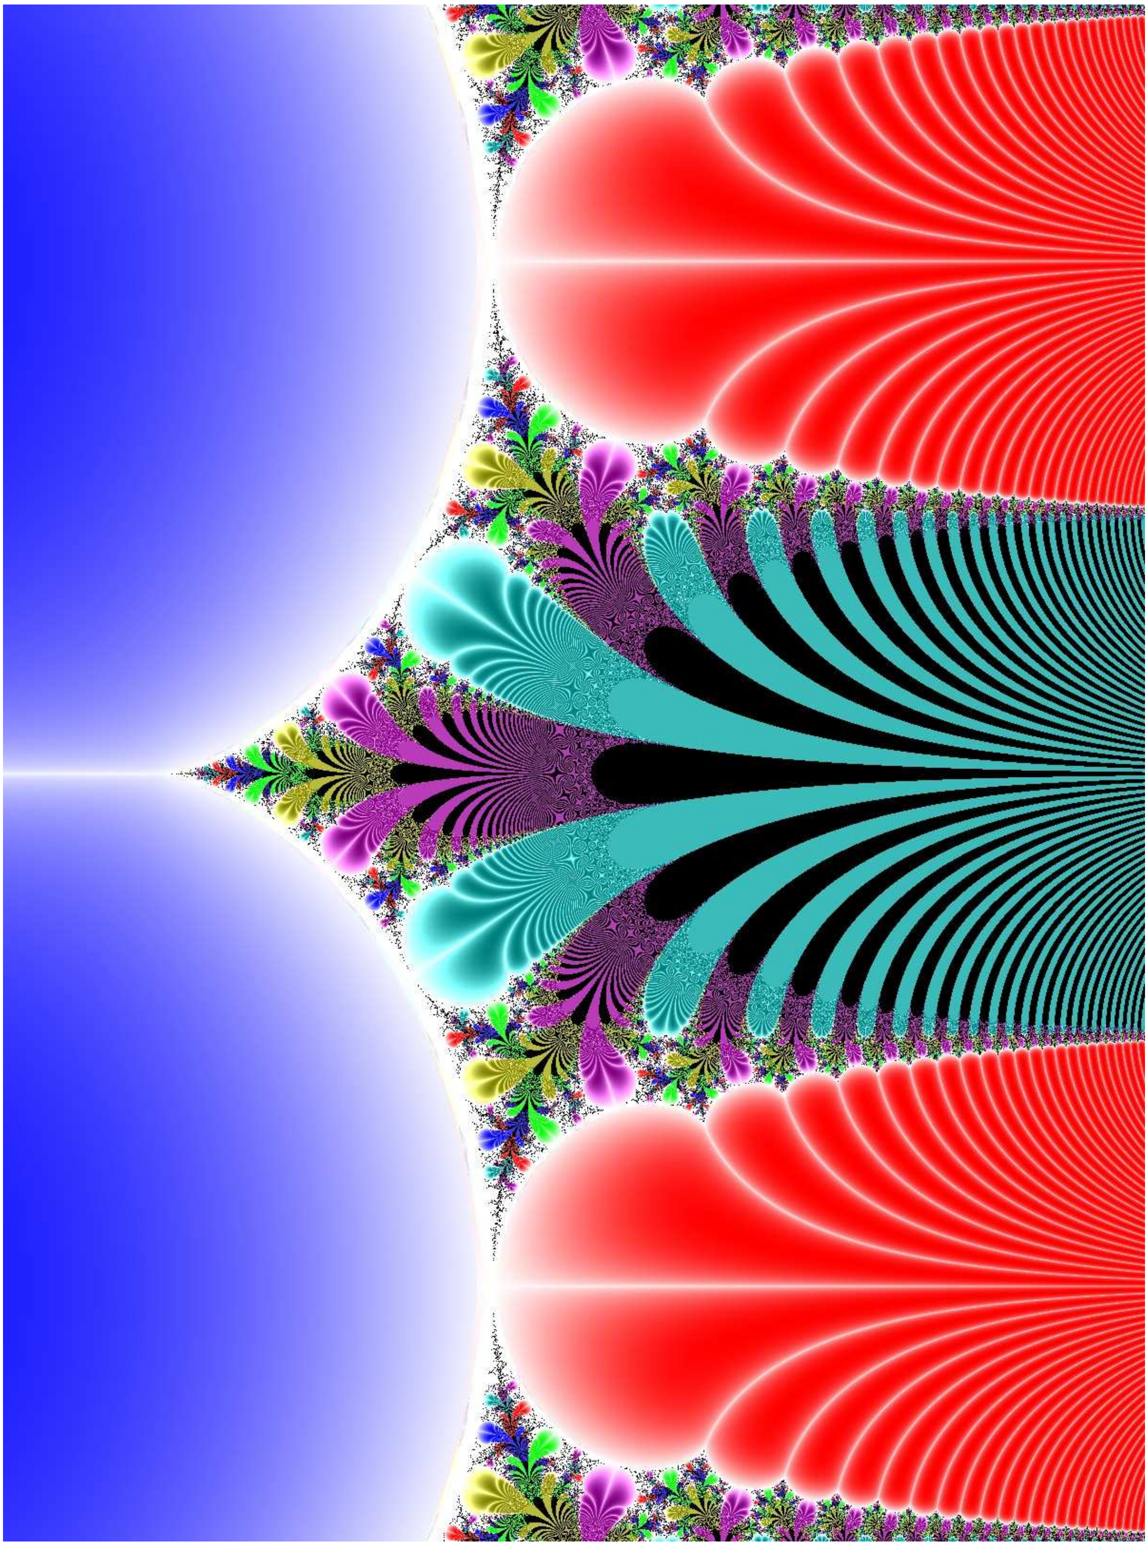

Figure 1.1: Parameter space for the exponential family.

marks at the end of Chapter 2. In the following, when quoting established results, we translate them into our parametrization.

In some sense, exponential dynamics was first studied by Euler [32], who determined for which real b the sequence $b, b^b, b^{b^b}, \dots$ converges. However, the study of iterated entire transcendental functions truly began with Fatou's mémoire [34], in which he conjectured, among other things, that the Julia set of exp is the complex plane. This was proved by Misiurewicz in 1981 [59].

The study of holomorphic iteration received a strong revival in the 1980's with Sullivan's proof of the nonexistence of wandering domains for rational functions [81] and Douady and Hubbard's celebrated study of the Mandelbrot set [27]. (The availability of stunning computer pictures also contributed to an immense rise in popularity.) These ideas were then also applied to the study of transcendental functions, with exponential maps being the natural one-parameter family to study in analogy to quadratic polynomials. This was first carried out by Baker and Rippon [5] and Eremenko and Lyubich [30, 31] (who considered the exponential family as an example of more general classes of entire functions). Subsequently, several phenomena where exposed which differ from those occurring for polynomials; see e.g. [50, 53]. There was also a large body of work done notably by Devaney and several coauthors (e.g. [17, 20, 23]) to understand the dynamics of E_κ for $\kappa \in \mathbb{R}$. Also, as early as in [21] (recently published as [11, 12]), the idea of considering exponential maps as a dynamical limit of unicritical polynomials was considered. However, there was little subsequent study of exponential parameter space or the dynamics of general exponential maps.

This changed in the late 1990's with work by Schleicher (in some parts as joined work with Zimmer) [73, 74, 77, 76] which aimed at understanding exponential dynamics in the combinatorial terms used to study the Mandelbrot set. For example, [77] gives a complete description of the set of escaping points of an arbitrary exponential map, and in [73] Schleicher proved two conjectures by Eremenko and Lyubich [30] on hyperbolic components of exponential maps. (Several parts of [73] were published and generalized in [74, 77, 76, 35], though large parts are still unpublished.)

This thesis continues the investigation of exponential maps in this spirit. We prove several new results, resolving in particular a question stated in [73, Section VI.6] on the landing of periodic external rays and a question by Herman, Baker and Rippon regarding the boundaries of Siegel disks [13]. We also give new and simplified proofs of the abovementioned theorems. (Since [73] was not widely circulated, this makes proofs of the conjectures from [30] available for the first time.) In the following, we describe our main results, which are for a large part contained in the manuscripts and preprints [64, 65, 66, 67, 68].

We first develop a model for the set of escaping points of an exponential map and construct a correspondence between this model and the set of escaping points of any exponential map. This construction yields a new, simpler proof of the classification theorem of escaping points mentioned above. To state this result, we need to introduce some minimal amount of symbolic dynamics for exponential maps. (See also Section 3.1.) We associate to (certain) orbits under E_κ an *external address*, i.e., a sequence $\underline{s} = s_1 s_2 s_3 \dots$ of entire numbers, which records the position of each orbit point with respect to a partition of the plane into horizontal strips of height 2π. More precisely, we say that a point z has external address $\underline{s}$ under E_κ if $\operatorname{Im}(E_\kappa^{k-1}(z)) \in \big((2s_k-1)\pi,(2s_k+1)\pi\big)$ for all $k \geq 1$. (Note that not every point has an external address.) Recall also that $I(E_\kappa)$ denotes the set of escaping points of E_κ; i.e. the set of all points z with $|E_\kappa^n(z)| \to \infty$.

Theorem 1.1 (Classification of Escaping Points [77])
Let $\kappa \in \mathbb{C}$. Let $\underline{s}$ be any external address for which there exists some $x > 0$ such that $|s_k| < \exp^{k-1}(x)$ for all $k \geq 1$. Then there exists a curve $g_{\underline{s}} := g_{\underline{s}}^\kappa : (0,\infty) \to I(E_\kappa)$ or $g_{\underline{s}} := g_{\underline{s}}^\kappa : [0,\infty) \to I(E_\kappa)$ (called the external ray *at address $\underline{s}$) with the following properties.*

- *$g_{\underline{s}}$ is a path connected component of $I(E_\kappa)$.*
- $\lim_{t\to\infty} \operatorname{Re}(g_{\underline{s}}(t)) = \infty$.
- *For large t, $g_{\underline{s}}(t)$ has external address $\underline{s}$.*

Conversely, if the singular value does not escape, then every path connected component of $I(E_\kappa)$ is such an external ray. If the singular value does escape, then every path connected component is either an external ray or is mapped into the external ray containing κ in finitely many steps.

Rather than just provide a new proof of a known theorem, our construction yields topological information on the set of escaping points. In particular, we prove the following result.

Theorem 1.2 (Canonical Correspondence of Escaping Points)
Let $\kappa_1, \kappa_2 \in \mathbb{C}$ be parameters for which the singular orbit does not escape. Then there exists a unique bijection

$$\Phi : I(E_{\kappa_1}) \to I(E_{\kappa_2})$$

such that $\Phi(E_{\kappa_1}(z)) = E_{\kappa_2}(\Phi(z))$ and $\big|\Phi(E_{\kappa_1}^n(z)) - E_{\kappa_2}^n(\Phi(z))\big| \to 0$ for all z and such that Φ maps $g_{\underline{s}}^{\kappa_1}$ to $g_{\underline{s}}^{\kappa_2}$ for all $\underline{s}$.

Furthermore, if R is large enough, then Φ is a homeomorphism from

$$I_R := \big\{z \in I(E_{\kappa_1}) : \operatorname{Re}(E^n(\kappa_1)) \geq R \text{ for all } n \geq 0\big\}$$

to $\Phi(I_R)$. Here R can be chosen to be of the order $\log(\max\{1,|\kappa_1|,|\kappa_2|\}) + O(1)$.

In fact, if there is any conjugacy between two exponential maps on their sets of escaping points which respects the combinatorial order of external rays (see Section 3.5), then this conjugacy is given by the above map Φ. In particular, we show the following theorems. (A *Misiurewicz* parameter is one for which the singular value is preperiodic; an *escaping* parameter is one for which the singular value escapes.)

Theorem 1.3 (Uniqueness of Conjugacy)
Suppose that E_{κ_1} and E_{κ_2} are conjugate by an orientation-preserving homeomorphism $\Psi : \mathbb{C} \to \mathbb{C}$. Then there exists $k \in \mathbb{Z}$ such that

$$\Psi(z) = \Phi(z + 2\pi i k)$$

for all $z \in I(E_{\kappa_1})$, where Φ is the map from the previous theorem.

Theorem 1.4 (No Conjugacy)
Suppose that κ_1 and κ_2 are escaping or Misiurewicz parameters such that E_{κ_1} and E_{κ_2} are topologically conjugate. Then there exists $k \in \mathbb{Z}$ such that $\kappa_1 = \kappa_2 + 2\pi i k$ or $\overline{\kappa_1} = \kappa_2 + 2\pi i k$.

This latter result is a generalization of a theorem by Douady and Goldberg [26] who show this result for real (escaping) parameters $\kappa_1, \kappa_2 > -1$.

In the case of an exponential map which has an attracting or parabolic orbit, on the other hand, our model becomes an analog of the "pinched disk model" [25] for polynomials. In particular, this implies the following.

Theorem 1.5 (Conjugacy for Attracting and Parabolic Dynamics)
Suppose that E_{κ_1} and E_{κ_2} both have an attracting or parabolic periodic orbit. Then the map Φ from Theorem 1.2 is a conjugacy on all of $I(E_{\kappa_1})$.

We should note that the topology of exponential maps with an attracting fixed point was described in [1], which was generalized to arbitrary periods in [10]. In particular, these articles proved the landing of all external rays for such parameters. However, the models depended on the parameter, thus these results do not imply Theorem 1.5.

Our model allows explicit computation of the speed of escape. For example, we show that, for any exponential map, there are escaping points which escape arbitrarily slowly. Furthermore, we also show the existence of nonlanding external rays for many exponential maps whose singular value is in the Julia set. (As usual, an external ray $g_{\underline{s}}$ *lands* if $\lim_{t \to 0} g_{\underline{s}}(t)$ exists.) For Misiurewicz parameters the existence of nonlanding rays was observed by Schleicher (personal communication), and later this was proved for real parameters $\kappa \in (-1, \infty)$ by Devaney and Jarque [22].

Theorem 1.6 (Nonlanding External Rays)
Suppose that E_κ is an exponential map for which some external ray accumulates at κ. Then there exists an external ray whose accumulation set contains an entire external ray.

In many cases — in particular if κ is escaping or Misiurewicz — we can even show that there is a ray which accumulates on itself.

By a result of Viana [84], external rays are C^∞ curves. Recall from Theorem 1.1 that some rays may have an endpoint $g_{\underline{s}}(0)$ which also escapes. It was previously not known which of these rays are differentiable in their endpoints, and whether this depends on the parameter. Using our construction, we are able to answer this question.

Theorem 1.7 (Differentiability in Endpoints)
Let $\kappa \in \mathbb{C}$, and suppose that $\underline{s}$ is an external address for which $g_{\underline{s}}$ has an escaping endpoint $z_0 := g_{\underline{s}}(0)$. Then the curve $g_{\underline{s}}$ is continuously differentiable in z_0 if and only if the sum

$$\sum_{k\geq 0} \arg(E_\kappa^k(z_0))$$

converges. Moreover, this condition depends only on $\underline{s}$.

We also study the combinatorics of attracting parameters more closely and succeed in giving a complete description of the bifurcation structure of hyperbolic components in exponential parameter space — i.e., components in which all parameters have an attracting periodic orbit. In particular, we develop algorithms for computing intermediate external addresses, characteristic external addresses, kneading sequences and internal addresses of hyperbolic components. As a corollary, we obtain a necessary and sufficient condition for an attracting exponential map to have infinitely many periodic points at which at least two external rays land (Corollary 5.9.8). A non-necessary sufficient condition was presented in [9]. Furthermore, we prove that there are infinitely many *bifurcation trees* of hyperbolic components (Corollary 5.8.7), which had been conjectured in [30]. We also give a simplified version of Schleicher's proof of the following important conjecture from [30].

Theorem 1.8 (Boundaries of Hyperbolic Components [73, Proposition V.6.4])
Let W be a hyperbolic component. Then $\partial W \subset \mathbb{C}$ is connected.

The results described in the previous paragraph — Section 4.5 and Sections 5.3 to 5.11, to be precise — are joint work with Dierk Schleicher.

Förster [35] extended work of Schleicher [73, Chapter II] to give a classification of those parameters for which the singular value lies on a ray (but is not an escaping endpoint of a ray). Using our results, we can complete this theorem to a complete classification of escaping parameters.

Theorem 1.9 (Parameter Rays)
Let $\underline{s}$ be an exponentially bounded external address. Then there exists a curve $\mathcal{G}_{\underline{s}} : (0,\infty) \to \mathbb{C}$ or $\mathcal{G}_{\underline{s}} : [0,\infty) \to \mathbb{C}$ such that, in the dynamical plane of $\kappa := \mathcal{G}_{\underline{s}}(t)$, $\kappa \in g_{\underline{s}}(t)$. $\mathcal{G}_{\underline{s}}$ is called the parameter ray *at address $\underline{s}$. Every escaping parameter lies on a (unique) parameter ray.*

Possibly the most important result of this thesis is an analog of Douady and Hubbard's landing theorem for periodic rays of polynomials [56, Theorem 18.10].

Theorem 1.10 (Landing of External Rays)
Let E_κ be an exponential map with nonescaping singular value. Let $\underline{s}$ be any periodic external address. Then $g_{\underline{s}}$ lands at a parabolic or repelling periodic point.

The proof of this theorem makes use of the theory of holomorphic motions [51] and a landing theorem for periodic *parameter* rays due to Schleicher [73, Theorem V.7.2]. It is rather unusual that the landing of parameter rays is used to prove the landing of dynamical rays.

Unfortunately, these methods cannot be generalized to higher dimensional parameter spaces, which is why we also give a direct dynamical proof of this theorem in the important case where the ray does not intersect the postsingular set (Theorem 3.9.1). While this thesis was being prepared, we noticed that a proof by Schleicher and Zimmer [76] of the landing theorem in the special case where the singular orbit is bounded can also be generalized to handle the case covered by Theorem 3.9.1.

Theorem 1.10 suggests the important question whether the converse is true, i.e. whether every repelling periodic point is the landing point of some external ray. Recall that this is the case for polynomials with connected Julia set by a theorem of Douady [56, Theorem 18.11]. (It was shown by Eremenko and Levin [29] that every repelling periodic point is accessible from the basin of ∞ even if the Julia set is disconnected.) Using Theorem 1.10, we give a partial answer in this direction, again utilizing results in parameter space.

Theorem 1.11 (Almost all Periodic Points are Landing Points)
Let E_κ be any exponential map with nonescaping singular value. Then, except for at most one periodic orbit, every periodic point is the landing point of a periodic external ray.

Because a nonrepelling periodic point can never be the landing point of a periodic external ray by the Snail Lemma [56, Lemma 16.2], this shows that for attracting or indifferent parameters all repelling periodic points are landing points of periodic rays. We believe that these are the only exceptions that can occur in Theorem 1.11, but this question is still open. This problem is discussed in Section 7.2.

Finally, we also give an answer to the question of Herman, Baker and Rippon [13] whether unbounded exponential Siegel disks need contain the singular value in their boundary.

Theorem 1.12 (Siegel Disks and Singular Values)
Suppose that E_κ has a Siegel disk U (of arbitrary period) such that $\kappa \notin \partial E_\kappa^n(U)$ for all n. Then U is bounded.

By a result of Herman [42], this implies that the singular value is on the boundary of the Siegel disk when the rotation number is diophantine (and, in fact, when it belongs to the more general class H [60]). This result is rather separate from the rest of the thesis, except that the existence of curves in the Julia set is used in the proof.

Because the topics covered in this thesis already cover many aspects of the study of iterated exponential maps, we have attempted to make it into not only a presentation of new results, but also as complete an account as possible of the current state of knowledge in the field. We thus hope that it will be useful as a self-contained overview of exponential dynamics.

The thesis is organized as follows. Chapter 3 contains the general construction of external rays which is fundamental for the rest of the thesis, as well as a discussion of several general results, most of which are of a topological nature. This includes a discussion of the dynamical landing theorem (Theorem 3.9.1) that we mentioned earlier.

Chapter 4 then turns its attention to attracting, parabolic, escaping and Misiurewicz parameters, which are those for which there is much knowledge about combinatorics. Here a special focus lies on attracting and parabolic parameters; in particular, we prove Theorem 1.5. Results from [76] for escaping and Misiurewicz parameters (which are from [76]) are reviewed rather quickly. The chapter is concluded by a discussion of combinatorial ideas introduced in [74] and extended in [68], which applies to all classes of parameters studied in the chapter.

In Chapter 5, we discuss exponential parameter space; in particular we examine the bifurcation structure of hyperbolic components. This chapter also contains the proofs of Theorems 1.10 and 1.11.

Chapter 6 reviews results about exponential dynamics which do not belong to the areas covered by the previous chapters. We prove Theorem 1.12, and discuss the dimension paradox discovered by Karpinska [45] and generalized by Schleicher and Zimmer [77]. Finally, Chapter 7 lists some interesting questions which remain open.

As this work is rather long, we have attempted to make the different chapters self-contained to some degree. The construction of Sections 3.2 and 3.3 is the basis of almost all considerations in the thesis, and the combinatorial definitions of Section 3.7 are almost equally important. All other sections of Chapter 3 are largely independent of each other. They are likewise not required for the understanding of the subsequent chapters, with the exception of the results (but not the methods) of Sections 3.9 and 3.10. Similarly, only Sections 4.1 and 4.5 are fundamentally important for the study of Chapter 5.

Acknowledgements. Foremost, I would like to thank my advisor, Walter Bergweiler, as well as Misha Lyubich and Dierk Schleicher for constant support, friendship and encouragement. I would like to thank the Fulbright Commission for the chance to spend a year in the United States, and the State University of New York where I received my M.A. degree during this time. Many thanks go to John Milnor and the Institute for Mathematical Sciences at Stony Brook for the continuing support and hospitality. Also, I would like to thank the University of Warwick, where this thesis was completed, for welcoming me under the Marie Curie Training Site program of the European Union. I thank Adam Epstein, Alex Eremenko, Markus Förster, Nuria Fagella, Lukas Geyer, John Hubbard, Oleg Kozlovski, John Milnor, Hossein Namazi, Rodrigo Perez, Phil Rippon, Juan Rivera-Letelier, Weixiao Shen, Bartlomiej Skorulski, Sebastian van Strien and Saeed Zakeri, as well as the audiences of the seminars at Cornell, Kiel, Nottingham, Orsay, Stony Brook and Warwick for many interesting discussions. Last but certainly not least I thank my girlfriend Emma as well as my family and friends for all their support and encouragement.

Chapter 2

Preliminaries

2.1 Notation

We assume that the reader is familiar with the basic notions of function theory, the Fatou-Julia theory of iterated holomorphic functions, and the theory of Riemann Surfaces. Excellent references are e.g. [7, 15, 56, 80, 8, 36].

Throughout the text, $\mathbb{C}$ will denote the complex plane, $\hat{\mathbb{C}}$ the Riemann sphere, $\mathbb{D}$ the open unit disk and $\mathbb{H}$ the left half plane. Sometimes we will also consider the cylinder $\widetilde{\mathbb{C}} := \mathbb{C}/2\pi i\mathbb{Z} = \mathbb{C}/\exp$ and the punctured disk $\mathbb{D}^* := \mathbb{D}\setminus\{0\}$. Finally, $\mathbb{D}_r(z_0)$ is the disk of radius r around z_0. If U is an open subset of the plane, we will use the notation $V \Subset U$ to describe the fact that V *is compactly contained in* U; i.e. that V is bounded and $\overline{V} \subset U$.

Let $f : \mathbb{C} \to \mathbb{C}$ be a nonconstant and nonlinear entire function. As usual, f^n denotes the n-th iterate of f. By $F(f)$ we denote the *Fatou set* of f, i.e., the largest set where the family (f^n) is normal; $J(f) := \mathbb{C} \setminus F(f)$ is called the *Julia set*. The set of *escaping points* of f is denoted by

$$I(f) := \left\{ z \in \mathbb{C} : \lim_{n\to\infty} |f^n(z)| = \infty \right\};$$

recall that $J(f) = \partial I(f)$ [28, Formula (1)]. When the function f is fixed, we also denote these sets simply by F, J and I.

Another fundamental object is the set of *singularities of* f^{-1}, or short *singular values*. A number $c \in \mathbb{C}$ is a singular value for f if it is either a critical or asymptotic value of f. The set of singular values of f is denoted by $\operatorname{sing}(f^{-1})$. The *postsingular set* of f is the set

$$\mathcal{P}(f) := \overline{\bigcup_{n\geq 0} f^n(\operatorname{sing}(f^{-1}))}.$$

Note that $f : \mathbb{C} \setminus f^{-1}(\mathcal{P}(f)) \to \mathbb{C} \setminus \mathcal{P}(f)$ is a covering map.

A point $z \in \mathbb{C}$ is called *periodic* under f if there is some n such that $f^n(z) = z$. The smallest such n is called the *period* of z. A periodic point is called *superattracting*, *attracting*, *repelling* or *indifferent*, depending on whether $\left|(f^n)'(z)\right|$ is $= 0$, < 1, > 1 or $= 1$, respectively. Suppose that z is an indifferent periodic point, and let $(f^n)'(z) = e^{2\pi i\alpha}$.

Then z is called *parabolic* if α is rational. Otherwise it is called either a *Siegel* or *Cremer* point, depending on whether $z \in F$ or $z \in J$.

Similarly, a component of the Fatou set U is called *periodic* if there is some n such that $f^n(U) \subset U$. A periodic Fatou component U is called an *attraction domain*, a *Böttcher domain* or a *parabolic domain* if the iterates $(f^{nk})_{k\in\mathbb{N}}$ on U converge to an attracting, superattracting or parabolic periodic point, respectively. It is called a *Siegel disk* if $f|U$ is conjugate to an irrational rotation of the disk, and a *Baker domain* if the iterates of f on U converge locally uniformly to ∞. It is well-known [8, Theorem 6] that every periodic component must be of one of these types (note that f cannot have a Herman ring by the maximum principle). A component of the Fatou set which is neither periodic nor preperiodic is called *wandering*.

Throughout this thesis, we frequently consider *external addresses*, i.e. sequences of integers, and denote the shift map of such sequences by σ, i.e. $\sigma(s_1 s_2 s_3 \dots) = s_2 s_3 \dots$. Also, the notation $\overline{s_1 s_2 \dots s_n}$ is used to denote the sequence in which the entries $s_1, \dots, s_n$ are repeated periodically. Often, the function $F(t) = \exp(t) - 1$ will be used as a model function for exponential growth. We conclude any proof (or any result which is an immediate corollary of previously proved facts) with the symbol $\blacksquare$. Results which are cited without proof are concluded by the symbol $\square$.

2.2 The Hyperbolic Metric

In this section we review those results about hyperbolic geometry of plane domains which we shall require. See [2] or [54] for more details.

Recall that the unique metric of constant negative curvature -1 on $\mathbb{D}$ is given by

$$ds = \frac{2|dz|}{1 - |z|^2}$$

and is called the *hyperbolic metric* on $\mathbb{D}$. This metric is invariant under Möbius transformations of the disk. Suppose that $U \subset \mathbb{C}$ is any domain with $|\mathbb{C} \setminus U| > 1$, and $\pi : \mathbb{D} \to U$ is a universal covering. Then there exists a unique metric on S whose pullback under ϕ is the hyperbolic metric of $\mathbb{D}$; this metric is called the hyperbolic metric of U. We will denote its density by $\rho_U : U \to (0, \infty)$; i.e.

$$ds = \rho_U(w)|dw| = \frac{2|dw|}{|\pi'(z)|(1 - |z|^2)},$$

where z is chosen such that $\pi(z) = w$. In the few cases where the universal covering π is explicitly known, one can thus compute the hyperbolic metric. Two important cases are given by the left half plane $\mathbb{H}$ and the punctured disk $\mathbb{D}^*$:

$$\rho_{\mathbb{H}}(w) = \frac{1}{-\operatorname{Re} w} \quad \text{and}$$
$$\rho_{\mathbb{D}^*}(w) = \frac{1}{\big|w \log |w|\big|}.$$

One of the most useful elementary facts about the hyperbolic metric is given by the Schwarz Lemma.

2.2.1 Theorem (Schwarz Lemma)
Let $f : S \to T$ be a holomorphic map between hyperbolic Riemann Surfaces S and T. Then either f is a covering map (in which case it is a local isometry) or it strictly decreases the hyperbolic metric. □

2.2.2 Corollary (Monotonicity of the Hyperbolic Metric)
Suppose $U \subsetneq V \subset \mathbb{C}$ and $|\mathbb{C} \setminus V| > 1$. Then the hyperbolic metrics $\rho_U(z)|dz|$ and $\rho_V(z)|dz|$ satisfy

$$\rho_U(z) > \rho_V(z)$$

for all $z \in U$. ■

It is important to have estimates on the hyperbolic metric of a plane domain. The following elementary estimate, obtained from the Schwarz Lemma and Koebe's theorem, often suffices.

2.2.3 Theorem (Estimates on the Hyperbolic Metric)
Let $U \subset \mathbb{C}$, $|\mathbb{C} \setminus U| > 1$. Then the hyperbolic metric $\rho_U(z)|dz|$ satisfies

$$\rho_U(z) \leq \frac{2}{\operatorname{dist}(z, \partial U)}$$

(where dist *denotes euclidean distance). If U is simply connected, also*

$$\frac{1}{2\operatorname{dist}(z, \partial U)} \leq \rho_U(z).$$

□

In the case of multiply connected domains, the bound of the previous theorem may be false, as the case of $\mathbb{D}^*$ shows. An exact asymptotic is given by [54, Theorem 2.3]. However, we shall only require the following lower bound for a certain multiply connected domain, which can be obtained by much more elementary means.

2.2.4 Lemma (Hyperbolic Metric of a Multiply Punctured Domain)
Let $U := \mathbb{C} \setminus \{2\pi i k : k \in \mathbb{Z}\}$ and $\mathcal{H} := \{z : \operatorname{Re} z > 0\}$. Then there exists $C > 0$ such that for all z with $\operatorname{Re} z > 1$,

$$\rho_U(z) > C \cdot \rho_{\mathcal{H}}(z) = \frac{C}{\operatorname{Re} z}.$$

Proof. Taking the image of both domains under the covering map exp, it suffices to show that the hyperbolic metrics of $\exp(U) = \mathbb{C} \setminus \{0, 1\}$ and $\exp(\mathcal{H}) = \mathbb{C} \setminus \overline{\mathbb{D}}$ satisfy such a relation for all z with $|z| > e$. However, this follows by classical estimates on the hyperbolic metric of $\mathbb{C} \setminus \{0, 1\}$ [2, Theorem 1-12]. ■

2.3 Iteration of Exponential Functions

In this section, we review some important results concerning the dynamics of the exponential family

$$E_\kappa : \mathbb{C} \to \mathbb{C}; z \mapsto \exp(z) + \kappa.$$

Note that many of these results have in fact been proved for the larger class S of functions with finitely many singular values [31].

Let us begin by proving that escaping points for exponential maps must lie in the Julia set. (Compare [31, Section 2].)

2.3.1 Lemma (Escaping Points in Julia Set)
Let $\kappa \in \mathbb{C}$. Then $I(E_\kappa) \subset J(E_\kappa)$.

PROOF. Suppose that there is $z_0 \in I(E_\kappa) \cap F(E_\kappa)$. Then there exists $\delta_0 > 0$ such that $f^n \to \infty$ uniformly on $\mathbb{D}_{\delta_0}(z_0)$. Then also $\operatorname{Re} f^n|_{\mathbb{D}_{\delta_0}(z_0)} \to \infty$ (because $|\exp z| = \exp(\operatorname{Re} z)$). Thus we can assume that $\operatorname{Re} E_\kappa^n(z) > \max\{\operatorname{Re}(\kappa), 2\}$ for all $z \in \mathbb{D}_{\delta_0}(z_0)$. Let

$$\delta_n := \operatorname{dist}\left(E_\kappa^n(z_0), \partial E_\kappa^n(\mathbb{D}_{\delta_0}(z_0))\right).$$

Then $\delta_n \leq \pi$, as otherwise $\mathbb{D}_{\delta_n}(E_\kappa^n(z_0))$ would intersect one of the lines $\{\operatorname{Im} z = (2k-1)\pi\}$, whose images under E_κ have real part $< \operatorname{Re}\kappa$. But then E_κ is injective on $\mathbb{D}_{\delta_n}(E_\kappa^n(z_0))$, and so by Koebe's theorem,

$$\delta_{n+1} \geq \frac{\left|E_\kappa'(E_\kappa^n(z_0))\right|}{4}\delta_n = \frac{\exp\left(\operatorname{Re} E_\kappa^n(z_0)\right)}{4}\delta_n \geq \frac{e^2}{4}\delta_n \geq \dots \geq \left(\frac{e^2}{4}\right)^{n+1} \cdot \delta_0 \to \infty.$$

This is a contradiction. ■

2.3.2 Theorem (Classification of Fatou Components)
Let $\kappa \in \mathbb{C}$. Then every component of the Fatou set of E_κ is simply connected, and is either an attraction domain, a parabolic domain, a Siegel disk or a preimage of such a domain.

PROOF. Sullivan's theorem that a rational function does not have wandering domains [81] has been generalized to the class of exponential maps by Baker and Rippon [5, Theorem 6]. Thus, every component of the Fatou set is periodic or preperiodic. Also, by the previous lemma there are no Baker domains. E_κ has no superattracting periodic points because it does not have critical points. Thus every Fatou component is of one of the listed types. The fact that all Fatou components are simply connected follows by the maximum principle. Indeed, suppose that γ is a simple closed curve which maps to an attraction domain, a parabolic domain or a Siegel disk under iteration. Then, on γ, the iterates E_κ^n are bounded, which implies by the maximum principle that they are also bounded on the domain enclosed by γ. However, by Montel's theorem this domain then lies in the Fatou set.[1] ■

[1] In fact, if f is any entire function, then every nonwandering component of the Fatou set is simply connected [3].

A famous theorem of Shishikura states that the number of nonrepelling periodic cycles of a rational function is bounded by the number of its critical points. Shishikura's argument was adopted by Eremenko and Lyubich [31] for functions with finitely many singular values. For exponential maps, this means that there is at most one nonrepelling cycle.

2.3.3 Theorem (Number of Nonrepelling Cycles)
An exponential map has at most one periodic cycle which is not repelling. □

2.3.4 Definition (Types of Exponential Parameters)
We call a parameter κ attracting *(or* hyperbolic*),* parabolic, Siegel *or* Cremer *if the map E_κ has a nonrepelling cycle of the corresponding type. We call a parameter* Misiurewicz *if the singular value of E_κ is preperiodic and* escaping *if $\kappa \in I(E_\kappa)$; i.e. if the singular orbit escapes.*

Remark. An escaping or Misiurewicz parameter has no nonrepelling cycle. Indeed, every cycle of immediate attracting or parabolic basins must contain a singular value, and both boundaries of Siegel disks and Cremer cycles must be contained in the postsingular set [56, Theorem 14.4].

To conclude, let us prove that, as noted in the introduction, the exponential family is the only family of transcendental entire maps with only one singular value. (See also [57, Appendix D].)

2.3.5 Theorem (Entire Maps with One Singular Value)
Suppose that $f : \mathbb{C} \to \mathbb{C}$ is an entire function with $\operatorname{sing}(f^{-1}) = \{0\}$. Then there exist $a \in \mathbb{C}^*$ and $b \in \mathbb{C}$ such that f is either of the form $z \mapsto \exp(az + b)$ or of the form $z \mapsto a(z - b)^d$ ($d \geq 2$).

Remark. Note that $z \mapsto \exp(az + b)$ is conjugate to $z \mapsto \exp(z) + b + \log a$. Similarly $z \mapsto a(z - b)^d$ is conjugate to a map of the form $z \mapsto z^d + c$.

Proof. $f : \mathbb{C} \setminus f^{-1}(0) \to \mathbb{C}^*$ is a covering map. In particular, $\mathbb{C} \setminus f^{-1}(0)$ is a parabolic surface, so $f^{-1}(0)$ consists of at most one point. If it is empty, then f is a universal covering of $\mathbb{C}^*$. Because exp is also a universal covering of $\mathbb{C}^*$ and any two such coverings are related by a conformal isomorphism of $\mathbb{C}$, there are $a \neq 0$ and b such that $f(z) = \exp(az + b)$.

Now suppose that $f^{-1}(0)$ is nonempty. Then it consists of a single point $b \in \mathbb{C}$ and $f : \mathbb{C} \setminus \{b\} \to \mathbb{C}^*$ is a covering map. Again, it follows that there is some $d > 0$ and a biholomorphic map $\phi : \mathbb{C} \setminus \{b\} \to \mathbb{C}^*$ such that $f(z) = \phi(z)^d$ for all z. Because $f(b) = 0$, ϕ extends to a conformal isomorphism of $\mathbb{C}$, i.e. a Möbius transformation of the form $z \mapsto a(z - b)$. ■

Remarks on Parametrization. The exponential family has most often been parametrized in the form

$$z \mapsto \lambda \exp(z),$$

with $\lambda \in \mathbb{C}^*$. (Baker and Rippon [5] use the parametrization $z \mapsto \exp(\lambda z)$, and Eremenko and Lyubich [31] use the same parametrization as we do.) Note that the map $\lambda \exp$ is conjugate to E_κ whenever $\lambda = \exp(\kappa)$. Although the choice of parametrization is largely a matter of taste, we feel that we should include an explanation why we prefer the parametrization as E_κ.

We are considering escaping points and external rays, and in particular the structure of the dynamical plane near ∞. Therefore we prefer a parametrization in which the behavior of all maps at ∞ is the same for all exponential maps. For the map $z \mapsto \lambda \exp(z)$ this is not the case, so that the asymptotics of external rays depends on the parameter. This seems somewhat awkward.

Another reason why this choice seems preferable to us is that it further stresses the analogy to the unicritical polynomial families, which are likewise usually parametrized as $z \mapsto z^d + c$, for which, again, the behavior at ∞ is normalized to be the same in the entire family.

Finally, this parametrization — as in that of the polynomial families — has the conceptual advantage that the structure of parameter space around a parameter corresponds to the structure of the dynamical plane around the singular value, and not around the image of the singular value as in the more traditional parametrization.

The only disadvantage we can see is the fact that the κ-plane is not a genuine parameter space, because two parameters which differ by an integer multiple of $2\pi i$ are conformally conjugate. However, this also has a positive side, as this way the same periodic structure exists in parameter space as in the dynamical plane. Also, the combinatorial structure of the parameter plane takes place in the same space as that of the dynamical plane. However, it will sometimes be useful to remember that we can restrict to parameters with $\operatorname{Im} \kappa \in (-\pi, \pi]$.

Chapter 3

Escaping Points of Exponential Maps

This chapter deals with general facts concerning external rays of exponential maps. The first sections contain the construction of external rays, including the proofs of Theorems 1.1 and 1.2. We then turn to several results regarding the topology of the set of escaping points. Section 3.4 contains the proof of Theorem 1.7 and Section 3.5 contains a proof of Theorem 1.3 as well as a short investigation of escaping set rigidity. Section 3.8 proves the existence of nonlanding rays for broad classes of exponential maps (Theorem 1.6). Section 3.9 contains a proof of the landing theorem for periodic rays which do not intersect the postsingular set, as mentioned in the introduction. Also, Section 3.7 contains several combinatorial notions which will be utilized in later chapters.

3.1 Symbolic Dynamics for Exponential Maps

One of the most powerful ideas for studying dynamical systems is to encode dynamical behavior by partitioning phase space into different regions and associating to orbits the sequence of regions that they visit. In this way, the dynamics is reduced to the dynamics of a shift map over a symbol space, making it possible to study dynamical features as combinatorial problems. Classical examples include the Smale horseshoe, unimodal maps and, of course, the study of external rays of polynomials introduced by Douady and Hubbard.

Symbolic dynamics has been part of the study of iterated exponential maps for a long time (see e.g. [23, 21, 20]). Let us first restrict to the case of a real parameter $\kappa \in (-\infty, -1)$. Then E_κ has an attracting fixed point which attracts the entire interval $\big(-\infty, 0]$. Thus the line segment $(-\infty, \kappa)$ and its preimage, $\{z : \exists k \in \mathbb{Z} : \operatorname{Im} z = (2k+1)\pi\}$, lie in the Fatou set. In other words, the Julia set is completely contained within the strips

$$S_k := \big\{z : \operatorname{Im} z \in \big((2k-1)\pi, (2k+1)\pi\big)\big\};$$

see Figure 3.1(a).

Therefore, we can associate to each point $z \in J(E_\kappa)$ a unique sequence $\underline{s} = s_1 s_2 s_3 \dots$ of integers such that $E_\kappa^{k-1}(z) \in S_{s_k}$ for all k. We call this sequence the *external address* of z. It was shown by Devaney (see [23]) that, for every address $\underline{s}$ for which the set of points

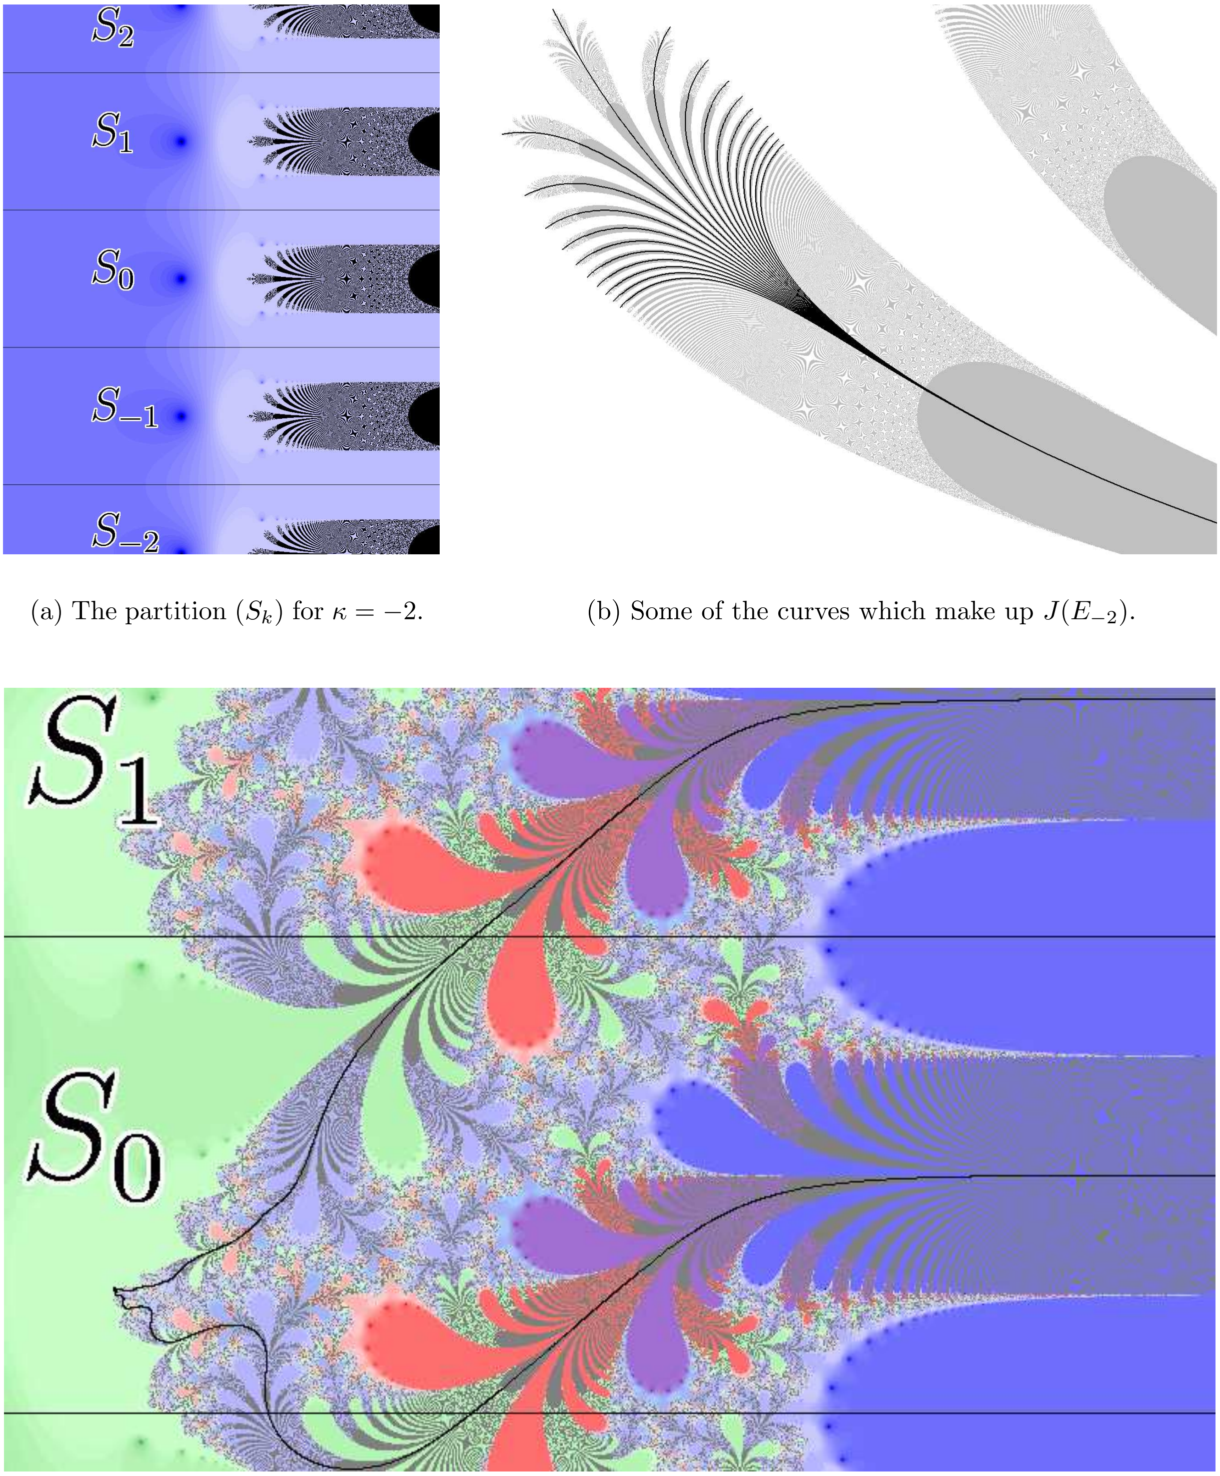

(a) The partition (S_k) for $\kappa = -2$.

(b) Some of the curves which make up $J(E_{-2})$.

(c) In general, external rays cross sector boundaries.

Figure 3.1: The definition of external addresses.

with address $\underline{s}$ is nonempty, this set consists of an injective curve $\gamma : [0, \infty) \to \mathbb{C}$ with $\gamma(t) \to \infty$ as $t \to \infty$; see Figure 3.1(b). For $t > 0$, $\gamma(t)$ is an escaping point, whereas $\gamma(0)$ may be or not be escaping.

The description of which sequences are realized used to depend on the parameter. It was recently observed by Schleicher [77] that this is not necessary, and that in fact the possible growth rates of orbits under exponential maps are independent of the parameter. These sequences are the so-called "exponentially bounded" sequences; see Section 3.6 for a definition.

For a general exponential map E_κ, the Julia set need no longer be contained in the strips S_k; see Figure 3.1(c). But if z is an escaping point, then $E_\kappa^n(z) \notin \{\kappa + t : t \in (-\infty, 0]\}$ for all n larger than some n_0. In particular, the orbit of $E_\kappa^{n_0}(z)$ never intersects the preimages of this line; i.e., $E_\kappa^{n_0}(z)$ has an external address.

Using this idea, Schleicher and Zimmer [77] first construct a "ray-end" for any exponentially bounded external address, and then use a pull-back argument to construct the full ray. Because this pull-back will never reach the endpoint of a ray, they have to use topological arguments to show that every escaping point is either on a ray or the endpoint of a ray, which concludes their proof of Theorem 1.1. However, this proof does not give much information about the behavior of escaping endpoints.

In the following, we modify this approach by using a model for the entire set of escaping points of an exponential map. This model allows us to make topological statements about escaping points, and in particular about the escape speed of escaping endpoints. We should note that our construction is independent of the previously known facts which we stated in this section, and, in particular, gives a new proof of these.

3.2 A Model for the Set of Escaping Points

We will now develop the promised model for the set of escaping points of an exponential map (with nonescaping singular orbit). There is considerable freedom in the definition, and the best choices may depend on the intended use. Here we have chosen a model which allows very explicit calculations. However, there is a downside to this, on which we will remark in the appropriate place.

Recall that a sequence $\underline{s} = s_1 s_2 \ldots$ of integers is called an *external address*. Our model will be situated in the space of all pairs $(\underline{s}, t)$, where $\underline{s}$ is an external address and $t \geq 0$. Note that the space of external addresses has a natural topological structure, namely the order topology of the lexicographic order on external addresses (i.e. the topology whose open sets are unions of open intervals). Thus the model space is equipped with the product topology of this topology and the usual topology of the real numbers. The first entry of the external address of a point should be thought of as indicating its imaginary part, while the second component indicates its real part. We thus define $Z(\underline{s}, t) := t + 2\pi i s_1$ and $|(\underline{s}, t)| := |Z(\underline{s}, t)|$. We shall also write T for the projection to the second component; i.e. $T(\underline{s}, t) = t$. In analogy to the potential-theoretic interpretation of external rays of polynomials, we will sometimes refer to t as the "potential" of the point $(\underline{s}, t)$.

Recall that $F(t) = \exp(t) - 1$. We define

$$\mathcal{F}(\underline{s}, t) := (\sigma(\underline{s}), F(t) - 2\pi|s_2|)),$$

where σ is the shift map, i.e. $\sigma(s_1 s_2 s_3 \dots) = s_2 s_3 \dots$. The key property of this model is the following: for any $\underline{s}$ and t with $T(\mathcal{F}(\underline{s}, t)) \geq 0$,

$$\frac{1}{\sqrt{2}} F(t) \leq |\mathcal{F}(\underline{s}, t)| \leq F(t).$$

This means that, as in exponential dynamics, the exponential of the real part of a point essentially determines the size of its image.

We now define

$$\begin{aligned} \overline{X} &:= \{(\underline{s}, t) : \forall n \geq 0 : T(\mathcal{F}^n(\underline{s}, t)) \geq 0\} \quad \text{and} \\ X &:= \{(\underline{s}, t) \in \overline{X} : T(\mathcal{F}^n(\underline{s}, t)) \to \infty\}. \end{aligned}$$

The dynamics of $\mathcal{F}$ on X will be our model of the set of escaping points; whereas $\overline{X}$ should be thought of as a model of the Julia set in the presence of an attracting fixed point, as discussed in Section 3.1.

3.2.1 Observation (Comb Structure of $\overline{X}$ and X)
For every external address $\underline{s}$, there exists a $t_{\underline{s}}$ such that

$$\overline{X} \cap \big(\{\underline{s}\} \times [0, \infty)\big) = \{(\underline{s}, t) : t \geq t_{\underline{s}}\};$$

this $t_{\underline{s}}$ depends lower semicontinuously on $\underline{s}$. Furthermore,

$$\{(\underline{s}, t) : t > t_{\underline{s}}\} \subset X.$$

Proof. Suppose that $(\underline{s}, t) \in \overline{X}$ and $t' = t + \delta$, $\delta > 0$. By the definition of $\mathcal{F}$, we have

$$\begin{aligned} T(\mathcal{F}(\underline{s}, t')) - T(\mathcal{F}(\underline{s}, t)) &= F(t') - F(t) \\ &= \exp(t + \delta) - \exp(t) \geq \exp(\delta) - 1 = F(\delta). \end{aligned}$$

By induction,

$$T(\mathcal{F}^n(\underline{s}, t')) \geq T(\mathcal{F}^n(\underline{s}, t')) - T(\mathcal{F}^n(\underline{s}, t)) \geq F^n(\delta) \to \infty.$$

This proves the first and last claim. To prove semicontinuity, note that $\overline{X}$ is a closed set. Therefore, for any $R > 0$ the set

$$\{\underline{s} : t_{\underline{s}} \leq R\} = \{\underline{s} : (\underline{s}, R) \in \overline{X}\}$$

is closed. ■

3.2.2 Lemma (Calculation of $t_{\underline{s}}$)
Define $t^0_{\underline{s}} := 0$ and, inductively,

$$t^{k+1}_{\underline{s}} := F^{-1}\left(t^k_{\sigma(\underline{s})} + 2\pi|s_2|\right).$$

Then $t^k_{\underline{s}} \to t_{\underline{s}}$ for every $\underline{s}$.

Proof. Note that, for every $\underline{s}$, the sequence $t^k_{\underline{s}}$ is nondecreasing, so $t'_{\underline{s}} := \lim t^k_{\underline{s}}$ exists. Also note that $T(\mathcal{F}^n(\underline{s}, t'_{\underline{s}})) = t'_{\sigma^n(\underline{s})} \geq 0$. This shows that $(\underline{s}, t'_{\underline{s}}) \in \overline{X}$, i.e., $t'_{\underline{s}} \geq t_{\underline{s}}$.

On the other hand, note that

$$T(\mathcal{F}^k(\underline{s}, t^k_{\underline{s}})) = t^0_{\sigma^k(\underline{s})} = 0$$

and thus $(\underline{s}, t^k_{\underline{s}}) \notin X$; i.e. $t^k_{\underline{s}} \leq t_{\underline{s}}$. Therefore $t'_{\underline{s}} = \sup t^k_{\underline{s}} \leq t_{\underline{s}}$. ■

3.3 Classification of Escaping Points

Throughout this section, let us fix some $\kappa \in \mathbb{C}$. Recall that a point z is said to have external address $\underline{s}$ if

$$\operatorname{Im}(E^{n-1}_\kappa(z)) \in S_{s_N}$$

for all n, where $S_k = \big\{z : \operatorname{Im} z \in \big((2k-1)\pi, (2k+1)\pi\big)\big\}$.

We will now construct a conjugacy (on a suitable subset of X). This is done by iterating forward in our model and then backwards in the dynamics of E_κ. To this end, we define the inverse branches L_k of E_κ by

$$L_k(w) := \operatorname{Log}(w - \kappa) + 2\pi i k,$$

where $\operatorname{Log} : \mathbb{C} \setminus (-\infty, 0] \to S_0$ is the principle branch of the logarithm. Thus $L_k(w)$ is defined and analytic whenever $w - \kappa \notin (-\infty, 0]$.

Define maps G_k inductively by $G_0(\underline{s}, t) := Z(\underline{s}, t)$ and

$$G_{k+1}(\underline{s}, t) := L_{s_1}(G_k(\mathcal{F}(\underline{s}, t)))$$

(wherever this is defined).

Fix $K > 2\pi + 6$ such that $|\kappa| \leq K$, and set $Q := \max\{\log(4(K + \pi + 3)), \pi + 2\}$. Consider the set

$$Y := Y_Q := \left\{(\underline{s}, t) \in \overline{X} : T(\mathcal{F}^n(\underline{s}, t)) \geq Q \text{ for all } n\right\}.$$

Note that Y contains the set of all $(\underline{s}, t)$ with $t \geq t_{\underline{s}} + Q$. Note also that, for every $x = (\underline{s}, t) \in X$, there exists some n so that $\mathcal{F}^n(x) \in Y$.

3.3.1 Lemma (Bound on G_k)
For all k, G_k is defined on Y and satisfies $|\operatorname{Re} G_k(\underline{s}, t) - t| < 2$; in particular,

$$|G_k(\underline{s}, t) - Z(\underline{s}, t)| < \pi + 2.$$

Proof. By induction. The case $k = 0$ is trivial. Let $k \geq 0$ such that the claim is true for k; we will show that it holds also for $k+1$. Let $(\underline{s}, t) \in Y$. By the induction hypothesis,

$$|G_k(\mathcal{F}(\underline{s}, t)) - Z(\mathcal{F}(\underline{s}, t))| \leq \pi + 2.$$

Therefore,

$$\operatorname{Re}\left(G_k(\mathcal{F}(\underline{s}, t))\right) \geq T(\mathcal{F}(\underline{s}, t)) - \pi - 2 \geq Q - \pi - 2 \geq 0.$$

Furthermore, we have

$$F(t) = \exp(t) - 1 \geq \exp(Q) - 1 \geq 4(K + \pi + 2)$$

In particular,

$$|G_k(\mathcal{F}(\underline{s}, t))| \geq |Z(\mathcal{F}(\underline{s}, t))| - \pi - 2 \geq \frac{F(t)}{\sqrt{2}} - \pi - 2 > 2K.$$

Thus $G_k(\mathcal{F}(\underline{s}, t)) - \kappa \notin (-\infty, 0]$, so $G_{k+1}(\underline{s}, t)$ is defined. Furthermore, we can write

$$G_k(\mathcal{F}(\underline{s}, t)) - \kappa = Z(\mathcal{F}(\underline{s}, t)) + \big(G_k(\mathcal{F}(\underline{s}, t)) - Z(\mathcal{F}(\underline{s}, t)) - \kappa\big),$$

and, by the definition of Q,

$$|G_k(\mathcal{F}(\underline{s}, t)) - Z(\mathcal{F}(\underline{s}, t)) - \kappa| \leq \pi + 2 + K \leq \frac{1}{4}\exp(Q) - 1 \leq \frac{1}{4}\exp(t) - 1.$$

Therefore

$$\begin{aligned}\operatorname{Re}(G_{k+1}(\underline{s}, t)) &= \log|G_k(\mathcal{F}(\underline{s}, t)) - \kappa| \geq \log\left(|Z(\mathcal{F}(\underline{s}, t))| + 1 - \frac{1}{4}\exp(t)\right)\\ &\geq \log\left(\frac{1}{\sqrt{2}}\exp(t) - \frac{1}{4}\exp(t)\right) > t - \log 4.\end{aligned}$$

Analogously $\operatorname{Re}(G_{k+1}(\underline{s}, t)) \leq t + \log 4$. Thus

$$|G_{k+1}(\underline{s}, t) - Z(\underline{s}, t)| \leq |\operatorname{Re}(G_{k+1}(\underline{s}, t) - t| + |\operatorname{Im}(G_{k+1}(\underline{s}, t)) - 2\pi s_1| \leq \log 4 + \pi < \pi + 2.$$

This completes the proof. ■

3.3.2 Theorem (Convergence to a Conjugacy)
On Y, the functions G_n converge uniformly (in $(\underline{s},t)$ and κ) to a function $G : Y \to J(E_\kappa)$ such that $G(\underline{s}, t)$ has external address $\underline{s}$ for each $(\underline{s}, t) \in Y$. This function satisfies $G \circ \mathcal{F} = E_\kappa \circ G$ and

$$\Big|G(\underline{s}, t) - \big(\operatorname{Log}(Z(\mathcal{F}(\underline{s}, t))) + 2\pi i s_1\big)\Big| \leq e^{-t} \cdot (2K + 2\pi + 4). \tag{3.1}$$

Furthermore, $G(\underline{s}, t) \in I$ if and only if $(\underline{s}, t) \in X$, G is a homeomorphism between Y and its image, and $G(\underline{s}, t)$ depends holomorphically on κ for fixed $(\underline{s}, t) \in Y$.

Remark. Note that (3.1) implies that

$$G(\underline{s},t) = t + 2\pi i s_1 + O(e^{-t}) \tag{3.2}$$

for every $\underline{s}$.

Proof. Recall from the previous proof that, for $n \geq 1$,

$$|G_k(\mathcal{F}^n(\underline{s},t))| > 2K$$

and

$$\operatorname{Re}\left(G_k(\mathcal{F}^n(\underline{s},t))\right) > 0.$$

Furthermore, the distance between $G_k(\mathcal{F}(\underline{s},t))$ and $G_{k+1}(\mathcal{F}(\underline{s},t))$ is at most $2\pi + 4$. Thus we can connect these two points by a straight line within the set

$$\{z \in \mathbb{C} : |z-\kappa| \geq 2 \text{ and } z - \kappa \notin (-\infty, 0]\}.$$

Since $L'_{s_1}(z) = \frac{1}{z}$, it follows that

$$|G_{k+1}(\underline{s},t) - G_{k+2}(\underline{s},t)| \leq \frac{1}{2}|G_k(\mathcal{F}(\underline{s},t)) - G_{k+1}(\mathcal{F}(\underline{s},t))|.$$

It follows by induction that

$$|G_{k+1}(\underline{s},t) - G_{k+2}(\underline{s},t)| \leq 2^{-(k+1)}(\pi+2),$$

so the G_n converge uniformly on Y. Because $E_\kappa \circ G_n = G_{n-1} \circ \mathcal{F}$, we see that G satisfies the functional relation

$$E_\kappa \circ G = G \circ \mathcal{F}.$$

To prove the asymptotics (3.1), note that

$$\left|G(\underline{s},t) - \big(\operatorname{Log}(Z(\mathcal{F}(\underline{s},t))) + 2\pi i s_1\big)\right| = \left|\operatorname{Log}\left(1 + \frac{G(\mathcal{F}(\underline{s},t)) - \kappa - Z(\mathcal{F}(\underline{s},t))}{Z(\mathcal{F}(\underline{s},t))}\right)\right|$$
$$\leq \sqrt{2}\left|\frac{G(\mathcal{F}(\underline{s},t)) - Z(\mathcal{F}(\underline{s},t)) - \kappa}{Z(\mathcal{F}(\underline{s},t))}\right| \leq \sqrt{2}\frac{\pi + 2 + K}{F(t)} \leq \frac{2\pi + 4 + 2K}{\exp(t)}.$$

(Note that $|\operatorname{Log}(1+z)| \leq \sqrt{2}|z|$ for $|z| < \frac{1}{4}$).

By Lemma 3.3.1, $G(\underline{s},t)$ escapes under iteration of E_κ if and only if $(\underline{s},t)$ escapes under iteration of $\mathcal{F}$. Clearly the point $G(\underline{s},t)$ has the correct external address (note that $\arg\big(G_{k-1}(\underline{s},t)\big)$ is bounded away from $\pm\pi$, so that the values $G_k(\underline{s},t)$ cannot converge to the strip boundaries). In particular, $G(\underline{s},t) \neq G(\underline{s}',t')$ whenever $\underline{s} \neq \underline{s}'$, because the points have different external addresses. On the other hand, under iteration of $\mathcal{F}$, the points $(\underline{s},t)$ and $(\underline{s},t')$, $t \neq t'$, will eventually be arbitrarily far apart, and therefore the same holds for $G(\underline{s},t)$ and $G(\underline{s},t')$ under E_κ. This proves injectivity.

The function G is continuous as uniform limit of continuous functions; for the same reason, $G(\underline{s},t)$ is analytic in κ. To prove that the inverse G^{-1} is continuous, note that we can compactify both Y and $G(Y)$ by adding a point at infinity. The extended map G is still continuous, and the inverse of a continuous bijective map on a compact space is continuous. ∎

REMARK. The asymptotic description of $G(\underline{s},t)$ in terms of

$$\begin{aligned}\operatorname{Log}(Z(\mathcal{F}(\underline{s},t))) + 2\pi i s_1 = &\log\sqrt{(F(t)-2\pi|s_2|)^2+(2\pi s_2)^2}+\\ &i\arg\big((F(t)-2\pi|s_2|+2\pi i s_2\big)+2\pi i s_1\end{aligned}$$

is somewhat awkward. Had we used, instead of $\mathcal{F}$, the map

$$\mathcal{F}'(\underline{s},t) := (\sigma(\underline{s}), \sqrt{F(t)^2-(2\pi s_2)^2}),$$

then the whole construction would have carried through analogously (with somewhat improved constants). For the map $G' : Y \to I$ that we obtain this way, we would correspondingly have the following asymptotics:

$$|\operatorname{Re}(G'(\underline{s},t)) - t| < e^{-t}\cdot(2K+2\pi+4).$$

However, in this thesis we will never use the asymptotics in any other form than (3.2), whereas we will rather often make direct calculations in the model. This is why we have opted to use the function $\mathcal{F}$ rather than $\mathcal{F}'$.

In order to extend G to a bijection $G : X \to I$, the main remaining problem is to decide whether a point is contained in $G(Y)$. The following Theorem is a counterpart to Theorem 3.3.2.

3.3.3 Theorem (Points in the Image of G)
Suppose that $z\in\mathbb{C}$ spends its entire orbit in the half plane $\{w\in\mathbb{C} : \operatorname{Re} w \geq Q+1\}$. Then z has an external address $\underline{s}$, and there exists t such that $(\underline{s},t)\in Y$ and $z = G(\underline{s},t)$.

PROOF. First note that, for $n \geq 1$,

$$|E_\kappa^n(z)-\kappa| = \exp(\operatorname{Re}(E_\kappa^{n-1}(z)) \geq \exp(Q+1) > 2K$$

and $\operatorname{Re} E_\kappa^n(z) > 0$, so $E_\kappa^n(z)-\kappa \notin (-\infty,0)$. Therefore, no iterate of z lies on the strip boundaries, and thus z has an external address $\underline{s}$.

Consider the sequence t_k of potentials uniquely defined by

$$\mathcal{F}^k(\underline{s},t_k) = (\sigma^k(\underline{s}), \operatorname{Re}(E_\kappa^k(z))).$$

By an analogous argument to that of Lemma 3.3.1, we see that

$$\left|T(\mathcal{F}^j(\underline{s},t_k)) - \operatorname{Re} E_\kappa^j(z)\right| \leq 1$$

for all $j \leq k$. Let t be any limit point of the sequence t_k. Then, by the above,

$$\left|T(F^j(\underline{s},t) - \operatorname{Re} E_\kappa^j(z)\right| \leq 1;$$

in particular, $(\underline{s}, t) \in Y$.

Since $G(\underline{s}, t)$ also has external address $\underline{s}$, it now follows that the distance between $E_\kappa^j(G(\underline{s}, t))$ and $E_\kappa^j(z)$ is bounded for all j. By the same contraction argument as in the proof of Theorem 3.3.2, they are equal. ■

We can now prove the existence of a global correspondence between X and $I(E_\kappa)$. The following result proves the classification of escaping points (Theorem 1.1) — except for the fact that external rays form the path-connected components of I, which will be shown in Section 3.8 — as well as Theorem 1.2.

3.3.4 Corollary (Global Correspondence)
Suppose that $\kappa \notin I$. Then $G|(Y \cap X)$ extends to a bijective function

$$G : X \to I$$

which satisfies $G(\mathcal{F}(\underline{s}, t)) = E_\kappa(G(\underline{s}, t))$.

Furthermore, G is a homeomorphism on every $\mathcal{F}^{-k}(Y \cap X)$ and for every $\underline{s}$ the function $t \mapsto G(\underline{s}, t)$ is continuous ("G is continuous along rays").

Proof. It is sufficient to show, by induction G extends to a homeomorphism

$$G : \mathcal{F}^{-k}(Y \cap X) \to E_\kappa^{-k}(G(Y \cap X))$$

for every $k \geq 1$. Indeed, the sets of definition clearly exhaust all of X, while the range exhausts I by Theorem 3.3.3. Continuity along rays also follows because every $(\underline{s}, t) \in X$ has a neighborhood on the ray that is completely contained in the same $\mathcal{F}^{-k}(Y \cap X)$.

So let us suppose that G has been extended to $\mathcal{F}^{-k}(Y \cap X)$. First note that we can extend G to $F^{-(k+1)}(Y \cap X)$ in such a way that the extension is continuous along rays. Indeed, for every $\underline{s}$, we can choose a branch L of E_κ^{-1} on the ray such that $L(G(\mathcal{F}(\underline{s}, t))) = G(\underline{s}, t)$ whenever $(\underline{s}, t) \in F^{-k}(Y \cap X)$. This extension is also continuous in both variables because the branch L varies continuously. ■

Remark. If $\kappa \in I$, then there is a similar statement; we will only need to exclude all preimages of the ray which contains κ. As the full statement becomes rather complicated, we will not discuss it here. However, clearly G still has a maximal extension, which can be described as follows. There exists some $(\underline{s}^0, t^0)$ with $G(\underline{s}^0, t^0) = \kappa$, and G is defined for all $(\underline{s}, t)$ except those for which $\mathcal{F}^n(\underline{s}, t) = (\underline{s}^0, t')$ for some $n \geq 1$ and $t' \leq t^0$.

When considering individual rays, it is often cumbersome to take into account the starting potential $t_{\underline{s}}$. For convenience, we make the following definition.

3.3.5 Definition (External Rays)
Let $\kappa \in \mathbb{C}$, and let $\underline{s}$ be any external address with $t_{\underline{s}} < \infty$. We define a curve $g_{\underline{s}}$ — the external ray at address $\underline{s}$ *by*

$$g_{\underline{s}}(t) = G(\underline{s}, t + t_{\underline{s}}).$$

If $g_{\underline{s}}$ is not defined for all $t > 0$ (i.e., if there exists $t_0 > t_{\underline{s}}$ such that $G(\mathcal{F}^n(\underline{s}, t_0)) = \kappa$), then we will sometimes call $g_{\underline{s}}$ a "broken ray". We say that an unbroken ray $g_{\underline{s}}$ lands *at a* *point* z_0 *if* $\lim_{t\to 0} g_{\underline{s}}(t) = z_0$. *Similarly, we say that* $g_{\underline{s}}(t)$ has an escaping endpoint *if* $g_{\underline{s}}(0)$ *is defined and escaping; i.e. if* $(\underline{s}, t_{\underline{s}}) \in X$.

3.3.6 Remark (Convergence of Rays)
Suppose that $\underline{s}^n$ *is a sequence of external addresses converging to a sequence* $\underline{s}^0$ *such that also* $t_{\underline{s}^n} \to t_{\underline{s}^0}$, *and let* $t_0 > 0$ *such that* $g_{\underline{s}^0}(t)$ *is defined for all* $t > t_0$. *Then*

$$g_{\underline{s}^n}|_{[t_0,\infty)} \to g_{\underline{s}^0}|_{[t_0,\infty)}$$

uniformly.

Proof. There exists k such that $\{(\underline{s}, t) : t - t_{\underline{s}} \geq t_0\} \subset \mathcal{F}^{-k}(Y_Q)$. The claim then follows from Corollary 3.3.4. ■

Remark. In the case where $g_{\underline{s}^0}$ is broken, we can say the following (with the same proof). Suppose that $\underline{s}^n > \underline{s}^0$ (or $< \underline{s}^0$) for all n. Then, under the assumptions of Remark 3.3.6 the rays $g_{\underline{s}^n}$ converge locally uniformly (in $\hat{\mathbb{C}}$) to a curve $\widetilde{g_{\underline{s}^0}} : (0, \infty) \to \hat{\mathbb{C}}$. This curve has $\widetilde{g_{\underline{s}^0}}(t) = \infty$ if and only if $G(\mathcal{F}^n(\underline{s}^0, t + t_{\underline{s}^0})) = \kappa$ for some $n \geq 1$ and coincides with $g_{\underline{s}^0}$ where the latter is defined. Thus if the ray which contains κ is periodic, ∞ is assumed infinitely many times on this curve. Also, one can see that this curve must accumulate everywhere on itself. A Theorem of Curry [16] can be used to show that the accumulation set of $\widetilde{g_{\underline{s}^0}}$ in $\mathbb{C}$ can be compactified to an indecomposable continuum. This was done for $\kappa > -1$ in [18] and for certain other parameters in [70].

In the following, we will sometimes write G^κ or $g^\kappa_{\underline{s}}$ for the objects constructed previously when the parameter is not fixed in the context.

There is an interesting corollary of Theorem 3.3.2. Note that Y_Q contains many points of $\overline{X} \setminus X$; in particular endpoints of periodic addresses. Which of these addresses lie in Y_Q depends on Q (and thus on the parameter); however, we can use this fact to give an elementary bound on those parameters for which we know that these rays cross sector boundaries or are not defined. This is the content of the following result, which will later be used to bound *parameter rays*.

3.3.7 Corollary (Bound on Parameter Rays)
Let $\underline{s}$ *be an external address with* $t_{\underline{s}} < \infty$. *Let* $t_0 := \inf_{n\geq 0} t_{\sigma^n(\underline{s})}$ *and suppose that* $t_0 > \pi+2$. *If* $\kappa \in \mathbb{C}$ *such that* $G(\mathcal{F}(\underline{s}, t)) = \kappa$ *for some* t, *then*

$$|\kappa| > \frac{1}{5}\exp(t_0).$$

In particular, if $\underline{s}$ *is periodic of period* n *and* $M := \max |s_k|$ *is large enough, then* $|\kappa| \geq \frac{1}{5}F^{-n}(2\pi M)$.

Similarly, suppose that $\underline{s}^1$ *and* $\underline{s}^2$ *are external addresses for which there is* $n \in \mathbb{N}$ *and some large enough* M *such that*

$$\max_{kn+1\leq m\leq (k+1)n} |\underline{s}^j_m| \geq M$$

for all $k \geq 0$ and $j \in \{1, 2\}$. If κ is a parameter such that $g^{\kappa}_{\underline{s}^1}$ and $g^{\kappa}_{\underline{s}^2}$ land together, then $|\kappa| \geq \frac{1}{5} F^{-n}(2\pi M)$.

PROOF. Recall that, for any address $\underline{s}^0$ such that $(\underline{s}^0, t_{\underline{s}^0}) \in Y_Q = Y_{Q(\kappa)}$, the ray $g^{\kappa}_{\underline{s}^0}$ lands, and in particular, $g^{\kappa}_{\sigma(\underline{s}^0)}$ does not contain κ. Also, no two such rays have the same landing point, because the landing points have the same external address as the rays. Thus the first claim is proved upon noting that $t_0 < Q(\kappa) = \max\{\log(4|\kappa| + \pi + 3), \pi + 2\}$ implies

$$|\kappa| > \frac{1}{4}\exp(t_0) - \pi - 3 > \frac{1}{5}\exp(t_0)$$

(for $t_0 > \pi + 2$).

The second and third claim follow from this by realizing that, among all addresses $\underline{s}$ one of whose entries $s_2, \dots, s_{n+1}$ is of size at least M, the value of $t_{\underline{s}}$ is minimized by the address

$$\underline{s} = 00\dots0M\overline{0}$$

(where the first block of 0s consists of $n - 1$ entries). For this $\underline{s}$, the value can be easily computed to be $F^{-n}(2\pi M)$. ■

3.4 Differentiability of Rays

Viana [84] proved (using a different parametrization) that the rays $g_{\underline{s}}$ are C^∞. His arguments also apply to the parametrization of the curves given by our construction. (Compare also the proof of Theorem 3.4.2 below.)

3.4.1 Theorem (Rays are Differentiable [84])
Let $\underline{s}$ be an exponentially bounded external address. Then $g_{\underline{s}} : (0, \infty) \to \mathbb{C}$ is C^∞. □

A proof of the differentiability of rays using the parametrization of rays constructed in [77] can also be found in [35], where this was carried out to obtain specific estimates on the first and second derivatives. However, there is no information about which rays with escaping endpoints are also differentiable in these endpoints. Using the results of the previous section, we can answer this question.

3.4.2 Theorem (Differentiability of Rays in Endpoints)
Let $\underline{s}$ be an exponentially bounded external address such that $(\underline{s}, t_{\underline{s}}) \in X$. Then the curve $g_{\underline{s}}\big([0, \infty)\big)$ is continuously differentiable in $g_{\underline{s}}(0)$ if and only if the sum

$$\sum_{j=0}^{\infty} \frac{2\pi s_{j+1}}{T(\mathcal{F}^j(\underline{s}, t_{\underline{s}}))} \tag{3.3}$$

converges.

Remark. By the formulation "the curve is continuously differentiable in $g_{\underline{s}}(0)$" we mean that it is continuously differentiable under a suitable parametrization (e.g., by arclength), *not* that the function $g_{\underline{s}}$ itself is necessarily differentiable in 0. If the convergence of the sum is absolute, then one can show that the function $g_{\underline{s}}$ itself is differentiable in 0.

Proof. Let Q be the number from the previous section for which all G_k are defined on the set Y_Q and converge uniformly to the function G there; it is clearly sufficient to prove the theorem for addresses for which $(\underline{s}, t_{\underline{s}}) \in Y_Q \cap X$.

By the definition of the functions G_k, their t-derivatives in any point $(\underline{s}, t) \in Y_Q$ are given by

$$\begin{aligned}\frac{\partial G_k}{\partial t}(\underline{s}, t) &= \frac{1}{G_{k-1}(\mathcal{F}(\underline{s}, t)) - \kappa} \cdot \frac{\partial G_{k-1}}{\partial t}(\mathcal{F}(\underline{s}, t)) \cdot \exp(t) = \dots \\ &= \prod_{j=1}^{k} \frac{\exp(T(\mathcal{F}^{j-1}(\underline{s}, t)))}{G_{k-j}(\mathcal{F}^j(\underline{s}, t)) - \kappa} \\ &= \left(\prod_{j=1}^{k} \frac{\exp(T(\mathcal{F}^{j-1}(\underline{s}, t)))}{G(\mathcal{F}^j(\underline{s}, t)) - \kappa}\right) \cdot \prod_{j=1}^{k} \left(1 + \frac{G(\mathcal{F}^j(\underline{s}, t)) - G_{k-j}(\mathcal{F}^j(\underline{s}, t))}{G_{k-j}(\mathcal{F}^j(\underline{s}, t)) - \kappa}\right).\end{aligned}$$

Recall that

$$\left|G(\mathcal{F}^j(\underline{s}, t)) - G_{k-j}(\mathcal{F}^j(\underline{s}, t))\right| \leq 2^{-(k-j)} \cdot (2\pi + 4),$$

so the last product converges uniformly for $t \geq t_{\underline{s}}$. It is not difficult to see that the first product converges locally uniformly (and is nonzero) for $t > t_{\underline{s}}$ (see e.g. [35]). Note that this proves that the ray without the endpoint is C^1.

Now the ray is continuously differentiable in its endpoint if and only if $\arg\left(\frac{\partial G}{\partial t}(\underline{s}, t)\right)$ has a limit as $t \to t_{\underline{s}}$. The above argument shows that this is equivalent to the question whether the function

$$\Theta(t) := \sum_{j=1}^{\infty} \arg\bigl(G(\mathcal{F}^j(\underline{s}, t)) - \kappa\bigr)$$

has a limit for $t \to t_{\underline{s}}$.

Claim The limit $\lim_{t \to t_{\underline{s}}} \Theta(t)$ exists if and only if the sum $\Theta(t_{\underline{s}})$ is convergent.

To prove this claim, fix some number $m \geq 3$ and define, for n large enough, $t_n > t_{\underline{s}}$ to be the unique number for which

$$T(\mathcal{F}^n(\underline{s}, t_n)) = T(\mathcal{F}^n(\underline{s}, t_{\underline{s}})) + \log m.$$

Note that $|G(\mathcal{F}^n(\underline{s}, t_n)) - G(\mathcal{F}^n(\underline{s}, t_{\underline{s}}))| \leq K := 2\pi + 4 + \log m$. It follows again by contraction that, for $k \leq n$,

$$|G(\mathcal{F}^k(\underline{s}, t_n)) - G(\mathcal{F}^k(\underline{s}, t_{\underline{s}}))| \leq 2^{-(n-k)} \cdot K.$$

Thus

$$\left|\sum_{k=1}^{n} \arg\bigl(G(\mathcal{F}^k(\underline{s}, t_{\underline{s}})) - \kappa\bigr) - \sum_{k=1}^{n} \arg\bigl(G(\mathcal{F}^k(\underline{s}, t_{\underline{s}})) - \kappa\bigr)\right| \leq \pi K \cdot \sum_{k=1}^{n} \frac{2^{-(n-k)}}{|G(\mathcal{F}^k(\underline{s}, t_{\underline{s}})) - \kappa|},$$

which is easily seen to converge to 0 as $n \to \infty$.

Also observe that, for $k \geq n+1$,

$$T(\mathcal{F}^k(\underline{s}, t_n)) \geq F^{k-n-1}((m-1) \cdot F(T(\mathcal{F}^n(\underline{s}, t_{\underline{s}}))))$$

and

$$2\pi s_{n-1} \leq F^{k-n-1}(F(T(\mathcal{F}^n(\underline{s}, t_{\underline{s}})))).$$

It easily follows that

$$\sum_{k=n+2}^{\infty} \left| \arg\bigl(G(\mathcal{F}^k(\underline{s}, t_n)) - \kappa\bigr) \right| \to 0$$

as $n \to \infty$. Similarly, for large enough n, the value $\left|\arg\bigl(G(\mathcal{F}^{n+1}(\underline{s}, t_n)) - \kappa\bigr)\right|$ is no larger than $\frac{2}{m-1} + \varepsilon$. If $\arg G(\mathcal{F}^{n+1}(\underline{s}, t_{\underline{s}}))$ tends to 0, then $\arg G(\mathcal{F}^{n+1}(\underline{s}, t_n))$ also does.

Now let us first consider the case that $\arg G(\mathcal{F}^j(\underline{s}, t_{\underline{s}}))$ does not converge to 0 (and thus in particular the sum $\Theta(t_{\underline{s}})$ is divergent). So let $\delta > 0$ and let n_k be a subsequence such that

$$\left| \arg\bigl(G(\mathcal{F}^{n_k}(\underline{s}, t_{\underline{s}})) - \kappa\bigr) \right| \geq \delta.$$

If m was chosen to be $1 + \frac{5}{\delta}$, then it follows from the above considerations that

$$|\Theta(t_{n_k - 1}) - \Theta(t_{n_k})| \geq \delta - \frac{4}{m-1} + o(1) > \frac{\delta}{5} + o(1)$$

(as $k \to \infty$). In particular, the sequence $\Theta(t_n)$ does not have a limit for $n \to \infty$. This proves the claim in this case.

So we can now suppose that $\arg G(\mathcal{F}^j(\underline{s}, t_{\underline{s}})) \to 0$. Then, by our observations,

$$\Bigl|\Theta(t_n) - \sum_{k=1}^{n} \arg(G(\mathcal{F}^k(\underline{s}, t_{\underline{s}})) - \kappa)\Bigr| \to 0.$$

Thus in particular the sequence $\Theta(t_n)$ has a limit if and only if the sum $\Theta(t_{\underline{s}})$ is convergent. It remains to show that this implies that Θ has a limit as $t \to t_{\underline{s}}$. However, it is easy to show that

$$\sup_{t \in [t_n, t_{n+1}]} |\Theta(t_n) - \Theta(t)| \to 0$$

as $n \to \infty$. Indeed, by the above observations,

$$\sum_{k=1}^{n-1} \bigl|\arg G(\mathcal{F}^k(\underline{s}, t)) - \arg G(\mathcal{F}^k(\underline{s}, t_n))\bigr|$$

is small, as is

$$\sum_{k=n+2}^{\infty} |\arg(G(\mathcal{F}^k(\underline{s}, t)) - \kappa)|.$$

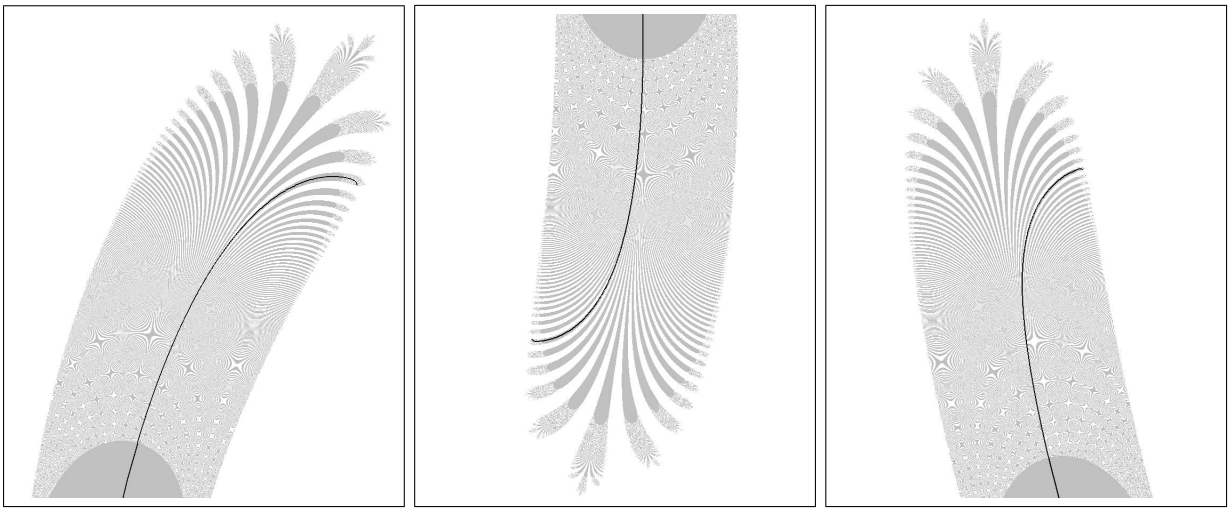

Figure 3.2: These pictures are intended to illustrate Theorem 3.4.2. Shown is the ray at address 01234 . . . for the parameter $\kappa = -2$ in three successive magnifications. The images demonstrate that the ray indeed spirals around its escaping endpoint.

The two entries that remain to be dealt with tend to 0 because $\arg G(\mathcal{F}^j(\underline{s}, t_{\underline{s}}))$ does. This proves the claim in the second case.

To conclude the proof, we need to show that the convergence of the sum $\Theta(t_{\underline{s}})$ is equivalent to the convergence in the sum (3.3) of the statement of the theorem. It is clear that the terms of (3.3) converge to 0 if and only if those of $\Theta(t_{\underline{s}})$ do. So we can suppose that $\arg G(\mathcal{F}^k(\underline{s}, t_{\underline{s}})) \to 0$. It is easy to see that then there exists x such that $|G(\mathcal{F}^k(\underline{s}, t_{\underline{s}}))| \geq F^k(x)$ for all large enough k. Because of this and since $|G(\mathcal{F}^k(\underline{s}, t_{\underline{s}})) - Z(\mathcal{F}^k(\underline{s}, t_{\underline{s}}))|$ is bounded by $2 + \pi$, we have

$$\sum_k^{k_0} \left| \arg(G(\mathcal{F}^k(\underline{s}, t_{\underline{s}})) - \kappa) - \sum_k^{k_0} \arg Z(\mathcal{F}^k(\underline{s}, t_{\underline{s}})) \right| \leq \pi \cdot (2 + \pi + |\kappa|) \cdot \sum_k^{k_0} \frac{1}{F^k(x)}.$$

The last sum is clearly absolutely convergent, and so the convergence of the sum $\Theta(t_{\underline{s}})$ and that of

$$\sum_k^{\infty} \arg Z(\mathcal{F}^k(\underline{s}, t_{\underline{s}}))$$

are equivalent. Similarly, one sees that the convergence of this last sum and the sum (3.3) are equivalent. ■

3.5 Canonical Correspondence and Escaping Set Rigidity

The bijection $G : X \to I(\kappa)$ constructed in the previous section — while having certain continuity properties — is, in general, quite far from being a conjugacy. The question poses itself whether one could, using another construction, create a different map which is a conjugacy, or at least is continuous on a larger set. We will now show that this is not the case. The reason is that otherwise points with different escape speeds must be identified under this map. On the other hand, the periodical structure of the dynamical plane would have to be preserved, so that the kind of stretching that is happening in the real direction cannot occur in the imaginary direction. Both of these directions interact under the map, and one derives a contradiction. This argument is essentially that of Douady and Goldberg [26] showing that for $\kappa_1, \kappa_2 \in (-1, \infty)$, the maps E_{κ_1} and E_{κ_2} are not conjugate on their Julia sets. In the terminology of our model, the proof will become somewhat simpler.

The second statement of the next theorem is somewhat complicated. The reader should simply think of it as a more precise version of the first part which makes the theorem more applicable.

3.5.1 Theorem (No Nontrivial Self-Conjugacies)
Let $Q > 0$ and suppose that $f : Y_Q \cap X \to X$ is a continuous map with $f \circ \mathcal{F} = \mathcal{F} \circ f$ and

$$f(\underline{r}, t) \in \{\underline{r}\} \times (0, \infty) \tag{3.4}$$

for all $\underline{r}$ and t. Then f is the identity.

More precisely, let $(\underline{s}, t_0) \in Y_Q \cap X$. Suppose that a function

$$f : \bigcup_{j \geq 0} \sigma^{-j}(\mathcal{F}(\underline{s}, t_0)) \cap Y_Q \to X$$

satisfies $f \circ \mathcal{F} = \mathcal{F} \circ f$ and (3.4). If $f(\underline{s}, t_0) \neq (\underline{s}, t_0)$, then f is not continuous in $(\underline{s}, t_0)$.

Proof. The first statement follows immediately from the second. So let $x := (\underline{s}, t_0)$ and f be as in the second part of the theorem, and suppose that $f(\underline{s}, t_0) = (\underline{s}, t_0')$ with $t_0 \neq t_0'$. If n is large enough, we can find $m(n) \in \mathbb{N}$ such that

$$0 \leq \log(2\pi m(n) + t_0 + 1) - T(\mathcal{F}^n(x)) < 1.$$

Let us define $y_n := \big(m(n)\, s_2\, s_3\, s_4 \ldots, t_0\big)$. Now pull back the points y_n along the orbit of x. More precisely, let z_n be the uniquely defined point with address $s_1 s_2 \ldots s_n m(n) s_2 s_3 \ldots$ such that $\mathcal{F}^{n+1}(z_n) = y_n$.

By the choice of $m(n)$,

$$0 \leq T(\mathcal{F}^n(z_n)) - T(\mathcal{F}^n(x)) < 1,$$

and thus

$$T(\mathcal{F}^j(z_n)) - T(\mathcal{F}^j(x)) < F^{-(n-j)}(1)$$

for every $j \leq n$. In particular, $z_n \in Y_Q$ and $z_n \to x$.

On the other hand, consider the image points $f(z_n)$. Because $\mathcal{F}^n(f(z_n)) = f(y_n) = (m(n) s_2 s_3 \ldots, t_0')$, we have, once more,

$$|T(f(z_n)) - T(z_n)| < F^{-(n+1)}\big(|t_0' - t_0|\big) \to 0.$$

Thus
$$\lim_{n\to\infty} f(z_n) = \lim_{n\to\infty} z_n = x \neq f(x) = f\left(\lim_{n\to\infty} z_n\right),$$
so f is not continuous in x. ■

An immediate consequence is that there is no other map with the same continuity properties as our map G. In fact, we can make a stronger statement.

3.5.2 Corollary (No Other Conjugacies)
Suppose that E_{κ_1} and E_{κ_2} have nonescaping singular values and are conjugate on their sets of escaping points by a conjugacy f that sends each external ray of E_{κ_1} to the corresponding external ray of E_{κ_2}. Then $f = G^{\kappa_2} \circ (G^{\kappa_1})^{-1}$.

Proof. Let
$$\Phi : I(E_{\kappa_1}) \to I(E_{\kappa_1})$$
be defined by $\Phi := f^{-1} \circ G^{\kappa_2} \circ (G^{\kappa_1})^{-1}$. Then if If Q is large enough, Φ is continuous on $G^{\kappa_1}(Y_Q)$. Suppose that $z = G(\underline{s}, t_0) \in I(E_{\kappa_1})$ such that $\Phi(z) \neq z$. We may assume (by possibly exchanging κ_1 and κ_2 and passing to a forward image of z if necessary) that $z \in Y_{Q+\varepsilon}$ and $T(G^{-1}(\Phi(z))) > t_0$.

Then the function
$$f : \bigcup_{j\geq 0} \sigma^{-j}(\mathcal{F}(\underline{s}, t_0)) \cap Y_Q \to Y_Q, (\underline{r}, t) \mapsto G^{-1}(\Phi(G(\underline{r}, t)))$$
is continuous. This contradicts Theorem 3.5.1. ■

One expects that the condition of every external ray being send to the corresponding ray in the limit dynamics is satisfied by every "reasonable" conjugacy (up to a possible relabeling of the combinatorics). In order to make this precise, let us introduce some additional notation. When discussing conjugacies between functions of the plane, one natural requirement is that the conjugacy is orientation-preserving. When we talk about conjugacies on the sets of escaping points, as we will in this section, then this notion should be replaced by the preservation of the order of external rays. More precisely, a continuous function from some subset of X to X is called order-preserving if the map on external addresses that it induces is order-preserving. Similarly, if $\kappa_1, \kappa_2 \in \mathbb{C}$ and $f : I(E_{\kappa_1}) \to I(E_{\kappa_2})$ is continuous, then we call f order-preserving if it maps external rays in such a way that the order of their external addresses is preserved. Note that any orientation-preserving conjugacy of two exponential maps will induce an order-preserving map on their sets of escaping points. (We implicitly use the fact that the lexicographic order of external addresses coincides with the vertical order of external rays, see Lemma 3.7.1.)

3.5.3 Lemma (Self-Conjugacies of the Shift)
Let f be an order-preserving homeomorphism of the space of external addresses such that $f \circ \sigma = \sigma \circ f$. Then there exists $j \in \mathbb{Z}$ such that
$$f(s_1 s_2 s_3 \dots) = (s_1 + j)(s_2 + j)(s_3 + j)\dots$$

for all external addresses $\underline{s} = s_1 s_2 s_3 \dots$.

PROOF. Recall that $\overline{j}$ denotes the constant address $jjj\dots$. The image of every constant address must also be constant, and without loss of generality, we may assume that $f(\overline{0}) = \overline{0}$. (Otherwise, we can consider the map $\tilde{f}(\underline{s}) = (\tilde{s}_1 - j)(\tilde{s}_2 - j)\dots$, where $\underline{\tilde{s}} = f(\underline{s})$ and $f(\overline{0}) = \overline{j}$).) Note that f must map every periodic address to an address of the same period; since periodic addresses are dense, it is sufficient to prove that f fixes these. Because f is order-preserving, it must fix every address of period 1. If $\underline{s}$ is such that $|s_k| \leq A$ for all k, then $\underline{s}$ and all its images under σ lie between the period 1 addresses $\overline{-A}$ and $\overline{A}$. Because f is order-preserving, all entries of $f(\underline{s})$ lie between $-A$ and A as well.

Now, fix some $A, N \in \mathbb{N}$. By the above, f is an order-preserving permutation of the space of all periodic addresses $\underline{s}$ of period $\leq N$ which satisfy $|s_k| \leq A$ for all k. However, this set is finite, so all its images must be fixed by f. ∎

3.5.4 Corollary (Order-Preserving Conjugacies)
Suppose that f is an order-preserving conjugacy between two maps E_{κ_1} and E_{κ_2} (with nonescaping singular orbit) on their escaping sets. Then there exists a parameter $\kappa_2' = \kappa_2 + 2\pi i k$ such that

$$G_{\kappa_2'} \circ G_{\kappa_1}^{-1} : I(E_{\kappa_1}) \to I(E_{\kappa_2})$$

is a conjugacy.

PROOF. The map f induces an order-preserving self-conjugacy of the shift. By Lemma 3.5.3 this map consists of shifting all labels by some number k. Let $\kappa_2' := \kappa_2 - 2\pi i k$. The maps E_{κ_2} and $E_{\kappa_2 - 2\pi i k}$ are conjugate by the map $z \mapsto z - 2\pi i k$, and the induced self-conjugacy of the shift consists of shifting all labels by $-k$. Thus the map $f' : z \mapsto f(z) - 2\pi i k$ is a conjugacy between E_{κ_1} and E_{κ_2} which preserves external rays. The claim follows by Corollary 3.5.2. ∎

The next theorem is a generalization of the result of Douady and Goldberg [26].

3.5.5 Theorem (No Conjugacy for Escaping Parameters)
Let $\underline{s}$ be an external address and let $t_1 \neq t_2$ with $(\underline{s}, t_1), (\underline{s}, t_2) \in X$. Suppose that κ_1 and κ_2 are parameters such that $G^{\kappa_1}(\underline{s}, t_1) = \kappa_1$ and $G^{\kappa_2}(\underline{s}, t_2) = \kappa_2$. Then E_{κ_1} and E_{κ_2} are not conjugate on $\mathbb{C}$.

PROOF. By contradiction, let $f : \mathbb{C} \to \mathbb{C}$ be a conjugacy between E_{κ_1} and E_{κ_2}. For some set $Q > 0$, the map

$$\alpha : Y_Q \to X; (\underline{s}, t) \mapsto (G^{\kappa_2})^{-1}\left(f(G^{\kappa_1}(\underline{s}, t))\right)$$

is defined. α is either order-preserving or order-reversing, depending on whether f is orientation-preserving or -reversing. Also $\alpha \neq \mathrm{id}$ because $\alpha(\mathcal{F}^n(\underline{s}, t_1)) = \mathcal{F}^n(\underline{s}, t_2) \neq \mathcal{F}^n(\underline{s}, t_1)$.

Note that $\alpha(\sigma^n(\underline{s})) = \sigma^n(\underline{s})$. Thus, if $\underline{s}$ is not constant, then α preserves external addresses (i.e. satisfies (3.4). As in the proof of Corollary 3.5.2, this contradicts Theorem 3.5.1.

So suppose that $\underline{s}$ is constant, say $\underline{s} = \overline{0}$, and that α is order-reversing. Then the map $\widetilde{\alpha}(s_1 s_2 \dots) := \alpha((-s_1)(-s_2)\dots)$ preserves external addresses, and we can again apply Theorem 3.5.1. ■

3.5.6 Theorem (Rigidity for Parameters with Different Combinatorics)
Suppose that $\underline{s}$ is an external address and κ is a nonescaping parameter for which the singular value is contained in the limit set of $g^{\kappa}_{\underline{s}}$. Suppose furthermore that κ' is another nonescaping parameter such that the limit set of $g^{\kappa'}_{\underline{s}}$ does not contain the singular value and is bounded. Then $G^{\kappa} \circ (G^{\kappa'})^{-1}$ is not continuous.

Proof. We first prove the theorem under the simplifying assumption that $g^{\kappa'}_{\underline{s}}$ lands, and comment at the end on the minor changes necessary if it does not. Let

$$f := G^{\kappa} \circ (G^{\kappa'})^{-1} : I(E_{\kappa'}) \to I(E_{\kappa}).$$

Let z_0 be the landing point of $g^{\kappa'}_{\underline{s}}$, and let $w \in I(E_{\kappa'})$. Then (because $w \in J(E_{\kappa'})$), there exists a sequence z_n of preimages of z_0 under iterates of $E_{\kappa'}$ which converges to w. For every n, there is a preimage $\underline{s}^n$ of $\underline{s}$ such that $g^{\kappa'}_{\underline{s}^n}$ lands at z_n. In particular, there exists a sequence t_n such that $g^{\kappa'}_{\underline{s}^n}|_{(0,t_n]} \to w$ uniformly.

On the other hand, the rays $g^{\kappa}_{\underline{s}^n}$ are all preimages of $g^{\kappa}_{\underline{s}}$, which accumulates on κ. Thus $g^{\kappa}_{\underline{s}^n}$ accumulates at ∞, and so we can find $t'_n < t_n$ such that $g^{\kappa}_{\underline{s}^n}(t'_n) \to \infty$. Thus

$$g^{\kappa'}_{\underline{s}^n}(t'_n) \to w,$$

but

$$f(g^{\kappa'}_{\underline{s}^n}(t'_n)) = g^{\kappa}_{\underline{s}^n}(t'_n) \not\to f(w').$$

This shows that f is not continuous.

If $g^{\kappa'}_{\underline{s}}$ does not land, let A denote its accumulation set. The proof will work analogously if we find a sequence of preimages of A which converges uniformly to w. This is easy to achieve as follows. First note that every component of the preimage of A is compact. Indeed, otherwise there is no continuous branch of E_{κ}^{-1} on A, which means that A separates κ from ∞. However, this is impossible: if $\kappa \in J(E_{\kappa})$, then there must be escaping points close to κ, which are connected to ∞ by an external ray, and if $\kappa \in F(E_{\kappa})$, then it is easy to see that there is a curve in the Fatou set which connects κ to ∞ (see Section 4.1). All preimages of A are translates of each other; let K denote the diameter of any of these preimages. For n sufficiently large, let U_n be a small neighborhood of w which is mapped biholomorphically to $\mathbb{D}_{2\pi+1}(E^n_{\kappa}(w))$ by E^n_{κ}. (It is easy to see that such a neighborhood exists for large enough n.) Chose among the preimages of A that one, call it A_n, which satisfies

$$|E^{n+1}(w) - \kappa| \leq |z - \kappa| \leq |E^{n+1}(w) - \kappa| + 2\pi + K$$

and for all $z \in A_n$. If n is so large that $|E^{n+1}(w) - \kappa| > 2\pi + K$, then

$$\log|z - \kappa| - \log|E^{n+1}(w) - \kappa| \leq \frac{2\pi + K}{|E^{n+1}(w) - \kappa|} < 1.$$

Thus if we take the pullback A'_n of A_n by the same branch of E_κ^{-1} that carries $E^{n+1}(w)$ to $E^n(w)$, then $A'_n \subset \mathbb{D}_{2\pi+1}(E_\kappa^n(w))$. We can then further pull back A'_n to U_n, and the diameters of the U_n shrink to 0 by Koebe's theorem. This concludes the proof. ■

The preceding result yields Theorem 1.4, which we state here in a slightly stronger form. The proof will require results from the following chapters, but we decided to include it here in order to keep the rigidity results in one place.

3.5.7 Corollary (Rigidity for Escaping and Misiurewicz Parameters)
Suppose that $\kappa_1 \neq \kappa_2$ are attracting, parabolic, Misiurewicz or escaping parameters, at least one of which is not attracting or parabolic. Suppose that $\operatorname{Im}\kappa_1, \operatorname{Im}\kappa_2 \in (-\pi, \pi]$. Then E_{κ_1} and E_{κ_2} are not conjugate on their sets of escaping points by an order-preserving conjugacy.

Proof. Clearly an escaping parameter cannot be conjugate to a nonescaping parameter. So let us first suppose that κ_1 and κ_2 are escaping parameters, and suppose the singular values lie on the rays at external addresses $\underline{s}^1$ and $\underline{s}^2$. Then by Remark 5.12.3, both addresses have first entry 0. Since the conjugacy must map the singular value of E_{κ_1} to that of E_{κ_2}, it follows by Lemma 3.5.3 that $\underline{s}^1 = \underline{s}^2$. As in the proof of Theorem 3.5.5, it follows that their potentials are equal as well. However, this contradicts Theorem 5.12.2.

Now suppose that both κ_1 and κ_2 are Misiurewicz. Assume that the preperiod of κ_1 is smaller or equal to that of κ_2. By Theorem 4.4.2, there exists a preperiodic address $\underline{s}$ such that $g_{\underline{s}}^{\kappa_1}$ lands at κ_1. Because of Theorem 3.9.1 (and because the preperiod of κ_2 is greater or equal to that of κ_1), $g_{\underline{s}}^{\kappa_2}$ lands at a preperiodic point. This point is $\neq \kappa_2$ by Theorem 4.4.4. Thus we can apply Theorem 3.5.6. The same argument works (without reference to Theorem 4.4.2) if κ_2 is parabolic or attracting. ■

It seems reasonable to conjecture that the escaping dynamics of exponential maps whose singular value lies in the Julia set is rigid. This would imply density of hyperbolicity. Indeed, a non-hyperbolic stable parameter would be (quasiconformally) conjugate to all nearby parameters, and in particular the maps would be conjugate on their sets of escaping points. (See Section 5.1.)

3.5.8 Conjecture (Escaping Set Rigidity)
Suppose that κ_1 is a parameter such that $\kappa_1 \in J(E_{\kappa_1})$, and let $\kappa_2 \notin \{\kappa_1 + 2\pi i k\}$. Then there exists no order-preserving conjugacy

$$f : I(E_{\kappa_1}) \to I(E_{\kappa_2})$$

between E_{κ_1} and E_{κ_2}.

3.6 Speed of Escape

Our model gives us very good control over the speed with which escaping points escape, as we shall see in this section. Let us first discuss some properties of external addresses which were considered in the construction of [77].

3.6.1 Definition (Properties of External Addresses)
Let $\underline{s}$ be an external address.

- $\underline{s}$ *is called* exponentially bounded *if there exists $x > 0$ such that, for all $k \geq 1$,*

$$2\pi|s_k| < F^{k-1}(x).$$

- *We say that $\underline{s}$ has* positive minimal potential, *if there exists an $x > 0$ such that*

$$2\pi|s_k| > F^{k-1}(x)$$

  for infinitely many k.

- *If S is a family of external addresses, then we call S* uniformly exponentially bounded *if all addresses in S are exponentially bounded with the same x.*

- *An external address $\underline{s}$ is called* slow *if there exists a subsequence $\sigma^{n_k}(\underline{s})$ of the iterates of $\underline{s}$ under the shift map which is uniformly exponentially bounded. Otherwise, $\underline{s}$ is called* fast.

It is known that escaping points exist exactly for exponentially bounded addresses, and that a ray of escaping points has an escaping endpoint if and only if its address is fast [77]. The results of Section 3.3 show that our model has the same properties; however, this can be easily proved by elementary means. Moreover, the last statement of the following lemma, concerning the speed of escape of the endpoints of addresses with positive minimal potential, is new.

3.6.2 Lemma (Speed of Escaping Points)
Let $\underline{s}$ be an external address. Then

- *$t_{\underline{s}} < \infty$ if and only if $\underline{s}$ is exponentially bounded.*

- *$\underline{s}$ is a fast external address if and only if $(\underline{s}, t_{\underline{s}}) \in X$.*

- *$\underline{s}$ has positive minimal potential if and only if there exists a positive x for which $T(\mathcal{F}^j(\underline{s}, t_{\underline{s}})) > F^j(x)$ for all j.*

Proof. Let $\underline{s}$ be an external address with $t_{\underline{s}} < \infty$. Then, by definition,

$$F(T(\mathcal{F}^{k-1}(\underline{s}, t_{\underline{s}}))) \geq 2\pi|s_{k+1}|.$$

Since $T(\mathcal{F}^{k-1}(\underline{s}, t_{\underline{s}})) \leq F^k(t_{\underline{s}})$, it follows that

$$2\pi|s_{k+1}| \leq F^k(t_{\underline{s}}).$$

Conversely, let $\underline{s}$ be an exponentially bounded external address, say $2\pi|s_k| < F^{k-1}(x)$ for all $k \geq 2$; we may suppose that $x \geq \log 2$. We show that $(\underline{s}, 2x) \in X$. More precisely, we will show by induction that $T(\mathcal{F}^k(\underline{s}, 2x)) \geq 2F^k(x)$. This is trivial for $k = 0$. Furthermore,

$$\begin{aligned} T(\mathcal{F}^{k+1}(\underline{s}, 2x)) &= F(T(\mathcal{F}^k(\underline{s}, 2x))) - 2\pi|s_{k+2}| \\ &\geq F(2F^k(x)) - F^{k+1}(x) = F^{k+1}(x)\left(\exp(F^k(x))\right) \geq 2F^{k+1}(x). \end{aligned}$$

Let us now prove the statement about fast addresses. By the above, a family of addresses $(\underline{s}^k)$ is uniformly exponentially bounded if and only if the entries (s_1^k) and the family $(t_{\underline{s}^k})$ are both uniformly bounded. Thus if $\underline{s}$ is a slow address, then some subsequence of $(t_{\sigma^n(\underline{s})})$ is bounded, and thus $t_{\sigma^n(\underline{s})} \not\to \infty$. Conversely, if a subsequence $t_{\sigma^{n_k}(\underline{s})}$ is uniformly bounded, then the sequence $t_{\sigma^{n_k+1}}$ is uniformly exponentially bounded.

Finally, suppose that $\underline{s}$ has positive minimal potential; say

$$2\pi|s_k| > F^{k-1}(t)$$

for infinitely many k. Let $x = (\underline{s}, t_{\underline{s}})$. Since

$$2\pi|s_k| \leq F(T(\mathcal{F}^{k-2}(x))),$$

we see that $T(F^{k-2}(x)) > F^{k-2}(t)$, and inductively,

$$T(\mathcal{F}^j(x)) > F^j(t) \tag{3.5}$$

for all $j < k - 1$. Since k was arbitrarily large, (3.5) is true for all j.

On the other hand, suppose that $\underline{s}$ does not have positive minimal potential. Let $a > 0$; then there exists an n such that $2\pi|s_k| < F^{k-1}(a)$ for all $k \geq n$. Then, if $k \geq n-1$ is so large that $F^k(a) > 3$,

$$T\left(\mathcal{F}\left(\sigma^k(\underline{s}), 2F^k(a)\right)\right) = F(2F^k(a)) - 2\pi|s_{k+2}| < 3F^{k+1}(a) - F^{k+1}a = 2F^{k+1}(a).$$

Thus $(\mathcal{F}(\sigma^k(\underline{s}), 2F^k(a))) \in X$, and therefore $t_{\sigma^k}(\underline{s}) \leq 2F^k(a)$. ■

The real part of any point $G(\underline{s}, t)$ with $t > t_{\underline{s}}$ grows like an iterated exponential. However, we have just seen that this is the case for the endpoint $G(\underline{s}, t_{\underline{s}})$ if and only if $\underline{s}$ has positive minimal potential in the sense of [77]. We will now show that the real parts of escaping points can grow arbitrarily slowly.

3.6.3 Theorem (Arbitrarily Slow Escape)
Suppose r_n is a sequence of positive real numbers such that $r_n \to \infty$ and $r_{n+1} \leq F(r_n - 5)$. Let $\kappa \in \mathbb{C}$. Then there exists a $z \in I$ and a $n_0 \in \mathbb{N}$ such that, for $n \geq n_0$,

$$|\operatorname{Re}(E^{n-n_0}(z)) - r_n| \leq \pi + 2.$$

Proof. It is sufficient to construct $(\underline{s}, t) \in X$ such that $|T(\mathcal{F}^{n-1}(\underline{s}, t)) - r_n| \leq \pi$. Let $t := r_1$; we will construct the entries of $\underline{s}$ inductively. (Remember that $T\big(\mathcal{F}^n(\underline{s}, t)\big)$ depends only on the first n entries of $\underline{s}$.) The construction is very simple: If $s_1, \dots s_n$ have been constructed, choose s_n such that

$$\big|F(T(F^{n-1}(\underline{s}, t))) - r_n - 2\pi|s_n|\big| \leq \pi.$$

The point $(\underline{s}, r_1)$ clearly has the desired property. ■

In [40], it was shown that, if the singular orbit escapes within some sector $\{z : |\operatorname{Im} z| < T \operatorname{Re} z\}$, the limit set of the orbit of almost every point is the postsingular set. It is clear that this "sector condition" is satisfied for $G(\underline{s}, t)$ if $t > t_{\underline{s}}$, and that it can be satisfied for the endpoint $G(\underline{s}, t_{\underline{s}})$ only if $\underline{s}$ has positive minimal potential. Using our explicit model, we can characterize exactly for which addresses the endpoint does satisfy this condition, and show in particular that there are many addresses with positive minimal potentials for which this is not the case. Note that the set of parameters that satisfies this condition has Hausdorff dimension two [62], which is an analog of the corresponding result by McMullen in the dynamical plane [53]. See Section 6.2 for a discussion.

3.6.4 Theorem (Endpoints that Escape in Sector)
Let $\underline{s}$ be an exponentially bounded external address, and define

$$b := \limsup_{n \to \infty} F^{-(n-1)}(2\pi|s_n|).$$

Set $t_n := T(\mathcal{F}^{n-1}(\underline{s}, t_{\underline{s}}))$. The following two conditions are equivalent:

a) There are K and n_0 such that $2\pi|s_n| < Kt_n$ for $n \geq n_0$.

b) There are C and n_0 such that $w\pi|s_n| < CF^{n-1}(b)$ for $n \geq n_0$.

Remark. b is the minimal potential as defined in [77].

Proof. Let us define $b_n := F^{-(n-1)}(2\pi|s_n|)$. As in the proof of the last statement of Lemma 3.6.2, we have for $k < n$:

$$t_k \geq F^{k-1}(b_n).$$

In particular, $t_k \geq F^{k-1}(b)$, which proves that condition b) implies condition a).

We may assume that for each $n_0 \in \mathbb{N}$, the set $\{b_n : n \geq n_0\}$ has a maximum (otherwise, condition b) is satisfied by definition). Define a sequence n_k inductively by

$$n_{k+1} := \min\{n > n_k : b_n = \max_{m > n_k} b_m\}.$$

Let us show that, for all $n > n_{k-1}$,

$$t_n \leq 2F^{n-1}(b_{n_k}). \tag{3.6}$$

Let $x_n := (\sigma^{n-1}(\underline{s}), 2F^{n-1}(b_{n_k}))$; we must prove that $x_n \in \overline{X}$. Indeed, it is shown by a standard induction that

$$T\big(\mathcal{F}(x_n)\big) \geq T(x_{n+1})$$

for all l, and the claim follows.

Taking the infimum (over k) on the right side of (3.6), it follows that

$$t_n \leq 2F^{n-1}(b),$$

and the claim follows. ■

We believe that these examples are sufficient to illustrate that our model is well suited for the computing of combinatorial conditions corresponding to growth conditions. We should note that, analogously to the previous theorem, one can derive a (nonempty) condition for an endpoint to eventually escape within every parabola $\operatorname{Re} z \geq |\operatorname{Im}(z)|^d$. (By the results of Karpinska [45], the set of such points has Hausdorff dimension 1. Compare the discussion in Section 6.2.)

3.7 Extended Combinatorics

In the following, we will often use the combinatorial information given by external rays (and other curves to ∞). In this section, we discuss the general notions and extend our combinatorial notation a little bit.

Suppose that $\mathcal{C}$ is any family of curves $\gamma : [0,\infty) \to \mathbb{C}$ for which $\operatorname{Re}\gamma(t) \to +\infty$. Suppose furthermore that for any two curves $\gamma_1, \gamma_2 \in \mathcal{C}$, the set of intersection points of γ_1 and γ_2 is bounded. Then $\mathcal{C}$ is equipped with a natural vertical order: of any two curves in $\mathcal{C}$, one is *above* the other. More precisely, define $\mathcal{H}_R := \{z \in \mathbb{C} : \operatorname{Re} z > R\}$ for $R > 0$. If $\gamma \in \mathcal{C}$ and R is large enough, then the set $\mathcal{H}_R \setminus \gamma$ has exactly two unbounded components, one above and one below γ, and any other curve of $\mathcal{C}$ must (eventually) tend to ∞ within one of these.

For the family of external rays, it is an immediate consequence of the construction that this order coincides with the lexicographic order on external addresses.

3.7.1 Lemma (Vertical Order of External Rays)
Let $\underline{s}$ and $\underline{\tilde{s}}$ be exponentially bounded external addresses. Then $g_{\underline{s}}$ is above $g_{\underline{\tilde{s}}}$ if and only if $\underline{s} > \underline{\tilde{s}}$.

Remark. We have already used this statement implicitly in Section 3.5.

Proof. Suppose that $\underline{s} > \underline{\tilde{s}}$. Let j be the first entry at which $\underline{s}$ and $\underline{\tilde{s}}$ differ. Then $s_j > \tilde{s}_j$ by the definition of the lexicographic order. Since $g_{\sigma^{j-1}(\underline{s})}$ tends to ∞ within the strip S_{s_j} and $g_{\sigma^{j-1}(\underline{\tilde{s}})}$ tends to ∞ within $S_{\tilde{s}_j}$, $g_{\sigma^{j-1}(\underline{s})}$ lies above $g_{\sigma^{j-1}(\underline{\tilde{s}})}$. Because the pullback of two curves under E_κ (into the same strip S_k) preserves vertical order, the claim follows. ■

We would now like to use the structure of external rays to assign combinatorics to other curves to infinity. However, the space of external addresses is not order-complete with respect to the lexicographic ordering; i.e., not every bounded subset has a supremum. Therefore we will now add the so-called *intermediate external addresses* to our repertoire. An intermediate external address is a finite sequence of the form

$$\underline{s} = s_1 s_2 \dots s_{n-1}\infty,$$

where $s_1, \dots, s_{n-2} \in \mathbb{Z}$ and $s_{n-1} \in \mathbb{Z} + \frac{1}{2}$. When we wish to make the distinction, we will refer to an external address in the original sense as an "infinite" external address. We will denote the space of all infinite and intermediate external addresses by $\mathcal{S}$.

To illustrate this definition, consider the following construction. Let

$$f : \mathbb{R} \setminus \bigl\{(2k-1)\pi : k \in \mathbb{Z}\bigr\} \to \mathbb{R}, t \mapsto \tan(2t).$$

Then to any (infinite) external address $\underline{s}$ we can associate a unique point x for which $f^{k-1}(x) \in \bigl((2s_k - 1)\pi, (2s_k + 1)\pi\bigr)$ for all k. However, this subset is missing the countably many preimages of ∞ under the iterates of f. Adding the intermediate external addresses corresponds to filling in these points. The resulting space is then isomorphic to $\mathbb{R}$ (and thus order-complete).

We will also often add ∞ as an (intermediate) external address. The set $\overline{\mathcal{S}} := \mathcal{S} \cup \{\infty\}$ carries a complete circular ordering (and is isomorphic to S^1). Sometimes we will also form a quotient by identifying all elements of $\mathcal{S}$ whose first entries differ by an integer. (We denote the elements of this space by replacing the first entry of $\underline{s}$ by an asterisk $*$.) The resulting space, denoted by $\tilde{\mathcal{S}}$, is then mapped isomorphically to $\overline{\mathcal{S}}$ under the shift. This should be thought of as analogous to E_κ being a conformal isomorphism from the cylinder $\widetilde{\mathbb{C}} = \mathbb{C}/2\pi i\mathbb{Z}$ to $\mathbb{C} \setminus \{\kappa\}$. To stress this analogy, we will refer to $\tilde{\mathcal{S}}$ as a "cylinder". Of course, $\tilde{\mathcal{S}}$ is topologically a circle, but we imagine each point having a ray attached to it.

Now let $\gamma : [0, \infty) \to \mathbb{C}$ be any curve with $\operatorname{Re}\gamma(t) \to \infty$ and suppose that for every $\underline{s}$, the set $\{t : g_{\underline{s}}(t) \in \gamma\}$ is bounded. Then we can assign an external address to this curve:

$$\begin{aligned} \operatorname{addr}(\gamma) &:= \sup\{\underline{s} : \gamma \text{ is above } g_{\underline{s}}\} \\ &\qquad \bigl(= \inf\{\underline{s} : \gamma \text{ is below } g_{\underline{s}}\}\ \bigr), \end{aligned}$$

where $\operatorname{addr}(\gamma) = \infty$ if the set in the definition is unbounded or empty. If γ is a curve for which $\gamma(t) \to -\infty$ or $|\operatorname{Im}\gamma(t)| \to \infty$, then we also set $\operatorname{addr}(\gamma) := \infty$.

As an example, consider the curve $\gamma : [0, \infty) \to \mathbb{C}, t \mapsto t + (2k+1)\pi$, $k \in \mathbb{Z}$. This curve lies between the strips S_k and S_{k+1}. As we would expect, $\operatorname{addr}(\gamma) = (k + \frac{1}{2})\infty$.

3.7.2 Remark

If γ *is a curve as above and* $\operatorname{addr}(\gamma) \neq \infty$, *then* $\operatorname{addr}(E_\kappa \circ \gamma) = \sigma(\operatorname{addr}(\gamma))$. ■

In the following, we shall sometimes encounter external rays which have the property that $\lim_{t\to 0} g_{\underline{s}}(t) = \infty$. These rays then have two external addresses: the ordinary external address $\underline{s}$, but also the address $\underline{s}'$ of the curve $t \mapsto g_{\underline{s}}(1/t)$. For brevity of notation, we shall denote this latter address by

$$\operatorname{addr}^-(\underline{s}) := \operatorname{addr}^-(g_{\underline{s}}) := \underline{s}'.$$

Finally, we will occasionally (particularly in the next section) need to deal with the case of curves which do not land at ∞, but do accumulate there. To such a curve γ, we want to associate an accumulation set of addresses. For a fixed $R > 0$, we shall say that a point z is *separated from γ in $\mathcal{H}_R$* if there exists some external ray g such that z and γ lie in different components of $\mathcal{H}_R \setminus g\big([1,\infty)]\big)$. We define

$$\mathrm{ADDR}(\gamma) := \bigcap_{R>0} \overline{\{\underline{s} : g_{\underline{s}}\big([1,\infty)\big) \text{ is not separated from } \gamma \text{ in } \mathcal{H}_R.\}},$$

where the closure is taken in $\overline{\mathcal{S}}$. The set $\mathrm{ADDR}(\gamma)$ is closed, and if $\mathrm{addr}(\gamma)$ is defined, then $\mathrm{ADDR}(\gamma) = \{\mathrm{addr}(\gamma)\}$. Again, in the case where an external ray accumulates at ∞ we shall denote the corresponding set by $\mathrm{ADDR}^-(g_{\underline{s}}) := \mathrm{ADDR}^-(\underline{s})$.

3.8 Limiting Behavior of External Rays

In this section, we will analyze the limiting behavior of rays more closely. In particular, we shall show that in many cases there exist rays which do not land, and in fact limit on themselves. However, let us first prove two results which essentially already appear in [77, Proof of Corollary 6.9] and show that a ray at a "slow" address cannot land at an escaping point. This proves that every ray is indeed a path connected component of I, and thus completes the proof of Theorem 1.1.

3.8.1 Lemma (Limit Set of Ray)
Let $g : (0,\infty) \to I$ be an external ray, and let

$$L := \bigcap_{t>0} \overline{g\big((0,t)\big)}$$

denote the limit set of L. If there is $(\underline{s}, t_0)$ such that $g_{\underline{s}}(t_0) \in L$, then

$$g_{\underline{s}}\big((0,t_0]\big) \subset L.$$

PROOF. Let us define addresses $\underline{s}^{n\pm}$ by

$$s_k^{n\pm} = \begin{cases} s_k \pm 1 & k = n \\ s_k & \text{otherwise.} \end{cases}$$

One sees easily that $t_{s^{n\pm}} \to t_{s^0}$ for $n \to \infty$.

Now pick some t, $0 < t < t^0$. Then

$$g_{\underline{s}^{n\pm}}\big([t,\infty)\big) \to g_{\underline{s}}\big([t,\infty)\big)$$

uniformly. Therefore any curve which does not intersect the $g_{\underline{s}^{n\pm}}$ and accumulates at $g_{\underline{s}}(t_0)$ must also accumulate at $g_s(t)$. ∎

3.8.2 Lemma (Addresses of Rays Landing Together)
Suppose that $\underline{s}$ and $\underline{s}'$ are two addresses such that the rays $g_{\underline{s}}$ and $g_{\underline{s}'}$ land at the same point. Then $|s_k - s'_k| \leq 1$ for all k.

PROOF. Assume, by contradiction, that $|s_k - s'_k| > 1$ for some k; by passing to a forward iterate if necessary we can assume that $k = 1$. Let S denote the union of $g_{\underline{s}}$, $g_{\underline{s}'}$ and their common landing point z_0, which is a Jordan arc tending to ∞ in both directions. Note that E_κ is injective on S. Indeed, E_κ is injective on every ray, and it is injective on $g_{\underline{s}} \cup g_{\underline{s}'}$ unless $\underline{s}$ and $\underline{s}'$ differ only in their first entries. However, in that case $g_{\underline{s}'}$ would be a translate of $g_{\underline{s}}$ by a multiple of $2\pi i$, which means that $g_{\underline{s}}$ and $g_{\underline{s}'}$ cannot land together. Finally, E_κ is injective on S as otherwise z_0 would lie on an external ray, which contradicts Lemma 3.8.1.

On the other hand, if $|s_1 - s'_1| > 1$, then the two ends of R tend to ∞ with a difference of more than 2π in their imaginary parts. This means that $R \cap (R + 2\pi i) \neq \emptyset$, and thus E_κ is not injective on R. This is a contradiction. ■

3.8.3 Corollary (No Landing at Escaping Points)
Suppose $\underline{s}$ is a slow address. Then $g_{\underline{s}}$ does not land at an escaping point.

PROOF. By Lemma 3.8.1, $g_{\underline{s}}$ could only land at the escaping endpoint of some ray $g_{\underline{s}'}$ with fast address $\underline{s}'$. However, because $\underline{s}$ is slow and $\underline{s}'$ is fast, there is a $k \in \mathbb{N}$ such that $s_k + 1 < s'_k$, which contradicts Lemma 3.8.2. ■

We now prove the main theorem of this section.

3.8.4 Theorem (Existence of Nonlanding Rays)
Suppose that there exists $\underline{s}$ such that the external ray $g_{\underline{s}}$ either lands at κ or contains κ. Then there is $\underline{r}$ such that the accumulation set of $g_{\underline{r}}$ contains $g_{\underline{r}}$.

REMARK. We will actually show the existence of uncountably many such $\underline{r}$. By a result of Curry [16], the accumulation set of such a ray either disconnects the plane into infinitely many components, has topological dimension 2 or is an indecomposable continuum. It van be shown that, under the hypotheses of Theorem 3.8.4, the limit set is always an indecomposable continuum.

The following lemma reduces Theorem 3.8.4 to the question of constructing an address $\underline{r}$ with $\underline{r} \in \mathrm{ADDR}^-(g_{\underline{r}})$.

3.8.5 Lemma (Accumulation Addresses)
Let $g_{\underline{r}}$ be an external ray and suppose that $\underline{s} \in \mathrm{ADDR}^-(\underline{r})$ is exponentially bounded. Then the accumulation set of $g_{\underline{r}}$ contains the ray $g_{\underline{s}}$.

PROOF. Analogous to that of Lemma 3.8.1. ■

PROOF OF THEOREM 3.8.4. We shall first restrict to the case where $\underline{s}$ is bounded. At the end of the proof, we will remark on the modifications required to deal with the case that $\underline{s}$ is unbounded.

Let $m \in \mathbb{Z}$ such that $m\underline{s} < \underline{s}$, and let $K \in \mathbb{Z}$ be such that $s_k < K$ for all k. We define $\underline{s}' := Km\underline{s}$; note that $g_{\underline{s}'}$ lands at ∞. The address $\underline{r}$ will consist of a concatenation of ever longer initial blocks of $\underline{s}'$; i.e. it will be of the form

$$\underline{r} = Kms_1s_2\dots s_{n_1}Kms_1s_2\dots s_{n_2}Kms_1s_2\dots s_{n_3}\dots\ . \tag{3.7}$$

Note first that there exists a such that $t_{\underline{r}} \leq a$ for any address $\underline{r}$ of the form (3.7); for example $a = 2t_{\underline{s}'}$ suffices.

We will define $n_1, n_2, \dots$ inductively. (We shall see that at each step, we have infinitely many choices, yielding uncountably many possible choices of r.) Suppose that $n_1, \dots, n_k$ have already been chosen. Consider the address

$$\underline{s}^k := Kms_1s_2\dots s_{n_1}\dots Kms_1s_2\dots s_{n_k}\underline{s}'.$$

The ray $g_{\underline{s}^k}$ is a preimage of $g_{\underline{s}}$ and thus lands at ∞. We claim that

$$\operatorname{addr}^-(\underline{s}^k) = Kms_1s_2\dots s_{n_1}\dots Kms_1s_2\dots s_{n_k}\left(K - \frac{1}{2}\right)\infty. \tag{3.8}$$

The proof proceeds by induction on the steps of pullbacks taken to obtain $\underline{s}^k$ from $m\underline{s}$ (which has $\operatorname{addr}^-(m\underline{s}) = \infty$). Recall that the preimages of $g_{\underline{s}}$ disconnect the plane. Thus, there is no preimage of $\underline{s}$ between $\operatorname{addr}^-(Km\underline{s})$ and $Km\underline{s}$. Since

$$K\underline{s} > Km\underline{s} > \left(K - \frac{1}{2}\right)\infty > (K-1)\underline{s},$$

it follows that

$$\operatorname{addr}^-(Km\underline{s}) = \left(K - \frac{1}{2}\right)\infty.$$

For the induction step, suppose we know that $\operatorname{addr}^-(\sigma^j(\underline{s}^k))$ has the correct address. Then, by the same reasoning as above, we only need to show that $\operatorname{addr}^-(\sigma^j(\underline{s}^k))$ and $\sigma^j(\underline{s}^k)$ do not surround $\underline{s}$. However, these two addresses agree on all entries up to the second but last entry of $\operatorname{addr}^-(\sigma^j(\underline{s}^k))$, in which one is K and the other $K - \frac{1}{2}$. Since all entries of $\underline{s}$ are smaller than K, the claim follows.

Choose some $t_n < \frac{1}{n}$ such that $|g_{\underline{s}^n}(t)|$ is larger than, say, k, and choose δ_n small enough (see below). By Remark 3.3.6, there exists N such that, for any address $\underline{r}$ of the form (*) with $n_{k+1} > N$, we have $|g_{\underline{r}}(t) - g_{\underline{s}^n}(t)| < \delta_n$ for $t \geq t_n$. We can then choose n_{k+1} such that $n_{k+1} > N$. This completes the inductive construction.

Clearly if the δ_n were chosen small enough, then $\underline{r} \in \operatorname{ADDR}^-(\underline{r})$. By Lemma 3.8.5, this completes the proof in the case where $\underline{s}$ is bounded.

If $\underline{s}$ is unbounded, we construct the address $\underline{r}$ in a slightly different way. Suppose without loss of generality that the entries of $\underline{s}$ are not bounded from above. We again choose m such that $m\underline{s} < \underline{s}$ and consider addresses of the form

$$\underline{r} = ms_1\dots s_{n_1}ms_1\dots s_{n_2}\dots,$$

where each n_k is chosen so that $s_n < s_{n_k}$ for $n < n_k$. The rest of the proof then proceeds analogously. ■

We will now adopt the proof to the case where the ray does not necessarily land at κ but accumulate there.

3.8.6 Theorem (Existence of Nonlanding Rays, II)
Suppose that $\underline{s}$ is an external address such that $g_{\underline{s}}$ accumulates at κ. Then there exist $\underline{r}$ and $\underline{r}'$ such that the accumulation set of $g_{\underline{r}}$ contains $g_{\underline{r}'}$.

Remark. We believe that one can always achieve $\underline{r} = \underline{r}'$. In most cases the proof directly gives this; however we are currently unable to show this in general.

The construction is largely the same, if somewhat more complicated because we need to deal with rays for which ADDR^- consists of more than one element. Let us first prove the following simple topological facts.

3.8.7 Lemma
Let $\underline{r}$ and $\underline{s}$ be exponentially bounded external addresses. Let $\underline{t}$ be any (infinite or intermediate) external address.

a) Suppose that $\inf \mathrm{Re}(g_{\underline{s}}(t)) = -\infty$ and $\underline{t} \in \mathrm{ADDR}^-(\underline{r})$. If $\underline{r} < \underline{s} < \underline{t}$ or $\underline{t} < \underline{s} < \underline{r}$, then $\inf \mathrm{Re}(g_{\underline{r}}(t)) = -\infty$.

b) Suppose that $\underline{t} \in \mathrm{ADDR}^-(\underline{r})$ and $\underline{t}' \in \mathrm{ADDR}^-(\underline{s})$. If $\underline{r} < \underline{s} < \underline{t} < t'$ or $\underline{t}' < \underline{t} < \underline{s} < \underline{r}$, then $\underline{t}' \in \mathrm{ADDR}^-(\underline{r})$.

c) Suppose that $\underline{t} \in \mathrm{ADDR}^-(\underline{s})$. If $\underline{t}^- := (t_1 - 1)t_2t_3 \cdots > \underline{s}$, then $\underline{t}^- \in \mathrm{ADDR}^-(\underline{s})$. Similarly, if $\underline{t}^+ := (t_1 + 1)t_2t_3 \cdots < \underline{s}$, then $\underline{t}^+ \in \mathrm{ADDR}^-(\underline{s})$.

Proof. To prove part (a), suppose by contradiction that $\underline{r} < \underline{s} < \underline{t}$ and that $M := \inf \mathrm{Re}(g_{\underline{r}}(t)) > -\infty$. Let $t_0 := \max\{t : \mathrm{Re}\, g_{\underline{s}}(t) = M$. Then $g_{\underline{s}}$ separates $g_{\underline{r}}$ from any address $> \underline{s}$, which is a contradiction.

Part b) is topologically the same as part a). To prove part c), note first that $\underline{t}^+ \in \mathrm{ADDR}^-(g_{\underline{s}} + 2\pi i)$. By b), also $\underline{t} \in \mathrm{ADDR}^-(g_{\underline{s}} + 2\pi i)$, and thus $\underline{t}^- \in \mathrm{ADDR}^-(\underline{s})$. ■

Sketch of the Proof of Theorem 3.8.6. The proof is essentially the same as that of Theorem 3.8.4. However, the assertion (3.8) will be replaced by showing that $\mathrm{ADDR}^-(\underline{s}^k)$ contains an address $\underline{t}^k$ whose entries do not differ by more than 1 from the entries of

$$Kms_1s_2 \dots s_{n_1} \dots Kms_1s_2 \dots s_{n_k} \left(K - \frac{1}{2}\right) \infty.$$

This is easily proved applying part (c) of Lemma 3.8.7 repeatedly in the pullback. We leave it to the reader to convince themselves that the rest of the construction of the proof can be modified so that $\mathrm{ADDR}^-(\underline{r})$ contains any limit address of the $\underline{t}^k$. Because all $\underline{t}^k$ are uniformly exponentially bounded, all their limit addresses are also infinite and exponentially bounded. Again, the claim then follows by Lemma 3.8.5. ■

To conclude this section, we give a negative answer to a question of Devaney and Jarque [22]. For simplicity, we restrict here to considering the map $E_0 = \exp$. The article asked whether, for any address $\underline{s}$ which contains arbitrarily long blocks of zeros, the accumulation set is an indecomposable continuum. We show that this is not the case.

3.8.8 Theorem (Landing Ray with Long Blocks of Zeroes)
Suppose that (n_k) is a sequence of positive natural numbers. Then there exists a sequence (m_k) such that the ray at address $\underline{s}$ lands, where $\underline{s}$ consists of m_1 entries of 1, n_1 entries of 0, m_2 entries of 1, n_2 entries of 0 and so on.

Sketch of proof. All addresses considered in the following are assumed to consist only of entries in $\{0,1\}$. Let us say that two addresses $\underline{s}$ and $\underline{s}'$ are *ε-close* if $|g_{\underline{s}}(t) - g_{\underline{s}'}(t)| < \varepsilon$ for all t. We claim that, for any n and any $\varepsilon > 0$, there exists an m such that if $\overline{1}$ and $\underline{s}$ are 1-close, then the addresses $\overline{1}$ and $1^m 0^n \underline{s}$ are ε-close.

Let us consider first the addresses $0^n\overline{1}$ and $0^n\underline{s}$. These addresses are K-close for some (large) K, because the derivative of Log^n is bounded from below on a neighborhood of radius 1 of $g_{\overline{1}}$. Similarly, continually pulling back a point into the strip $\{\mathrm{Im}\, z \in (\pi, 3\pi)\}$ contracts the Euclidean metric. Thus if m is large enough, then $1^m 0^n \underline{s}$ and $1^m 0^n \overline{1}$ are, say, $\frac{\varepsilon}{2}$-close. It is therefore sufficient to show that $\underline{r} := 1^m 0^n \overline{1}$ and $\overline{1}$ are arbitrarily close if m is large enough. Set $\underline{r}' := 0^n\overline{1}$.

Fix some $t_0 > 0$, and let t_1 be such that $E_\kappa^m(g_{\underline{r}}(t_1)) = g_{\underline{r}'}$. It is easy to see that, if m is large enough, then

$$|g_{\underline{r}}(t) - g_{\overline{1}}(t)| \leq \varepsilon$$

for $t \geq t_1$. On the other hand, the diameter of $g_{\underline{r}'}\big((0, t_0]\big)$ is bounded. By the previous contraction argument, if m was large enough, then

$$|g_{\underline{r}}(t) - g_{\overline{1}}(0)| < \frac{\varepsilon}{2}$$

for $0 < t < t_1$, and

$$|g_{\overline{1}}(t) - g_{\overline{1}}(0)| < \frac{eps}{2}.$$

(Here $g_{\overline{1}}(0)$ denotes the landing point of the ray $g_{\overline{1}}$.)

Now, for each k choose such a number m_k for $n = n_k$ and $\varepsilon = \frac{1}{2^k}$. Then the addresses

$$\underline{s}^k := 1^{m_1} 0^{n_1} 1^{m_2} 0^{n_2} \ldots 0^{n_k} \overline{1}$$

converge to an address $\underline{s}$ of the form described in the theorem. By construction, it follows that $\underline{s}$ is 1-close to $\overline{1}$. In particular, the accumulation set of $g_{\underline{s}}$ is bounded. Only slightly more care is required to show that, for each k, $\underline{s}^k$ is $\frac{1}{2^k}$-close to $\underline{s}^{k-1}$, which shows that $g_{\underline{s}}$ lands, as claimed. ∎

3.9 A Landing Theorem for Periodic Rays

In this section, we will discuss landing properties of periodic rays of periodic external rays. While we will eventually show that all periodic rays of exponential maps with nonescaping singular orbit land, the proof uses deep results on parameter space. Also, the proof actually uses the fact that the theorem is known for certain parameters, such as those for which we prove it in this section.

Our proof is an adaptation of the proof that periodic rays land for polynomials with connected Julia set. Such a polynomial is a covering map of the set of escaping points to itself, and therefore locally preserves the hyperbolic metric. So the hyperbolic length of the piece of a fixed ray between any point and its image is constant (and is in fact exactly $\log d$, where d is the degree of the map). Since the ray accumulates on the Julia set, where the hyperbolic metric blows up, this means that the *euclidean* distance between a point on the ray and its image tends to 0. Therefore all accumulation points are periodic. Since the limit set of the ray is connected, it lands at a single periodic point. (For more details, see [56, Theorem 18.10].)

The main obstruction to carrying this proof over to the exponential case is that the set of escaping points is no longer open. However, we shall remove this difficulty by considering an open set that is an approximation to the set of escaping points: All points whose orbits become large at some point in the future and stay so for at least a certain number of iterations. Now if we can show that the ray's accumulation set is contained in the boundary of this domain, then we can use the same proof as before. However, we will have to take out the singular orbit. Thus the proof, as presented here, requires the ray does not intersect the postsingular set at low potentials. This condition, although it includes the previously known cases, is still rather special. Indeed, for a generic set in the bifurcation locus, the singular orbit is dense in the plane (Theorem 5.1.6). As the example of rays at addresses which do not contain zeros shows below, it may nonetheless be possible to generalize the proof to larger classes of parameters.

In the following, fix $\kappa \in \mathbb{C}$ and $E := E_\kappa$. Let $g := g_{\underline{s}} : (0, \infty) \to \mathbb{C}$ be a periodic external ray of period p whose orbit does not contain the singular value. Let $f : (0, \infty) \to (0, \infty)$ be the unique function satisfying $g(f(t)) = E_\kappa^p(g(t))$. (In fact, $f(t) = T(\mathcal{F}^p(\underline{s}, t + t_{\underline{s}})) - t_{\underline{s}}$.) Recall that $\mathcal{P}$ denotes the postsingular set of E_κ.

3.9.1 Theorem (Periodic Rays Land)
If there is some $T_0 > 0$ such that $g\big((0, f(T_0)]\big) \cap \mathcal{P} = \emptyset$, then there exists a periodic point z such that $\lim_{t \to 0} g(t) = z$.

Remark. The landing point must be repelling or parabolic by the Snail lemma [56, Lemma 16.2], just as in the polynomial case.

We shall begin with some preliminary observations. Let $A := (2 \max |s_k| + 1)\pi$. Then, if t is large enough, $\big|\mathrm{Im}\big(E^k(g(t))\big)\big| < A$ for all $k \geq 0$. For an arbitrary $R > 0$, consider the set

$$S_R := \{z : \mathrm{Re}\, z > R \text{ and } |\,\mathrm{Im}\, z| < A\}.$$

We define

$$U_R := \{z : \exists k : E^{kp+1}(z), \dots, E^{(k+1)p}(z) \in S_R\} \setminus \mathcal{P}.$$

U_R is open, and by the above it contains $g\big((0, f(T_0)]\big)$. Furthermore, $E^{-p}(U_R) \subset U_R$, and $E^p : E^{-p}(U_R) \to U_R$ is a covering map.

Consider the component V'_R of $E^{-p}(U_R)$ that contains $g\big((0, T_0]\big)$, and the image component $V_R := E^p(V'_R)$ of U_R. Then $E^p : V'_R \to V_R$ is still a covering map, and thus expands the hyperbolic metric of V_R. In particular, let ℓ be the hyperbolic length of the piece $g\big([T_0, f(T_0)]\big)$. Then, for any $t \leq T_0$, the length of $g\big([t, f(t)]\big)$ is at most ℓ.

We are now ready to prove Theorem 3.9.1. Let

$$G := \bigcap_{t>0} \overline{g\big((0, t]\big)}$$

denote the accumulation set of g. The key step is proving the following lemma.

3.9.2 Lemma (Limit Set on Boundary)
There exists an $R > 0$ such that $G \subset \partial V_R$.

PROOF. Let w be a periodic point of E with $\mathrm{Re}(w) > 0$ and set $R_0 := \max_k \mathrm{Re}(E^k(w))$. Note that $V_{R_0} \subset \widetilde{U} := \mathbb{C} \setminus \{w + 2\pi i k : k \in \mathbb{Z}\}$. By the monotonicity of the hyperbolic metric, the hyperbolic metric of V_{R_0} is larger than that of $\widetilde{U}$. Thus, by Lemma 2.2.4, the hyperbolic metric on V_{R_0} satisfies

$$\rho_{V_{R_0}}(z) \geq \frac{C}{\mathrm{Re}(z) - \mathrm{Re}(w)} \geq \frac{C}{\mathrm{Re}\, z}$$

whenever $\mathrm{Re}\, z > \mathrm{Re}(w) + 1$.

Now fix some large R and suppose that $G \cap V_R \neq \emptyset$. Because G is forward invariant under E^p, this means that there is $z_0 \in G$ such that $E^k(z_0) \in S_R$ for $k = 1, \dots, p$. Thus if we choose $z_1 = g(t)$, $t < T_0$, close enough to z_0, then $E^k(z_1) \in S_R$ for $k = 1, \dots, p$.

If R was chosen large enough, then, whenever $\mathrm{Re}\, z > R$,

$$|E(z)| \geq \exp(\mathrm{Re}(z) - 1) + A + 1.$$

In particular, if $|\mathrm{Im}\, E(z)| < A$ and $\mathrm{Re}(E(z)) > 0$, then

$$\mathrm{Re}\,(E(z)) - 1 \geq |E(z)| - |\mathrm{Im}\, E(z)| - 1 \geq \exp(\mathrm{Re}(z) - 1).$$

By induction, we thus obtain

$$\mathrm{Re}\, E^p(z_1) \geq \exp^p(\mathrm{Re}(z_1) - 1).$$

Denote the hyperbolic distance on V_{R_0} by dist_h. Let $r := \mathrm{Re}(z_1)$. Then

$$\begin{aligned} \ell \geq \mathrm{dist}_h(z, E^p(z)) &\geq C \int_r^{\exp^p(r-1)} \frac{dx}{x} \\ &= C\left(\exp^{p-1}(r) - \log(r)\right). \end{aligned}$$

The last term is certainly unbounded as R (and thus r) tends to ∞. This is a contradiction. ■

Proof of Theorem 3.9.1. Let R be such that $G \subset \partial V_R$. We will first show that every point of G is fixed by E^p. Because the set of fixed points of E^p is discrete, this proves that G lands either at a periodic point or at ∞.

The proof is the same as in the polynomial case: Suppose that $g(t_n) \to z$, where $t_n \to 0$. Because $z \in \partial U_R$, and $\operatorname{dist}_h\big(g(t_n), g(f(t_n))\big) < \ell$, it follows that the euclidean distance $\operatorname{dist}(g(t_n), g(f(t_n)))$ tends to 0. Because $g(F^p(t_n)) \to E(z)$, this implies $E(z) = z$.

It remains to show that g cannot land at ∞; so suppose this is the case. Then all rays in the cycle of g also land at ∞, and therefore $\lim_{t\to 0} \operatorname{Re} g(t) = +\infty$. Since E is injective on g, g cannot intersect any of its translates; the same goes for its images $E^k \circ g$. It follows that g and its images have bounded imaginary parts. As in the proof of Lemma 3.9.2 this implies that we can choose R large enough such that if $r = \operatorname{Re}(g(t)) > R$, then $\operatorname{Re}\big(g(f^{-1}(t))\big) < \exp^{-p}(r) + 1$. Thus it is impossible for the backward orbit $g\big(f^{-k}(t)\big)$ to converge to ∞. This is a contradiction. ■

In the following, suppose that $\operatorname{Im}\kappa \in [-\pi, \pi]$. Then we can adapt the proof to show that every periodic ray whose address has only nonzero entries lands. This has been known for a long time, and can be shown in many different ways. In fact, *any* external ray with such an address lands (see Section 3.11 for a discussion). However, we wish to present this argument as an indication that the proof of Theorem 3.9.1 might be generalized even to cases where the postsingular orbit is dense.

3.9.3 Theorem (Nonzero Periodic Rays Land [21])
Let $\kappa \in \mathbb{C}$, $\operatorname{Im}\kappa \in [-\pi, \pi]$. Then every periodic ray whose address contains no zeros lands.

Sketch of proof. The proof relies on the fact that these rays can never cross the lines $\{\operatorname{Im} z = (2k-1)\pi\}$, and in particular never enter the strip which contains κ. The proof proceeds as above; however, instead of U_R we use the sets

$$\widetilde{U}_R := \{z : \exists k : E^{kp+1}(z), \dots, E^{(k+1)p}(z) \in S_R\} \setminus \{\kappa + x : x \in (-\infty, 0]\}.$$

Again, let $\widetilde{V}_R$ be the component of $\widetilde{U}_R$ which contains the ray g, and let $\widetilde{V}'_R$ be the corresponding component of $E_\kappa^{-p}(\widetilde{V}_R)$. Then, by definition $\widetilde{V}'_R$ does not intersect the set $\{\operatorname{Im} z \in [-\pi, \pi]\}$, and thus $\widetilde{V}'_R \subset \widetilde{V}_R$. The rest of the proof proceeds as above. ■

We wish to emphasize that the main difficulty for using this method to prove the landing of a periodic ray in a more general setting appears to consist of first separating the ray from κ. Indeed, if the ray accumulated on κ and the orbit of κ accumulates on the ray at arbitrarily low potentials, then it is impossible to find a domain U such that $g \in U$ and $g : U \to E(U) \supset U$ is a covering map. The general landing theorem (Theorem 5.13.1) shows that this case never occurs, but we do not currently know of any intrinsic dynamical reason for this.

3.10 Holomorphic Motion of External Rays

In this section we present a rather peculiar method to prove the landing of external rays, which will allow us to prove theorem 1.10 in Section 5.13. It is a direct consequence of the "λ-lemma" in the theory of holomorphic motions [51]. For completeness, we give the simple proof here, which is largely the same as for the general λ-lemma.

3.10.1 Lemma (Holomorphic Motion of External Rays)
Let $\underline{s}$ be an exponentially bounded external address. Suppose that $U \subset \mathbb{C}$ is a domain such that $\kappa \notin g^{\kappa}_{\sigma^n(\underline{s})}$ for all $\kappa \in U$ and $n \geq 1$. If there exists some $\kappa_0 \in U$ such that $g^{\kappa_0}_{\underline{s}}$ lands, then $g^{\kappa}_{\underline{s}}$ lands for all $\kappa \in U$.

Proof. Let us use the notation $h_t(\kappa) := g^{\kappa}_{\underline{s}}(t)$. Note that the family $(h_t)_{t\in(0,1)}$ of holomorphic functions is normal by Montel's theorem, as it omits the holomorphic functions $\kappa \mapsto \kappa$ and $\kappa \mapsto h_1(\kappa)$. Let h_0 be a limit function of this family as $t \to 0$. Note that $h_0(\kappa_0) = z_0$ and, by Hurwitz's theorem, $h_0(\kappa) \neq h_t(\kappa)$ for all $\kappa \in W$ and $t > 0$.

Now if $\widetilde{h}$ is some other limit function of h_t as $t \to 0$, then $\widetilde{h}(\kappa_0) = z_0 = h_0(\kappa_0)$. Since the h_t omit the holomorphic function h_0, we can apply Hurwitz's theorem once more to see that $\widetilde{h} = h_0$. In other words, $h_t \to h_0$ as $t \to 0$. In particular, $g_\kappa(t) \to h_0(\kappa)$ for every κ. ■

3.11 Rays at Nonzero Addresses

In this section, we shortly remark on the case of external addresses containing only nonzero entries (under the assumption that $\operatorname{Im}\kappa \in [-\pi, \pi]$). These addresses were the main focus of attention in [21], where they were called *regular*. However, despite the connection with the position of the singular value which is the origin of this terminology, we feel that rays at addresses which contain zeros do not deserve the title "singular". In fact, these are the rays that contain combinatorial information about the map, as only rays of this kind can have a common landing point. Milnor has suggested calling them "unreal" (in the context of the traditional parametrization $z \mapsto \lambda \exp(z)$) because they do not intersect the real axis. Since the real axis has no special significance in our parametrization, however, we prefer to simply call them "nonzero".

As we have noticed before, rays at nonzero addresses can never cross the lines $\{\operatorname{Im} z = (2k-1)\pi\}$, which makes them rather special. In Section 3.9, we have already proved the result from [21] that nonzero periodic rays always land. For completeness, we will give here the short (and standard) proof of a converse result.

3.11.1 Theorem (Periodic Points Outside the Singular Strip)
Suppose that $\operatorname{Im}\kappa \in [-\pi, \pi]$. If z_0 is a periodic point whose orbit is contained in $\{z : |\operatorname{Im} z| > \pi\}$, then z_0 is the landing point of some nonzero periodic ray.

PROOF. Let $\underline{s} = \overline{s_1 s_2 \dots s_n}$ be the external address of z_0. Then the ray $g_{\underline{s}}$ has a periodic landing point z_1, and the orbits of both points lie in the same strips $S_k = \{z : \operatorname{Im} z \in ((2k-1)\pi, (2k+1)\pi)\}$ of the partition.

There exists a unique component U of $E_\kappa^{-n}(S_{s_1})$ such that $E_\kappa^{k-1}(U) \subset S_{s_k}$ for $k = 1, \dots, n$. Then $z_0, z_1 \in U$. Now $E_\kappa^n : U \to S_k$ is a covering map, and thus expands the hyperbolic metric of S_k. However, this is only possible if $z_0 = z_1$. ■

Finally, we should note that, in fact, *all* external rays at nonzero addresses land. An elegant and natural proof of this fact can be found in [77]. However, for completeness, we will give a short proof using the result of the previous section.

3.11.2 Theorem (All Nonzero Rays Land [77, Proposition 6.11])
Suppose that $\operatorname{Im} \kappa \in (-\pi, \pi)$, *and suppose that* $\underline{s}$ *is a nonzero address. Then* $g_{\underline{s}}$ *lands.*

PROOF. For such κ, neither $g_{\underline{s}}$ nor any of its image rays can contain κ, because these rays do not intersect the partition boundaries. Because $g_{\underline{s}}$ lands for e.g. $\kappa = -2$, as mentioned in Section 3.1 and proved in Section 4.2, the theorem follows from Lemma 3.10.1. ■

Chapter 4

Combinatorics of Exponential Maps

In this chapter, we shift our focus from considering general exponential maps to those parameters whose dynamics allows us to gain more combinatorial information. The largest part of the chapter deals with attracting or parabolic parameters. In particular, we prove Theorem 1.5 in Section 4.2 and show that topological renormalization fails even for attracting parameters (Section 4.3).

The rest of the chapter consists of a combinatorial discussion of *itineraries* and *orbit portraits* which is taken from [68] and is an extension of concepts introduced by Schleicher in [74]. These concepts will be of great importance in our study of exponential parameter space (Chapter 5). In Section 4.4 we also discuss the combinatorics of escaping and Misiurewicz parameters, as introduced in [76], allowing us to incorporate these into the combinatorial discussion which follows. (In this, our description slightly differs from that in [68], which is restricted to attracting maps for simplicity.)

For many of the results from [74] which we cite, we have at least included a sketch of the proof. We hope that this will give the reader a better understanding of the essential ideas, especially since many of the proofs convey interesting additional information on the combinatorics which is being considered.

4.1 Symbolic Dynamics of Attracting Fatou Components

We shall now examine one of the striking features of attracting and parabolic exponential dynamics which is not present in the polynomial case: the Fatou components are unbounded, which allows us to incorporate their dynamics into the combinatorial model of external addresses.

Throughout this section, let $\kappa \in \mathbb{C}$ such that E_κ has an attracting or parabolic periodic point. Then the singular value κ is contained in some periodic Fatou component which we call the *characteristic Fatou component*. Let $U_0 \mapsto U_1 \mapsto \ldots \mapsto U_n = U_0$ be the cycle of Fatou components, labeled such that U_1 is the characteristic component. Since U_1 contains a neighborhood of the singular value, U_0 contains an entire left half plane. In particular,

U_0 contains a horizontal curve along which $\mathrm{Re}(z) \to -\infty$. This curve is unique up to homotopy. Its pullback to U_1 under E_κ^{n-1} is also a curve to ∞, and all countably many preimages of this curve are curves to ∞ in U_0. By taking pullbacks of these curves, we see that every Fatou component U contains infinitely many homotopy classes of curves to ∞ (i.e., infinitely many prime ends for which ∞ is the only principal point.) In fact these prime ends are dense in the prime end topology on ∂U, compare [4]. Also, every prime end contains ∞ in its impression [6]. (See [20] for the case $n = 1$.)

A homotopy class as above will be called an *access to infinity* within U. We will also call the homotopy class in U_0 which contains all horizontal curves to $-\infty$ the "access to $-\infty$". Accesses to infinity fit naturally into our combinatorial description: if γ is a curve to ∞ within any Fatou component, then, as in Section 3.7, it has an associated infinite or intermediate external address $\mathrm{addr}(\gamma)$, and this address depends only on the access to ∞ represented by γ.

Note that $\mathrm{addr}(\gamma) = \infty$ if and only if γ represents the access to $-\infty$ in U_0. Note that any curve γ to ∞ in the Fatou set for which $\mathrm{addr}(\gamma)$ is intermediate will eventually map to a curve in U_0 with address ∞. In other words, γ can be obtained, up to homotopy, by a pullback of the horizontal curve to $-\infty$ in U_0. Because this pullback is unique as long as it does not pass through U_1, there is a distinguished access to ∞ within each Fatou component:

4.1.1 Definition and Lemma (Principal Address of Fatou Components)
Let $\kappa \in \mathbb{C}$ be an attracting or parabolic parameter, and let U be a Fatou component. If $m \in \mathbb{N}$ is minimal with the property that $E^{m-1}(U) = U_0$, then there is a unique access to infinity in U whose address is intermediate and has length m; it corresponds to the access to $-\infty$ in U_0 under the (bijective) map $E_\kappa^{m-1} : U \to U_0$. The address of this access is denoted by $\mathrm{addr}(U)$. *We also define the* (intermediate) external address of κ *to be* $\mathrm{addr}(\kappa) := \mathrm{addr}(U_1)$. ■

As we shall see in the next section, $\mathrm{addr}(\kappa)$ completely determines the dynamics of E_κ on the Julia set. Because the previous definition is somewhat abstract and it is vitally important for the understanding of much that follows, let us explain it in a slightly more hands-on manner. Let U be a component as in the previous Lemma. If T is large and negative, then the curve $\gamma_m : [0, \infty) \to \mathbb{C}, t \mapsto T - t + \mathrm{Im}\,\kappa$ lies in U_0. Let $\gamma_{m-1}, \dots, \gamma_1$ be successive pullbacks of this curve such that γ_1 lies in U. Then

$$\mathrm{addr}(U) = \mathrm{addr}(\gamma_1) = s_1 \dots s_{m-1}\infty$$

can be interpreted in the following way. The curve γ_m tends to $-\infty$ in U_n, which is why the corresponding entry is ∞. The curve γ_{m-1} has constant imaginary part $2\pi s_{m-1}$; in other words, it lies on the boundary between the strips $S_{s_{m-1}+\frac{1}{2}}$ and $S_{s_{m-1}-\frac{1}{2}}$. The curves γ_k for $k < m-1$ tend to ∞ in S_{s_k}; in fact their imaginary part is asymptotic to $2\pi s_k$.

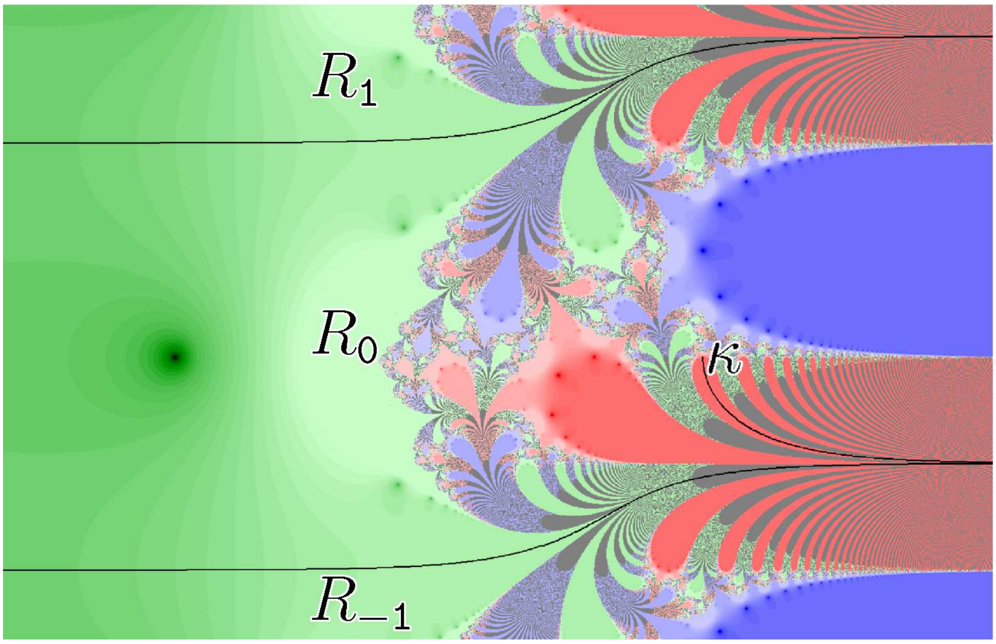

(a) The dynamical partition for a period 3 parameter.

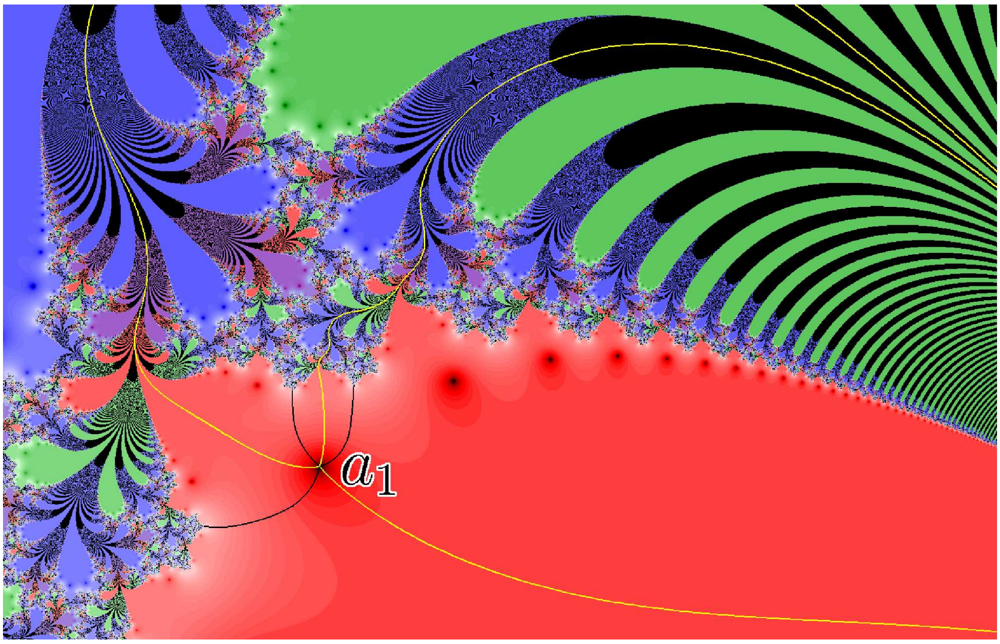

(b) A period 5 parameter whose multiplier has angle 1/3. The picture shows the three broken attracting rays and three unbroken rays landing at the distinguished boundary cycles.

Figure 4.1: Illustrations of the definition of itineraries and Lemma 4.1.7.

Itineraries

The above construction allows us to define a partition of $\mathbb{C} \setminus U_0$ that is dynamically more meaningful than the static partition into horizontal strips used in the definition of external addresses. Select some curve $\gamma \subset U_1$ connecting κ to ∞ with $\operatorname{addr}(\gamma) = \operatorname{addr}(U_1)$. Its pullback under E_κ consists of countably many curves connecting $-\infty$ to $+\infty$ within U_0 (compare Figure 4.1(a)); all of these curves are translates of each other (by $2\pi i\mathbb{Z}$). Thus, they partition the dynamical plane into "regions" R_j, where the labeling is chosen so that R_j contains all z with $\operatorname{Im}(z) = (2j+1)\pi$ and $\operatorname{Re}(z)$ large. (This labeling has the same ambiguity as the labeling of the static partition.) To any point z in the Julia set, we can then associate an itinerary, recording into which strips its orbit maps.

4.1.2 Definition ((Dynamical) Itinerary)
If $z \in J(f)$, then the itinerary *of z is defined to be* $\operatorname{itin}(z) := \mathtt{u}_1\mathtt{u}_2\dots$, *where* $\mathtt{u}_k = \mathtt{j}$ *iff* $E_\kappa^{k-1}(z) \in R_{\mathtt{j}}$. *If $z \in F(f)$, then we similarly define* $\operatorname{itin}(z) = \mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}*$, *where n is the smallest number such that $E_\kappa^{n-1}(z) \in U_0$.*

Remark. If an orbit enters U_0 but does not fall on the boundary between the sectors R_j, we could define an entry for its itinerary at this step. However, this entry would depend on the noncanonical choice of γ.

Analogously, one can define a combinatorial notion of itineraries for external addresses roughly as follows (the precise definition is Definition 4.1.3 below). Let $\underline{s} = \operatorname{addr}(\kappa)$. We want the itinerary of any external address $\underline{r}$ to record the intervals of

$$\mathcal{S} \setminus \sigma^{-1}(\underline{s}) = \mathcal{S} \setminus \{j\underline{s} : j \in \mathbb{Z}\}$$

containing the orbit points of $\underline{r}$ under the shift. If we wish to assign itineraries to every address of $\mathcal{S}$, we will also need to assign labels to the strip boundaries $j\underline{s}$; this is done by labelling them with a "boundary symbol" ${}^{\mathtt{j}}_{\mathtt{j}-1}$. (This convention means simply that the address lies between the strips labelled j and $j-1$.) In Section 4.4, we will also define itineraries for other types of parameters; this time using external rays rather than curves in Fatou components to connect the singular value to ∞. Therefore we make the definition here in all generality, i.e. for addresses $\underline{s}$ which are not necessarily intermediate.

4.1.3 Definition ((Combinatorial) Itinerary)
Let $\underline{s} \in \mathcal{S}$ and $\underline{r} \in \overline{\mathcal{S}}$. Then the itinerary of $\underline{r}$ with respect to $\underline{s}$ *is* $\operatorname{itin}_{\underline{s}}(\underline{r}) = \mathtt{u}_1\mathtt{u}_2\dots$, *where*

$$\begin{cases} \mathtt{u}_k = \mathtt{j} & \text{if } \ \mathtt{j}\underline{s} < \sigma^{k-1}(\underline{r}) < (\mathtt{j}+1)\underline{s} \\ \mathtt{u}_k = {}^{\mathtt{j}}_{\mathtt{j}-1} & \text{if } \ \sigma^{k-1}(\underline{r}) = \mathtt{j}\underline{s} \\ \mathtt{u}_k = * & \text{if } \ \sigma^{k-1}(\underline{r}) = \infty \end{cases} .$$

We also define $\operatorname{itin}^+_{\underline{s}}(\underline{r})$ *and* $\operatorname{itin}^-_{\underline{s}}(\underline{r})$ *to be the address obtained by replacing each boundary symbol ${}^{\mathtt{j}}_{\mathtt{j}-1}$ by $\mathtt{j}$ or $\mathtt{j}-1$, respectively.*

When κ is a fixed hyperbolic or parabolic parameter, we shall usually abbreviate $\operatorname{itin}(\underline{r}) := \operatorname{itin}_{\operatorname{addr}(\kappa)}(\underline{r})$.

REMARK.

(a) If z lies in the closure of the external ray $g_{\underline{r}}$, then $\mathrm{itin}(z) = \mathrm{itin}(\underline{r})$. If z lies in a Fatou component U, then $\mathrm{itin}(z) = \mathrm{itin}(\mathrm{addr}(U))$.

(b) In the case $\underline{s} = \infty$, we can define itineraries analogously, but the addresses $\mathtt{j}\underline{s}$ and $(\mathtt{j}+1)\underline{s}$ in the definition will have to be replaced by $(\mathtt{j} - \frac{1}{2})\infty$ and $(\mathtt{j} + \frac{1}{2})\infty$. With this definition, $\mathrm{itin}_\infty(\underline{r}) = \underline{r}$ for all infinite external addresses $\underline{r}$.

4.1.4 Definition (Kneading Sequence)
Let $\underline{s}$ be an intermediate external address. Then

$$\mathbb{K}(\underline{s}) := \mathrm{itin}_{\underline{s}}(\underline{s})$$

is called the kneading sequence *of $\underline{s}$. We again set $\mathbb{K}^{\pm}(\underline{s}) := \mathrm{itin}^{\pm}_{\underline{s}}(\underline{s})$ and $\mathbb{K}(\kappa) = \mathbb{K}(\mathrm{addr}(\kappa)$.*

Note that the n-th entry $\mathtt{u}_n$ of the kneading sequence $K(\underline{s})$ is locally constant (when defined) as a function of $\underline{s}$, unless $\mathtt{u}_n$ is a boundary symbol ${}^{\mathtt{j}}_{\mathtt{j}-1}$ (which happens exactly when $\sigma^n(\underline{s}) = \underline{s}$) or $\mathtt{u}_n = *$. In the former case, it will be $\mathtt{j}$ for addresses slightly above $\underline{s}$, and $\mathtt{j}-1$ slightly below $\underline{s}$. The latter case occurs when $\underline{s}$ is an intermediate external address; i.e. $\underline{s} = s_1 \dots s_{n-2}(k+\frac{1}{2})\infty$ and $\mathbb{K}(\underline{s}) = \mathtt{u}_1 \dots \mathtt{u}_{n-1}*$. Then the addresses $\overline{s_1 \dots s_{n-1} k \mathtt{j}}$ approximate $\underline{s}$ from below (for $\mathtt{j} \to \infty$) and have kneading sequence $\overline{\mathtt{u}_1 \dots \mathtt{u}_{n-1}{}^{\mathtt{j}}_{\mathtt{j}-1}}$ for large j. Similarly, $\overline{s_1 \dots s_{n-1}(k+1)\mathtt{j}}$, $\mathtt{j} \to -\infty$, have kneading sequence $\overline{\mathtt{u}_1 \dots \mathtt{u}_{n-1}{}^{\mathtt{j}}_{\mathtt{j}-1}}$ when $\mathtt{j}$ is small.

Finally let us make the simple observation that every itinerary is realized (with a possible exception).

4.1.5 Lemma (Existence of Itineraries)
Let $\underline{\mathtt{u}}$ be any sequence of entire numbers, and let $\underline{s} \in \mathcal{S}$. Then there exists an (infinite) external address $\underline{t}$ such that $\mathrm{itin}_{\underline{s}}(\underline{t}) = \underline{\mathtt{u}}$, unless $\sigma^n(\mathtt{u}) \in \{\mathbb{K}^+(\underline{s}), \mathbb{K}^-(\underline{s})\}$ for some $n \geq 1$. If $\underline{\mathtt{u}}$ is periodic, then $\underline{t}$ can be chosen to be periodic.

REMARK. Note that the resulting address $\underline{t}$ might not be exponentially bounded. However, this is the case if and only if $\mathtt{u}$ is not exponentially bounded. Also, note that it may happen that $\mathbb{K}(\underline{s})$ is periodic but there is no periodic address $\underline{t}$ with $\mathrm{itin}_{\underline{s}}(\underline{t}) = \mathbb{K}(\underline{s})$.

PROOF. Let us denote

$$\mathcal{S}_k := [k\underline{s}, (k+1)\underline{s}] = \{\underline{t} \in \mathcal{S} : k\underline{s} \leq \underline{t} \leq (k+1)\underline{s}\}.$$

and define

$$M_k := \{\underline{t} \in \mathcal{S} : \sigma^{j-1}(\underline{t}) \in \mathcal{S}_{\mathtt{u}_j} \text{ for all } j \leq k\}.$$

Then each M_k is compact and $M_{k+1} \subset M_k$. One also easily checks that M_k is nonempty. Thus $\bigcap M_k \neq \emptyset$; let $\underline{t}$ be any element of this intersection. Then $\underline{t}$ clearly has itinerary $\underline{\mathtt{u}}$, unless, for some $n \geq 1$, $\sigma^n(\underline{t}) = \underline{s}$. In this case, choose a monotone sequence $\underline{t}_k \in M_k$ with $\underline{t}_k \to \underline{t}$ and $\sigma^j(\underline{t}_k) \neq \underline{s}$ for all $j \leq k+1$. Then $\operatorname{itin}_{\underline{s}}(\underline{t}_k)$ begins with $\mathtt{u}_1 \dots \mathtt{u}_k$. To fix our ideas, suppose that $\underline{t}_k$ converges to $\underline{t}$ from above. It follows that $\operatorname{itin}^+_{\underline{s}}(\underline{t}) = \mathtt{u}$; or in other words $\sigma^n(\mathtt{u}) = \mathbb{K}^+(\underline{s})$.

For the last statement, see [76, Lemma 5.2]. ■

Attracting Dynamic Rays

For the following considerations, suppose that E_κ has an attracting orbit, which we label $a_0 \mapsto a_1 \mapsto \dots \mapsto a_n = a_0$ such that $a_i \in U_i$. We can construct a curve representing $\operatorname{addr}(\kappa)$ using the linearizable dynamics on the Fatou set: connect the attracting point a_1 to κ by a straight line in linearizing coordinates and consider the pullback under the first return map of U_1. The result is a curve γ connecting a_1 to ∞ with $\operatorname{addr}(\gamma) = \operatorname{addr}(\kappa)$.

To generalize this construction, consider the extended Koenigs map

$$\phi : U_1 \to \mathbb{C},$$
$$\phi(E_\kappa^n(z)) = \mu\phi(z).$$

(Here μ is the multiplier of the attracting orbit.) For simplicity, we normalize ϕ in such a way that $\phi(\kappa) = 1$. Now we can define dynamic rays inside the Fatou component U_1:

4.1.6 Definition (Attracting Dynamic Ray)
Let $\theta \in \mathbb{R}/\mathbb{Z}$, and $R_\theta(t) = te^{2\pi i\theta}$ (for $t \geq 0$). Then the component γ of $\phi^{-1}\big(R_\theta\big([0,\infty)\big)\big)$ which contains a_1 is called the attracting dynamic ray *at angle θ. If $e^{2\pi i\theta} = \mu^{-n}$ for some $n > 0$, the ray γ is called* broken, *and every component of $\phi^{-1}\big(R_\theta(|\mu|^{-n},\infty)\big)$ is called a* ray piece *at angle θ.*

The attracting dynamic ray at angle $-\frac{\arg(\mu)}{2\pi}$ is called the principal attracting ray *of E_κ.*

We can apply the same procedure to obtain dynamic rays inside $U_2, \dots, U_n$.

Remark. The principal attracting ray is just the curve considered above, which connects a_1 to ∞ with external address $\operatorname{addr}(\kappa)$.

Now suppose that $\mu = e^{2\pi i \frac{p}{q}}$, where p and q are relatively prime. Among the rays starting at a_1, there is a cycle of q broken rays, completely containing the singular orbit in U_1. Rays outside this cycle are periodic with period q and unbroken, see Figure 4.1(b). In fact, we can say even more:

4.1.7 Definition and Lemma (Distinguished Boundary Orbit)
Consider an attracting exponential map of period n whose multiplier μ has rational angle $\frac{p}{q}$. Then on ∂U_1, there is an orbit (under E_κ^n) of period q such that every unbroken attracting ray starting at a_1 lands at a point of this orbit. It is called the distinguished boundary orbit *on ∂U_1.*

Remark. For $q = 1$ this is [74, Lemma 6.1].

Proof. The proof is analogous to the last part of the proof of Theorem 3.9.1: The attracting dynamic ray accumulates on the boundary, and thus every accumulation point is a periodic point or ∞. The ray cannot land at ∞, because the imaginary parts of points on the ray (and its finitely many images) is bounded; thus large points on the ray would have exponentially growing real parts and thus escape. ■

4.2 Topology of the Julia Set for Attracting and Parabolic Parameters

We will now completely describe the Julia sets of attracting and parabolic exponential maps (and their dynamics thereon) as a quotient of our model X. In particular, any two attracting exponential maps are conjugate on their sets of escaping points. We will give the complete construction for attracting parameters, and remark on the parabolic case later.

So let $\kappa \in \mathbb{C}$ such that E_κ has an attracting cycle $a_0 \mapsto a_1 \mapsto \dots \mapsto a_n = a_0$ and corresponding Fatou components $U_0 \mapsto U_1 \mapsto \dots \mapsto U_n = U_0$, again labelled such that $\kappa \in U_1$. By the Koenigs linearization theorem, we can find open neighborhoods U_j of a_j such that $\kappa \in U_1$, $E_\kappa(U_j) \Subset U_{j+1}$ and $E_\kappa(U_n) \Subset U_1$. Define $U_0 := E_\kappa^{-1}(U_1)$ and consider the set

$$W := \mathbb{C} \setminus \left(\bigcup_{i=0}^{n-1} \overline{U_i} \right).$$

Then $E_\kappa^{-1}(W) \subset W$, and $E_\kappa : E_\kappa^{-1}(W) \to W$ is a covering map.

As in the definition of itineraries, let $\gamma \subset U_1$ connect κ to ∞ with $\operatorname{addr}(\gamma) = \operatorname{addr}(\kappa)$, and set $V := \mathbb{C} \setminus \gamma$. Recall that $E_\kappa^{-1}(V)$ consists of the regions R_k used in the definition of itineraries; let

$$\widetilde{L}_k : V \to R_k$$

denote the corresponding branch of E_κ^{-1}. These differ from the branches L_k considered in Chapter 3. Note that $\widetilde{L}_k$ is well-defined everywhere on the Julia set.

Choose some $A > 0$ and $B < 0$ such that the map $G = G^\kappa : X \to I(E_\kappa)$ satisfies $|\operatorname{Re} G(\underline{s}, t) - t| < 2 + \pi$ on Y_A and such that

$$\mathcal{H} := \{z \in \mathbb{C} : \operatorname{Re} z > A - 2\} \subset W \subset \{z \in \mathbb{C} : \operatorname{Re} z > B\}.$$

Now let us define functions $H_k : \overline{X} \to J(E_\kappa)$ by

$$H_0(\underline{s}, t) := G(\underline{s}, t + A) \text{ and}$$
$$H_{k+1}(\underline{s}, t) := \widetilde{L}_{\mathtt{u}_1}(H_k(\mathcal{F}(\underline{s}, t))),$$

where $\mathtt{u}_1$ is the first entry of $\operatorname{itin}_{\operatorname{addr}(\kappa)}(\underline{s})$. Note that $H_k(\underline{s}, t)$ always lies on the external ray $g_{\underline{s}}$.

4.2.1 Theorem (Conjugacy for Attracting Parameters)
In the hyperbolic metric of W, the functions H_k converge uniformly to a continuous, surjective function

$$H : \overline{X} \to J$$

with $H \circ \mathcal{F} = E_\kappa \circ F$. Furthermore, $H|_X : X \to I$ is a conjugacy.

Remark. By Corollary 3.5.2, $H_{|X}$ must be equal to G^κ.

Proof. Since $E_\kappa : E_\kappa^{-1}(W) \to W$ is a covering map, it expands the hyperbolic metric of W. In fact, there exists $K > 1$ such that $\| \mathrm{D}E_\kappa(z)\|_{\text{hyp}} \geq K$ for all $z \in E_\kappa^{-1}(W)$.

To prove this, let W' be any component of $E_\kappa^{-1}(W)$, and let

$$\widetilde{W} := \big\{ z \in W : z + 2\pi i k \in W \text{ for all } k \in \mathbb{Z} \big\}.$$

(Thus $\widetilde{W}$ is obtained from W by obtaining all translates of the sets $\overline{U_i}$.) Then $W' \subset \widetilde{W} \subset W$. Because $\rho_{\widetilde{W}} \geq \rho_W$ by the monotonicity of the hyperbolic metric, it is sufficient to show that, for every $z \in W'$,

$$\frac{\rho_{W'}(z)}{\rho_{\widetilde{W}}(z)} \geq K > 1. \tag{4.1}$$

However, recall that U_0 contains a left half plane, so that, for some $R_0 > 0$, the set $\mathbb{C} \setminus W'$ contains the curves $\{R + (2k+1)\pi : R \geq R_0\}$. By Theorem 2.2.3, this implies that $\rho_{W'}(z)$ is bounded from below for large real parts of z. On the other hand, $\widetilde{W}$ contains the right half plane $\mathcal{H}$, and therefore $\rho_{\widetilde{W}}(z) \to 0$ as $\operatorname{Re} z \to +\infty$. Thus

$$\lim_{\operatorname{Re} z \to +\infty} \frac{\rho_{W'}(z)}{\rho_W(z)} = \infty.$$

Because the expression $\frac{\rho W'}{\rho_{\widetilde{W}}}$ is $2\pi i$-periodic and $\partial W' \subset \widetilde{W}$, the claim (4.1) follows.

For an arbitrary $(\underline{s}, t) \in X$, consider the two points $z_1 := E_\kappa(H_0(\underline{s}, t)) = G(\mathcal{F}(\underline{s}, t + A))$ and $z_2 := E_\kappa(H_1(\underline{s}, t)) = H_0(\mathcal{F}(\underline{s}, t))$. Both points have real parts greater than $A - 2$ and thus can be connected by a straight line g_0 in W. Note that g_0 is homotopic (in W) to the piece of the ray $g_{\underline{s}}$ between z_1 and z_2, as this piece is also contained in the half plane $\mathcal{H}$. Thus we can pull back g_0 and obtain a curve g_1 between $H_0(\underline{s}, t)$ and $H_1(\underline{s}, t)$. We claim that the (euclidean) length of g_1 is uniformly bounded (independent of $\underline{s}$ and t).

To prove this claim, recall that

$$\operatorname{Re}(z_1) \leq T(\mathcal{F}(\underline{s}, t + A)) + 2 = \exp(t + A) - 2\pi|s_2| + 1$$

and

$$\operatorname{Re}(z_2) \geq T(\mathcal{F}(\underline{s}, t)) - 2 = \exp(t) - 2\pi|s_2| - 3.$$

It follows that the euclidean length of g_0 satisfies

$$\ell(g_0) \leq \exp(t + A) - \exp(t) + 4 + 2\pi = O(\exp(t)).$$

Because all points of g_0 have absolute value at least $|z_2| - 2\pi \geq \frac{\exp(t)}{\sqrt{2}} - 2 - 3\pi$, we see that

$$\ell(g_1) \leq \frac{1}{|z_2| - 2\pi} \cdot \ell(g_0) = O(1).$$

Recall that, since $\rho_W(z) \to 0$ as $\operatorname{Re} z \to \infty$, the function ρ_W is uniformly bounded on W'. Thus the hyperbolic length of g_1 in the hyperbolic metric of W is also bounded by some constant C. Now, taking pullbacks inductively, we see that the hyperbolic distance between $H_k(\underline{s}, t)$ and $H_{k+1}(\underline{s}, t)$ is bounded by $\frac{C}{K^k}$. Thus the H_k converge uniformly. The functional equation is satisfied by construction.

To show surjectivity of H, it is sufficient to see that $H(\overline{X})$ is dense in J (note that because the hyperbolic distance between $H(x)$ and $H_0(x)$ is uniformly bounded, H is again continuous as a map $\overline{X} \cup \{\infty\} \to J \cup \{\infty\}$). However, density of the image is trivial because $E_\kappa^{-1}(H(\overline{X})) \subset H(\overline{X})$, and backward orbits of any point (except κ) accumulate on the whole Julia set. Injectivity of H on X follows by the same argument as before. ■

An immediate corollary is the following.

4.2.2 Proposition (External rays landing at a common point)
Let κ be an attracting parameter. Then every non-escaping point in $J(E_\kappa)$ is the landing point of at least one external ray. Two external rays land at the same point if and only if they have identical itineraries.

PROOF. The only thing we still need to check is that, whenever $\operatorname{itin}(\underline{s}) = \operatorname{itin}(\underline{s}')$, then $H(\underline{s}, t_{\underline{s}}) = H(\underline{s}', t_{\underline{s}'})$. However, for any n, the n-th entries of $\underline{s}$ and $\underline{s}'$ differ by at most 1. It follows easily (for example by Lemma 3.2.2) that $|t_{\sigma^n(\underline{s})} - t_{\sigma^n(\underline{s}')}|$ is bounded independently of n. Thus the distance between the points $H_0(\sigma^n(\underline{s}), t_{\sigma^n(\underline{s})})$ and $H_0(\sigma^n(\underline{s}'), t_{\sigma^n(\underline{s}')})$ is uniformly bounded as well. The claim now follows by the contraction argument from the previous proof. ■

Note that the proof of 4.2.2 for periodic addresses is much easier, compare [76] and also Section 4.4.

4.2.3 Corollary ("Pinched Cantor Bouquet")
Let κ be an attracting parameter and $\underline{r} = \operatorname{addr}(\kappa)$. Form the quotient $\widetilde{X}$ of X by identifying all point $(\underline{s}, t_{\underline{s}})$ and $(\underline{s}', t_{\underline{s}'})$ for which

$$\operatorname{itin}_{\underline{r}}(\underline{s}) = \operatorname{itin}_{\underline{r}}(\underline{s}').$$

Then $\mathcal{F}$ projects to a map $\mathcal{F} : \widetilde{X} \to \widetilde{X}$ which is conjugate to $E_\kappa : J(E_\kappa) \to J(E_\kappa)$. ■

We have only constructed the Julia set as a quotient of our model space, and it is not quite obvious a priori that this quotient has a natural embedding in the plane, and that the homeomorphism extends continuously. While it is possible to do this, we will contend ourselves here to discuss the boundaries of Fatou components directly.

4.2.4 Theorem (Points on ∂U_1)
Let κ be an attracting parameter, and let U_1 be its characteristic Fatou component. Let $\mathtt{u}_1 \dots \mathtt{u}_{n-1}*$ be the kneading sequence of κ. Then a point $z \in J$ lies on ∂U_1 if and only if its itinerary $\tilde{\mathtt{u}}_1\tilde{\mathtt{u}}_2 \dots$ satisfies

$$\tilde{\mathtt{u}}_{kn+j} = \mathtt{u}_j \tag{4.2}$$

for every $k \geq 0$ and $1 \leq j < n$.

Proof. Clearly any point on ∂U_1 must have an itinerary of the form (4.2). Conversely, let z satisfy this condition. Then it is easy to see that the hyperbolic distance between $E_\kappa^{kn+j}(z)$ and ∂U_j is universally bounded, and the claim follows by contraction. ■

All the preceding theorems remain true also for parabolic parameters. Clearly, the main issue is to prove an analog of Theorem 4.2.1, whose proof gets somewhat more complicated by the fact that the hyperbolic metric is no longer uniformly expanded. This can be dealt with by considering the euclidean metric in a neighborhood of the parabolic point (using the asymptotic behavior in the repelling directions), and the hyperbolic metric outside this neighborhood. This is the same method as in the original proof that the Julia sets of parabolic quadratic polynomials are locally connected [27]. Those arguments, however, are somewhat technical and hardly very enlightening. Furthermore, in recent work by Haissinsky [39], parabolic rational maps were constructed from hyperbolic maps by using Guy David's transquasiconformal surgery. In particular, the resulting parabolic map is topologically conjugate to the hyperbolic function it originated from. It seems likely that these methods generalize to the space of exponential maps and thus yield a natural proof of the conjugacy of parabolic exponential maps to attracting exponential maps with the same combinatorics (i.e., with the same intermediate external address). In view of these issues, we have decided against a presentation of rigorous proofs of the above theorems in the parabolic case. Note also that we will, in the remainder of this thesis, apply these results only for periodic rays or periodic points, for which the proofs are much simpler.

Let us finally discuss the question of accessible points on ∂U_1. Clearly only the landing point of an external ray can be accessible from U_1. It is not difficult to prove the converse, compare [10]. (Note that [10] used a different, more complicated combinatorial description; however this is not substantial to the proof of accessibility.)

4.2.5 Theorem (Accessible Points)
Let κ be an attracting parameter and let $\mathbb{K}(\kappa) = \mathtt{u}_1\mathtt{u}_2 \dots \mathtt{u}_{n-1}*$. Then a point $z \in \partial U_1$ is accessible if and only if it is the landing point of an external ray. □

The following theorem was proved for the case of $\kappa \in (-\infty, -1)$ in [20]; the general case is not much different (but has, to our knowledge, not previously been published). Compare also [6].

4.2.6 Theorem (Prime Ends of U_1)
For an attracting parameter κ, let $\phi : \mathbb{D} \to U_1$ be a Riemann map of U_1. Then for every $\theta \in \mathbb{R}/\mathbb{Z}$, the curve

$$\gamma_\theta : [0, 1) \to U_1, \gamma_\theta(r) = \phi(re^{2\pi i\theta})$$

has a limit point for $r \to 1$.

In fact, every prime end impression of U_1 consists either of ∞ alone or of ∞ together with one or two rays and a (common) landing point. □

REMARK. Recall that the *prime end impression* of γ_θ is the set

$$\bigcap_{\varepsilon>0} \overline{\phi\big(\mathbb{D}_\varepsilon(e^{2\pi i\theta}) \cap \mathbb{D}\big)}.$$

PROOF. We may assume that ϕ is normalized so that γ_0 lies in the same homotopy class of U_1 as the attracting dynamic ray. (I.e., $E_\kappa^n(\gamma_0)$ is a curve to κ.) It is not difficult to see [20] that the map $\phi^{-1} \circ E_\kappa \circ \phi$ extends continuously to $\partial\mathbb{D} \setminus \{1\}$ and that the preimages of 1 are dense in ∂D. Let $\psi : (\mathbb{R}/\mathbb{Z}) \setminus \{0\} \to \mathbb{R}/\mathbb{Z}$ be the corresponding map on angles.

If $\theta \in K := \bigcup \psi^{-n}(0)$, then clearly γ_θ tends to ∞ for $r \to 1$. Furthermore, the prime end impression of γ_θ is $\{\infty\}$ because this is true for γ_0. Thus let $\theta \notin K$. We can choose sequences θ_n^+ and θ_n^- in K which converge to θ from above and below, respectively. Let

$$\mathrm{addr}_n^\pm := \mathrm{addr}(\gamma_{\theta_n^\pm}).$$

These sequences are (eventually) monotone decreasing resp. increasing, and thus converge to two addresses $\underline{s}^+$ and $\underline{s}^-$. It is easy to check that these addresses have a common itinerary. (Otherwise there would be a curve $\gamma_{\theta'}$ with $\theta' \in K$ between them, which is a contradiction.) Clearly every point in the impression of γ_θ must then have this same itinerary.

If the addresses $\underline{s}^+$ and $\underline{s}^-$ are not exponentially bounded, then it follows that the impression of γ_θ is $\{\infty\}$ (and indeed $\underline{s}^+ = \underline{s}^- = \mathrm{addr}(\gamma_\theta)$). Otherwise, the two rays $g_{\underline{s}^+}$ and $g_{\underline{s}^-}$ land together, and the impression consists of these rays together with ∞ and their common landing point. Furthermore, no other $\gamma_{\theta'}$ can accumulate at these points. By Theorem 4.2.5, the landing point z of the rays is accessible; thus it is the single principle point of this prime end and $\lim_{r\to 1} \gamma_\theta(r) = z$. ■

A corresponding statement can be made for attracting dynamic rays (cf Section 4.1). To the best of our knowledge, this result has not previously been published, so we again include a sketch of its proof.

4.2.7 Theorem (Attracting Dynamic Rays Land)
Let κ be an attracting parameter. Then all unbroken attracting dynamic rays land in $\mathbb{C}$.

SKETCH OF PROOF. In the case where the multiplier is rational, this is Lemma 4.1.7. So suppose that $\mu = re^{2\pi i\alpha}$ with $\theta \notin \mathbb{Q}$. Let $\Phi : U_1 \to \mathbb{C}$ be the extended linearizing map, normalized as usual so that $\Phi(\kappa) = 1$. Denote by g_θ the attracting dynamic ray at angle θ; i.e. the component of

$$\Phi^{-1}\big(\{re^{2\pi i\theta} : r > 0\}\big)$$

which contains a_1.

Let $\theta \in \mathbb{R}/\mathbb{Z}$ with $n\alpha + \theta \neq 0 \pmod 1$. Then we can find $n_k^\pm$ so that

$$-n_k^\pm \cdot \alpha \to \theta \pmod 1$$

from above and below, respectively. The claim now follows similarly to that of Theorem 4.2.6, by considering the attracting rays $g_n^\pm := g_{-n_k^{pm} \cdot \alpha}$. The landing point cannot be ∞ because its itinerary is bounded. ■

4.3 Invalidity of Renormalization

Renormalization plays an important part in the study of polynomial dynamics. While there is no natural generalization of the original Douady-Hubbard concept to the study of exponential maps, it was hoped that some concept of renormalization does exist in this family, and that in particular, the renormalization principle is topologically valid. Let us now describe a consequence that this principle would have, and prove that this is false even for attracting parameters.

Suppose that E_κ is any hyperbolic exponential map of period $n > 1$, let μ be the multiplier of its attracting orbit and let U_0 be the Fatou component containing a left half plane. $E_\kappa^n|_{U_0}$ is conformally conjugate to $E_{\kappa_0}|_{F(E_{\kappa_0})}$ where κ_0 is such that E_{κ_0} has an attracting fixed point with multiplier κ_0 (in fact, $\kappa_0 = \log\mu - \mu$). This can be proved either by constructing the conjugacy directly using the linearizing coordinates of E_κ and E_{κ_0}, or by conjugating these maps to a normal form as in [73, Section III.4] or [20]. Let

$$\Phi : U_0 \to F(E_{\kappa_0})$$

be this conjugacy, and note that $\Phi(z + 2\pi i) = \Phi(z) + 2\pi i$. For simplicity, we will assume that $\kappa_0 = -2$ (in fact, this can always be achieved by following Φ by a quasiconformal homeomorphism, see Theorem 5.2.3). The principle of topological renormalization would imply that this map extends continuously to $\overline{U_0}$. We show that this cannot be the case. The reason is similar to the argument from Section 3.5: Orbits of E_κ^n can grow much faster than those of E_{κ_0}, but imaginary parts of two points which correspond under Φ differ by no more than a fixed constant because the $2\pi i$-periodic structure must be preserved.

4.3.1 Theorem (No Topological Renormalization)
Let E_κ and Φ be as above. Then Φ does not have a continuous extension to the boundary of U_0.

PROOF. Assume, by contradiction, that Φ does extend continuously to ∂E_κ.

Again, cut the plane into the strips R_k used in the definition of itineraries; these strips have bounded imaginary parts. We may assume that they are numbered so that R_0 contains 0. Let

$$R := R_0 \cup \bigcup_{1 \leq j \leq n-1} R_{\mathtt{u}_j},$$

where $\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}*$ is the kneading sequence of $\underline{s}$. Thus the orbit of any point $z \in \overline{U_0}$ lies in $\overline{U_0} \cup R$. Define $A := \sup_{z\in R} |\operatorname{Im}(z)|$, and note that $A \geq \pi$.

Choose K large enough such that, whenever $|E_\kappa(z)| \geq K$, then

$$|E_\kappa(z)| - A - 1 \geq \exp(\operatorname{Re}(z) - 1) \tag{4.3}$$

and

$$|E_\kappa(z)| + A + 1 \leq \exp(\operatorname{Re}(z) + 1). \tag{4.4}$$

Let

$$M := \max\left\{|\Phi(z)| : z \in U_0 \cap \overline{R_0} \text{ and } \operatorname{Re}(z) \in [K-1, K+1]\right\}$$

and choose k so large that $\exp^k(K - \pi - A) > M$.

Pick any point $z_1 \in \overline{U_0}$ with $\operatorname{Re}(z_1) = 0$ and

$$\operatorname{Im}(z_1) \in \left[\exp^{kn}(K) - \pi, \exp^{kn}(K) + \pi\right].$$

There exists a unique point $z_0 \in R_0 \cap \overline{U_0}$ with $E_\kappa^{kn}(z_0) = z_1$ and $E_\kappa^{jm}(z_0) \in R_0$ for $j < n$. By (4.3), we see that

$$\operatorname{Re}\bigl(E_\kappa^{kn-1}(z_0)\bigr) - 1 \leq \log(|z_1| - A - 1) \leq \exp^{kn-1}(K)$$

and, similarly, by (4.4,

$$\operatorname{Re}\bigl(E_\kappa^{kn-1}(z_0)\bigr) + 1 \geq \exp^{kn-1}(K).$$

In particular, $\operatorname{Re}\bigl(E_\kappa^{kn-1}(z_0)\bigr) \geq K$. Repeating this argument inductively, it follows that

$$\operatorname{Re}(z_0) \in [K-1, K+1].$$

Because $\Phi(z + 2\pi i) = \Phi(z) + 2\pi i$, we can estimate that

$$\operatorname{Im}(\Phi(z_1)) \geq \exp^{kn}(K) - \pi - A$$

On the other hand, $\operatorname{Re}(\Phi(z_0)) \leq M$, and thus

$$\begin{aligned}|\Phi(z_1)| = |E_{-2}^k(\Phi(z_0))| &< \exp^k(\operatorname{Re}(\Phi(z_0))) \leq \exp^k(M) < \exp^{2k}(K - \pi - A)\\ &< \exp^{kn-1}(K) - \pi - A.\end{aligned}$$

This is a contradiction. ■

4.4 Escaping and Misiurewicz Parameters

Let κ be an escaping parameter; say $\kappa = g_{\underline{s}}^\kappa(t_0)$. The preimages of $g_{\underline{s}}\bigl((t_0, \infty)\bigr)$ form a partition of the plane, see Figure 4.2(a). Thus we can assign each point in $\mathbb{C}$ an itinerary, completely analogous to Definition 4.1.2. (Note that the itinerary of any point whose forward orbit intersects the ray $g_{\underline{s}}$ will contain boundary symbols.)

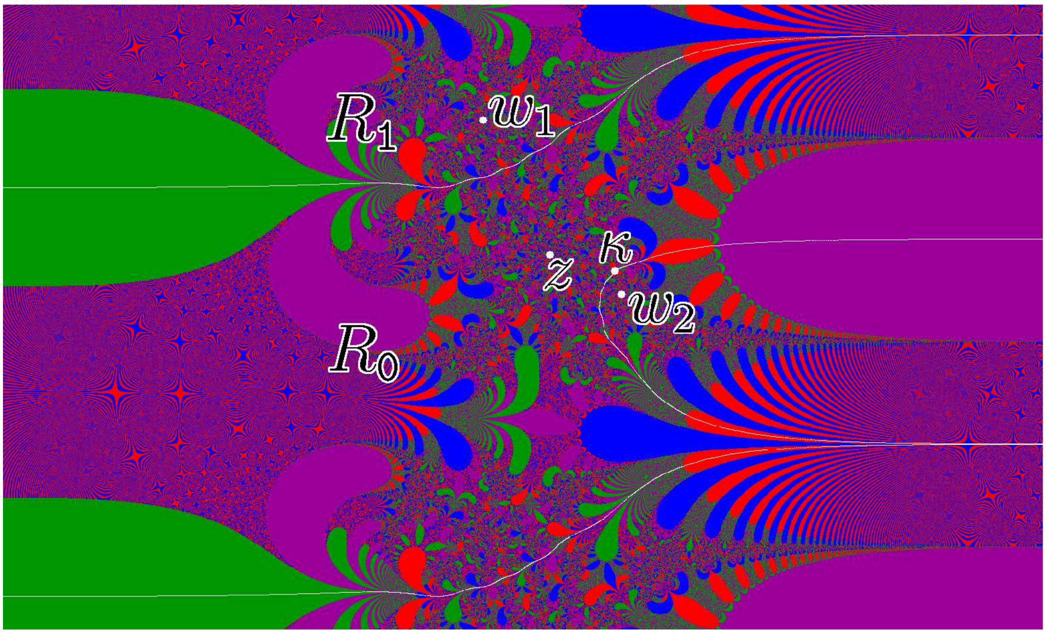

(a) The dynamical partition for an escaping parameter with $\mathrm{addr}(\kappa) = \overline{01}$. There are no periodic rays landing at either the fixed point z or the cycle $w_1 \mapsto w_2 \mapsto w_1$.

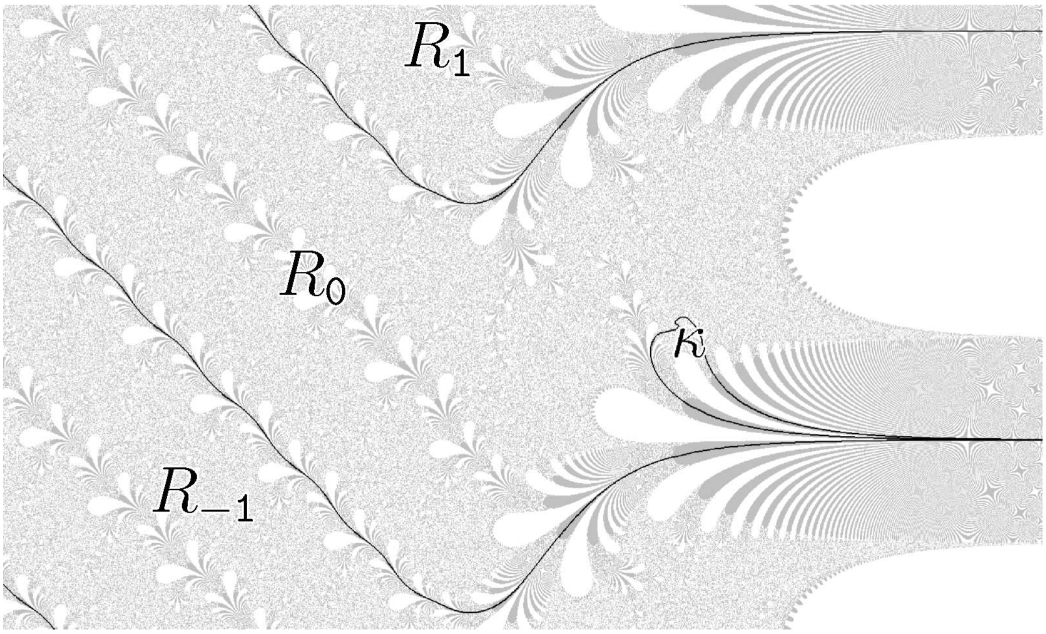

(b) The dynamical partition for a Misiurewicz parameter with $\mathrm{addr}(\kappa) = 01\overline{10}$. There are two preperiodic rays landing at the singular value.

Figure 4.2: Illustration of the definition of itineraries.

4.4.1 Theorem (Periodic Points are Landing Points [76])
Let κ be an escaping parameter. Then no two periodic or preperiodic points of E_κ have the same itinerary. A periodic external ray lands at a given periodic point if and only if both have the same itinerary.

Furthermore, every (pre)-periodic point z is the landing point of at least one (pre)-periodic ray, unless there exists $n \geq 1$ such that $\mathrm{itin}(E_\kappa^n(z)) \in \{\mathbb{K}^+(\underline{s}), \mathbb{K}^-(\underline{s})\}$.

Sketch of proof. The first statement is proved by the usual hyperbolic contraction argument: The map E_κ expands the hyperbolic metric of

$$U := \mathbb{C} \setminus \bigcup_{k \geq 0} E_\kappa^k \left(g_{\underline{s}}([t_0, \infty)) \right).$$

The hyperbolic length of any curve in U connecting two periodic points with the same itinerary thus shrinks to 0 under the corresponding pullbacks. Since every periodic ray (which does not map to $g_{\underline{s}}$ under iteration of E_κ) lands at a periodic point with the same itinerary 3.9.1, the second claim is also proved. The last statement now follows from Lemma 4.1.5. □

Now let κ be a Misiurewicz parameter; i.e., $E^n(\kappa)$ is periodic for some $n \geq 1$. Using the theory of "spiders", a method of iteration in Teichmüller spaces, Schleicher and Zimmer construct a ray that lands at the singular value.

4.4.2 Theorem (Ray at Singular Orbit [76, Theorem 4.3])
Let κ be a Misiurewicz parameter. Then there exists a preperiodic address $\underline{s}$ such that the ray $g_{\underline{s}}^\kappa$ lands at the singular value of E_κ. □

Remark. We shall write $\mathrm{addr}(\kappa) = \underline{s}$ in analogy with the attracting and escaping cases. One should bear in mind, however, that the address $\underline{s}$ need not be uniquely determined.

Again, one can now define itineraries, using the preimages of $g_{\underline{s}}$ as partition boundaries; see Figure 4.2(b). It should not be surprising that this suffices to prove that periodic points are landing points for Misiurewicz parameters.

4.4.3 Theorem (Periodic Points are Landing Points [76])
No two periodic or preperiodic points of E_κ have the same itinerary. A (pre)-periodic external ray lands at a given (pre)-periodic point if and only if both have the same itinerary. Furthermore, every (pre)-periodic point z is the landing point of a (pre)-periodic external ray.

Sketch of proof. The only significant difference to the escape case consists of showing that any periodic point z whose itinerary coincides with that of $w := E_\kappa^n(\kappa)$ must be equal to w. This is done by connecting z to a linearizing neighborhood of w and then repeating the contraction argument. □

Schleicher and Shishikura [75] extended Thurston's classification theorem to postsingularly finite exponential maps, which yields the following theorem.

4.4.4 Theorem (Existence of Misiurewicz Parameters [75])
For every preperiodic address $\underline{s}$, there exists a unique parameter $\kappa \in \mathbb{C}$ for which $g_{\underline{s}}^{\kappa}$ lands at κ. □

4.5 Orbit Portraits and Characteristic Rays

Let $\kappa \in \mathbb{C}$. Following Milnor [58], we define the notion of *orbit portraits* to encode the dynamics of periodic rays landing at common points.

4.5.1 Definition (Orbit Portrait)
Let $(z_1, \dots, z_n)$ be a repelling or parabolic periodic orbit for E_κ. Define

$$A_k := \{\underline{r} \in \mathcal{S} : \underline{r} \text{ is periodic and } g_{\underline{r}} \text{ lands at } z_k\}.$$

Then $\mathcal{O} := \{A_1, \dots, A_p\}$ is called the orbit portrait *of (z_n). The orbit (and the orbit portrait) is called* essential *if $|A_k| > 1$ for any k.*

4.5.2 Lemma (Basic Properties of Orbit Portraits)
Let $\mathcal{O} = \{A_1, \dots, A_p\}$ be an orbit portrait. Then all A_k are finite, and all addresses in the portrait share the same period qn. qn is called the ray period *of the orbit. The shift map carries A_k bijectively onto A_{k+1}.*

Proof. The proof that all rays share the same period and are transformed bijectively by the shift is completely analogous to the polynomial setting [56, Lemma 18.12]. Let n be the common period of the rays in $\mathcal{O}$ and let $\underline{s} = \overline{s_1 s_2 \dots s_n} \in A_1$. By Lemma 3.8.2, it follows that

$$A_1 \subset \left\{ \overline{s'_1 s'_2 \dots s'_n} : |s'_j - s_j| \leq 1 \right\},$$

thus A_1 is finite. ■

An orbit portrait is a combinatorial object, and we will often suppress the actual choice of parameter present in its definition.

4.5.3 Definition (Characteristic Rays)
Let (z_k) be a periodic orbit as before, and let $\mathcal{O}$ be its orbit portrait. Suppose there exist j and $\underline{r}, \underline{\tilde{r}} \in A_j$ such that the rays $g_{\underline{r}}$ and $g_{\underline{\tilde{r}}}$, together with their common landing point z_j, separate the singular value from all other rays of the orbit portrait.

Then $g_{\underline{r}}$ and $g_{\underline{\tilde{r}}}$ are called the characteristic rays *(and $\underline{r}$ and $\underline{\tilde{r}}$ the* characteristic addresses*) of the orbit (z_k). The interval in $\mathcal{S}$ bounded by $\underline{r}$ and $\underline{\tilde{r}}$ is called the* characteristic sector *of $\mathcal{O}$.*

4.5.4 Lemma (Permutation of Dynamic Rays [74, Lemma 5.2])
Every essential periodic orbit has exactly two characteristic rays; their addresses depend on $\mathcal{O}$, but not on κ. Furthermore, if $|A_k| > 2$, then the rays landing at each z_j are permuted cyclically by the first return map. □

As in the Mandelbrot set, it will turn out that orbit portraits are "born" in bifurcations of hyperbolic components. Therefore we will now turn our attention once more to the case of attracting or parabolic parameters. Let κ be an attracting parameter, and let $U_0 \mapsto U_1 \mapsto U_2 \mapsto \dots \mapsto U_n$ denote the cycle of Fatou components, with the usual convention that $U_n = U_0$ contains a left half plane. We are interested in points on ∂U_1 which are fixed by the first return map E^n of U_1 (see Figure 4.3(a)). Let us first state how these are combinatorially distinguished.

4.5.5 Lemma (Boundary Fixed Points)
Let $\kappa \in \mathbb{C}$ be an attracting or parabolic parameter of period n, and let $\underline{\mathtt{u}} = \mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}* = \mathbb{K}(\kappa)$. For every $\mathtt{m} \in \mathbb{Z}$, there exists a unique periodic point $z_{\mathtt{m}} \in \partial U_1$ (of period dividing n) with $\operatorname{itin}(z_{\mathtt{m}}) = \overline{\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}\mathtt{m}}$. There are no other fixed points of E_κ^n on ∂U_1. Furthermore, if $|\mathtt{m}|$ is large enough, then $z_{\mathtt{m}}$ is the landing point of the external ray at address $\overline{s_1 s_2 \dots s_{n-2}(s_{n-1} + \frac{1}{2})\mathtt{m}}$ (if $\mathtt{m} < 0$) or $\overline{s_1 s_2 \dots s_{n-2}(s_{n-1} - \frac{1}{2})(\mathtt{m}+1)}$ (if $\mathtt{m} > 0$).

Proof. The fact that each itinerary is realized follows by Lemma 4.1.5, and for large $|\mathtt{m}|$ the given addresses have the correct itinerary. The rest of the claim follows from Theorem 4.2.4. ■

For hyperbolic quadratic polynomials, there is always a pair of periodic rays landing together at the boundary of a periodic attracting Fatou component. This is the main motor of the combinatorial description of their dynamics. Therefore it is important to know that the same is the case for exponential maps.

4.5.6 Theorem and Definition (Dynamic Root [74, Theorem 6.2])
Let κ be an attracting or parabolic parameter, and let $n \geq 2$ be the period of U_1. Then exactly one fixed point of the first return map $E_\kappa^n|_{\partial U_1}$ is the landing point of two or more rays.

This point is called the dynamic root *of U_1. The characteristic rays of its orbit portrait are called the* characteristic rays *of κ.*

Proof. Note that it is clear that there can be at most one such fixed point. Indeed, if $\mathcal{O}$ is *any* nonessential orbit portrait, then the two characteristic rays of $\mathcal{O}$ will enclose κ, and thus all of ∂U_1. Because rays do not intersect, this means that any other pair of rays landing together on ∂U_1 separates κ from all rays in the orbit portrait $\mathcal{O}$. In particular, there can be at most one such pair.

It thus suffices to prove existence. We will give a proof which is a variation on that in [74] and shows how the orbit of the characteristic rays relates to the corresponding Fatou components.

Let $\underline{s} = \operatorname{addr}(\kappa)$ and let $\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}* = \mathbb{K}(\underline{s})$. For abbreviation, we define

$$\operatorname{ADDR}(U_j) := \{\operatorname{addr}(\gamma) : \gamma : [0,\infty) \to U_j, |\gamma(t)| \to \infty\}.$$

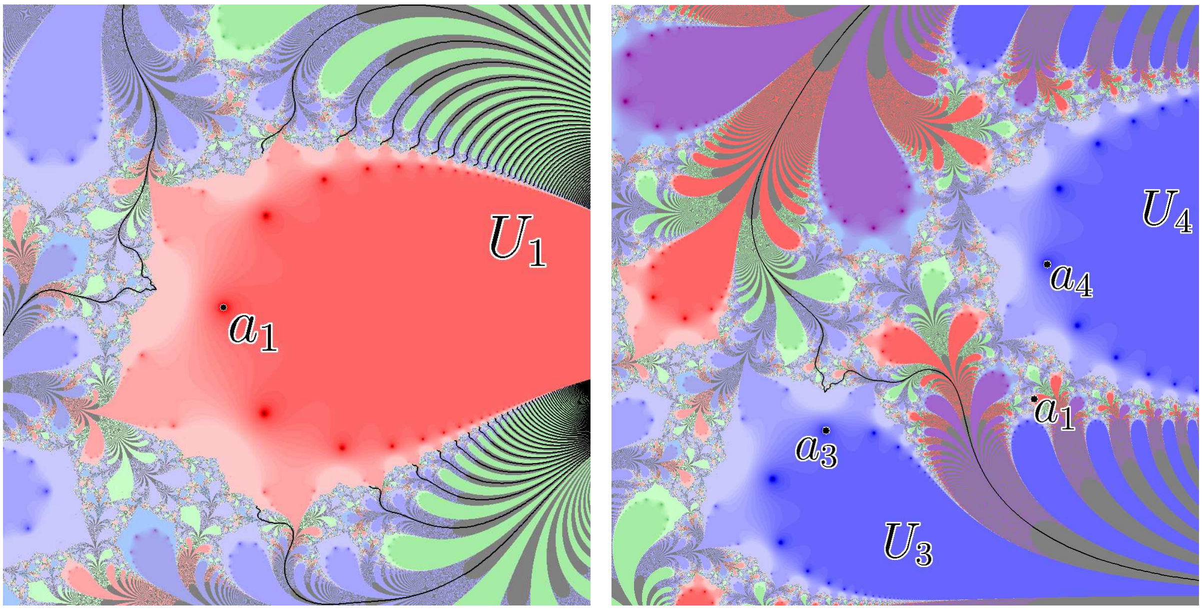

(a) Periodic rays landing at boundary fixed points.

(b) U_3 surrounds the rays $g_{\sigma^2(\underline{t}^{\pm})}$ (as well as U_1 and U_4)

Figure 4.3: External rays for an attracting parameter with address $020\frac{1}{2}\infty$.

We define two external addresses by

$$\underline{t}^+ := \sup \mathrm{ADDR}(U_1) \qquad \text{and}$$
$$\underline{t}^- := \inf \mathrm{ADDR}(U_1),$$

and claim that these addresses are periodic of period n and have a common itinerary of the form $\mathtt{u}_1\mathtt{u}_2\ldots\mathtt{u}_{n-1}\mathtt{m}$ (which will conclude the proof). Note that $\underline{t}^{\pm} \neq s$; indeed, U_0 contains curves both above and below U_{n-1}, and these map to curves in U_1 below and above any curve representing $\underline{s}$.

To prove this claim, we show inductively that, for $1 \leq j \leq n+1$, the pairs $(\underline{t}_j^-, \underline{t}_j^+) := (\sigma^{j-1}(\underline{t}^-), \sigma^{j-1}(\underline{t}^+))$ have the following properties.

a) $\underline{t}_j^- = \inf\{\underline{r} \in \mathrm{ADDR}(U_j) : \underline{r} \geq \underline{t}_j^-\}$ and $\underline{t}_j^+ = \sup\{\underline{r} \in \mathrm{ADDR}(U_j) : \underline{r} \leq \underline{t}_j^+\}$;

b) There exists k such that $k\underline{s} \leq \underline{t}_j^+, \underline{t}_j^- \leq (k+1)\underline{s}$;

c) For every address $\underline{r} \in \mathrm{ADDR}(U_{j+1})$,

$$\underline{t}_j^- \leq \underline{r} \leq \underline{t}_j^+ \tag{4.5}$$

in the circular ordering of $\overline{\mathcal{S}}$.

Note that, for $j \neq n$, b) follows from a) (with $k = \mathtt{u}_j$). Let us now show a) to c) by induction on j; the proof of the theorem will then be complete: for $j = n+1$, a) and c) imply that $\underline{t}_j^+ = \underline{t}^+$ and $\underline{t}_j^- = \underline{t}^-$, and b) ensures that the addresses have the correct itinerary entries.

Clearly the claims a) to c) are true for $j = 1$. Now suppose they is satisfied for some j, $1 \leq j \leq n$. Then it follows from b) that (4.5) is also true for the projections of the addresses to the cylinder $\tilde{\mathcal{S}}$. Because the shift is order-preserving as a map from the cylinder $\tilde{\mathcal{S}}$ to $\overline{\mathcal{S}}$, claims a) and c) follow for $j+1$. In the case where $j \neq n-1$, this also proves b) for $j+1$.

It thus remains to show b) for $j = n$ if we know that a) and c) hold. Let $k := \max k : k\underline{s} \leq \underline{t}_n^+$. Because $(k+1)\underline{s} \in \mathrm{ADDR}(U_0)$, it follows that $k\underline{s} \leq \underline{t}_n^+ \leq \underline{t}_n^- \leq (k+1)\underline{s}$, which proves b). ■

Remark. Note that it is possible that $\sigma^j(\underline{t}^+) < \sigma^j(\underline{t}^-)$ for $j < n-1$, in which case $\mathrm{ADDR}(U_{j+1})$ lies *outside* of this ray pair in the (linear) ordering of $\mathcal{S}$. Compare Figure 4.3(b). It is easy to see that in this case the ray pair must enclose one of the components $U_1, \dots, U_j$. In particular, $\mathtt{u}_n \in \{\mathtt{u}_1, \dots, \mathtt{u}_{n-1}\}$.

It is apparent from the proof of the previous theorem that the relationship between $\underline{s}$ and the characteristic addresses of κ is purely combinatorial. (One needs to check that the sets $\mathrm{ADDR}(U_j)$ depend only on $\underline{s}$, but this is easily done; compare Theorems 4.2.4 and 4.2.6.) In fact, this relationship can be expressed in the following form:

4.5.7 Definition (Intermediate Address Associated to an Orbit Portrait)
Let $\underline{s}$ be an intermediate external address of length $n \geq 2$, and let $\underline{r}$ and $\underline{\tilde{r}}$ be periodic of period n with $\underline{r} < \underline{s} < \underline{\tilde{r}}$. Then we say that $\underline{r}$ and $\underline{\tilde{r}}$ are characteristic addresses for $\underline{s}$ *if* $\mathrm{itin}_{\underline{s}}(\underline{r}) = \mathrm{itin}_{\underline{s}}(\underline{\tilde{r}})$, *the first $n-1$ entries of this itinerary agree with those of $\mathbb{K}(\underline{s})$ and*

$$\big(\{\sigma^j(\underline{r}) : j \geq 0\} \cup \{\sigma^j(\underline{\tilde{r}}) : j \geq 0\}\big) \cap (\underline{r}, \underline{\tilde{r}}) = \emptyset.$$

Conversely, we say that $\underline{s}$ is associated to an orbit portrait $\mathcal{O}$ *if the characteristic addresses for $\mathcal{O}$ are characteristic for $\underline{s}$ in the above sense.*

Remark. If κ is an attracting or parabolic parameter and $\underline{s} = \mathrm{addr}(\kappa)$, then the addresses of the characteristic rays of E_κ are characteristic addresses for $\underline{s}$ in the above sense. Thus, using this terminology, we could formulate a purely combinatorial version of Theorem 4.5.6 as follows: Every intermediate external address has exactly one pair of characteristic addresses. However, it is not immediate how to efficiently construct these addresses from $\underline{s}$; an algorithm for this will be given in Section 5.7. On the other hand, it is very easy to reconstruct $\underline{s}$ from its characteristic addresses, compare Lemma 4.5.12. (In particular, every orbit portrait has exactly one associated intermediate external address.)

One of the fundamental questions for the study of parameter space that will follow is the following: When does an exponential map realize a given orbit portrait? We can give a (combinatorial) answer to this question using the following observations.

4.5.8 Lemma (Behavior of Itineraries under Change of Partition)
Let $\mathcal{O}$ be an essential orbit portrait, and let $\underline{s} \in \mathcal{S}$. Then the characteristic addresses of $\mathcal{O}$ have common itinerary with respect to $\underline{s}$ if and only if $\underline{s}$ lies in the characteristic sector of $\mathcal{O}$.

If this is the case, then this common itinerary does not depend on the choice of $\underline{s}$, nor does the itinerary of any intermediate address associated to $\mathcal{O}$.

Proof. When we refer to "strips" in the following, we mean the intervals $\big(m\underline{s}, (m+1)\underline{s}\big) \subset \mathcal{S}$, which are used to define itineraries. Let $\underline{r} \leq \underline{\tilde{r}}$ be the characteristic addresses of $\mathcal{O}$. We again denote $\underline{r}_j := \sigma^{j-1}(\underline{r})$.

First suppose that $\underline{s}$ does not belong to the characteristic sector of $\mathcal{O}$. The shift map takes the interval between $\underline{\tilde{r}}_{n-1}$ and $\underline{r}_{n-1}$ bijectively to $\overline{\mathcal{S}} \setminus [\underline{r}, \underline{\tilde{r}}]$, where $[\underline{r}, \underline{\tilde{r}}]$ is the characteristic sector of $\mathcal{O}$. It follows that there exists some preimage of $\underline{s}$ between $\underline{\tilde{r}}_n$ and $\underline{r}_n$. In other words, these two addresses lie in different strips, and thus their itineraries differ. Conversely suppose that $\underline{s}$ does lie between $\underline{r}$ and $\underline{\tilde{r}}$. Then, in analogy to the previous argument, the entire interval $[\underline{\tilde{r}}_n, \underline{r}_n]$ is contained in a single strip. We claim that, for every $j < n$, the two addresses $\underline{r}_j$ and $\underline{\tilde{r}}_j$ lie in one of the translates of that interval. (By this we mean the set obtained by adding some integer to the first entry of all addresses in the interval.) Indeed, otherwise they would be mapped into the characteristic sector of $\mathcal{O}$ by the same argument as before. This is impossible by the definition of characteristic addresses. Thus $\underline{r}_j$ and $\underline{\tilde{r}}_j$ lie in the same strip for every j, which concludes the proof of the first claim.

To prove the second statement, first consider any $\underline{s}$ and $\underline{t}$ with $\sigma(\underline{t}) \neq \underline{s}$. Then, if $\underline{s}$ varies continuously without moving through $\sigma(\underline{t})$, the first entry of $\operatorname{itin}_{\underline{s}}(\underline{t})$ does not change. The addresses $\underline{r}$ and $\underline{\tilde{r}}$ never map into the characteristic sector under a positive iterate of the shift. Thus every itinerary entry of these addresses is locally constant, and therefore constant, when $\underline{s}$ is varied in the characteristic sector.

It is also true that any intermediate external address $\underline{t}$ associated to $\mathcal{O}$ never maps into the characteristic sector. One way to see this is the following: by [74, Theorem 3.5], there exists some attracting parameter κ with $\operatorname{addr}(\kappa) = \underline{t}$; compare Theorem 5.2.5. Then $\underline{r}$ and $\underline{\tilde{r}}$ are the characteristic addresses of κ, and there exists some curve γ in U_1 with $\operatorname{addr}(\gamma) = \underline{t}$. The forward images of γ are separated from γ by the rays $g_{\underline{r}}$, $g_{\underline{\tilde{r}}}$ and their common landing point, and the claim follows. (It is also possible to prove this by purely combinatorial means.) ∎

4.5.9 Corollary (Characteristic Rays determine Orbit Portrait)
Let $\mathcal{O}$ be an essential orbit portrait. Let κ be some attracting, parabolic, Misiurewicz or escaping parameter with $\operatorname{addr}(\kappa) = \underline{s}$*. Then the following are equivalent:*

1. *E_κ has an orbit with portrait $\mathcal{O}$.*
2. *$\underline{s}$ lies in the characteristic sector of $\mathcal{O}$.*
3. *The characteristic rays of $\mathcal{O}$ land together under E_κ.*

Proof. Assume that E_κ has an orbit with portrait $\mathcal{O}$. Then the characteristic rays of $\mathcal{O}$ enclose the singular value κ, and therefore also the curve γ ending at κ used in the definition of itineraries. Thus $\underline{s} = \operatorname{addr}(\gamma) = \operatorname{addr}(\kappa)$ lies between the addresses of these rays. So 1. implies 2. Also, 2. is equivalent to 3. by Lemma 4.5.8.

Now suppose that 3. holds, i.e. that the characteristic rays of $\mathcal{O}$ land together for E_κ. Then the period of their common landing point z is equal to that of its itinerary, which does not depend on $\underline{s}$ by Lemma 4.5.8. Thus, the orbit portrait of z is an extension of $\mathcal{O}$ (by which we mean that the set A_k of rays landing at $E_\kappa^{k-1}(z)$ contains at least all rays required by $\mathcal{O}$). It remains to show that this orbit portrait cannot contain any additional orbits of rays. But otherwise there would be at least three rays landing at z, while the orbit portrait contains more than one cycle of rays. This contradicts Lemma 4.5.4. ■

Recall that, in the Mandelbrot set, the *wake* of a hyperbolic component is the set of parameters for which the characteristic rays of this component land together, or equivalently the region enclosed by its two characteristic *parameter rays*. The analog of this fact for exponential maps will have to wait until Section 5.13; however the previous Lemma shows that at least a *combinatorial* variant of this notion exists.

4.5.10 Definition (Combinatorial Wake)
Let $\underline{s}$ be an intermediate external address, and $\underline{r} < \underline{\tilde{r}}$ its characteristic addresses. The (combinatorial) wake *of $\underline{s}$ is defined to be*

$$\mathcal{W}(\underline{s}) := \{\underline{t} \in \overline{\mathcal{S}} : \underline{r} < \underline{t} < \underline{\tilde{r}}\}.$$

Remark. This object is, strictly speaking, just the characteristic sector of $\mathcal{O}$. However, one should think of the orbit portrait as situated in "combinatorial dynamical space", while the wake lies in "combinatorial parameter space". (This is the same distinction as for "itineraries" and "kneading sequences".)

4.5.11 Definition (Forbidden Kneading Sequence)
Let $\underline{t}_1, \underline{t}_2$ be the characteristic addresses of an orbit portrait, and let $\underline{s}$ be an associated intermediate external address. Then $\mathtt{u} = \operatorname{itin}_{\underline{s}}(\underline{t}_1) = \operatorname{itin}_{\underline{s}}(\underline{t}_2)$ is called the forbidden kneading sequence *of $\underline{s}$ and denoted by $\mathbb{K}^*(\underline{s})$.*

Remark. The reason for this terminology is the following: If $\underline{r} \in \mathcal{W}(\underline{s})$ is any external address, then $\mathbb{K}(\underline{s})$ cannot begin with $\mathtt{u}_1\mathtt{u}_2\ldots\mathtt{u}_n$. The relevance of this fact (which is also proved there) will be seen in Section 5.9.

To conclude, we now give an algorithm for finding the (unique) intermediate address associated to an orbit portrait.

4.5.12 Lemma (Unique Address for Orbit Portraits)
Every essential orbit portrait $\mathcal{O}$ has exactly one associated intermediate external address. This address can be computed by a simple algorithm.

Proof. Let $\underline{r} < \underline{\tilde{r}}$ be the characteristic addresses of $\mathcal{O}$. Choose any $\underline{s}'$ between $\underline{r}$ and $\underline{\tilde{r}}$. Let us construct an address $\underline{s}$ whose itinerary with respect to $\underline{s}'$ agrees with that of the characteristic rays on the first $n-1$ entries. We define the entries s_k of $\underline{s}$ inductively, starting with $s_{n-1} \in \mathbb{Z} + \frac{1}{2}$.

We have to choose s_{n-1} such that $\underline{s}_{n-1} = s_{n-1}\infty$ is not separated from $\underline{r}_{n-1}$ and $\underline{\tilde{r}}_{n-1}$ by the partition $*\underline{s}'$. This obviously determines s_{n-1} uniquely; namely $s_{n-1} = k + \frac{1}{2}$ where $k\underline{s}' < \underline{r}_{n-1} < (k+1)\underline{s}'$. When s_j is known for $j > k$, the same argument yields s_k: it is chosen such that $s_k s_{k+1} \dots s_{n-1}\infty$ is the unique address of this form in the strip containing $\big(\mathtt{u}_k\underline{s}, (\mathtt{u}_k + 1)\underline{s}\big)$.

This construction, together with the second observation in Lemma 4.5.8, proves uniqueness of the associated intermediate address. To prove that the constructed address is associated to the orbit portrait, it only remains to see that it lies in the characteristic sector. However, $\underline{r}_n < \infty < \underline{\tilde{r}}_n$ in the circular ordering of $\overline{\mathcal{S}}$, and this order is preserved under repeated pullbacks through the same strips. ■

Chapter 5

Parameter Space

After having studied the dynamical plane of exponential maps in general and several important classes of maps in particular, we now turn to an investigation of the κ-plane. The first main result of this chapter is the fact that boundaries of hyperbolic components are connected (Theorem 1.8). The proof (in Section 5.3) will use an important tool called the "Squeezing Lemma" (Theorem 5.3.5), which is proved in the subsequent sections. Along the way, we also give a complete description of the bifurcation structure of hyperbolic components, including several new combinatorial algorithms, e.g. to calculate the characteristic rays associated to an intermediate external address.

In Section 5.12, we also discuss parameter rays. Using our results in Chapter 3 and the Squeezing Lemma, we extend the work of Schleicher [73, Chapter 2] and Förster [35] to a complete classification of escaping parameters.

Finally, with all the tools in place, we give proofs of Theorems 1.10 and 1.11 in Sections 5.13 and 5.14, using arguments in the parameter plane.

5.1 Structural Stability

5.1.1 Definition (Structural Stability and Bifurcation Locus)
A parameter $\kappa_0 \in \mathbb{C}$ is called structurally stable *if there exists a neighborhood U of κ_0 such that, for every $\kappa \in U$, the map E_κ is conjugate to E_{κ_0}. The complement of the set of structurally stable parameters is called the* bifurcation locus *and is denoted by $\mathcal{B}$.*

The theorem of Mañé, Sad and Sullivan [51] that structural stability is dense in the space of rational maps has been generalized by Eremenko and Lyubich [31] to parameter spaces of entire functions with only finitely many singularities; in particular to the exponential family.

5.1.2 Theorem (Structural Stability is Dense [31])
The structurally stable parameters form an open and dense subset of $\mathbb{C}$. □

As in the case of rational maps, there are several equivalent definitions of structural stability. We mention some of these here; for further details the reader is referred to [85].

5.1.3 Theorem (Equivalent Definitions of Structural Stability [31, 85])
Let $\kappa_0 \in \mathbb{C}$. Then the following properties are equivalent.

- κ_0 *is structurally stable.*
- κ_0 *is* J-stable*; i.e. there exists a neighborhood of* κ_0 *in which all maps are conjugate on their Julia sets.*
- *There exists a neighborhood* U *of* κ_0 *such that either* $J(E_\kappa) = \mathbb{C}$ *for all* $\kappa \in U$ *or* $J(E_\kappa) \neq \mathbb{C}$ *for all* $\kappa \in U$.
- *There is a neighborhood of* κ_0 *which contains no indifferent parameters.*
- *The family of functions* $\kappa \mapsto E_\kappa^n(\kappa)$ *is normal in* κ_0. □

It was proved by Devaney [17] that $E_0 = \exp$ is not structurally stable. This result has been generalized to arbitrary escaping parameters by Ye [85]. (This can also be proved using a perturbation argument as in Section 5.4.) Note also that Misiurewicz parameters cannot be structurally stable, because the equation $E_\kappa^m(\kappa) = E_\kappa^{m+n}(\kappa)$ has a discrete set of solutions by the identity theorem.

5.1.4 Theorem (Structurally Unstable Parameters)
Escaping and Misiurewicz parameters, as well as parameters with an indifferent periodic orbit, are structurally unstable. □

5.1.5 Lemma (Parameters near an unstable parameter)
Suppose that $\kappa_0 \in \mathcal{B}$. *Then there are Misiurewicz, attracting and escaping parameters arbitrarily close to* κ_0.

Proof. Let a_0 be a repelling periodic point of E_{κ_0}, say of period $n \geq 3$. Let U be a (small) neighborhood of κ_0. By the implicit mapping theorem, there exists a holomorphic function $\kappa \mapsto a(\kappa)$ with $a(\kappa_0) = a_0$ and $E_\kappa^n(a(\kappa)) = a(\kappa)$. By choosing U sufficiently small, we may assume that a is defined on all of U. By Theorem 5.1.3 and Montel's theorem, there exist m and $\kappa \in U$ such that $E_\kappa^m(\kappa) \in \{a(\kappa), E_\kappa(a(\kappa)), E_\kappa^2(a(\kappa))\}$. Thus κ is a Misiurewicz point.

Furthermore, by Theorem 5.1.3, indifferent parameters are dense in the bifurcation locus. Every indifferent parameter can be perturbed so that the indifferent orbit becomes attracting.

Finally, consider any open set U which contains no escaping parameters. Then for every $(\underline{s}, t) \in X$, the point $G^\kappa(\underline{s}, t)$ depends holomorphically on κ in U. In other words, $I(E_\kappa)$ moves holomorphically, and by the λ-lemma, so does $J(E_\kappa) = \overline{I(E_\kappa)}$. ■

Let us take this opportunity to remark on the fact that the landing theorem 3.9.1 fails on a very large set of parameters. Indeed, the set of parameters whose postsingular set is all of $\mathbb{C}$ is generic in $\mathcal{B}$, as is the set of points for which all periodic external rays land.

5.1.6 Theorem (Genericity of $\mathcal{P} = \mathbb{C}$)
The following two conditions are both generic in the bifurcation locus of exponential maps:

1. *All periodic external rays land at repelling periodic points.*
2. $\mathcal{P} = \mathbb{C}$.

Proof. Fix any periodic address $\underline{s}$ of period n, and consider the set $V_{\underline{s}}$ of parameters where the ray $g_{\underline{s}}$ lands at a repelling periodic point. This set is open by a routine perturbation argument (well known from the polynomial setting [38, Lemma B.1]), which we sketch in the following. Suppose that for some parameter κ_0, $g_{\underline{s}}$ lands at a repelling periodic point w_0. Choose a linearizing neighborhood U of w_0, and some $t, t' > 0$ such that $E_\kappa^n(g_{\underline{s}}(t)) = g_{\underline{s}}(t')$ and $g_{\underline{s}}\big([t,t']\big) \subset U$ (in the dynamical plane of E_{κ_0}). For κ close to κ_0, U is still a linearizing neighborhood for a repelling periodic point $w(\kappa)$ of E_κ. Furthermore, $g_{\underline{s}}\big([t,t']\big)$ depends continuously on κ, so $g_{\underline{s}}\big([t,t']\big) \subset U$ for κ close to κ_0. Since the backward iterates $E_\kappa^{-kp} : U \to U$ converge uniformly to $w(\kappa)$, it follows that the ray $g_{\underline{s}}$ lands at $w(\kappa)$. Thus $\kappa \in V_{\underline{s}}$.

For any Misiurewicz parameter, all periodic rays land; thus $V_{\underline{s}} \cap \mathcal{B}$ is dense in $\mathcal{B}$. Therefore the intersection of all $V_{\underline{s}}$ is generic in $\mathcal{B}$.

Let $U \subset \mathbb{C}$ be open. The set

$$V(U) := \{\kappa \in \mathbb{C} : \exists n : E_\kappa^n(\kappa) \in U\}$$

is open. We will show that $V(U) \cap \mathcal{B}$ is dense in the bifurcation locus. Let $\kappa_0 \in \mathcal{B}$; we may assume that $J(\kappa_0) = \mathbb{C}$. Choose any two repelling periodic points of E_{κ_0} in U. These points move holomorphically (and stay in U) under a small perturbation of κ_0. Thus, by Montel's theorem, there exists κ arbitrarily close to κ_0 such that, for some n, $E_\kappa^n(\kappa)$ is one of these periodic points. Thus $\kappa \in V(U)$; furthermore $\kappa \in \mathcal{B}$ because κ is a Misiurewicz parameter.

Thus $V(U) \cap \mathcal{B}$ is open and dense in $\mathcal{B}$ for every U. Consequently,

$$\{\kappa \in \mathcal{B} : \mathcal{P}(\kappa) = \mathbb{C}\} = \mathcal{B} \cap \bigcap_{\substack{p,q \in \mathbb{Q} \\ \varepsilon \in \mathbb{Q}^+}} V(\mathbb{D}_\varepsilon(p + iq)$$

is generic in $\mathcal{B}$. ■

5.2 Hyperbolic Components

5.2.1 Definition (Stable Components)
Suppose W is a component of $\mathbb{C} \setminus \mathcal{B}$. Then W is called hyperbolic *if all parameters in W are attracting. Otherwise W is called* nonhyperbolic *or* queer.

It is conjectured that queer components do not exist. Note that all parameters within the same hyperbolic component W have the same period, because the attracting orbit moves holomorphically throughout W. We can say even more: the combinatorics of all parameters within the same component are the same.

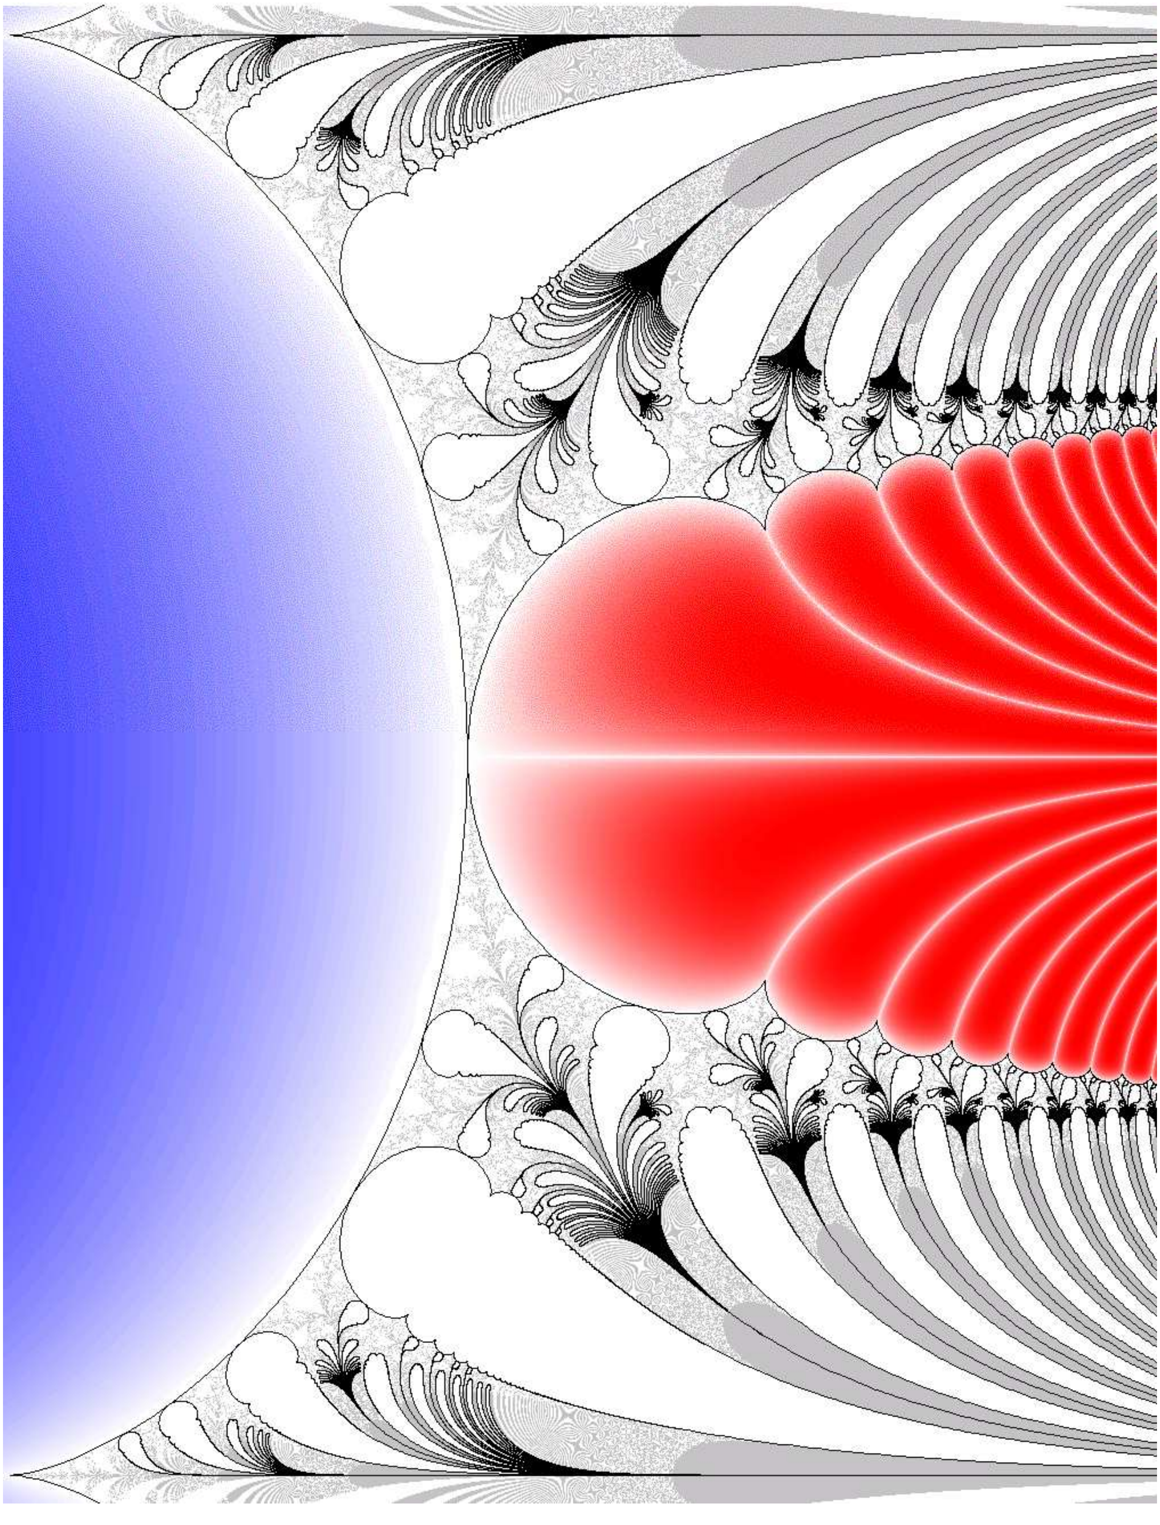

Figure 5.1: Several hyperbolic components in the strip $\{\operatorname{Im} \kappa \in [0, 2\pi]\}$. Within the period two component $W_{\frac{1}{2}\infty}$, parameters with positive real multiplier are colored white.

5.2.2 Definition and Lemma (External Addresses of Hyperbolic Components)
Let W be a hyperbolic component. Then all parameters $\kappa \in W$ have the same intermediate external address $\operatorname{addr}(\kappa)$. *This address is called the* intermediate external address of W *and denoted by* $\operatorname{addr}(W)$.

Proof. The attracting dynamic ray which connects the periodic point a_1 to κ moves continuously throughout W because the linearizing coordinates on the attractive basin depend holomorphically on the parameter. The branch of E_κ^{-n} which takes this curve back into U_1 also depends continuously, and thus the external address of this pullback is constant.

(A different proof can be furnished as follows: If $\kappa, \kappa' \in W$, then the two maps are topologically conjugate on their Julia sets, see Theorem 5.2.3 below. In particular, they have the same characteristic rays, which implies by Lemma 4.5.12 that they have the same intermediate external address.) ∎

Similarly, the kneading sequences, characteristic external addresses etc. only depend on W. Correspondingly, we will also denote these objects by $\mathbb{K}(W)$, $\mathbb{K}^*(W)$ etc.

Hyperbolic components of exponential maps were first considered in [5] and [31]. It was proved in [31] that, for every hyperbolic component, the multiplier map is a universal covering map.

5.2.3 Theorem (Multiplier Map [31])
Let W be a hyperbolic component. Then W is simply connected and the map $\mu : W \to \mathbb{D}^*$ which maps each parameter to the multiplier of its unique attracting cycle is a universal covering. Furthermore, if $\kappa, \kappa' \in W$, then E_κ and $E_{\kappa'}$ are quasiconformally conjugate. □

Note that hyperbolic components in exponential parameter space do not have "centers" because the multiplier can never be 0. In some sense, their center is at ∞:

5.2.4 Lemma (Small Multipliers)
Let W be a hyperbolic component of period n. If (κ_i) is a sequence of parameters in a hyperbolic component W such that $\mu(\kappa_i) \to 0$, then $\kappa_i \to \infty$. In particular, every hyperbolic component is unbounded.

Proof. Let $\kappa \in W$ and let $a_1, \dots, a_n$ denote its attracting orbit as usual. Recall that

$$\mu(\kappa_i) = \prod_{j=1}^{n} \exp(a_j). \tag{5.1}$$

Since $|\mu(\kappa_i)| \leq 1$, there is some j such that $\operatorname{Re} a_j \leq 0$. This implies that $|a_{j+1}| < |\kappa| + 1$. Denote $f(t) = \exp(t) + |\kappa|$, then it follows that

$$|a_k| < f^{n-1}(|\kappa| + 1)$$

for all k. Thus the product (5.1) is bounded from below in terms of $|\kappa|$, which concludes the proof. ∎

Hyperbolic components are completely classified by the following theorem, which was the main result of [74]. (A proof of the existence part of this theorem has also been published in [19]).

5.2.5 Theorem and Definition (Classification of Hyperbolic Components [74]) *For every intermediate external address $\underline{s}$, there exists exactly one hyperbolic component W with $\operatorname{addr}(W) = \underline{s}$, which we denote by $W_{\underline{s}}$. The vertical order of hyperbolic components coincides with the lexicographic order of their external addresses.*

REMARK. To explain the last statement of the theorem, note that for any hyperbolic component W and any curve $\gamma : [0, \infty) \to W$ such that $\lim \mu(\gamma(t)) = 0$, Lemma 5.2.4 shows that $|\gamma(t)| \to \infty$. By Theorem 5.2.3 there is a unique homotopy class which contains all such curves. We shall refer to this as the *preferred homotopy class of* W. As in Section 3.7, these homotopy classes have a vertical order, and this is the order referred to in the theorem. Note, however, that we cannot yet exclude the possibility that there are other homotopy classes of curves to ∞ within hyperbolic components which could have a different vertical ordering.

SKETCH OF PROOF. The proof of the existence part of the statement is essentially the argument used in Section 5.4 to prove part of the Squeezing Lemma. (More precisely, the argument in Section 5.4 is an extension of the original argument in [74].)

To prove uniqueness, suppose that κ_1, κ_2 are two parameters with the same intermediate external address and the same positive real multiplier $\mu \in (0, 1)$. Suppose furthermore that for both κ_1 and κ_2 the dynamic root is the distinguished boundary fixed point; it is easy to show that every hyperbolic component contains a parameter of this kind (see Section 5.5). Under these conditions, Schleicher uses quasiconformal methods to show that κ_1 and κ_2 are conformally conjugate.

The essence of the argument is as follows. In a neighborhood of the attracting periodic orbits, the maps E_{κ_1} and E_{κ_2} are conformally conjugate. This conformal conjugacy can be extended to a quasiconformal homeomorphism which sends the curves on the orbit of the principal attracting rays to curves in the corresponding homotopy class for E_{κ_2}. By propagating this map using the dynamics of E_{κ_1} and E_{κ_2}, one then obtains a sequence of quasiconformal homeomorphisms with uniformly bounded dilatation which are conformal conjugacies on larger and larger subsets of the Fatou set. In the limit, one obtains a quasiconformal conjugacy between the two maps which is conformal on the Fatou set. Because the Julia set has measure zero, this conjugacy is conformal on the plane. □

Let us also note the following elementary facts about hyperbolic components of periods one and two [5]; see Figure 5.1.

5.2.6 Lemma (Hyperbolic Components of Period 1 and 2) *There is a unique component W_∞ of period 1, which contains a left half plane. W_∞ is the biholomorphic image of the left half plane $\mathbb{H} = \{\operatorname{Re} z < 0\}$ under the map*

$$\mathbb{H} \to \mathbb{C}; \mu \mapsto \mu - \exp(\mu).$$

For every $k \in \mathbb{Z} + \frac{1}{2}$, the period 2 component $W_{k\infty}$ asymptotically contains the strip

$$\left\{\kappa : \operatorname{Im}\kappa \in \left(2\pi(k - \frac{1}{4}), 2\pi(k + \frac{1}{4})\right)\right\}.$$

□

5.3 Boundaries of Hyperbolic Components

We now turn our attention to the boundaries of hyperbolic components.

5.3.1 Lemma (Indifferent Parameters)
Let W be a hyperbolic component of period n and $\kappa_0 \in \partial W$. Then κ_0 has an indifferent cycle of period dividing n. Furthermore, as $\kappa \to \kappa_0$ in $\overline{W}$, the nonrepelling cycles of E_κ converge to the indifferent cycle of E_{κ_0}.

Proof. Let $\kappa_i \to \kappa_0$ in $\overline{W}$. Since the multiplier of the nonrepelling orbit of κ_i is the product of E'_{κ_i} along the cycle, there exists at least one point z_i on this orbit such that $|\exp(z_i)| = |E'_{\kappa_i}(z_i)| \leq 1$. Because the κ_i stay bounded, the sequence $E_{\kappa_i}(z_i) = \exp(z_i) + \kappa_i$ is bounded and thus has a limit point z, which is a nonrepelling fixed point of $E^n_{\kappa_0}$. Since κ_0 has at most one nonrepelling orbit, and since z can clearly not be attracting because $\kappa_0 \in \partial W$, the claim follows. ■

Similarly to the definition of external (and internal) rays for polynomials, the foliation of the punctured disk by radial rays gives rise to a foliation of the hyperbolic component by *internal rays*. These rays are of a natural interest when studying the boundary of hyperbolic components. In fact it is not difficult to show — and we will do so in this section — that the question whether the boundary of a hyperbolic component is connected reduces to that whether an internal ray can land at ∞.

5.3.2 Definition (Internal Rays)
Let W be a hyperbolic component, and let $\mu : W \to \mathbb{D}^*$ be the multiplier map. If $\theta \in \mathbb{R}/\mathbb{Z}$, then any component of

$$\mu^{-1}\left(\{re^{2\pi i\theta} : r \in (0,1)\}\right)$$

is called an internal ray *of angle θ. Internal rays are usually parametrized as $\gamma : (0,1) \to \mathbb{C}$ such that $\mu(\gamma(t)) = te^{2\pi i\theta}$. We say that an internal ray γ* lands *at a point z if $z = \lim_{t\to 1}\gamma(t)$.*

Remark. By Lemma 5.2.4, $\lim_{t\to 0}\gamma(t) = \infty$.

5.3.3 Lemma (Boundary is Locally Connected)
Let W be a hyperbolic component. Then $\partial W \subset \hat{\mathbb{C}}$ is locally connected. In particular, every internal ray of W lands in $\hat{\mathbb{C}}$.

Proof. Let $\mathbb{H}$ denote the left half plane as usual, and let $\Phi : H \to W$ be a conformal isomorphism such that $\mu \circ \Phi = \exp$. Such a map exists because both exp and μ are universal covers of $\mathbb{D}^*$.

We will show that Φ extends continuously to a map $\overline{\mathbb{H}} \cup \{\infty\} \to \hat{\mathbb{C}}$. Let $\theta \in \mathbb{R}$ and let $L \subset \hat{\mathbb{C}}$ denote the set of limit points of $\Phi(z)$ as $z \to i\theta$. We must show that L consists of a single point. Let $\kappa_0 \in L \cap \mathbb{C}$; then by Lemma 5.3.1, E_{κ_0} has an indifferent periodic point a with multiplier $\mu = \exp(i\theta)$. Note that the set of such parameters κ is discrete. Indeed, suppose there is a sequence $\kappa_n \to \kappa_0$ of such parameters. By the argument of Lemma 5.3.1, the indifferent orbits of the κ_n orbits accumulate at the indifferent orbit of κ_0. Pick an indifferent periodic point a_n of E_{κ_n} such that the a_n converge to a. Then the points (κ_n, a_n) lie in the zero set of the function of two complex variables given by

$$f(\kappa, a) := (E^n_\kappa(a) - a, E^n_\kappa)'(a) - \mu).$$

However, because f is analytic, this implies that for all κ in a neighborhood of κ_0 there is a solution of $f(\kappa, z) = 0$. In other words, all parameters in a neighborhood of κ are indifferent, which is absurd. Thus the limit set L is connected, but is contained in a discrete set. It follows that $|L| = 1$, as required.

Finally, let us show that $\Phi(z)$ has no accumulation points in $\mathbb{C}$ as $z \to \infty$. Suppose then that $\gamma \subset \mathbb{H}$ is a curve to ∞ with a finite accumulation point κ_0. Then by Lemma 5.3.1, κ_0 has an indifferent orbit of multiplier, say, $e^{2\pi i\theta}$. We can continue the multiplier of this orbit to an analytic function on a finite sheeted cover of a neighborhood U of κ_0. In particular, there are finitely many connected subsets of U in which the orbit becomes attracting, which means that (if U was chosen small enough) the set $\Phi^{-1}(U)$ consists of finitely many bounded components. Thus $\gamma(t) \not\to \infty$. ■

5.3.4 Corollary (Landing Points of Internal Rays)
Every boundary point of W in $\mathbb{C}$ is the landing point of a unique internal ray; in particular every component of $\partial W \cap \mathbb{C}$ is a Jordan arc extending to ∞ in both directions.

Proof. The fact that every boundary point is the landing point of an internal ray follows immediately from Lemma 5.3.3. We thus need to show that no two internal rays can land at the same boundary point κ_0. Let $\Phi : \mathbb{H} \to W$ be a conformal isomorphism as before, and suppose that $\Phi(i\theta) = \Phi(i\theta') =: \kappa_0$ for some $\theta < \theta'$. (Recall that by Lemma 5.3.3, Φ extends continuously to the boundary.)

Connect $i\theta$ and $i\theta'$ by a curve in $\mathbb{H}$; the image of this curve under Φ is then a simple closed curve γ in $\overline{W}$ which intersects ∂W only in κ_0. By the F. and M. Riesz theorem [56, Theorem A.3], there exists some parameter $\kappa_1 \in \Phi(i[\theta, \theta']) \setminus \{\kappa_0\}$. The indifferent parameter κ_1 then lies in the bifurcation locus $\mathcal{B}$ and is separated from ∞ by the curve γ. Because Misiurewicz parameters are dense in $\mathcal{B}$, there exists some Misiurewicz parameter κ_2 which is also enclosed by γ. By Lemma 5.1.5, we can finally find some attracting parameter κ with $|\kappa_2 - \kappa| < \operatorname{dist}(\kappa, \partial W)$. It follows that the hyperbolic component containing κ is separated from ∞ by γ, which contradicts Lemma 5.2.4. ■

The only question that remains is thus whether an internal ray can land at ∞, disconnecting the boundary of W in $\mathbb{C}$. In order to study this question, we will return to combinatorial considerations. Recall that hyperbolic components have a natural vertical order, given by the preferred homotopy classes; i.e. the homotopy classes of curves along which the multiplier tends to 0. If $\gamma : [0, \infty) \to \mathbb{C}$ is a curve to ∞ which does not contain indifferent parameters, then, as in Section 3.7, we can associate to it an address

$$\operatorname{addr}(\gamma) := \inf\{\underline{s} : W_{\underline{s}} \text{ is above } \gamma\}.$$

The following fundamental result, which we call the "Squeezing Lemma", shows that there are only two possibilities for such a curve.

5.3.5 Theorem (Squeezing Lemma)
Let $\gamma : [0, \infty) \to \mathbb{C}$ be a curve in parameter space with $|\gamma(t)| \to \infty$. Suppose that γ contains no indifferent parameters. Then

(a) $\underline{s} := \operatorname{addr}(\gamma)$ is either intermediate or exponentially bounded.

(b) If $\underline{s}$ is exponentially bounded, then $\gamma(t)$ is escaping for large t. More precisely, $\gamma(t) = g_{\underline{s}}^{\gamma(t)}(r)$, where $r \to \infty$ as $t \to \infty$.

(c) If $\underline{s}$ is intermediate, then γ is contained in $W_{\underline{s}}$ and is homotopic to a curve in $W_{\underline{s}}$ along which the multiplier tends to 0.

Remark. Part (b) of the Squeezing Lemma is essentially [73, Proposition II.9.2], although our proof is different. (In fact, [73, Proposition II.9.2] was restricted to bounded external addresses, but this restriction is not essential.) Also, part (a) and (c) of the theorem were proved in [73] for parameter rays and internal rays of hyperbolic components [73, Proposition V.6.1], which suffices to prove Theorem 1.8. Proofs of these facts using simplified proofs were contained in [68]. The general form of the squeezing lemma as above is new.

Part (b) of the Squeezing Lemma can be proved in an elementary way, which will be done in Section 5.4. Parts (c) and (a) require a more intimate understanding of the bifurcation structure of hyperbolic components and will be proved as Theorem 5.8.8 and 5.10.1, respectively.

With the help of the Squeezing Lemma, let us now prove Theorem 1.8.

5.3.6 Theorem (Boundary of Hyperbolic Components [73, Proposition V.6.4])
Let W be a hyperbolic component. Then the boundary of W in $\mathbb{C}$ is a Jordan arc which tends to ∞ in both directions.

Proof. By Corollary 5.3.4, it remains only to show that no internal ray lands at ∞. So suppose by contradiction that γ is an internal ray which lands at ∞. Then by the Squeezing Lemma, both ends of γ lie in the same homotopy class. This is again impossible by the F. and M. Riesz theorem. ∎

As we have mentioned before (in Section 4.2), there are methods available for rational functions to construct parabolic parameters through transquasiconformal surgery. It seems likely that these methods can be transferred to the exponential family to show the landing of internal rays at rational angles in a more canonical way. For irrational angles, an analytical method to construct landing points by surgery is not available, and the difficulties in developing such a surgery, even for polynomials, are formidable. Another way to find a more analytical proof Theorem 5.3.6 would be to obtain estimates on internal rays in terms of the combinatorics of the component (and the *sector* of the internal ray, see Section 5.5). However, currently no such results are known.

Finally, let us use the Squeezing Lemma to deduce a fact about nonhyperbolic components. Compare also the discussion in Section 7.2.

5.3.7 Corollary (Nonaccessibility of ∞ in Queer Components)
Suppose that W is a nonhyperbolic stable component. Then ∞ is not accessible from W.

PROOF. W cannot contain attracting or escaping parameters; thus the claim follows directly from the Squeezing Lemma. ■

5.4 Proof of Part (b) of the Squeezing Lemma

5.4.1 Theorem (Squeezing Lemma, Part (b))
Let $\gamma : [0,\infty) \to \mathbb{C}$ be a curve in parameter space which does not contain any indifferent parameters. Suppose that $\gamma(t) \to \infty$ and that $\underline{s} := \operatorname{addr}(\gamma)$ is exponentially bounded. Then, if t is large enough, $\gamma(t)$ is escaping; in fact $\gamma(t) = g_{\underline{s}}^{\gamma(t)}(r)$ with $r \to \infty$ as $t \to \infty$.

This fact was originally proved — for bounded external addresses — by approximating the *parameter ray* at a given address (see Section 5.12) above and below by other parameter rays [73, Proposition II.9.2]. This is the origin of the name "Squeezing Lemma". Our proof will instead squeeze from above and below using hyperbolic components. It relies on the same ideas as the proof of the existence of hyperbolic components in [74]; the essence of the argument goes back as far as [5, Section 7] (which, however, did not use a combinatorial description to distinguish these components).

The first step is to obtain a condition which allows us to identify parameters in a given hyperbolic component. This is the content of the following lemma, which is a generalization of [74, Lemma 3.4].

5.4.2 Lemma (Identifying Hyperbolic Components)
Let $x \geq 2$ be arbitrary. Let $\kappa_0 \in \mathbb{C}$ with $\operatorname{Re}\kappa_0 > 2x+1$, and let $\kappa_0 =: z_1 \mapsto z_2 \mapsto \dots$ be the singular orbit of E_{κ_0}. Suppose there exists $n \in \mathbb{N}$ such that, for all $1 \leq k \leq n-1$, $\operatorname{Re}(z_k) > 0$ and $|\operatorname{Im}(z_k)| < F^{k-1}(x)$. Then

$$\operatorname{Re}(z_k) \geq F^{k-1}(\operatorname{Re}(\kappa_0) - 1). \tag{5.2}$$

Furthermore, consider the map $f : \kappa \mapsto E_\kappa^{n-2}(\kappa)$. Then

$$|f'(\kappa_0)| \geq F^{n-2}(\mathrm{Re}(\kappa_0) - 1). \tag{5.3}$$

If $\mathrm{Im}(z_{n-1}) \in \pi + 2\pi\mathbb{Z}$, then E_{κ_0} has an attracting periodic orbit of exact period n. The multiplier of this orbit tends to 0 as (for fixed x) either $\mathrm{Re}(\kappa_0)$ or n tend to ∞. Furthermore, $\mathrm{addr}(\kappa_0) = s_1 \dots s_{n-1}\infty$, where $\mathrm{Im}\, z_k \in \big((2s_k - 1)\pi, (2s_k + 1)\pi\big)$ for $k < n-1$ and $\mathrm{Im}\, z_{n-1} = 2s_{n-1}\pi$.

Proof. First note that, for $t \geq 2$, the following inequalities hold:

$$\exp(t) \geq 2t + 1 \tag{5.4}$$

$$\exp(2t - 1) \geq 2\exp(t) \tag{5.5}$$

$$\exp(2t) - 3\exp(t) \geq \exp(2t - 1). \tag{5.6}$$

We shall show by induction that

$$\mathrm{Re}(z_{k+1}) - 1 \geq F(\mathrm{Re}(z_k) - 1); \tag{5.7}$$

this proves (5.2). In the following, let us fix $t_0 := \frac{\mathrm{Re}(z_1)}{2} > x$.

Suppose that $k \geq 0$ and that (5.7) holds for all smaller values of k. Note that this implies by (5.5) that $\mathrm{Re}\, z_k \geq 2F^{k-1}(t_0)$.

Because $\mathrm{Re}(z_{k+1}) \geq |z_{k+1}| - F^k(x)$, we can estimate

$$\begin{aligned}
\mathrm{Re}(z_{k+1}) - 1 &\geq |z_{k+1}| - F^k(x) - 1 \\
&\geq \exp(\mathrm{Re}(z_k)) - 1 - |\kappa_0| - F^k(x) - 1 \\
&\geq F(\mathrm{Re}(z_k)) - \mathrm{Re}(\kappa_0) - |\mathrm{Im}(\kappa_0)| - F^k(x) \\
&\geq F(\mathrm{Re}(z_k)) - 2t_0 - x - F^k(x) \\
&\geq F(\mathrm{Re}(z_k)) - 3F^k(t_0) \geq F(\mathrm{Re}(z_k) - 1),
\end{aligned}$$

where the last inequality follows from (5.6).

Now consider the maps $E^k(\kappa) := E_\kappa^k(\kappa)$. Then

$$\frac{\partial E^{k+1}}{\partial \kappa}(\kappa_0) = \exp(z_{k+1}) \frac{\partial E^k}{\partial \kappa}(\kappa_0) + 1.$$

Therefore, by the above

$$\left| \frac{\partial E^{k+1}}{\partial \kappa}(\kappa_0) \right| \geq |\exp(z_{k+1})| - 1 \geq F^{k+1}(\mathrm{Re}(\kappa_0) - 1).$$

The fact that E_{κ_0} is attracting (with exact period n) was proved in [74, Lemma 3.4] under the weaker assumption that $\mathrm{Im}\, z_k$ is bounded. The proof remains essentially the same, so we shall only sketch it without working out the precise estimates. Using similar estimates as above, it is shown that the left half plane $\{z \in \mathbb{C} : \operatorname{Re} z < \operatorname{Re} z_n + 1\}$ is mapped by $E^n_{\kappa_0}$ into a very small circle around z_n (whose size goes to 0 as either $\operatorname{Re} z_1$ or n tend to ∞). This shows that E_{κ_0} has an attracting orbit of period n, and the statement about the multiplier follows by the Schwarz lemma.

To see that κ_0 has the correct external address, connect z_n to $-\infty$ by a curve γ_n at constant imaginary parts, and consider the pullbacks γ_k of this curve along the orbit (z_k). Now for $1 < k \leq n-1$, the curve γ_k contains no points at real parts $\leq \operatorname{Re} \kappa_0$. Indeed any point on γ_k maps, in $n-k$ iterations, to a point with real part less than $\operatorname{Re} \kappa_n < -F^{n-1}(\operatorname{Re} \kappa_0 - 1)$. It is easy to see that this is impossible for a point with real part less than $\operatorname{Re} \kappa_0$. Thus the curve γ_{k-1} is completely contained in the strip $\{\operatorname{Im} z \in \big((2s_{k-1} - 1)\pi, (2s_{k-1} + 1)\pi\big)\}$, which completes the proof. ■

Now that we have a way to recognize points in a given hyperbolic component, we wish to construct these. The next lemma shows that we can construct parameters satisfying the conditions of Lemma 5.4.2 for which the point z_{n-1} is at any prescribed position (subject to the obvious condition that the real part has to be sufficiently large).

5.4.3 Lemma (Parameters with Prescribed Singular Orbit)
Let $\underline{s}$ be any exponentially bounded address, say $2\pi(|s_k| + 1) < F^{k-1}(x)$ with $x \geq 2$. Set $R := 2x + 2$. Let $n \in \mathbb{N}$ be arbitrary and define $\underline{s}^+ := s_1 \dots (s_{n-1} + \frac{1}{2})\infty$ and $\underline{s}^- := s_1 \dots (s_{n-1} - \frac{1}{2})\infty$. Then there exists R' with $F^{-(n-2)}(R') = R + O(1)$ and a conformal map Φ from $K_{R'} := \{z : \operatorname{Re}(z) \geq R', \operatorname{Im}(z) \in [(2s_{n-1} - 1)\pi, (2s_{n-1} + 1)\pi]\}$ into parameter space such that

1. *For any $z \in K_{R'}$, $\Phi(z)$ satisfies the hypothesis of Lemma 5.4.2, and, for $1 \leq k < n-1$,*
$$\operatorname{Im}(z_k) \in \big((2s_k - 1)\pi, (2s_k + 1)\pi\big).$$
2. *$E^{n-2}_{\Phi(z)}(\Phi(z)) = z$.*
3. *$\operatorname{Re}\Big(\Phi\big(R' + (2s_{n-1} - 1)\pi\big)\Big) = R$.*

Remark. By Lemma 5.4.2, this proves the existence of hyperbolic components with an arbitrary intermediate external address.

Proof. For all κ, set $t_\kappa := \max\{2x+1, \log(2(|\kappa|+2))\}$. Let us define $\phi_k : K_{F^{n-2}(t_\kappa)} \to \mathbb{C}$ by $\phi_{n-1}(z) = z$ and $\phi_k(z) = L_{s_k}(\phi_{k+1}(z))$, where L_{s_k} is the branch of E_κ^{-1} that takes values in the strip $\operatorname{Im} z \in [(2s_k - 1)\pi, (2s_k + 1)\pi]$. A simple induction (quite similar to that of Lemma 3.3.1) shows that ϕ_k is defined for $1 \leq k \leq n-1$, and $|\operatorname{Re}(\phi_k(z)) - F^{k-n+1}(\operatorname{Re} z)| < 1$. Define $\phi_\kappa := \phi_1$.

Now let us return to the parameter plane. Note that, if κ is any parameter with $\operatorname{Re} \kappa = R$ and $\operatorname{Im} \kappa \in \big((2s_1 - 1)\pi, (2s_2 + 1)\pi\big)$, then $t_\kappa < R - 1$. Now move κ horizontally from $R + i(2s_1 - 1)\pi$ to $R + (2s_1 + 1)\pi$. Because κ starts out below the image of ϕ_κ but ends up above it, and because ϕ depends continuously on κ, there is some value of κ such that $\kappa = \phi_\kappa(r + 2s_{n-1}\pi)$ for some r. Set $R' := r$.

We claim that we can extend this solution of the implicit equation $\phi_\kappa(z) = \kappa$ to an analytic function of z in all of $K_{R'}$. Because $K_{R'}$ is simply connected, it is sufficient to show that this solution can be continued analytically along every curve $\gamma : [0,1] \to K_{R'}$ with $\gamma(0) = r + 2s_{n-1}\pi$. Let $I = [0,t_0]$ or $I = [0,t_0)$ be the maximum interval through which the solution can be continued. First note that any κ which solves this equation for some $z \in K_{R'}$ satisfies the conditions of Lemma 5.4.2. By the derivative estimate (5.3) and the inverse function theorem, we can extend every solution locally, so the set I is open.

It thus remains to show only that $t_0 \in I$. is closed. Let κ be a limit point of $\Phi(\gamma(t))$ as $t \to t_0$. Then, by continuity of ϕ, $\phi_\kappa(\gamma(t_0)) = \kappa$. Again, by (5.3), the set of such κ is discrete, and thus $\Phi(\gamma(t)) \to \kappa$ as $t \to t_0$. ■

Proof of Theorem 5.4.1. Set $\underline{s} := \operatorname{addr}(\gamma)$, and let $x \geq 2$ be such that $2\pi(|s_k| + 1) \leq F^{k-1}(x)$. For every n, let Φ_n and R'_n be as in Lemma 5.4.3, and let

$$\underline{s}^{\pm} := s_1 s_2 \dots s_{n-2}(s_{n-1} \pm \frac{1}{2})\infty.$$

By Lemma 5.4.2, $\Phi_n(z) \in W_{\underline{s}^{\pm}}$ whenever $\operatorname{Im} z = (2s_{n-1} \pm 1)\pi$. Note also that, by the estimate (5.3),

$$\operatorname{diam} \Phi\Big(\big\{R'_n + bi : b \in [(2s_{n-1} - 1)\pi, (2s_{n-1} + 1)\pi]\big\}\Big) \leq \frac{2\pi}{F^{n-2}(2x)}.$$

Recall that γ tends to ∞ between the components $W_{\underline{s}^-}$ and $W_{\underline{s}^+}$, and thus between the curves $\{\Phi(z) : \operatorname{Im} z = (2s_{n-1} - 1)\pi\}$ and $\{\Phi(z) : \operatorname{Im} z = (2s_{n-1} + 1)\pi\}$. It follows that, for large t,

$$\gamma(t) \in \bigcap_n \Phi_n(K_{R'_n}).$$

By (5.2) and the construction of Φ, this means that $\gamma(t)$ is escaping and has external address $\underline{s}$. It follows easily from the estimates of Lemma 5.4.2 that the potential of $\gamma(t)$ goes to ∞ as $t \to \infty$. ■

In fact, this proof actually *constructs* a curve with external address $\underline{s}$ as the intersection of the images of the Φ_n. This is a rather curious way of constructing parameter rays, which we will record here for use in Section 5.12.

5.4.4 Corollary (Parameter Ray Tails)

Let $\underline{s}$ be any exponentially bounded external address. Then, for some $t_0 > 0$, there exists a curve

$$H_{\underline{s}} : [t_0, \infty) \to \mathbb{C}$$

with $\operatorname{addr}(H_{\underline{s}}) = \underline{s}$ such that $g_{\underline{s}}^{H_{\underline{s}}(t)}(t) = H_{\underline{s}}(t)$ for all $t \geq t_0$. For each $t \geq t_0$, $H_{\underline{s}}(t)$ is a simple zero of the function

$$\kappa \mapsto g_{\underline{s}}^{\kappa}(t) - \kappa.$$

Furthermore, if $\gamma : [0,\infty) \to \mathbb{C}$ is any curve which does not contain indifferent parameters and which has external address $\underline{s}$, then $\gamma(t) \in H_{\underline{s}}[t_0, \infty)$ whenever t is large enough.

Proof. It is easy to show that the curves

$$\gamma_n^+ : [F^{-(n-2)}(R'_n), \infty) \to \mathbb{C}, t \mapsto \Phi_n(F^{n-2}(t) + (2s_{n-1} + 1)\pi i)$$

and

$$\gamma_n^+ : [F^{-(n-2)}(R'_n), \infty) \to \mathbb{C}, t \mapsto \Phi_n(F^{n-2}(t) + (2s_{n-1} + 1)\pi i)$$

converge uniformly to a curve which, after reparametrization, has the required properties. To show that the given parameter is a simple solution, estimate the derivative $\frac{\partial}{\partial\kappa} g_{\underline{s}}^{\kappa}(t)$ using the estimates in Lemma 5.4.2. Details are left to the reader. ■

5.5 Sectors of Hyperbolic Components

Recall that the multiplier map of a hyperbolic component W is a universal covering. In particular, W can be decomposed into countably many fundamental domains, which we call *sectors*. The appearance of these sectors is well-known from the setting of Multibrot sets, where every hyperbolic component has $d - 1$ sectors (if d is the degree). In order to describe the combinatorics of components bifurcating from W, we have to study the structure of these sectors more carefully.

5.5.1 Definition (Sectors of a Hyperbolic Component)
Let W be a hyperbolic component, and let $\mu : W \to \mathbb{D}^*$ be the multiplier map. Then any component of $\mu^{-1}(\mathbb{D} \setminus [0, 1))$ is called a sector *of W.*

Remark. Every sector is bounded by two internal rays of angle 0 and is mapped biholomorphically to $\mathbb{D} \setminus [0, 1)$ by the multiplier map.

We will now describe how to distinguish the sectors of a hyperbolic component combinatorially.

5.5.2 Definition (Sector Numbers)
Let $W = W_{\underline{s}}$ be a hyperbolic component of period $n \geq 2$, and let $\kappa \in W$ be a parameter with $\mu := \mu(\kappa) \notin (0, 1)$. Let γ be the attracting dynamic ray connecting a_0 to $-\infty$, and let γ' be the component of $E_\kappa^{-n}(\gamma)$ which starts at a_0. Then $\operatorname{addr}(\gamma')$ is of the form $s_*\underline{s}$ with $s_* \in \mathbb{Z}$. The entry $s_* = s_*(\kappa)$ is called the sector number *of κ.*

The justification for this terminology is given by the following observation, see Figure 5.2(a).

5.5.3 Theorem (Labeling Map)
The map $\kappa \mapsto s_*(\kappa)$ is constant on sectors of W. Furthermore, when κ passes a sector boundary so that μ passes through $(0, 1)$ in positive orientation, then $s_*(\kappa)$ increases exactly by 1. In particular the induced map from sectors to indices in $\mathbb{Z}$ is bijective.

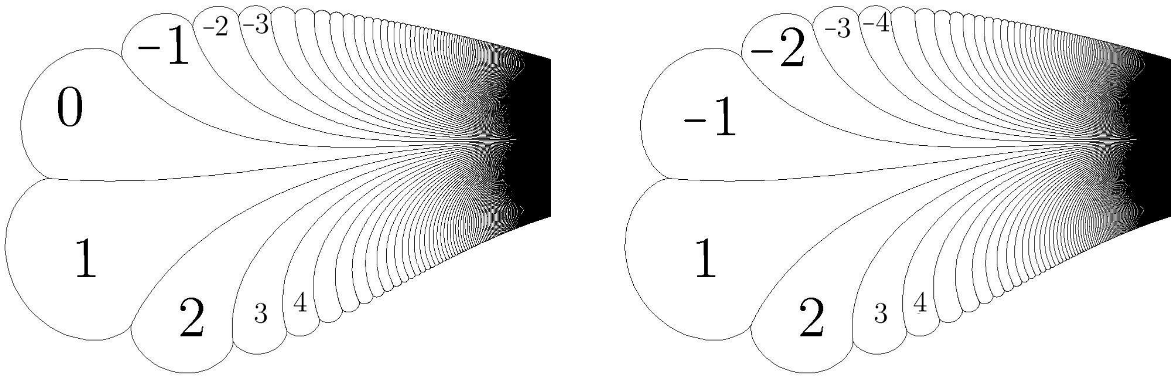

(a) Sector Labels (b) Kneading Sequences

Figure 5.2: Sectors of the period 5 component $W = W_{020\frac{1}{2}\infty}$, with $\mathbb{K}^*(W) = \overline{\texttt{01000}}$. The illustrations show the value s_* and the 5-th entry of the sector kneading sequence, respectively.

Proof. In the following, let γ^κ denote the principal attracting ray (which connects a_1 to ∞ with $\operatorname{addr}(\gamma^\kappa) = \operatorname{addr}(W)$).

As in the proof of Lemma 5.2.2, the branch of the pullback which takes γ^κ to a curve starting at a_0 varies continuously as long as γ^κ does not move through the singular value; i.e. as long as κ does not move through an internal ray at angle 0. Thus the entry s_* is constant on sectors.

Conversely, consider a parameter $\kappa_0 \in W$ with positive real multiplier μ. In the dynamical plane of E_{κ_0}, the singular value κ_0 lies on γ^{κ_0}, cutting it into two pieces γ_0, from a_1 to κ_0, and γ_1, from κ_0 to ∞. Let $L : U_1 \setminus \gamma_1 \to U_0$ be the branch of $E_{\kappa_0}^{-1} = \log(z - \kappa_0)$ that maps a_1 to a_0. The image of L is a strip bounded by two preimages of γ_1, with external addresses $s^*\underline{s}$ and $(s^*-1)\underline{s}$ for some $s^* \in \mathbb{Z}$ (Figure 5.3(a)). Note that $L(\gamma_0)$ is a broken attracting ray tending to $-\infty$.

Now change γ^κ in a neighborhood of κ_0 to a curve $\widetilde{\gamma}$ which avoids κ_0 below γ^κ (Figure 5.3(c)). Then evidently $\operatorname{addr}(L \circ \widetilde{\gamma}) = s^*\underline{s}$. Now, move κ so that $\mu(\kappa)$ becomes small and positive. We simultaneously move this curve continuously — using linearizing coordinates — and obtain a family of curves $\widetilde{\gamma}^\kappa$, which avoid κ and lie in the same homotopy class of $U_1(\kappa) \setminus \{\kappa\}$ as γ_κ. Then the pullback of $\widetilde{\gamma}^\kappa$ that starts at $a_0(\kappa)$ has external address $s^*\underline{s}$, so $s_*(\kappa) = s^*$. In the same way, one shows that $s_*(\kappa) = s^* - 1$ if $\mu(\kappa)$ becomes small and negative. ∎

Remark. The proof also shows how the sector boundaries can be distinguished dynamically. Consider a parameter with positive real multiplier. By Lemma 4.1.7 all unbroken attracting dynamic rays from a_1 must land at a common fixed point of the first return map of U_1. For parameters on the boundary between the sectors $\{\underline{s}_* = j\}$ and $\{\underline{s}_* = j+1\}$, we see that this fixed point must lie between the two preimages of γ_1 which have external addresses $j\underline{s}$ and $(j+1)\underline{s}$, so the distinguished fixed point has itinerary $\overline{\mathtt{u}_1\mathtt{u}_2\ldots\mathtt{u}_{n-1}j}$, where $\mathbb{K}(W)=\mathtt{u}_1\ldots\mathtt{u}_{n-1}*$. In particular, there is exactly one internal ray of angle 0 consisting of parameters for which the distinguished boundary fixed point is the dynamic root. This internal ray is called the *central ray* of the component W.

5.5.4 Definition (Parameter (co)-roots)
If the central internal ray of a hyperbolic component W lands, its landing point is called the root *of W. A landing point of any other internal ray of angle* 0 *is called a* co-root *of W.*

As we will see later, W can touch a component of smaller period only at its root point.

5.5.5 Definition (Kneading Sequences of Sectors)
Let W be a hyperbolic component of period $n\geq 2$ with $\mathbb{K}^*(W)=\overline{\mathtt{u}_1\ldots\mathtt{u}_n}$, and consider the sector of W numbered s_*.

Let $r_1<r_2$ be the n-th entries of the characteristic addresses of W. Let $\mathtt{m}\in\mathbb{Z}$ be defined as

$$\mathtt{m}:=\begin{cases}s_*-r_2 & \text{if } s_*\leq r_2\\ s_*-r_1 & \text{if } s_*\geq r_1\end{cases}.$$

Then the kneading sequence *of this sector is defined to be $\overline{\mathtt{u}_1\ldots\mathtt{u}_{n-1}(\mathtt{u}_n+\mathtt{m})}$.*

The n-th entry of the sector itinerary jumps from $\mathtt{u}_n-1$ to $\mathtt{u}_n+1$ (i.e. omits the forbidden kneading sequence $\mathbb{K}^*(W)$) when s_* passes the central internal ray. Compare Figure 5.2(b).

We will explain this definition without proofs in two different ways. Suppose that κ is a parameter in the sector with kneading sequence $\overline{\mathtt{u}_1\ldots\mathtt{u}_{n-1}\mathtt{m}}$. Then, by Theorem 4.2.7, the attracting dynamic ray γ from a_1 through κ either lands at ∞ (if the ray is broken) or at the landing point of some dynamic ray. In either case, we can associate to κ an external address $\underline{s}$; namely the address of γ itself in the former case, or the address of the corresponding external address in the latter. As it turns out, the kneading sequence of this address either equals the kneading sequence of the sector (if $\underline{s}$ is infinite) or is of the form

$$(\mathtt{u}_1\ldots\mathtt{u}_{n-1}\mathtt{m})^{km}\mathtt{u}_1\ldots\mathtt{u}_{n-1}*$$

(if $\underline{s}$ is intermediate). In the important case where $\underline{s}$ is intermediate (i.e., where the multiplier has rational angle), this fact will be proved in Section 5.7.

The second interpretation of the sector kneading sequence also uses the above address $\underline{s}$. However, now we will associate to $\underline{s}$ a kneading sequence based on a quite different partition of the sequence space, namely by the addresses of the period n-rays which land on ∂U_0. (Compare Section 4.5.) This process yields the same kneading sequence as above. This explains why the entry jumps from $\mathtt{u}_n-1$ to $\mathtt{u}_n+1$: The region with entry $\mathtt{u}_n$ is bounded by the $n-1$st images of the characteristic rays of E_κ, which land together on the boundary of U_0. Thus these rays, together with their landing point, separate this region from U_0, and the curve constructed above can never lie in this region.

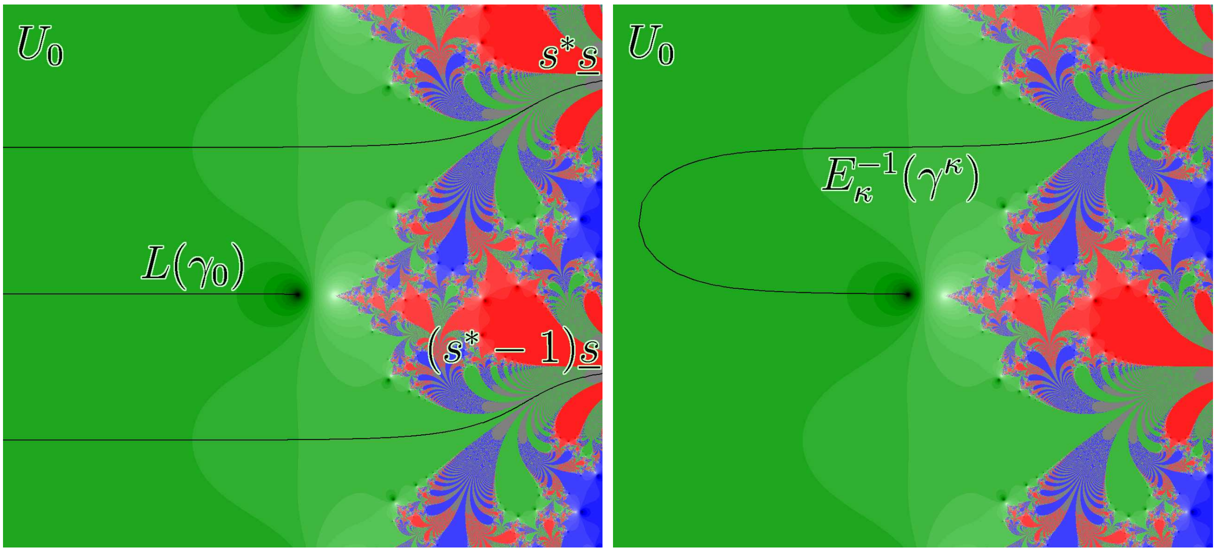

(a) The broken ray $L(\gamma_0)$ and the two pieces of $E_\kappa^{-1}(\gamma_1)$ which bound the image of L

(b) The pullback of γ^κ after a perturbation under which $\arg(\mu)$ becomes positive.

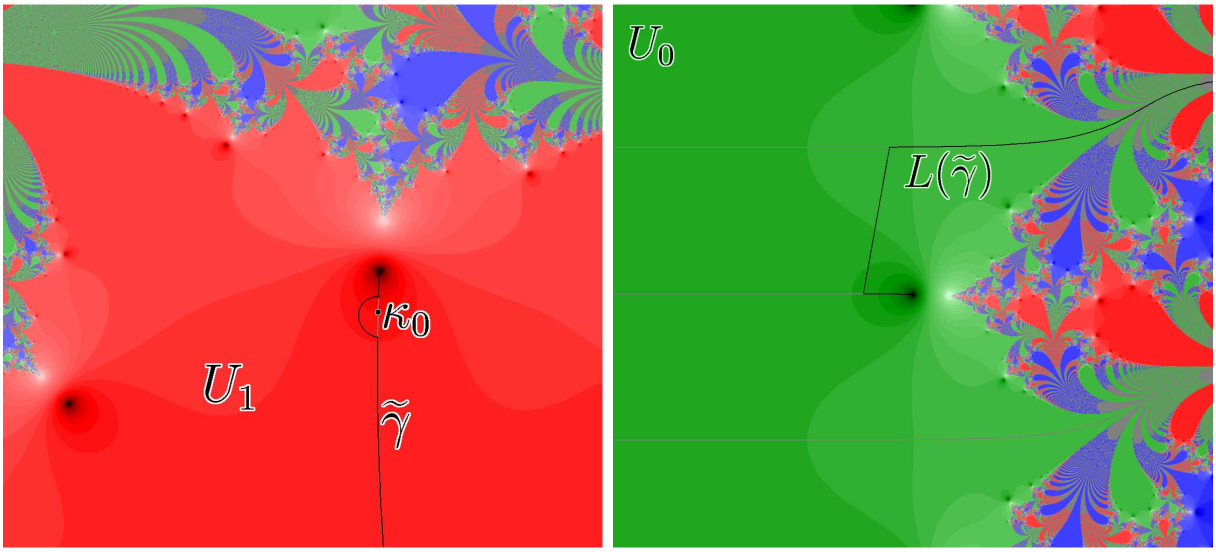

(c) The modified curve $\widetilde{\gamma}$ for the parameter κ_0

(d) The pullback $L(\widetilde{(\gamma)})$ has address $s^*\underline{s}$

Figure 5.3: Attracting dynamic rays, illustrating the proof of Theorem 5.5.3. (a) shows the broken attracting rays in U_0 for the parameter κ_0 with positive real multiplier, and (b) shows the situation after a small perturbation. (c) and (d) illustrate the construction of the perturbed curve $\widetilde{\gamma}$ in the dynamical plane of E_{κ_0}.

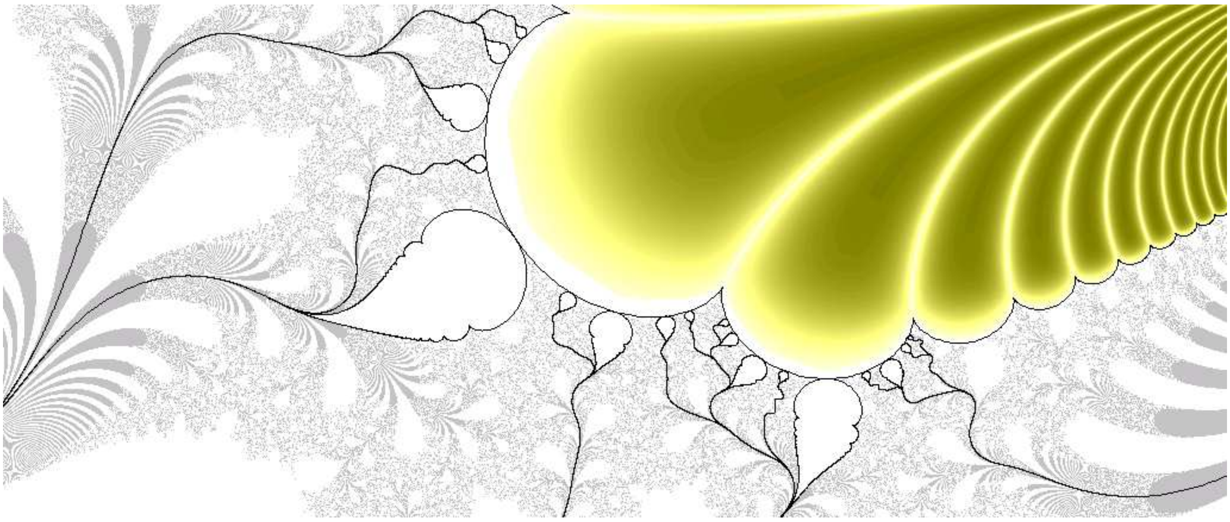

Figure 5.4: Some child components of a period 5 hyperbolic component.

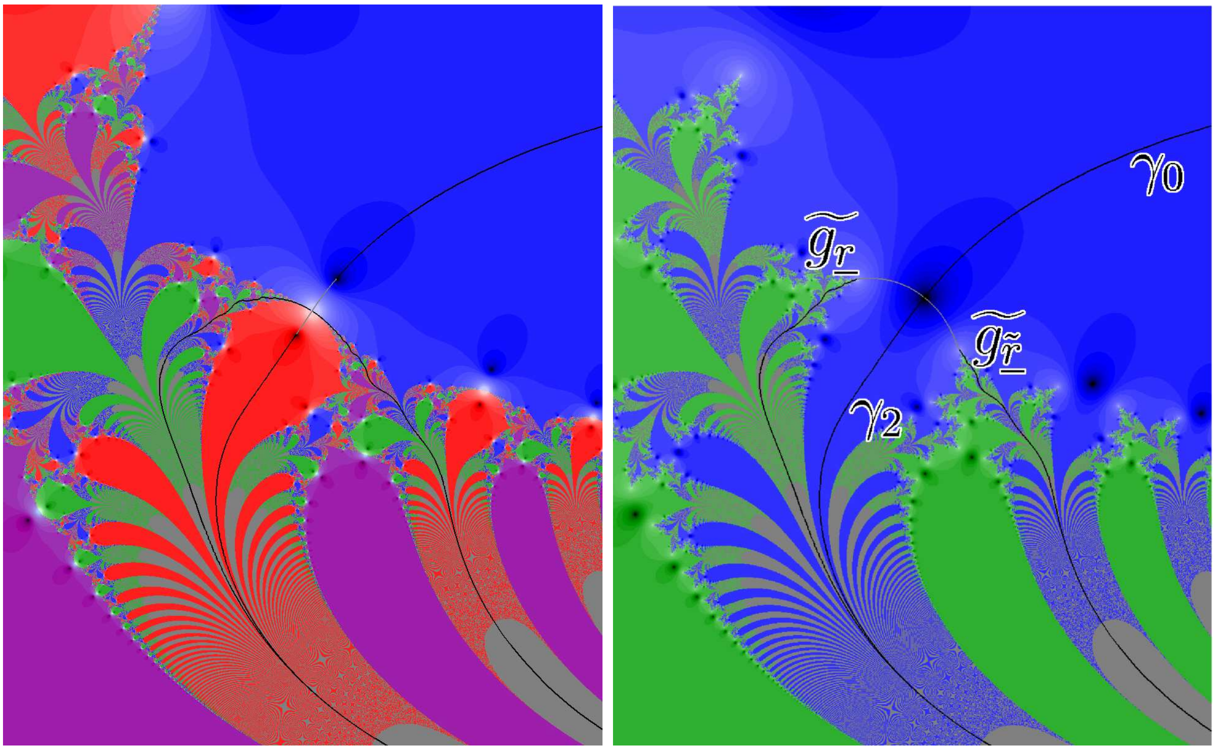

(a) Child Component

(b) Parent Component

Figure 5.5: Attracting rays in the child resp. parent component of a satellite parabolic point

To conclude this section, we describe the sector partition in the case of the period one component W_∞. Here the sector boundaries are all $2\pi k$-translates of the real interval $(-\infty, 1)$. The numbers $s_* \in \mathbb{Z} + \frac{1}{2}$ are defined analogously to the case $n \geq 2$. Theorem 5.5.3 trivially holds also for $n = 1$, since moving to another sector means relabeling the map. However, there is no distinguished sector boundary as before (just like there is no dynamic root point for a period-1 map). The itinerary of the sector $\{z \in \mathcal{W}_\infty : \mathrm{Im}(z) \in (2\mathtt{m}\pi, 2(\mathtt{m}+1)\pi)\}$ is defined to be $\overline{\mathtt{m}}$.

5.6 Bifurcation Structure at a Parabolic Point

In this section, we will describe, given a parabolic parameter κ_0, the combinatorics of the components which contain this parameter on its boundary. To this end, we will prove the following important bifurcation result. (Recall that the *ray period nq* of a parabolic parameter of period n is the period of the rays in its orbit portrait; equivalently, q is the number of repelling petals at each point of the parabolic orbit).

5.6.1 Theorem (Bifurcation at a Parabolic Parameter)
Let κ_0 be a parabolic parameter of period n, with ray period $qn > 1$ and intermediate external address $\underline{s}$. Let $\underline{\mathtt{u}} = \mathbb{K}^*(\kappa)$ and $\underline{\tilde{\mathtt{u}}} = \overline{\mathtt{u}_1\mathtt{u}_2\ldots\mathtt{u}_{n-1}\tilde{\mathtt{u}}_n}$ be the itinerary of the parabolic orbit.

Then κ_0 is a (co-)root of the hyperbolic component $W_{\underline{s}}$ (which is called the child component *of κ_0). It is the root point of $W_{\underline{s}}$ if and only if $\underline{\mathtt{u}} = \underline{\tilde{\mathtt{u}}}$, or equivalently if the dynamic root lies on the parabolic orbit. Furthermore:*

- ***Co-root and primitive cases.*** *If $q = 1$, then no other hyperbolic components touch $W_{\underline{s}}$ at κ_0.*
- ***Satellite case.*** *If $q \geq 2$, then κ_0 is the root of $W_{\underline{s}}$ and there is exactly one component $W_{\underline{s}'}$ (the* parent component*) touching $W_{\underline{s}}$ at κ_0. This component has period n and intermediate external address $\underline{s}' = \sigma^{(q-1)n}(\underline{s})$. κ_0 is the landing point of exactly one internal ray in $W_{\underline{s}'}$, which lies in the sector of $W_{\underline{s}'}$ with sector number $s_{(q-1)n}$.*

In order to prove Theorem 5.6.1, we will use some of the well-known properties of the analytic structure at a bifurcation. These are beautifully exposed, and proved using elementary complex analysis, in [58, Section 4].

5.6.2 Proposition (Perturbation of Parabolic Orbits)
Let κ_0 be a parabolic parameter of period n, with ray period qn.

- ***(Primitive Case)*** *If $q = 1$ (so the multiplier of the parabolic orbit is 1), then, under perturbation, the parabolic orbit splits up into two orbits of period n that can be defined as holomorphic functions of a two-sheeted cover around κ_0.*

  Any hyperbolic component touching κ_0 must correspond to one of these orbits becoming attracting (and therefore have period n).

- ***(Satellite Case)*** *If $q \geq 2$, then, under perturbation, the parabolic orbit splits into one orbit of period n and one of period nq. The period n orbit can be defined as a holomorphic function in a neighborhood of κ_0, as can the multiplier of the period nq-orbit. The nq-orbit itself can be defined on a q-sheeted covering around κ_0.*

  Any hyperbolic component touching κ_0 must correspond to one of these orbits becoming attracting (and therefore have period n or qn). □

Any hyperbolic component of period nq that touches κ_0 is called a *child component*; note that at least one such component always exists. In the satellite case, any period n component touching κ is called a *parent component.*

5.6.3 Proposition (Orbit Stability under Perturbation)
Under perturbation of a parabolic parameter into a child component, all periodic points retain the same orbit portraits.

Under perturbation into a parent component, the rays landing at the parabolic orbit are split up and land at the newly created repelling orbit. All repelling periodic points retain their orbit portraits. □

We are now ready to analyze the combinatorial behavior at a bifurcation. The following two theorems describe the change of combinatorics under perturbation into a child and parent component, respectively. These are the main information necessary to complete the proof of Theorem 5.6.1.

5.6.4 Theorem (Combinatorics in a Child Component)
Let κ_0 be a parabolic parameter of period n and ray period nq, and let W be a child component. Then $\operatorname{addr}(W) = \operatorname{addr}(\kappa_0)$.

Furthermore, for points on the internal ray of W landing at κ_0, the repelling point created in the bifurcation is the distinguished boundary fixed point. Therefore κ_0 is the root point of W if and only if its dynamic root lies on the parabolic orbit.

5.6.5 Theorem (Combinatorics in a Parent Component)
Let κ_0 be a satellite parabolic parameter of period n and multiplier $e^{2\pi i \frac{p}{q}}$, and let $W_{\underline{s}}$ be a parent component. If κ is a parameter on the $\frac{p}{q}$-internal ray in W which lands at κ_0, then the preferred boundary orbit for κ is the repelling period qn orbit created from the parabolic point.

Furthermore, for such a parameter κ, let $\underline{s}'$ be the address of the attracting dynamic ray that contains the singular value. Then $\underline{s}' = \operatorname{addr}(\kappa_0)$. In particular, $\underline{s} = \sigma^{(q-1)n}(\operatorname{addr}(\kappa_0))$, and the number s_* of the sector of W containing κ is given by the $[(q-1)n-1]$th entry of $\operatorname{addr}(\kappa_0)$.

Proof of Theorem 5.6.4. If $nq = 1$, then κ has period 1 and thus $W = W_\infty$ is the unique component of period 1. Now suppose that $nq > 1$. Choose a parameter $\kappa \in W$ on the internal ray landing at κ_0. The characteristic sector of a parabolic or hyperbolic parameter can be characterized as the minimal characteristic sector of all essential orbit

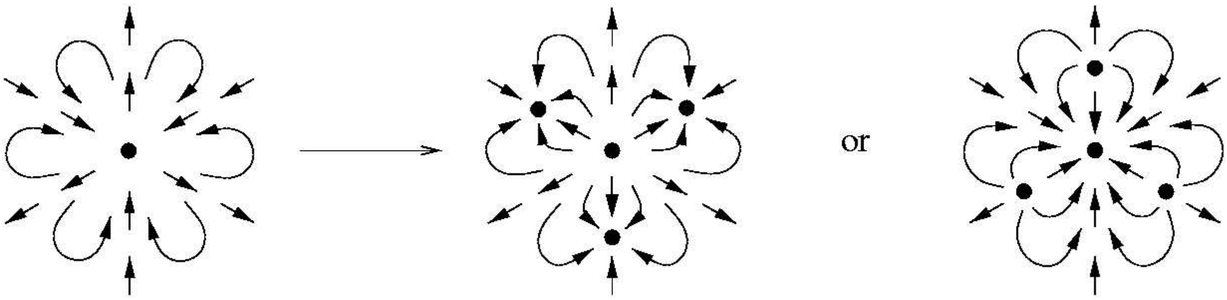

Figure 5.6: Local dynamics at a parabolic point and after deformation into a child resp. parent component. (This picture is courtesy of Saeed Zakeri.)

portraits. By Proposition 5.6.3, κ_0 and κ have the same orbit portraits and thus they have the same characteristic addresses. Lemma 4.5.12 then yields $\mathrm{addr}(\kappa_0) = \mathrm{addr}(\kappa)$.

Let us now show that the newly created repelling point w is the distinguished boundary fixed point for κ. First choose some small wedge-shaped attracting petals for the parabolic orbit of E_{κ_0}. If κ is close enough to κ_0 (on the internal ray landing at κ_0), then, for κ, the corresponding petals based at the orbit of w are still forward invariant and contain the newly created attracting period nq points (see [58, Section 4]).

Now suppose that w is not the distinguished boundary fixed point for κ. Let α be the piece of the principal attracting ray of E_κ which connects κ to ∞. Recall that all boundary fixed point differ in their nq-th itinerary entries, so their $nq-1$th images are separated from each other by the preimages of α. Because a_1 is connected to the distinguished boundary fixed point by an attracting dynamic ray and because attracting dynamic rays never intersect, this means that $E_\kappa^{n-1}(w)$ is separated from a_0 by a component α' of $E_\kappa^{-1}(\alpha)$. Since one of the attracting petals connects a_0 to $E_\kappa^{nq-1}(w)$, the curve α' must intersect this petal. Because the petals are forward invariant, it then follows that the piece $\gamma := E_\kappa^{nq}(\alpha)$ of the principal attracting ray between κ and $E_\kappa^{nq}(\kappa)$ intersects the petal at w which contains a_1.

Now let $\Phi : U_1 \to \mathbb{C}$ be the linearizing coordinate for E_κ, normalized so that $\Phi(\kappa) = 1$. Let $V \subset U_1$ be the component of the preimage of $\Phi^{-1}(\mathbb{D}(0, \frac{1}{\mu}))$ that contains a_0. Thus V contains the piece γ, and by definition, $\Phi(\gamma) = (\mu, 1)$. In the hyperbolic metric on $\mathbb{D}(0, \frac{1}{\mu})$, the length of $(\mu, 1)$ stays bounded as $\mu \to 1$. It follows that the hyperbolic length of this piece in V, and thus in U_1, is bounded as $\kappa \to \kappa_0$. By the standard estimates on the hyperbolic metric, this piece stays away from w (in the Euclidean metric). On the other hand, for κ close to κ_0, we can choose the attracting petal arbitrarily small. This is a contradiction. ■

Proof of Theorem 5.6.5. To prove the first statement, choose a cycle of repelling petals for the parabolic orbit of κ_0, which will contain the newly created repelling period nq point under a sufficiently small perturbation along an internal ray in the parent component. Choose any attracting dynamic ray which approaches a_1 through one of these petals. Recall that this ray is periodic under E_κ^n, of period q. Pulling back, it must land at the unique fixed point of the return map of this petal, which is on the repelling orbit created in the bifurcation.

Let $\underline{r}, \underline{\tilde{r}}$ be the characteristic addresses of κ_0. We will now find the combinatorial features of κ_0 within the attracting dynamics of κ, using attracting dynamic rays. (Compare Figure 5.5(b).) Let γ_0 be the attracting dynamic ray containing the singular value; so $\underline{s}' = \operatorname{addr}(\gamma_0)$. Then the attracting rays $\gamma_j := E_\kappa^{nj}(\gamma_0)$, $j = 0, \dots, q-1$ completely contain the singular orbit in U_1. By Proposition 5.6.3, the rays $g_{\underline{r}}$ and $g_{\underline{\tilde{r}}}$ do not land together for κ, but rather land separately on two points of the distinguished boundary cycle of κ. As we have seen in the first part of the proof, we can connect these two landing points to a_1 by two nonsingular attracting rays, staying within the chosen repelling petals. Let $\widetilde{g_{\underline{r}}}$ and $\widetilde{g_{\underline{\tilde{r}}}}$ denote the curves obtained by extending $g_{\underline{r}}$ and $g_{\underline{\tilde{r}}}$ to a_1 by these attracting rays.

Now consider the preimages of that part of γ_0 which connects κ to ∞; these are attracting ray pieces in U_0, connecting $-\infty$ to $+\infty$ with external addresses of the form $m\underline{s}'$. Since attracting dynamic rays never intersect, it follows that the images of $\widetilde{g_{\underline{r}}}$, $\widetilde{g_{\underline{\tilde{r}}}}$ and γ_0 always lie in the same strips of this partition.

Consequently, $\operatorname{itin}_{\underline{s}'}(\underline{r}) = \operatorname{itin}_{\underline{s}'}(\underline{\tilde{r}})$ is an itinerary (of period n) which agrees with $\mathbb{K}(\underline{s}')$ on its first $nq - 1$ entries. Therefore, $\underline{r}$ and $\underline{\tilde{r}}$ are characteristic addresses for $\underline{s}'$ (note that, $\underline{s}$ lies between $\underline{r}$ and $\underline{\tilde{r}}$ by Lemma 4.5.8; this also easily follows from the construction of $\widetilde{g_{\underline{r}}}$ and $\widetilde{g_{\underline{\tilde{r}}}}$). By uniqueness (Lemma 4.5.12), this implies $\underline{\tilde{s}} = \operatorname{addr}(\kappa_0)$. ■

Proof of Theorem 5.6.1. Let κ_0 be a parabolic parameter, and let $\underline{s}$ be its intermediate external address. By Theorem 5.6.4, every child component of κ_0 has intermediate external address $\underline{s}$. By the classification of hyperbolic components (Theorem 5.2.5), this means that κ_0 has exactly one child component, namely $W_{\underline{s}}$. Similarly, if κ_0 is a satellite parameter, then by Theorem 5.6.5, it has exactly one parent component, namely $W_{\sigma^{(q-1)n}(\underline{s})}$. Finally note that there are by definition at least q rays landing at each point of the parabolic orbit. Thus if $q \geq 2$, then by definition the dynamic root lies on the parabolic orbit. ■

5.7 Bifurcation from the parent component

In the previous section, we saw how to determine the combinatorics of both the parent and the child component of a parabolic parameter κ_0 when the combinatorics of this parameter are known. In this section, we will investigate how one determines the combinatorics of the child component (and thus of κ_0) given only the combinatorics of the parent component.

We will lead the discussion on a somewhat more combinatorial level. By Theorem 5.6.5, given a parameter on a rational internal ray, there is only one possible candidate for the component bifurcating from the landing point of this internal ray, irrelevant of whether this landing point exists or not. Let us first make this notion precise.

5.7.1 Definition (Combinatorial Bifurcation)
Let $W = W_{\underline{s}}$ be a hyperbolic component of period n, and let κ be a parameter in W whose multiplier has rational angle $\frac{p}{q}$ $(q > 1)$. Let A be the sector of W which contains κ, and

let $\underline{s}'$ be the external address of the (broken) attracting dynamic ray of E_κ which contains the singular value. Then we say that $W_{\underline{s}'}$ bifurcates combinatorially *from W (with angle θ and in sector A). We denote the address of this component by*

$$\operatorname{addr}(\underline{s}, A, \frac{p}{q}) := \operatorname{addr}(\underline{s}, s_*(A), \frac{p}{q}) := \operatorname{addr}(\underline{s}, \mathtt{m}, \frac{p}{q}) := \underline{s}'$$

where $\mathtt{m}$ is the n-th entry of $\mathbb{K}(A)$. The wake $\mathcal{W}(\underline{s}')$ of this component is called a subwake *of W.*

Remark. Note that $s_*(A)$ and $\mathtt{m}$ are, strictly speaking, both integers, so that the definition contains an ambiguity. However, in practice it will always be clear from the context which is meant.

5.7.2 Lemma (Bifurcation Formula for Kneading Sequences)
Let W be a hyperbolic component of period n, and let V, of period qn, bifurcate combinatorially from W. Let $\mathtt{u} = \overline{\mathtt{u}_1 \dots \mathtt{u}_{n-1}\mathtt{m}}$ be the kneading sequence of the sector of W from which V bifurcates. Then

$$\mathbb{K}(V) = (\mathtt{u}_1 \dots \mathtt{u}_{n-1}\mathtt{m})^{q-1}\mathtt{u}_1 \dots \mathtt{u}_{n-1} * .$$

Proof. Here and in the proofs throughout this section we will assume that $n > 1$; the simple modifications necessary for the case $n = 1$ are left to the reader. Let κ, $\underline{s}$ and $\underline{s}'$ be as in Definition 5.7.1. Also let $\mathtt{u}_n$ denote the n-th entry of the forbidden kneading sequence $\mathbb{K}^*(W) = \overline{\mathtt{u}_1 \dots \mathtt{u}_n}$. We shall require the following fact.

5.7.3 Lemma (Pullback Increases Addresses)
Let γ be any broken attracting dynamic ray of E_κ starting at a_1. If $\mathtt{m} < \mathtt{u}_n$, then γ lies above the principal attracting ray. Equivalently, $*\underline{s} \leq * \operatorname{addr}(\gamma) < * + \frac{1}{2}$ in the cylinder $\tilde{\mathcal{S}}$. Conversely, if $\mathtt{m} > \mathtt{u}_n$, then γ lies below the principal attracting ray.

Proof. Assume $\mathtt{m} < \mathtt{u}_n$; the other case is completely analogous. The curve γ can be obtained by pullbacks of the principal attracting ray along the attracting orbit (a_i). Therefore, we can prove the lemma inductively, taking further and further pullbacks. Assume then that γ is an attracting dynamic ray starting at a_1 and lies above (or is equal to) the principal attracting ray. Furthermore, assume that $\kappa \notin \gamma$ (otherwise, there are no further pullbacks).

Since $*\underline{s} \leq * \operatorname{addr}\gamma < * + \frac{1}{2}$ in the circular order of $\tilde{\mathcal{S}}$, the address of the unique pullback of γ that starts at a_0 must lie between $s_*\underline{s}$ and $s_* + \frac{1}{2}$. In particular, this pullback is below the n-th images of the characteristic external rays. Because the next $n-1$ pullbacks are univalent, the circular order of this configuration is preserved, and the claim follows. ■

Proof of Lemma 5.7.2, continued. Let γ be the attracting dynamic ray which contains κ (and has external address $\operatorname{addr}(\gamma) = \underline{s}'$). Because the first $nq - 1$ images of γ cannot intersect any preimages of γ, it is clear that

$$\underline{\tilde{\mathtt{u}}} := \mathbb{K}(\underline{s}') = (\mathtt{u}_1 \dots \mathtt{u}_{n-1}\tilde{\mathtt{m}})^{q-1}\mathtt{u}_1 \dots \mathtt{u}_{n-1}*$$

for some $\tilde{\mathtt{m}} \in \mathbb{Z}$. By the previous lemma, $k\underline{s} \le k\underline{s}'$ for every k. By the definition of $\mathtt{m}$, there are therefore exactly $\mathtt{m}$ addresses of the form $k\underline{s}'$ between $s_*\underline{s} = \sigma^{(q-1)n-1}(\underline{s}')$ and the n-th images of the characteristic external addresses. Thus $\tilde{\mathtt{m}} = \mathtt{m}$. ■

We now turn to the question of determining $\underline{s}'$ as well as its characteristic addresses. First note that $\underline{s}'$ can be easily computed:

5.7.4 Lemma (Bifurcation Algorithm for External Addresses)
Let $\underline{s}$ be an intermediate address. Then, for any s_* and $\frac{p}{q}$, the address

$$\underline{s}' = \mathrm{addr}(\underline{s}, s_*, \frac{p}{q})$$

can be determined from $\underline{s}$, s_* and $\frac{p}{q}$ by a simple algorithm.

Proof. Assume again that $n > 1$ and $\mathtt{m} < \mathtt{u}_n$. We know that $\sigma^{(q-1)n-1}(\underline{s}') = s_*\underline{s}$, so the last $n+1$ entries of $\underline{s}'$ are already known. We demonstrate how to determine $\underline{s}'_j$ from $\underline{s}'_{j+1}$ if $1 \le j < (q-1)n$. (Here $\underline{s}_j = s_j s_{j+1} \dots$ as usual.)

If j is not a multiple of n, the procedure is obvious. $\underline{s}'_j$ is the unique preimage of $\underline{s}'_{j+1}$ that, with respect to the partition $*\underline{s}$, lies in the same strip as $\underline{s}_{(k-1) \bmod n}$.

If j is a multiple of n, $\underline{\tilde{s}}_j$ must lie between $(s_*-1)\underline{s}$ and $(s_*+1)\underline{s}$, so there are two choices, depending on whether $\underline{\tilde{s}}_j$ falls above or below $s_*\underline{s}$. (In fact, we see from Lemma 5.7.3 that the two choices are $s_*\underline{\tilde{s}}_j$ or $(s_*-1)\underline{\tilde{s}}_j$.) Since σ^n permutes the addresses $\underline{\tilde{s}}_{kn}$ with rotation number $\alpha := \frac{p}{q}$, the information of which choice to make is coded in α as follows. For $k = 1, \dots q-2$ let i_k equal 1 if $k\alpha \pmod 1 \in [1-\alpha, 1]$, and 0 otherwise. Then $\underline{\tilde{s}}_{kn-1}$ falls above $s_*\underline{s}$ if and only if $i_k = 1$. ■

Recall that there are two addresses $\underline{r}^0 = \overline{r_1^0 r_2^0 \dots r_n^0}$ and $\underline{r}^1 = \overline{r_1^1 r_2^1 \dots r_n^1}$, such that $\underline{r}^0 < \underline{r}^1$ and $\mathrm{itin}_{\underline{s}}(\underline{r}^i) = \mathtt{u}_1 \dots \mathtt{u}_{n-1}(\mathtt{m}+i)$. These play a prominent role in Section 5.9, as *combinatorial sector boundaries*. For simplicity of notation, denote by r^0 and r^1 the n-tuples $r_1^0 \dots r_n^0$ and $r_1^1 \dots r_n^1$. We give a simple formula for $\underline{s}'$ and its wake, which is completely analogous to the corresponding formula for Mandel- and Multibrot sets. In fact, it is a special case of the general combinatorial tuning formula that we give later. The proof of the general case is not much different, but the setup is a bit more complicated. Since we will not require the general statement in the following, it has been relegated to Section 5.11.

5.7.5 Theorem (Subwake Formula)
Let $W_{\underline{s}'}$ bifurcate combinatorially from $W_{\underline{s}}$ at angle $\alpha = \frac{p}{q}$. Define

$$i_k := \begin{cases} 1 & k\alpha \pmod 1 \in [1-\alpha, 1] \\ 0 & \textit{otherwise} \end{cases}$$

and let $\underline{r}^0, \underline{r}^1$ as above. Then

$$\underline{s}' = r^{i_1} r^{i_2} \dots r^{i_{k-2}} r^q \underline{s}, \tag{5.8}$$

where q is 0 if $\mathtt{m} > \mathtt{u}_n$ and 1 if $\mathtt{m} < \mathtt{u}_n$. (If $n = 1$, then in the above formula r^q must be replaced by $(\mathtt{m} + \frac{1}{2})$.) Furthermore, the characteristic addresses of $\underline{s}'$ are

$$\overline{r^{i_1} r^{i_2} \dots r^{i_{k-2}} r^1 r^0} \text{ and}$$
$$\overline{r^{i_1} r^{i_2} \dots r^{i_{k-2}} r^0 r^1}.$$

Remark. The pairs $(\overline{i_1 \dots i_{q-2} 10}, \overline{i_1 \dots i_{q-2} 01})$ are exactly the external angles of the ray pairs landing at the main cardioid of the Mandelbrot set.

Proof. Suppose again for simplicity that $n > 1$ and $\mathtt{m} > \mathtt{u}_n$ (so that all considered addresses lie below $\underline{s}$). For any $\underline{t} \in \mathcal{W}(W)$, $\underline{t} \leq \underline{s}$, we have

$$\operatorname{itin}_{\underline{s}}(r^i \underline{t}) = \mathtt{u}_1 \dots \mathtt{u}_{n-1} (\mathtt{m} + i) \operatorname{itin}_{\underline{s}}(\underline{t}).$$

(To see this, note that $\underline{s}$ and its successive images never lie between $\underline{t}$ and $\underline{r}^i$, so the corresponding pullbacks will not be separated by $*\underline{s}$.)

From the proof of Lemma 5.7.4, we know that the right-hand side in (5.8) has the correct itinerary (with respect to $\underline{s}$), which defines $\underline{s}'$ uniquely. One easily checks that the itinerary for the characteristic addresses is also correct. Since an exponential map has no essential orbit portraits in its characteristic sector, these addresses are uniquely determined by their itinerary under $\underline{s}$. ■

5.7.6 Corollary (Subwakes Fill Wake)
Let W be a hyperbolic component of period n. Then the following hold.

(a) *Suppose $\underline{r} \in \mathcal{W}(W)$. Then there exist $\underline{s} < \underline{r} < \underline{s}'$ such that $W_{\underline{s}}$ and $W_{\underline{s}'}$ bifurcate combinatorially from W.*

(b) *Let $q > 1$. Then any component of period less than $(q-1)n + 2$ which lies in the wake of W must be contained in some $\frac{p}{q'}$ subwake of W, $q' \leq q$.*

(c) *Suppose $\underline{t} \in \mathcal{W}(W) \setminus \{\operatorname{addr}(W)\}$. Then $\underline{t}$ lies in the closure of some subwake of W, unless*
$$\mathbb{K}(\underline{t}) = \overline{\mathtt{u}_1 \mathtt{u}_2 \dots \mathtt{u}_{n-1} \mathtt{m}},$$
where $\overline{\mathtt{u}_1 \dots \mathtt{u}_n} = \mathbb{K}^*(W)$ and either $\mathtt{m} \in \mathbb{Z} \setminus \{\mathtt{u}_n\}$ or $\mathtt{m} = {}^{\mathtt{j}}_{\mathtt{j}-1}$, $\mathtt{j} \in \mathbb{Z}$.

Proof. To prove (b), suppose that V bifurcates combinatorially from W at angle $\frac{p}{q+1}$. Then by the subwake formula, the first $(q-1)n$ entries of the characteristic rays of V agree, so that $\operatorname{addr}(V)$ must be of length at least $(q-1)n+2$. The first statement follows similarly from the subwake formula, as we can choose the bifurcating addresses to lie arbitrarily close to the characteristic addresses of W.

Let us now prove the final claim. By (b), we know that the set of addresses in $\mathcal{W}(W)$ which do not lie in a subwake of W does not contain any intervals. If $\underline{t}$ is an address which does not lie in the closure of any subwake, we can thus find a sequence $W_{\underline{s}^n}$ of components which bifurcate combinatorially from W such that $\underline{s}^n \to \underline{t}$. The claim about $\mathbb{K}(\underline{t})$ then easily follows from Lemma 5.7.2. ■

To conclude this section, let us note that Theorem 5.7.5 gives a simple algorithm to compute the characteristic external addresses of $\underline{s}$.

5.7.7 Algorithm (Computing Characteristic Addresses)
Let $\underline{s}$ be an intermediate external address of length > 1 and let $\mathtt{u} := \mathbb{K}^*(\underline{s})$. Using Lemma 5.7.4, compute the addresses $\operatorname{addr}(\underline{s}, \mathtt{u}_n - 1, \frac{1}{2})$ and $\operatorname{addr}(\underline{s}, \mathtt{u}_n + 1, \frac{1}{2})$. The characteristic addresses of $\underline{s}$ are the periodic continuations of the first n entries of these addresses. □

Remark. This algorithm assumes that $\mathtt{u}_n$ is known. There are, however, only finitely many possibilities because $\mathtt{u}_n \in \{\mathtt{u}_1, \dots, \mathtt{u}_{n-1}\}$. A direct way to compute $\mathtt{u}_n$ from $\underline{s}$ will be shown in Section 5.9.

5.8 Proof of Part (c) of the Squeezing Lemma

In this section, we will prove part (c) of the Squeezing Lemma (Theorem 5.3.5). We also prove a variant of the squeezing lemma which allows to prove the landing in $\mathbb{C}$ of all internal rays of a hyperbolic component except its central internal ray. This already proves Theorem 5.3.6 for all satellite components. The idea of the proof is to "cut off" a curve from the desired direction at ∞ by connecting bifurcating hyperbolic components. This is easier when we know that the curve starts in a given hyperbolic component; otherwise we need a tool to make sure that the curves we use to cut can be chosen to lie far to the right. This tool is given by Corollary 3.3.7.

If W is a hyperbolic component and $\mu : W \to \mathbb{D}$ is its multiplier map, recall that there exists a conformal isomorphism $\Phi : \mathbb{H} \to W$ with $\mu \circ \Phi = \exp$. We can normalize this map in such a way that $\Phi\big((-\infty, 0)\big)$ is the central internal ray of W; we shall refer to this as the *preferred parametrization* of W. Recall that Φ extends to a continuous map $\Phi : \overline{\mathbb{H}} \to \hat{\mathbb{C}}$.

As we want to use bifurcating components to cut, we will need to make sure that these exist. Luckily there are sufficiently many of them:

5.8.1 Proposition (Bifurcation Angles are Dense)
Let W be a hyperbolic component, and let $\Phi : \mathbb{H} \to W$ be its preferred parametrization. Then the set

$$\{\theta \in \mathbb{Q} : \Phi(2\pi i\theta) \in \mathbb{C}\}$$

is dense in $\mathbb{R}$.

Proof. The set

$$\{\theta \in \mathbb{R} : \Phi(2\pi i\theta) \in \mathbb{C}\}$$

is open by continuity of Φ, and dense by the F. and M. Riesz theorem. ■

5.8.2 Proposition (Curves in the Wake of a Hyperbolic Component)
Let $\gamma : [0, \infty) \to \mathbb{C}$ be a curve to ∞ which contains no indifferent parameters. If W is a hyperbolic component with $\mathrm{addr}(\gamma) \in \mathcal{W}(W)$, then there exists a curve which separates γ from all hyperbolic components $W_{\underline{t}}$ with $\underline{t} \notin \mathcal{W}(W)$. This curve consists only of parameters in W, parameters in a child component of W and one parabolic parameter on ∂W.

Proof. Without loss of generality, we may assume that γ tends to ∞ either in or below the preferred homotopy class of W. Let Φ be the preferred parametrization of W. By Proposition 5.8.1 we can find an arbitrary small angle $\theta \in \mathbb{Q} \cap (0, \infty)$ for which $\kappa_0 := \Phi(2\pi i\theta) \in \mathbb{C}$. Let V be the child component which bifurcates from W at this point. By Corollary 5.7.6, if θ was small enough, then $\mathrm{addr}(V) < \mathrm{addr}(\gamma)$; because γ can intersect at most one hyperbolic component, we can also assume that it does not intersect V. Then the desired curve is obtained by following the central internal ray of V to κ_0 and then connecting κ_0 to ∞ above γ in the preferred homotopy class of W. ■

5.8.3 Theorem (Landing of Non-Central Internal Rays)
Let W be a hyperbolic component and let Φ be its preferred parametrization. Then $\Phi(2\pi i\theta) \in \mathbb{C}$ for all $\theta \neq 0$.

Proof. Suppose that $\theta \neq 0$ and $\Phi(2\pi i\theta) = \infty$. We may assume without loss of generality that $\theta > 0$. Let

$$\gamma : (-\infty, 0) \to \mathbb{C}, t \mapsto \Phi(t + 2\pi i\theta).$$

Then for $t \to -\infty$, γ has external address $\underline{t} := \mathrm{addr}(W)$; let $\underline{s}$ be the external address of γ for $t \to 0$.

By Proposition 5.8.1, there exists some $\theta_0 \in \mathbb{Q}$ between 0 and θ with $\Phi(2\pi i\theta_0) \in \mathbb{C}$. As in the proof of Proposition 5.8.2, we can surround γ by a curve in W and the child component of W bifurcating from $\Phi(2\pi i\theta_0)$. It follows that $\underline{s} \in \mathcal{W}(W)$ and $\underline{s} \leq \underline{t}$.

Similarly, there exists a rational $\theta_1 > \theta$ for which $\kappa_1 := \Phi(t + 2\pi i\theta_1) \in \mathbb{C}$. Since γ surrounds κ_1, it also surrounds the child component W' bifurcating from W at κ_1. This shows that $\underline{s} \neq \underline{t}$.

By part (b) of the Squeezing Lemma (Theorem 5.4.1), $\underline{s}$ is not exponentially bounded. By Corollary 5.7.6, there exists some component V which bifurcates combinatorially from W such that $\underline{s} \in \mathcal{W}(V)$. By Proposition 5.8.2, there exists a curve which is disjoint from W and separates the piece $\gamma([-1, 0))$ from W. This is a contradiction. ■

5.8.4 Corollary (Boundaries of Satellite Components)
Suppose that W is a satellite component; i.e., a component which bifurcates combinatorially from some other hyperbolic component. Then the boundary of W in $\mathbb{C}$ is connected.

Remark. In particular, W bifurcates not only combinatorially, but also in the usual sense from a parent component.

Proof. It only remains to show that the central internal ray of W lands, or equivalently that W has a root point. Let W' be the component from which W bifurcates combinatorially, and let γ be the corresponding internal ray in W'. Then by Theorem 5.8.3, this internal ray lands at a parabolic boundary point, which is necessarily the root point of W. ■

5.8.5 Proposition (Analytical Boundary)
Let W be a hyperbolic component and Φ its preferred parametrization. Then the function $\theta \mapsto \Phi(2\pi i\theta)$ is analytic in $\mathbb{R} \setminus \mathbb{Z}$. For $\theta \in \mathbb{Z} \setminus \{0\}$, ∂W has a cusp in $\Phi(2\pi i\theta)$.

If $\kappa_0 := \Phi(0) \in \mathbb{C}$, then ∂W is analytic in κ_0 or has a cusp in κ_0 depending on whether W is a satellite or primitive component, respectively.

Proof. The statements about roots and co-roots follow easily from the previous discussions and the local structure of the multiplier map in these points.

For $\theta \notin \mathbb{Z}$, the multiplier map μ extends analytically to a neighborhood of $\kappa := \Phi(2\pi i\theta)$. Thus the claim follows by the implicit function theorem, unless κ is a critical point of the multiplier map μ. So suppose that this is the case. Then for any small neighborhood U of κ, $\mu^{-1}(\mathbb{D}^* \cap U)$ has at least two components which contain κ on the boundary. By Corollary 5.3.4, only one of these lies in W, so there exists some other period n component V with $\kappa \in \partial V$. By Corollary 5.7.6, V lies in a subwake of W. By Proposition 5.8.2, V can be separated from W by a curve which consists of a parabolic parameter and attracting parameters of period $> n$. This is a contradiction. ■

5.8.6 Corollary (Closures Intersect at Parabolic Points)
Let W be a hyperbolic component and let $\kappa \in \partial W$ be an indifferent parameter. Then no other hyperbolic component contains κ on its boundary.

Proof. The multiplier map μ of W extends to a a holomorphic function in a neighborhood of κ. By Lemma 5.3.1, any hyperbolic component whose boundary contains κ must contain a component of $\mu^{-1}(\mathbb{D}^* \cap U)$. By the proof of Proposition 5.8.5, there is only one such component. ■

The *bifurcation forest* of hyperbolic components is the (infinite) graph on hyperbolic components in which two components are adjacent if and only if their closures intersect. A *bifurcation tree* is any component of this graph. Note that every bifurcation tree is indeed a tree in the graph-theoretical sense. Indeed, every hyperbolic component W can touch at most one whose period is not larger than that of W. (In fact, this can be proved in a much more elementary manner using the same proof as Corollary 5.3.4.) It was conjectured in [30] that there are infinitely many bifurcation trees. Using the previous results, we can now prove this fact.

5.8.7 Corollary (Infinitely Many Bifurcation Trees)
There are infinitely many bifurcation trees.

PROOF. Any primitive component is the root, i.e. component of smallest period, of a bifurcation tree. However, there are infinitely many primitive components, for example the components $W_{0(k+\frac{1}{2})\infty}$ with $k > 0$. ■

We will now prove the general form of part (c) of the squeezing lemma.

5.8.8 Theorem (Squeezing Lemma, Part (c))
Suppose that $\gamma : [0, \infty) \to \mathbb{C}$ is a curve to ∞ which does not contain indifferent parameters. If $\underline{s} := \operatorname{addr}(\gamma)$ is intermediate, then γ lies in the preferred homotopy class of $W_{\underline{s}}$.

PROOF. For simplicity, we may assume that γ does not intersect $W := W_{\underline{s}}$ (otherwise we can use the same proof as for Theorem 5.8.3). Suppose, by contradiction, that γ does not lie in the preferred homotopy class of W. We may then assume without loss of generality that γ tends to ∞ below this homotopy class.

Let Φ be the preferred parametrization of W and let θ_j be a sequence of rational numbers with $\theta_j \to \infty$. For every j, let $\kappa_j := \Phi(2\pi i \theta_j)$. Denote by g_j^1 the internal ray of W landing at κ_j and by g_j^2 the central internal ray of the child component V_j bifurcating from κ_j. Then the curve

$$g_j := g_j^1 \cup \{\kappa_j\} \cup g_j^2$$

surrounds γ. We will derive a contradiction by showing that

$$\lim_{j\to\infty} \min\{|\kappa| : \kappa \in g_j\} = \infty.$$

Because Φ is continuous in ∞, it is clear that

$$\lim_{j\to\infty} \inf\{|\kappa| : \kappa \in g_j^1\} = \infty.$$

It is thus sufficient to concentrate on the curves g_j^2. Denote the characteristic addresses of V_j by $\underline{r}_j$ and $\tilde{\underline{r}}_j$. Note that the dynamic rays at these addresses land together for every parameter on g_j^2. By Theorem 5.7.5, the n-th entries of $\underline{r}_j$ and $\tilde{\underline{r}}_j$ tend to ∞ as $j \to \infty$. By Corollary 3.3.7, this implies that

$$\inf\{|\kappa| : \kappa \in g_j^2\}$$

does too. ■

5.9 Internal Addresses

Fix a hyperbolic component W of period n, $\mathbb{K}^*(W) = \overline{\mathtt{u}_1 \ldots \mathtt{u}_n}$. Consider all external addresses in $\mathcal{W}(W)$ whose corresponding rays land at one of the boundary fixed points of U_1. In other words, these are the characteristic addresses of W and the addresses with itinerary $\overline{\mathtt{u}_1 \ldots \mathtt{u}_{n-1}\mathtt{j}}$, $\mathtt{j} \neq \mathtt{u}_n$. It is easy to show that for every $m \in \mathbb{Z}$, exactly one of these addresses has m as its nth entry. We can thus label these rays as $\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}$ such that $\mathbb{K}(\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}) = \overline{\mathtt{u}_1 \ldots \mathtt{u}_{n-1}{}^{\mathtt{j}}_{\mathtt{j}-1}}$.

The following lemma generalizes Corollary 4.5.9 and suggests a definition of sector wakes that we shall give below.

5.9.1 Lemma (Itineraries of Sector Boundaries)
Let $\underline{t} \in \mathcal{W}(W)$, $\mathtt{j} \in \mathbb{Z}$. Then

$$\operatorname{itin}_{\underline{t}}\left(\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}\right) = \begin{cases} \mathbb{K}^{+}\left(\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}\right) & \text{if } \underline{t} \geq \underline{r}^{\mathtt{j}}_{\mathtt{j}-1} \\ \mathbb{K}\left(\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}\right) & \text{if } \underline{t} = \underline{r}^{\mathtt{j}}_{\mathtt{j}-1} \\ \mathbb{K}^{-}\left(\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}\right) & \text{if } \underline{t} \leq \underline{r}^{\mathtt{j}}_{\mathtt{j}-1} \end{cases}$$

Proof. The proof is analogous to that of Lemma 4.5.8. ■

5.9.2 Definition (Wakes of Sectors)
Let A be the sector of W with kneading sequence $\overline{\mathtt{u}_1 \ldots \mathtt{u}_{n-1}\mathtt{j}}$, $\mathtt{j} \neq \mathtt{u}_n$. Then

$$\mathcal{W}(A) := \mathcal{W}_{\mathtt{j}}(W) := \{\underline{t} \in \mathcal{W}(W) : \underline{r}^{\mathtt{j}+1}_{j} < \underline{t} < \underline{r}^{\mathtt{j}}_{\mathtt{j}-1}\}$$

is called the wake of A. The two addresses bounding it are called the sector boundaries *or* characteristic addresses *of A.* □

5.9.3 Definition (Bifurcation Order and Combinatorial Arcs)
Let A and B be hyperbolic components or sectors. Then we say that $A \prec B$ if $\mathcal{W}(B) \subset \mathcal{W}(A)$.

If $A \prec B$, the combinatorial arc *$[A, B]$ is defined to be the collection of all C satisfying $A \prec C \prec B$. Similarly, when $\underline{s} \in \overline{\mathcal{W}(A)}$, we define $[A, \underline{s}]$ to be the set of all C satisfying $A \prec C$ and $\underline{s} \in \overline{\mathcal{W}(C)}$.* □

5.9.4 Lemma (Kneading Sequences Differ)
If W' is a hyperbolic component with $W \prec W'$, then $\mathbb{K}(W') \neq \mathbb{K}(W)$. Furthermore, $\operatorname{itin}_{\underline{s}}(\underline{t}) \neq \operatorname{itin}_{\underline{s}}(\underline{r}^{\mathtt{j}}_{\mathtt{j}-1})$ for all $\underline{s} \in \mathcal{W}(W)$, $\mathtt{j} \in \mathbb{Z}$ and $\underline{t} \in \mathcal{S} \setminus \{\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}\}$.

Proof. Suppose that $\mathbb{K}(W') = \mathbb{K}(W)$. Then there exists a sector boundary $\underline{t}$ of W' which has itinerary $\mathbb{K}^{*}(W)$ under $\underline{s}$. For parameters in W', this means that the rays at the characteristic addresses of W and the ray $g_{\underline{t}}$ land together. This is impossible by Corollary 4.5.9.

It is sufficient to prove the second statement for a dense set of addresses $\underline{s}$, so we can assume that $\underline{s}$ is in fact an intermediate external address. Thus $\underline{t}$ and $\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}$ belong to an essential orbit portrait, which must have an associated intermediate address $\underline{s}'$, $\mathcal{W}(\underline{s}') \in \mathcal{W}(W)$. Since the images of $\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}$ lie outside of $\mathcal{W}(W)$, $\underline{r}^{\mathtt{j}}_{\mathtt{j}-1}$ must be a characteristic address for $\underline{s}'$. However, this implies $\mathbb{K}(\underline{s}') = \mathbb{K}(W)$, which contradicts the first part of this lemma. ■

5.9.5 Theorem (Determining Bifurcation Structure)
Let A be a sector of some hyperbolic component and $\underline{s} \in \mathcal{W}(A)$. Let j be the first index at which $\mathbb{K}(A) = \mathtt{u}_1\mathtt{u}_2\dots$ and $\mathbb{K}(\underline{s}) = \tilde{\mathtt{u}}_1\tilde{\mathtt{u}}_2\dots$ differ. Then there are no components of period smaller than j on the combinatorial arc $[A, \underline{s}]$. Furthermore there is a unique component $W' \in [A, \underline{s}]$ of period j such that $\mathbb{K}^*(W') = \overline{\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_j}$. If $\tilde{\mathtt{u}}_j \in \mathbb{Z}$, then $\underline{s} \in \mathcal{W}_{\tilde{\mathtt{u}}_j}(W')$.

Remark. In particular it follows that every periodic address is a sector boundary of some hyperbolic component.

Proof. We prove the theorem by induction in j; note that we do not assume A (or the hyperbolic component to which A belongs) to be fixed throughout the induction. The theorem is trivial for $j = 1$.

So let $j > 1$ and assume the theorem is true for all smaller values. Let us first show that, for any address $\underline{s}$ as in the statement of the theorem, there can be no components of period smaller than j on $[A, \underline{s}]$. Indeed, suppose there is a component W' of minimal period $m < j$ in $[A, \underline{s}]$. Applying the induction hypothesis for m to the characteristic rays of this component, we see that $\mathbb{K}^*(W') = \overline{\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_m}$. Now let B be the sector of W' containing $\underline{s}$. Then $\mathbb{K}(B)$ and $\mathbb{K}(\underline{s})$ differ at the m-th position, so we can apply the induction hypothesis for B and $\underline{s}$. In other words, there is a component in $\mathcal{W}(B)$ of period m and kneading sequence $\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{m-1}* = \mathbb{K}(W)$. This is impossible by Lemma 5.9.4.

We will now locate hyperbolic components of period j by explicitly constructing their sector boundaries. To this end, we repeatedly use the following observations:

Observation 1 If $\underline{s} \in \mathcal{W}(A)$ satisfies $\mathbb{K}^+(\underline{s}) = \mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{j-1}\tilde{\mathtt{u}}_j\dots$ with $\tilde{\mathtt{u}}_j \neq \mathtt{u}_j$, then

$$\underline{r}^-(\underline{s}) := \sup\{\underline{r} \le \underline{s} : \mathbb{K}(\underline{r}) = \mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{j-1}\mathtt{m}\dots, \text{where } \mathtt{m} \neq \tilde{\mathtt{u}}_j\}$$

has kneading sequence $\mathbb{K}(\underline{r}^-(\underline{s})) = \overline{\mathtt{u}_1\dots\mathtt{u}_{j-1}{}^{\tilde{\mathtt{u}}_j}_{\tilde{\mathtt{u}}_j-1}}$. Similarly, if $\mathbb{K}^-(\underline{s}) = \mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{j-1}\tilde{\mathtt{u}}_j\dots$, then

$$\underline{r}^+(\underline{s}) := \sup\{\underline{r} \ge \underline{s} : \mathbb{K}(\underline{r}) = \mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{j-1}\mathtt{m}\dots, \text{where } \mathtt{m} \neq \tilde{\mathtt{u}}_j\}$$

has kneading sequence $\mathbb{K}(\underline{r}^+(\underline{s})) = \overline{\mathtt{u}_1\dots\mathtt{u}_{j-1}{}^{\tilde{\mathtt{u}}_j+1}_{\tilde{\mathtt{u}}_j}}$.

Proof of Observation 1. Let $\underline{r}^- := \underline{r}^-(\underline{s})$. Note that $\mathbb{K}^-(\underline{r}^-)$ must start with $\mathtt{u}_1\dots\mathtt{u}_{j-1}\mathtt{m}$, $\mathtt{m} \neq \tilde{\mathtt{u}}_j$, and $\underline{r}^-$ must be periodic with period at most j. Suppose the period was $k < j$. Then by the induction hypothesis, $\underline{r}^-$ is a sector boundary of a period k component V with $\mathbb{K}^*(V) = \overline{\mathtt{u}_1\dots\mathtt{u}_k}$. Thus, $\underline{r}^-$ would actually be the lower characteristic address of this component. Because $\underline{s} \notin \mathcal{W}(V)$ and $\underline{r}^- < \underline{s}$, all of $\mathcal{W}(V)$ must lie below $\underline{s}$ as well. However, addresses closely above $\mathcal{W}(V)$ have kneading sequences close to $\mathbb{K}^*(V) = \mathbb{K}^-(\underline{r}^-)$, and thus starting with $\mathtt{u}_1\dots\mathtt{u}_{j-1}\mathtt{m}$. This contradicts the definition of $\underline{r}^-$. Thus $\underline{r}^-$ has period j. Because addresses slightly above $\underline{r}^-$ have kneading sequences starting with $\mathtt{u}_1\dots\mathtt{u}_{j-1}\mathtt{m}+1$, it follows from the definition of $\underline{r}^-$ that $\mathbb{K}(\underline{r}^-)$ is as claimed.

Observation 2 Any component $V \prec A$ with $\mathbb{K}(V) = \mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{j-1}*$ has forbidden kneading sequence $\mathbb{K}^*(V) = \overline{\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{j-1}\mathtt{u}_j}$.

Proof of Observation 2. To fix our ideas, suppose there are period j components in $\mathcal{W}(W)$ with forbidden kneading sequence $\overline{\mathtt{u}_1 \ldots \mathtt{u}_{j-1}\mathtt{m}}$, $\mathtt{m} < \mathtt{u}_j$. Let W' be the lowest of these components; i.e. that whose wake is closest to the lower sector boundary of A. The lower characteristic address $\underline{s}$ of W' has kneading sequence $\overline{\mathtt{u}_1 \ldots \mathtt{u}_{j-1}{}^{\mathtt{m}+1}_{\mathtt{m}}}$. Let us apply Observation 1 repeatedly to obtain a decreasing sequence of periodic addresses $r^-(\underline{s}) \leq r^-(r^-(\underline{s})) \leq \ldots$, with kneading sequences $\overline{\mathtt{u}_1 \ldots \mathtt{u}_{j-1}{}^{\mathtt{m}}_{\mathtt{m}-1}}$, $\overline{\mathtt{u}_1 \ldots \mathtt{u}_{j-1}{}^{\mathtt{m}-1}_{\mathtt{m}-2}}$ etc. They must converge to some intermediate external address $\underline{t}$ with kneading sequence $\mathbb{K}(\underline{t}) = \mathtt{u}_1 \ldots \mathtt{u}_{j-1}*$. Because $\mathcal{W}(W')$ and $\mathcal{W}(\underline{t})$ are disjoint by Lemma 5.9.4, it follows by construction that $\mathbb{K}^*(\underline{t}) = \mathtt{u}_1 \ldots \mathtt{u}_{j-1}\tilde{\mathtt{m}}$, $\tilde{\mathtt{m}} \leq \mathtt{m} < \mathtt{u}_j$. This contradicts the definition of W' and proves Observation 2.

We can now finish the proof of the induction step. Let $\underline{s}$ be as in the statement of the theorem, and suppose without loss of generality that $\tilde{\mathtt{u}}_j \in \mathbb{Z}$, $\tilde{\mathtt{u}}_j < \mathtt{u}_j$. As above, we consider the sequence $\underline{r}^{\tilde{\mathtt{u}}_j+1}_{\tilde{\mathtt{u}}_j} := \underline{r}^+(\underline{s}), \underline{r}^{\tilde{\mathtt{u}}_j}_{\tilde{\mathtt{u}}_j-1} := \underline{r}^-(\underline{s}), \underline{r}^{\tilde{\mathtt{u}}_j-1}_{\tilde{\mathtt{u}}_j-2} := \underline{r}^-(\underline{r}_{\tilde{\mathtt{u}}_j}) \leq \ldots$ and the intermediate external address $\underline{t}$ to which they converge. By Observation 2, $\mathbb{K}^*(\underline{t}) = \overline{\mathtt{u}_1 \ldots \mathtt{u}_j}$. Since $\tilde{\mathtt{u}}_j < \mathtt{u}_j$, it follows that $\underline{r}^{\tilde{\mathtt{u}}_j}_{\tilde{\mathtt{u}}_j-1}$ and $\underline{r}^{\tilde{\mathtt{u}}_j+1}_{\tilde{\mathtt{u}}_j}$ are sector boundaries of $\mathcal{W}(\underline{t})$. Thus $\underline{s} \in \mathcal{W}_{\tilde{\mathtt{u}}_j}(\underline{t})$. ■

5.9.6 Definition (Internal Addresses)
Let $\underline{s} \in \mathcal{S}$. Consider all hyperbolic components $W \in [W_\infty, \underline{s}]$ with the property that the combinatorial arc $[W, \underline{s}]$ contains no component of smaller period than W. The periods of these components, ordered according to the combinatorial order $\prec$, form a strictly increasing (finite or infinite) sequence of integers starting with 1. This sequence is called the internal address *of $\underline{s}$, and denoted $1 \mapsto n_2 \mapsto n_3 \mapsto \ldots$.*

Theorem 5.9.5 immediately yields the following:

5.9.7 Algorithm (Determining Internal Addresses from Kneading Sequences)
Let $\underline{s}$ be an intermediate external address of length n. Then the internal address of $\underline{s}$, $1 \mapsto n_2 \mapsto n_3 \mapsto \ldots \mapsto n$ can be determined inductively from $\mathbb{K}(\underline{s}) = \underline{\mathtt{u}}$ by the following algorithm:

To compute n_{i+1} from n_i, continue the first n_i entries of $\mathtt{u}$ periodically, and compare this kneading sequence $\mathbb{K}_i$ to $\mathbb{K}(\underline{s})$. n_{i+1} is then the first index at which these sequences differ. ■

Remark. $\mathbb{K}_i$ is the kneading sequence of the period n_i component represented by the i-th entry in the internal address.

In particular, this algorithm allows us to calculate $\mathbb{K}^*(\underline{s})$, which consists of the first n entries of $\mathbb{K}_{n-1}$, continued periodically.

5.9.8 Corollary (Infinitely Many Essential Orbits)
Let κ be an attracting parameter. Then E_κ has infinitely many essential orbits if and only if its internal address does not have the form $1 = n_1 \mapsto n_2 \mapsto n_3 \ldots \mapsto n_k$, where $n_j | n_{j+1}$ for all $j \in \{1, \ldots, n_{k-1}\}$.

Remark. Using Algorithm 5.9.7, it is easy to rephrase this condition in terms of kneading sequences.

Proof. A map has only finitely many essential orbits if and only if it is contained in only finitely many wakes. By Corollary 5.7.6, this just means that the component can be reached by a series of bifurcations from the period 1 component, which is just what the statement about internal addresses means. ■

Angled Internal Addresses

Internal addresses do not label hyperbolic components uniquely. For completeness, we will now discuss a way of decorating internal addresses to restore uniqueness, called "angled internal addresses". However, since we shall not need uniqueness in the following, we do not give an independent combinatorial proof here. Rather, we shall deduce the result from the corresponding result for Multibrot sets, which was proved in [47].

If W is a hyperbolic component of period n, then we denote by $\mathcal{W}_{\left(\mathtt{m};\frac{p}{q}\right)}$ the subwake of angle $\frac{p}{q}$ of W in the sector whose nth kneading entry is $\mathtt{m}$.

Let $\underline{s} \in \mathcal{S}$ with internal address $1 \mapsto n_2 \mapsto n_3 \mapsto \dots$. Let W_j be the corresponding period n_j-components. Then the *angled internal address* of $\underline{s}$ is

$$1_{\left(\mathtt{m}_1;\frac{p_1}{q_1}\right)} \mapsto n_{2\left(\mathtt{m}_2;\frac{p_2}{q_2}\right)} \mapsto n_{3\left(\mathtt{m}_3;\frac{p_3}{q_3}\right)} \mapsto \dots$$

where $W_{j+1} \subset \mathcal{W}_{\left(\mathtt{m}_j;\frac{p_j}{q_j}\right)}$.

5.9.9 Theorem (Uniqueness of Angled Internal Addresses)
No two hyperbolic components share the same angled internal address.

Proof. Note that, in a Multibrot set of high enough degree, the characteristic addresses of all hyperbolic components appearing in the internal address will still be characteristic addresses for some corresponding hyperbolic components in the Multibrot set. The formulae for determining kneading sequences, internal addresses, and subwakes are exactly the same in the polynomial case. Thus, for any hyperbolic component in exponential parameter space, there is one with the same characteristic addresses and the same angled internal address in all Multibrot sets of high enough degree. Thus the theorem follows from the corresponding result for Multibrot sets. ■

5.10 Proof of Part (a) of the Squeezing Lemma

We will now complete the proof of the Squeezing Lemma, and thus of Theorem 5.3.6.

5.10.1 Theorem (Squeezing Lemma, Part (a))
Suppose that $\gamma : [0,\infty) \to \mathbb{C}$ is a curve to ∞ which does not contain indifferent parameters. Then $\operatorname{addr}(\gamma)$ *is either intermediate or exponentially bounded.*

Proof. Suppose, by contradiction, that $\underline{s} := \mathrm{addr}(\gamma)$ is infinite but not exponentially bounded. Set $\mathtt{u} := \mathbb{K}(\underline{s})$; then the sequence $\mathtt{u}$ is also not exponentially bounded, and in particular not bounded.

We can choose a subsequence $(\mathtt{u}_{n_k})$ of entries such that $|\mathtt{u}_j| < |\mathtt{u}_{n_k}|$ whenever $j < n_k$ and such that

$$F^{-n_k}(|\mathtt{u}_{n_k}|) \to \infty.$$

Then, by Algorithm 5.9.7, the internal address of $\underline{s}$ contains an entry n_k for every k. Let W_k denote the hyperbolic component of period n_k corresponding to this entry. Then

$$\mathcal{W}(W_1) \supset \mathcal{W}(W_2) \supset \dots \ni \underline{s}.$$

Now, for every k, let V_k be the child component of W_k with $\underline{s} \in \mathcal{W}(V)$. By Proposition 5.8.2, there is a curve γ_k which tends to ∞ in both directions, consists only of parameters in V and a child component of V together with a common boundary point and which surrounds γ. Now let $\underline{r}_k$ and $\underline{\tilde{r}}_k$ denote the characteristic addresses of V. Because $\mathbb{K}^*(V_k) = \overline{\mathtt{u}_1 \dots \mathtt{u}_{n_k}}$, the jn_k-th entries of $\underline{r}_k$ and $\underline{\tilde{r}}_k$ are of size at least $|\mathtt{u}_{n_k}| - 1$. By Corollary 3.3.7, it follows that

$$\gamma_n \subset \{z : |z| > \frac{1}{5}\mathcal{F}^{-n_k+1}(|\mathtt{u}_{n_k}| - 1)\}$$

However, this is a contradiction because $F^{-n_k+1}(|\mathtt{u}_{n_k}| - 1) \to \infty$ and γ_n surrounds γ. ■

5.11 Combinatorial Tuning

Let $\underline{s}$ be an intermediate external address. We will give an analog of the concept of tuning for polynomials, on the combinatorial level. For every i, let r_{i-1}^i consist of the first n entries of the sector boundary $\underline{r}_{i-1}^i$, as before.

A map $\tau : \overline{\mathcal{S}} \to \mathcal{W}(\underline{s})$ is called a tuning map for $\underline{s}$, if $\tau(-\infty) = \underline{s}$ and

$$\tau(k\underline{r}) = \begin{cases} r_{\mathtt{u}_n+k-1}^{\mathtt{u}_n+k}\tau(\underline{r}) & \tau(\underline{r}) > \underline{s} \\ r_{\mathtt{u}_n+k}^{\mathtt{u}_n+k+1}\tau(\underline{r}) & \tau(\underline{r}) < \underline{s} \\ r_{\mathtt{u}_n+k-\frac{1}{2}}^{\mathtt{u}_n+k+\frac{1}{2}}\tau(\underline{r}) & \tau(\underline{r}) = \underline{s}. \end{cases}.$$

There are exactly two such maps, which are uniquely defined by choosing $\tau(\overline{0})$ to be either $\overline{r_{\mathtt{u}_n}^{u_n+1}}$ or $\overline{r_{\mathtt{u}_n-1}^{\mathtt{u}_n}}$. (Note that under the tuning map, addresses which are not exponentially bounded may be mapped to addresses which are exponentially bounded, see Section 4.3.)

Before we prove that this map transforms combinatorics in the appropriate way, let us first make the following observation.

5.11.1 Lemma (Monotonicity of Tuning)
The map τ is increasing when restricted either to addresses $> \overline{0}$ or addresses $< \overline{0}$.

PROOF. This is easily verified from the definition. ■

5.11.2 Theorem (Tuning Theorem)
If the internal address of $\underline{r}$ *is* $1 \mapsto n_2 \mapsto n_3 \mapsto \dots$, *and the internal address of* $\underline{s}$ *is* $1 \mapsto m_2 \mapsto \dots \mapsto n$, *then the internal address of* $\tau(\underline{r})$ *is*

$$m_1 \mapsto m_2 \mapsto \dots \mapsto n \mapsto n * n_2 \mapsto n_3 \mapsto \dots .$$

The kneading sequence of $\tau(\underline{r})$ *can be computed from that of* $\underline{r}$ *by replacing each symbol* $\mathtt{u}$ *by the* n *symbols* $\mathtt{u}_1 \dots \mathtt{u}_{n-1}(\mathtt{u}_n + \mathtt{u})$ *if* $\underline{r} < \overline{0}$, *or* $\mathtt{u}_1 \dots \mathtt{u}_{n-1}(\mathtt{u}_n + \mathtt{u} + 1)$, *if* $\underline{r} > \overline{0}$.

PROOF. Of course it is sufficient to prove the statement about the kneading sequence. To fix our ideas, let us suppose that $\underline{r} < \overline{0}$. Let $\underline{t}$ be any external address, and suppose the first itinerary of $\underline{t}$ with respect to $\underline{s}$ is $\mathtt{u}$. We will show that $\sigma^{n-1}(\tau(\underline{s}))$ lies in the strip $\mathtt{u}_n + \mathtt{u}$ of the partition $*\tau(\underline{r})$.

Note that $\mathtt{u} = t_1$ if $\sigma(\underline{t}) > \underline{r}$ and $\mathtt{u} = t_1 - 1$ if $\sigma(t) < \underline{r}$. Now let us differentiate two cases. If $\tau(\sigma(\underline{t})) < \underline{s}$, then $\sigma(\underline{t}) \geq \overline{0} > \underline{r}$, and therefore $\mathtt{u} = t_1$. Furthermore, by the definition of τ, the n-th entry of $\tau(\underline{t})$ is $u_n + t_1 + 1$, and $\tau(\underline{t}) < \underline{s} < \tau(\underline{s})$. This shows that the n-th itinerary entry of $\tau(\underline{t})$ with respect to $\tau(\underline{r})$ is $u_n + t_1 = \mathtt{u}_n + \mathtt{u}$.

Now consider the case when $\tau(\sigma(\underline{t})) > \underline{s}$. In this case, the n-th entry of $\tau(\underline{t})$ is $u_n + t_1$. By Lemma 5.11.1, $\tau(\sigma(\underline{t})) > \tau(\underline{r})$ if and only if $\sigma(\underline{t}) > \underline{r}$. This completes the proof. ■

5.12 Parameter Rays

In [35], it was shown — generalizing results from [73, Chapter II] — that for every exponentially bounded $\underline{s}$ and every $t > 0$, there exists a unique point κ with $g_{\underline{s}}^{\kappa}(t) = \kappa$. If we denote this parameter by $\mathcal{G}_{\underline{s}}(t)$, then $\mathcal{G}_{\underline{s}}$ is a continuously differentiable curve in parameter space. Using the powerful methods developed in the previous sections, we can give a different proof of this fact, and extend this classification also to escaping endpoints of rays, to which the method of [35] does not generalize. We should note that the larger part of the following proof is similar to that of [35], except that we use the squeezing lemma to ensure that curves do not end at infinity and that we know that escaping endpoints of dynamic rays depend holomorphically on the parameter.

First, let us mention the following consequence of Hurwitz's theorem. It is rather intuitive, but we do not know of the proof being written elsewhere.

5.12.1 Lemma
Let U *be open, and suppose that, for* $t \in (-\delta, \delta)$,

$$f_t : U \to \mathbb{C}$$

is nonconstant and holomorphic in U *and continuous in* t. *If* $z_0 \in U$ *such that* $f_0(z_0) = 0$, *then, for some* $\varepsilon \leq \delta$, *there exists a continuous function*

$$g : (-\varepsilon, \varepsilon) \to U$$

with $f_t(g(t)) = 0$ *and* $g(0) = z_0$.

Remark. Without the demand for continuity of g, this is simply Hurwitz's theorem. Note that it is possible that the set of t for which the multiplicity of $g(t)$ as a zero of f_t is greater than one need not be discrete. Consider for example

$$f_t(z) := \begin{cases} (z-t)^2 & \text{if } t \geq 0 \\ z^2 + t^2 & \text{if } t < 0. \end{cases}$$

Sketch of proof. Let m be the multiplicity of z_0 as a zero of f_0. By Hurwitz's theorem, we can choose a neighborhood V of z_0 and an interval $I := (-\varepsilon, \varepsilon)$ such that, for $t \in (-\varepsilon, \varepsilon)$, f_t has exactly m zeros in V, counted by multiplicity. We will prove the claim by induction on m. Note that it follows by the inverse function theorem if $m = 1$. So suppose that $m > 1$. Then the set

$$A := \{t \in I : f_t \text{ has a zero of multiplicity} < m \text{ in } V\}$$

is open (again by Hurwitz's theorem). Let t_0 be any element of A, and let I' be the maximal subinterval of A containing t_0 on which we can define a continuous function $g : I' \to V$ with $f_t(g(t)) = 0$. Then I' is open by the induction hypothesis and closed by continuity, so it is a component of A. By choosing one such function g for every component of A, we can construct a continuous function $g : A \to V$ with $f_t(g(t)) = 0$. We can extend this function continuously to I by letting, for $t \notin A$, $g(t)$ denote the unique zero (of multiplicity m) of f_t in V. ■

5.12.2 Theorem (Classification of Escaping Parameters)
There exists a bijective function $\mathcal{G}$ *from* X *to the set of escaping parameters such that* $G^{\mathcal{G}(\underline{s},t)}(\underline{s}, t) = \mathcal{G}(\underline{s}, t)$ *for all* $(\underline{s}, t) \in X$. *For every* ε, $\mathcal{G}$ *is continuous on the set*

$$\{(\underline{s}, t) : t > t_{\underline{s}} + \varepsilon\}.$$

Remark. As before, we will call

$$\mathcal{G}_{\underline{s}}(t) := \mathcal{G}(\underline{s}, t_{\underline{s}} + t)$$

the *parameter ray* at address $\underline{s}$.

Proof. Let $\underline{s}$ be any exponentially bounded external address, and consider the set

$$A := \{\kappa : \kappa \in g_{\underline{s}}^{\kappa}\}.$$

We claim first that for every $\kappa_0 \in A$, say $\kappa_0 = g_{\underline{s}}^{\kappa_0}(t_0)$, there exists T with $t_0 < T \leq \infty$ and a continuous curve $H : [t_0, T) \to \mathbb{C}$ with $H(t) = g_{\underline{s}}^{H(t)}(t_0)$ and $\lim_{t \to T} H(t) = \infty$. Note first that κ_0 is an isolated zero of the function $f_{t_0} : \kappa \mapsto g_{\underline{s}}^{\kappa}(t_0) - \kappa$ because otherwise κ_0

would be a J-stable parameter, which contradicts theorem 5.1.4. By Lemma 5.12.1, we can choose a continuous function $H' : [t_0, t_0 + \varepsilon] \to \mathbb{C}$ with the required properties. (Note that, if $t_0 = 0$, the function f_t is not defined for $t < 0$. However, this makes no difference to the proof of the lemma; alternatively we could extend the definition of f_t by setting $f_t := f_0$ for $t < 0$.)

Now let $H : [t_0, T) \to \mathbb{C}$ be a maximal continuation of this piece. We need to show that $\lim_{t\to T} H(t) = \infty$. First suppose that $T < \infty$ and let κ_1 be any finite limit point of H as $t \to T$. We will show that $g_{\underline{s}}^{\kappa_1}(T) = \kappa_1$, which contradicts the maximality of H. Let $x := (\underline{s}, T + t_{\underline{s}})$. Let Q be so large that G^κ is defined on Y_Q for all κ in a neighborhood of κ_1. For some large enough n we must have $\mathcal{F}^n(x) \in Y_Q$, and then by continuity of G, we have $G^{\kappa_1}(\mathcal{F}^n(x)) = E_{\kappa_1}^n(\kappa_1)$. Because pullbacks of the ray $G_{\sigma^n(\underline{s})}$ vary continuously (until they pass through κ), it follows easily that $G^{\kappa_1}(x) = \kappa_1$.

If, on the other hand, $T = \infty$, then $\lim_{t\to\infty} H(t) = \infty$ follows easily from the construction of escaping points. This proves our claim that every point in A can be connected to ∞ by a curve H as above.

Now by Theorem 5.3.5, $\underline{s}' = \operatorname{addr}(H)$ is exponentially bounded. By Corollary 5.4.4, for some a (which does not depend on H) $H(t)$ lies on the curve $H_{\underline{s}}$ from Corollary 5.4.4 for all $t > a$. Thus for $t > a$, there exists exactly one parameter $\kappa = H_{\underline{s}}(t)$ with $g_{\underline{s}}^{\kappa}(t) = \kappa$, and κ is a simple zero of f_t.

Now let $H : (t_0, \infty) \to \mathbb{C}$ or $H : [t_0, \infty)$ be some maximal curve in H with $f_t(H(t)) = 0$ for all t. We show that, for every t, the value $H(t)$ is a simple zero of f_t. Suppose that the claim is false, and let $t_1 \geq a$ be maximal with the property that $H(t_1)$ is a multiple zero of f_t. Note that, for $t > t_1$, $H(t)$ must then be the unique zero of f_t in all of $\mathbb{C}$. By Hurwitz's theorem, for $t = t_1 + \varepsilon$, f_t has the same number of zeros near $H(t_1)$ as f_{t_1}, counting multiplicities. This is a contradiction.

In particular, we have shown that A consists only of the curve H. If the domain of H is $[t_0, \infty)$, then $t_0 = 0$. Indeed, otherwise there are zeros of $f_{t_0 - \varepsilon}$ close to $H(t_0)$, which contradicts the maximality of H. It then remains to show that, if the domain of H is (t_0, ∞), then $(\underline{s}, t_0 + t_{\underline{s}}) \notin X$ (and in particular $t_0 = 0$). This follows by the same argument as above: Otherwise, every limit point $\kappa_1 \in \mathbb{C}$ of H would satisfy $\kappa_1 = g^{\kappa_1}(t_0)$. Thus we only need to exclude that $\lim_{t\to t_0} H(t) = \infty$, which again follows by Theorem 5.3.5.

Thus we have constructed a bijective map $\mathcal{G}$ as in the statement of the theorem. The continuity of $\mathcal{G}$ on the mentioned sets follows from the continuous dependence of G on κ, using Hurwitz's theorem once more. ■

The construction of parameter rays has the following curious consequence.

5.12.3 Remark

Suppose that κ is an escaping parameter with $\operatorname{Im}\kappa \in (-\pi, \pi)$. Then the first entry of $\operatorname{addr}(\kappa)$ is 0.

Proof. By the previous theorem, κ lies on a parameter ray, whose address must lie between $-\frac{1}{2}\infty$ and $\frac{1}{2}\infty$. ■

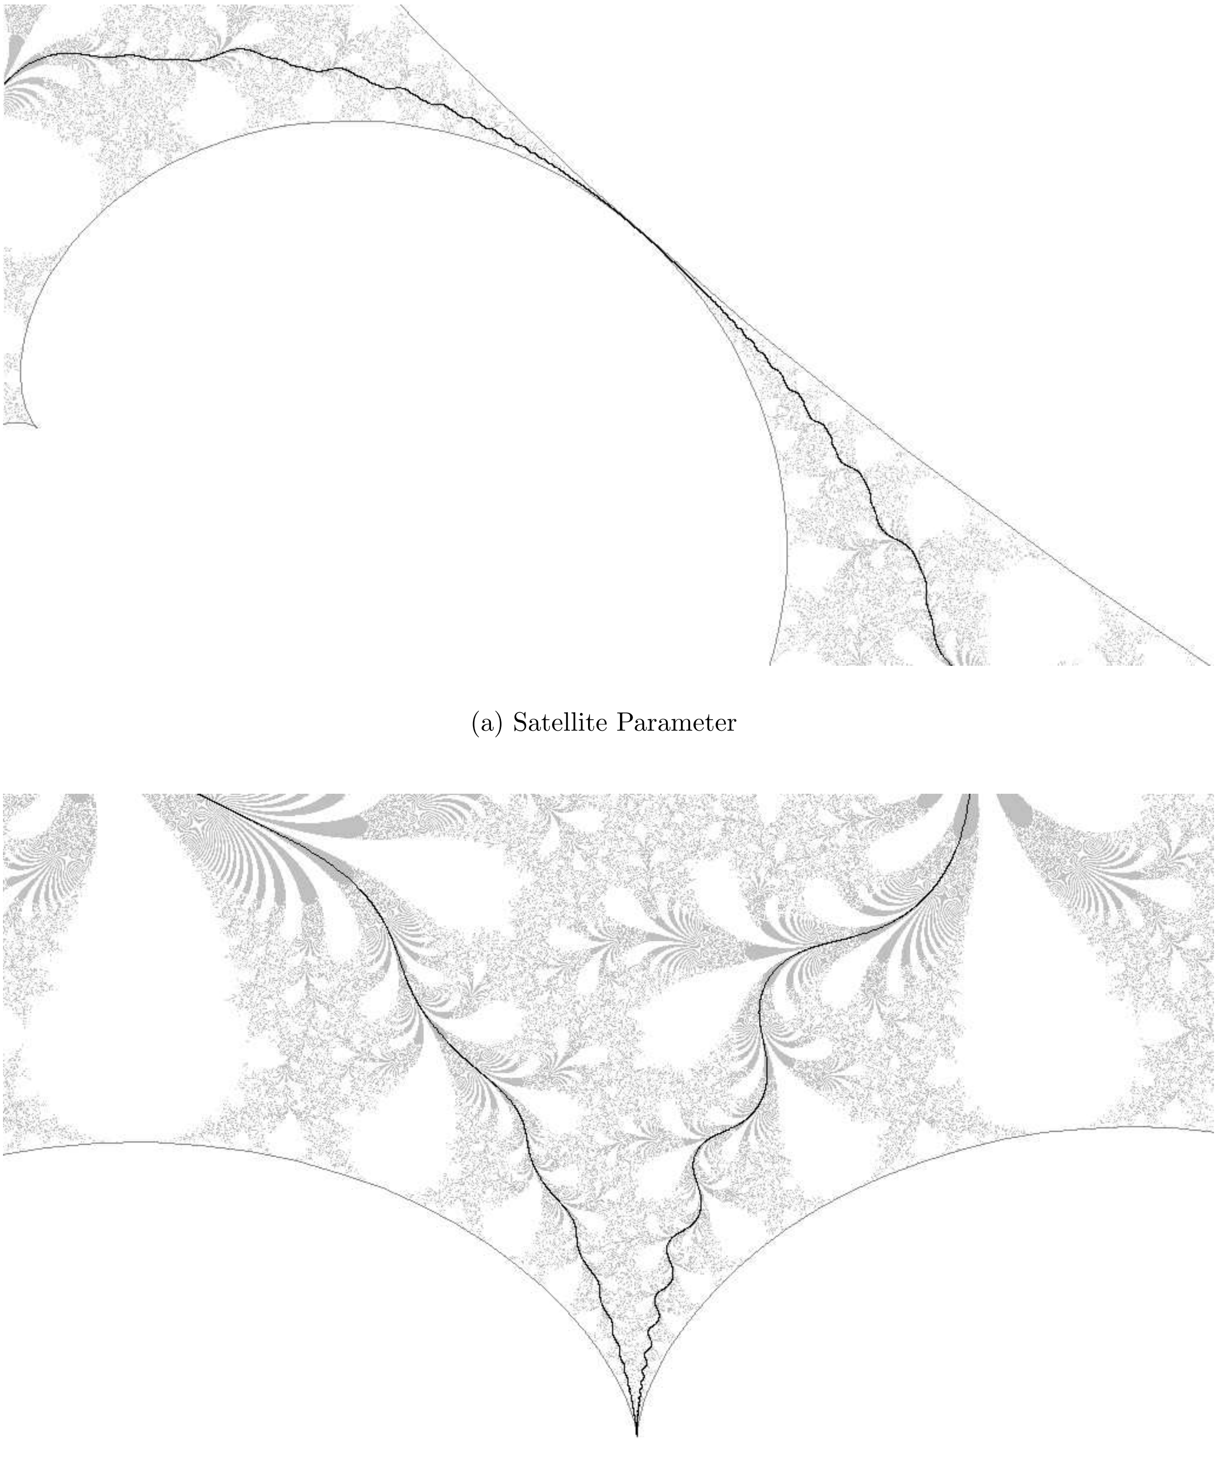

(a) Satellite Parameter

(b) Primitive Parameter

Figure 5.7: Parameter rays at root points of hyperbolic components.

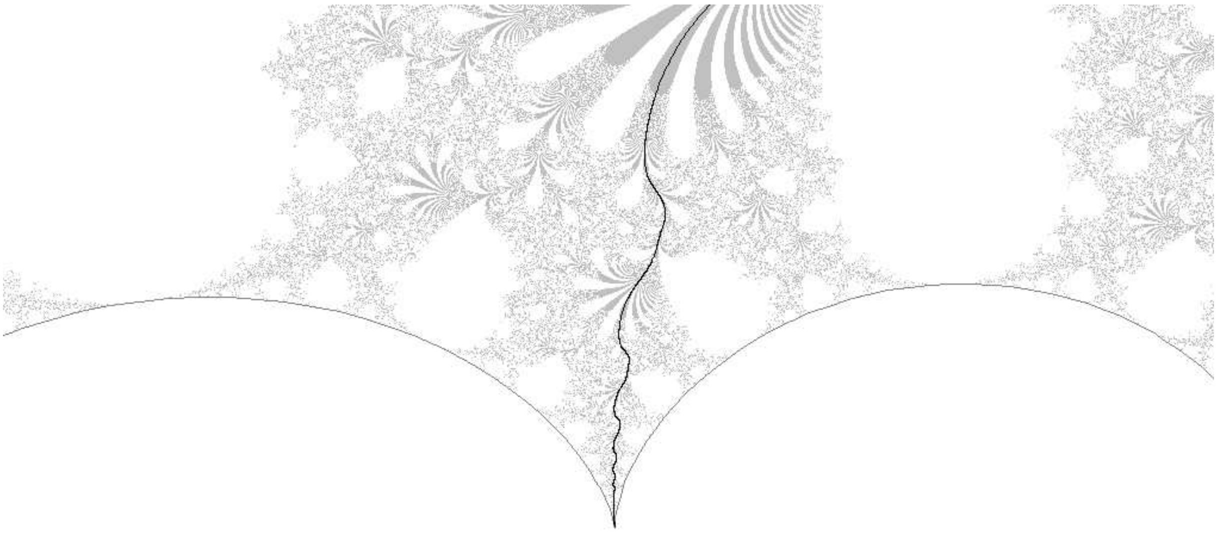

(a) Co-root

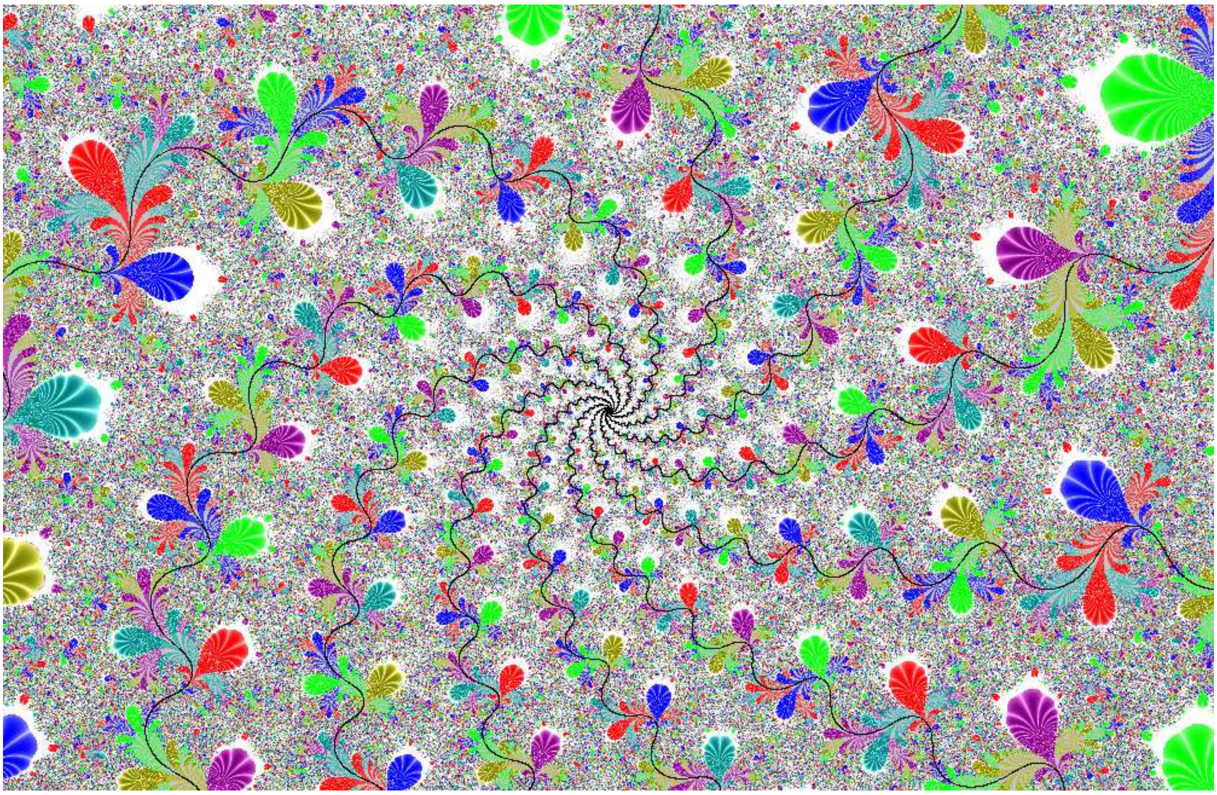

(b) Misiurewicz Parameter

Figure 5.8: Parameter rays at a co-root and at a Misiurewicz parameter.

In [73], Schleicher was able to prove that all parameter rays at periodic addresses land by showing that parabolic points are landing points of periodic parameter rays [73, Corollaries IV.4.4 and IV.5.2] and using the existence of all roots and co-roots of hyperbolic components to ensure that for every periodic ray there exists a suitable parabolic parameter.

5.12.4 Theorem (Parabolic Points are Landing Points [73])
Let κ_0 be a parabolic parameter. First suppose that the orbit portrait $\mathcal{O}$ of the parabolic periodic point of κ_0 is essential (i.e. that κ_0 is a root point), and let $\underline{s}_1$ and $\underline{s}_2$ denote its characteristic addresses. Then κ_0 is the landing point of the parameter rays $\mathcal{G}_{\underline{s}_1}$ and $\mathcal{G}_{\underline{s}_2}$.

If $\mathcal{O}$ is not essential, let $\underline{s}$ be the address of the unique periodic external ray which lands at the principal point of the parabolic orbit. Then κ_0 is the landing point of the parameter ray $\mathcal{G}_{\underline{s}}$. □

5.12.5 Corollary (Periodic Parameter Rays Land [73, Theorem V.7.2])
Let $\underline{s}$ be any periodic external address. Then $\mathcal{G}_{\underline{s}}$ lands at a parabolic parameter.

Proof. There exists a hyperbolic component W such that $\underline{s}$ is a characteristic address or sector boundary of W. If $\underline{s}$ is a characteristic address of W, then let κ_0 be the root of W. By the previous theorem, $G_{\underline{s}}$ lands at κ_0. Similarly, if $\underline{s}$ is a sector boundary, then $\mathcal{G}_{\underline{s}}$ lands at a suitable co-root of W. ■

We can now give a definition of the wake of W as a subset of parameter space.

5.12.6 Definition (Wake)
Let W be a hyperbolic component. Let κ_0 be the root point of W, and let $\underline{s}_1$ and $\underline{s}_2$ be the characteristic addresses of W. Then $\mathcal{W}(W)$, the wake *of W, is defined to be the component of*

$$\mathbb{C} \setminus \left(\kappa_0 \cup \mathcal{G}_{\underline{s}_1} \cup \mathcal{G}_{\underline{s}_2}\right)$$

containing W.

For completeness, let us note that, using the results of [75], Schleicher also proved the landing of all preperiodic rays at Misiurewicz parameters.

5.12.7 Theorem (Misiurewicz Points are Landing Points [73, Theorem IV.6.1])
Let κ_0 be a Misiurewicz parameter, and let $\underline{s}_1 < \underline{s}_2 < \cdots < \underline{s}_n$ be the addresses of the preperiodic dynamic rays landing at the singular value for E_{κ_0}. Then κ_0 is the landing point of the parameter rays $\mathcal{G}_{\underline{s}_1}, \ldots, \mathcal{G}_{\underline{s}_n}$.

Proof. Let $\underline{s} \in \{\underline{s}_1, \ldots, \underline{s}_n\}$. The orbit portraits of repelling periodic points are stable under small perturbations. Thus the ray $g_{\underline{s}}$ and its landing point move holomorphically in a neighborhood of κ_0. Using Hurwitz's theorem as in the proof of Theorem 5.12.2, we can find a point $\gamma(t)$ (for small t) near κ_0 with $g_{\underline{s}}^{\gamma(t)}(t) = \gamma(t)$, and $\gamma(t) \to \kappa_0$ for $t \to 0$. By Theorem 5.12.2, $\gamma(t) = \mathcal{G}_{\underline{s}}(t)$. ■

5.12.8 Theorem (Preperiodic Parameter Rays Land [73, Theorem IV.6.2])
Let $\underline{s}$ be a preperiodic external address. Then $\mathcal{G}_{\underline{s}}$ lands.

PROOF. This is an immediate consequence of Theorems 4.4.4 and 5.12.7. ■

5.13 All Periodic Dynamic Rays Land

We will now prove Theorem 1.10, that, for every nonescaping parameter, all periodic rays land.

5.13.1 Theorem (Periodic Rays Land)
Let $\underline{s}$ be any periodic external address, and let

$$\kappa \notin \bigcup_{n \geq 0} \mathcal{G}_{\sigma^n(\underline{s})}.$$

Then $g_{\underline{s}}^{\kappa}$ lands at a periodic point of E_κ.

PROOF. If κ is parabolic or hyperbolic, then $g_{\underline{s}}$ lands by Theorem 3.9.1. By Theorem 5.12.5, the set

$$K := \overline{\bigcup_n \mathcal{G}_{\sigma^n(\underline{s})}}$$

consists of the rays $\mathcal{G}_{\sigma^n(\underline{s})}$ together with finitely many parabolic parameters (for which all periodic rays land by Theorem 3.9.1). Let U be the component of $\mathbb{C} \setminus K$ which contains κ. By Lemma 3.10.1, it suffices to show that there is some hyperbolic parameter in U. However, this is clear because there are hyperbolic components arbitrarily close below and above every parameter ray. ■

With this theorem, we can finally characterize the wakes of hyperbolic components in a more natural way.

5.13.2 Corollary (Orbit Forcing)
Let W be a hyperbolic component, let s_1 and s_2 be the characteristic addresses of W and let $\mathcal{O}$ be their orbit portrait (for parameters in W). Then, for any $\kappa \in \mathbb{C}$, the following are equivalent.

a) $\kappa \in \mathcal{W}(W)$,

b) $g_{\underline{s}_1}^{\kappa}$ and $g_{\underline{s}_2}^{\kappa}$ land together,

c) $\mathcal{O}$ is realized by κ.

PROOF. The orbit of $g_{\underline{s}_1}$, $g_{\underline{s}_2}$ and their common landing point moves holomorphically in $\mathcal{W}(W)$. ■

5.14 Periodic Points are Landing Points

In this section, we will use the landing theorem for external rays to prove Theorem 1.11. The key observation is the following.

5.14.1 Lemma (Periodic Points which are not landing points)
Suppose that an exponential map E_κ with nonescaping singular value has some periodic orbit of period n whose points are not the landing points of periodic external rays. Then there exists a period n hyperbolic component W with $\kappa \in \mathcal{W}(W)$, but $\kappa \notin \mathcal{W}(W')$ for every child component W' of W.

Proof. Again, we can assume that κ is not parabolic. If κ is not contained in the wake of any period n component, then let $W = W_\infty$; otherwise let W be the unique period n component with $\kappa \in \mathcal{W}(W)$ but $\kappa \notin \mathcal{W}(W')$ for any period n component W' with $W \prec W'$.

Consider the set U obtained by removing from $\mathcal{W}(W)$ all period n parameter rays and their landing points (i.e., roots and co-roots of hyperbolic components of period n). This set is open by Corollary 3.3.7. Let V be the component of U containing κ and W. Let $\kappa_0 \in W$. Then all period n points of E_{κ_0} move holomorphically throughout U. We claim that the ray portraits of the repelling period n points also move holomorphically. This is clear if the ray portrait is not essential, as then the period of the rays is n, and U does not contain any parameter rays of period n.

On the other hand, if (z_k) is a periodic orbit of E_{κ_0} and its orbit portrait $\mathcal{O}$ is essential, then the characteristic rays of $\mathcal{O}$ are the characteristic rays of some period n hyperbolic component W'. The ray portrait of (z_k) then actually moves holomorphically throughout all of $\mathcal{W}(W') \supset U$.

Thus if (z_k) is a period n orbit whose points are not the landing points of periodic external rays, this orbit must be the analytic continuation of the attracting period n orbit of κ_0 (in particular, $W \neq W_\infty$). However, if W' is any child component of W, then this orbit becomes the characteristic orbit of parameters in W'. For any parameter in $\mathcal{W}(W')$, the characteristic point of this orbit is thus the landing point of the external dynamic rays whose addresses are the characteristic addresses of W'. This completes the proof. ■

Proof of Theorem 1.11. Suppose that E_κ has two periodic orbits of periods n and m, both of which consist of points which are not the landing points of periodic rays. Then there exist period n and m components $W_1 \prec W_2$ such that $\kappa \in \mathcal{W}(W_2)$, but κ does not lie in any subwake of W_1. In particular, $\mathcal{W}(W_2)$ is not contained in any subwake of W_1, which contradicts Corollary 5.7.6. ■

Chapter 6

Other Results on Exponential Dynamics

In this chapter, we will discuss some results concerning the exponential family which do not fit into the framework of the previous sections. Section 6.1 contains results concerning irrationally indifferent cycles and the proof of Theorem 1.12. In Section 6.2, we discuss the stunning "dimension paradox" discovered by Karpińska, as well as some other topological phenomena which appear in the exponential family. Section 6.3 discusses ergodic aspects of exponential dynamics, and finally we attempt to give an indication of work done in related classes of entire functions in Section 6.4.

6.1 Siegel Disks

Let $\theta \in \mathbb{R}$ and

$$\kappa(\theta) := 2\pi i\theta - e^{2\pi i\theta}.$$

Then E_κ has a fixed point of multiplier $\lambda = e^{2\pi i\theta}$ at $z_0 = 2\pi i\theta$. For the purposes of this section, the family is more conveniently parametrized as $f_\lambda : z \mapsto \lambda(\exp(z) - 1)$, for which the fixed point lies at the origin.

The linearization theorem of Bryuno shows that f_λ has a Siegel disk whenever θ is a Bryuno number, and it is easy to show that, for generic λ, f_λ has a Cremer point. For the families of unicritical polynomials, Geyer [37] showed that Bryuno's condition is necessary and sufficient for linearization. However, it is still open whether Bryuno's condition is necessary for linearization in the exponential family.

Herman [42] showed that every polynomial Siegel disk with diophantine rotation number has a critical point on its boundary provided it is injective on its boundary. In particular, this implies for the families of unicritical polynomials that every diophantine Siegel disk does have a critical point on the boundary. In fact, using unpublished work of Yoccoz on the optimal condition $\mathcal{H}$ for linearizability of analytic circle diffeomorphisms (see [60]), the same method shows that these facts are true for the (larger) class $\mathcal{H}$ of rotation numbers.

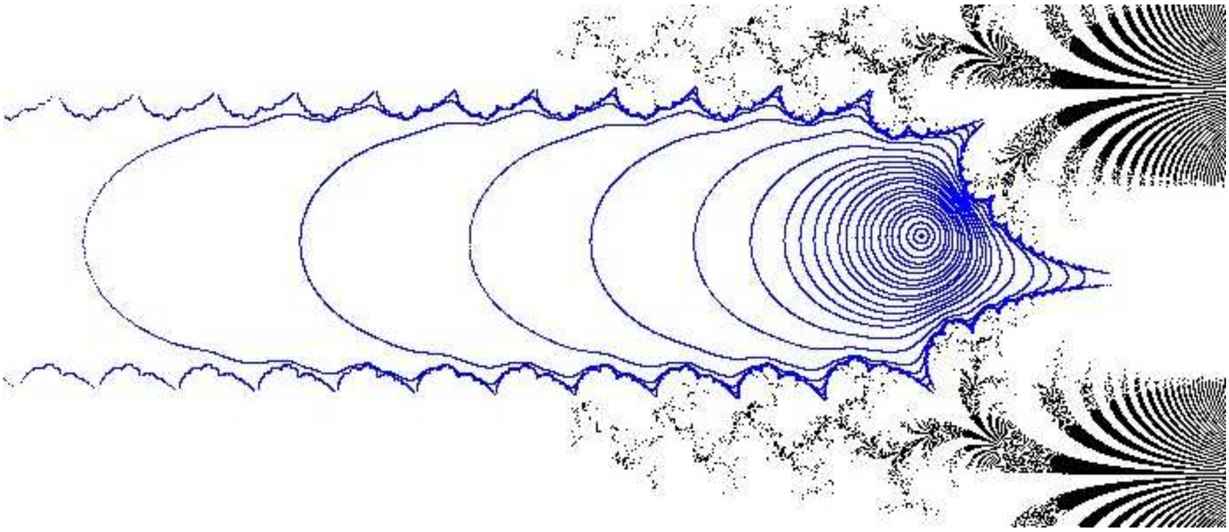

Figure 6.1: The Siegel disk of $E_{\kappa(\theta)}$ where θ is the golden mean.

Herman also applied his methods to the exponential family, but his methods only succeeded to show that the Siegel disk must be unbounded in the case where the rotation number is diophantine (or, more generally, in $\mathcal{H}$). This left open the question whether there are exponential Siegel disks which contain the singular value on their boundary. This question was answered by Rippon [69], who generalized a method of Carleson and Jones [15, page 86] to give an elegant elementary proof that $-\lambda \in \partial U$ for almost every λ.

However, it was until now unknown whether, for example, this is the case for the golden mean (see Figure 6.1). Thus in the collection [13] of research problems in complex analysis, the following question (Problem 2.86) was posed by Herman, Baker and Rippon.

"Let the function $f_\lambda(z) = \lambda(e^z - 1)$, $|\lambda| = 1$, have a Siegel disk U that contains 0.

(a) Prove that there exists some number λ such that U is bounded in $\mathbb{C}$.

(b) If U is unbounded in $\mathbb{C}$, does the singular value $-\lambda$ belong to ∂U?"

As noted in [69], part a) can in fact be done for a dense set of λ using a method by E. Ghys mentioned in [24]. We will shortly explain this method at the end of this section. Let us now answer part (b).

6.1.1 Theorem (Singular Value on Boundary)
Let $\kappa \in \mathbb{C}$ and suppose that E_κ has an unbounded Siegel disk U. Then there is j such that $\kappa \in \partial E_\kappa^j(U)$.

Fix a $\kappa \in \mathbb{C}$ for which E_κ has a Siegel disk U. Recall that κ belongs to the Julia set of E_κ. Let us suppose that $\kappa \notin \partial E_\kappa^j(U)$ for every $j \in \mathbb{N}$; we wish to show that U is bounded. Choose $\delta > 0$ such that

$$\overline{\mathbb{D}_\delta}(\kappa) := \{z \in \mathbb{C} : |z - \kappa| \le \delta\} \subset \mathbb{C} \setminus \bigcup_j E_\kappa^j(U).$$

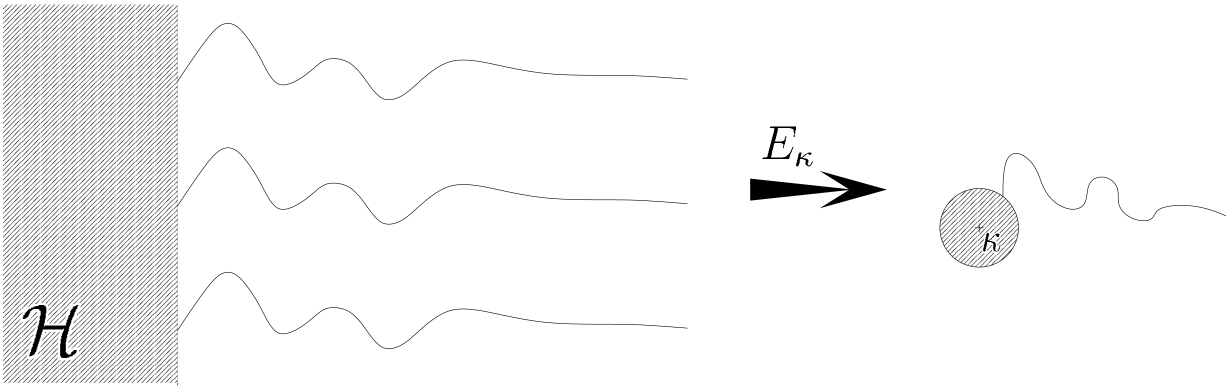

Figure 6.2: The set K and its image, illustrating the proof of Proposition 6.1.2.

Note that this implies that, for every j, $E_\kappa^j(U) \subset \mathbb{C} \setminus \overline{\mathcal{H}}$, where $\mathcal{H} = E_\kappa^{-1}\big(\mathbb{D}_\delta(\kappa)\big) = \{z \in \mathbb{C} : \operatorname{Re} z < \log \delta\}$.

The key is to prove that U and its images under the iterates of E_κ are bounded above and below:

6.1.2 Proposition (Bounded Imaginary Part)
For every $j \in \mathbb{N}$, the set $E_\kappa^j(U)$ has bounded imaginary part.

PROOF. Since $\kappa \in J(E_\kappa)$, we can find an escaping point $z_0 \in I$ with $|z_0 - \kappa| < \delta$. Now $z_0 = g_{\underline{s}}(t_0)$ for some address $\underline{s}$ and $t_0 \geq 0$. Choose the largest $t > 0$ with $|g_{\underline{s}}(t) - \kappa| = \delta$ and consider the set

$$K := E_\kappa^{-1}\left(\overline{\mathbb{D}_\delta(\kappa)} \cup g_{\underline{s}}\big([t_0, \infty)\big)\right).$$

This set consists of $\overline{\mathcal{H}}$ together with the preimages of $g_{\underline{s}}\big([t_0,\infty)\big)$ (see Figure 6.2). Each of these preimages is asymptotic to a line $\{\operatorname{Im} z = 2k\pi\}$, and thus has bounded imaginary part. Therefore every component of $\mathbb{C} \setminus K$ has bounded imaginary part. Since $E_\kappa^j(U) \subset \mathbb{C} \setminus K$, this proves the claim. ■

PROOF OF THEOREM 6.1.1. By the previous proposition, we can find $S > 0$ with $|\operatorname{Im} z| < S$ for all $j \in \mathbb{N}$ and $z \in E_\kappa^j(U)$. Choose $R > 1$ large enough such that $\exp(R-1) > -\log\delta$ and $\exp(R) \geq \exp(R-1) + S + 1 + |\kappa|$. Then, if $z \in \mathbb{C}$ with $r := \operatorname{Re} z \geq R$ and $|\operatorname{Im} E_\kappa(z)| < S$,

$$|\operatorname{Re} E_\kappa(z)| > |E_\kappa(z)| - S \geq \exp(r) - |\kappa| - S \geq \exp(r-1) + 1.$$

If also $\operatorname{Re} E_\kappa(z) > \log\delta$, then in particular $\operatorname{Re}\big(E_\kappa(z)\big) - 1 > \exp(r-1)$.

Now suppose that there was a $z \in U$ with $r = \operatorname{Re}(z) \geq R$. It then follows by induction that

$$\operatorname{Re}\big(E_\kappa^n(z)\big) - 1 > \exp^n(r-1),$$

which is a contradiction since points in a Siegel disk do not escape to ∞. ■

6.1.3 Corollary (Diophantine Siegel Disks)
Suppose that θ is diophantine (or, more generally, $\theta \in \mathcal{H}$). Then the Siegel disk $E_{\kappa(\theta)}$ contains the singular value on its boundary. ■

Let us now shortly explain how one constructs bounded Siegel disks in the exponential family. (I am indebted to Lukas Geyer who introduced me to this method.)

6.1.4 Theorem (Bounded Siegel Disks)
There exists $\theta \in \mathbb{R}/\mathbb{Z}$ such that $E_{\kappa(\theta)}$ has a bounded Siegel disk. The boundary of this Siegel disk is a quasicircle.

Sketch of proof. Let θ be diophantine and let $\kappa := \kappa(\theta)$. Then E_κ has a Siegel disk U on which E_κ is conjugate to an irrational rotation by a biholomorphic map

$$\phi : \mathbb{D} \to U.$$

Let $K := \phi(\overline{\mathbb{D}_r})$ for some $r < 1$. Then, by the Riemann mapping theorem, there is a biholomorphic map

$$\Phi : \mathbb{C} \setminus K \; to \mathbb{C} \setminus \overline{\mathbb{D}}$$

which can be extended to a quasiconformal homeomorphism $\Phi : \mathbb{C} \to \mathbb{C}$. Consider the maps

$$g_\lambda : \mathbb{C} \to \mathbb{C}, z \mapsto \lambda \cdot \Phi(E_\kappa(\Phi^{-1}(z))),$$

where $\lambda \in \partial\mathbb{D}$. By an unpublished theorem of Herman [43], there is some $\lambda \in \partial\mathbb{D}$ for which the map $g_\lambda|_{\partial\mathbb{D}}$ is quasisymmetrically but not C^2 conjugate to a rotation $R_{\theta'} : z \mapsto \theta' z$. (In fact, the set of such λ is dense.)

We can extend this quasisymmetrical conjugacy to a quasiconformal homeomorphism $\psi : \overline{\mathbb{D}} \to \overline{\mathbb{D}}$ and modify g_λ on $\mathbb{D}$ to the map

$$\tilde{g}_\lambda : z \mapsto \Big\{ g_\lambda(z) : z \notin \mathbb{D} \psi(R_\theta(\psi^{-1}(z))) : z \in \mathbb{D}.$$

Taking pullbacks of the conformal structure given to $\mathbb{D}$ by the conformal map $g_\lambda|_{\mathbb{C}\setminus\mathbb{D}}$, we obtain an invariant conformal structure for $\tilde{g}_\lambda$. Using the measurable Riemann mapping theorem, we see that $\tilde{g}_\lambda$ is conjugate to an entire function, which is easily seen to be an exponential map. ■

6.2 The Dimension Paradox and Other Topological Anomalies

McMullen [53] showed that for any exponential map the set of escaping points has Hausdorff dimension two. In fact, he showed that the set of points which escape in a sector $|\operatorname{Im} z| < \operatorname{Re} z$ has Hausdorff dimension two. The idea of the proof is the following: consider some rectangular box B of size of height 2π. Then the image of B under E_κ will be a large (slit) annulus, and we can fill about a quarter of this annulus with similar boxes, whose preimages will correspondingly fill about a quarter of B. Repeating this process, one obtains a nested intersection of compact sets which consists of points escaping in the sector. Simple estimates on the Hausdorff dimension then yield the desired result. Note that, by contrast, the area of I is zero by a general result of Eremenko and Lyubich [31, Theorem 7]. In summary:

6.2.1 Theorem (Hausdorff Dimension and Area of I [31, 53])
Let E_κ be an exponential map. Then $\dim_H(I(E_\kappa)) = 2$ *and* $\text{area}(I(E_\kappa)) = 0$. □

In [45], Karpinska showed the following stunning result for maps with an attracting fixed point. The result was later generalized to all exponential maps by Schleicher and Zimmer [77].

6.2.2 Theorem (The Dimension Paradox [45, 77])
Let E_κ be an exponential map. Then the set

$$I(E_\kappa) \setminus \{g_{\underline{s}}(0) : \underline{s} \text{ is fast}\}$$

has Hausdorff dimension one. □

In other word, all the Hausdorff dimension of $I(E_\kappa)$ lies in the escaping endpoints of E_κ: the rays with their endpoints removed only have dimension one. This is highly surprising because every element of the dimension two set of endpoints is connected to infinity by an entire ray. This indicates just how complicated the structure of I in the plane is.

The core of the argument rests on Karpinska's observation that the set of parameters which escape in a parabola of the form $|imz| < (\operatorname{Re} z)^{\frac{1}{p}}$ has Hausdorff dimension less than $1 + \frac{1}{p}$. It is easy to see from the construction of external rays that points which are not endpoints eventually escape in every one of these parabolas, which proves the theorem.

Note also that by McMullen's result even the set of endpoints which escape in some sector, and in particular the set of endpoints $g_{\underline{s}}(0)$ where $\underline{s}$ has positive minimal potential, has Hausdorff Dimension two. As far as I know, nothing is known about the dimension of the complementary sets of endpoints.

McMullen's result about the Hausdorff dimension of I has been transferred to the parameter plane by Qiu [62].

6.2.3 Theorem (Hausdorff Dimension of Escaping Parameters)
The set $\{\kappa \in \mathbb{C} : \kappa \in I(E_\kappa)\}$ has Hausdorff dimension two. □

Again, Qiu in fact considers parameters whose singular orbits escape within some sector. It seems likely that Karpinska's result can also be carried over to the parameter plane. However, to my knowledge this has not yet been carefully carried out.

Finally, let us mention another result illustrating the complicated topological structure of the set of endpoints.

6.2.4 Theorem (∞ is an Explosion Point [52])
Let κ be a parameter for which E_κ has an attracting fixed point. Then the set E of all landing points of external rays of E_κ is totally disconnected, but $E \cup \{\infty\}$ is connected. □

6.3 Ergodic Theory

In [48] and [63], Lyubich and Rees independently showed that the map exp is not recurrent in the following sense. Recall that the ω-limit set of a point z under E_κ is defined to be

$$\omega(z) = \bigcap_{n_0 \in \mathbb{N}} \overline{\bigcup_{n \geq n_0} E_\kappa^n(z)}.$$

6.3.1 Theorem (Behavior of Typical Orbits [48, 63])
Let $\kappa = 0$. Then for almost every z,

$$\omega(z) = \mathcal{P}(E_\kappa).$$

□

M. Hemke has generalized this theorem to parameters whose singular orbit escapes within a sector [40] and recently to larger classes of parameters and other entire functions [41].

Lyubich [50, 49] proved that exp is not ergodic with respect to Lebesgue measure.

Recent results of Urbanski and Zdunik, however, suggest that the Julia set may not be the correct object to ask questions of ergodic theory about. In [82], they consider, for an attracting exponential map E_κ, the *recurrent Julia set*

$$J_r(E_\kappa) := J(E_\kappa) \setminus I(E_\kappa).$$

They show that this set has Hausdorff dimension less than two and construct an ergodic measure on this set. In a sequel [83], they prove that the Hausdorff dimension of J_r varies real analytically with κ.

Urbanski and Zdunik also recently announced results which treat escaping parameters for which the singular orbit escapes in a sector in the same spirit [86]. In particular, they prove that for such a parameter the set J_r of orbits with

$$\omega(z) \neq \mathcal{P}(E_\kappa)$$

has Hausdorff dimension less than two.

6.4 Beyond the Exponential Family

Investigations in the same spirit as for the exponential family have been carried out in several other families of transcendental functions with finitely many singular values. Let us note here in particular the family $z \mapsto \lambda z e^z$, which is of interest as a degenerate case of the complex standard family [33], and the cosine family $z \mapsto ae^z + be^{-z}$ (where $a, b \in \mathbb{C}^*$ and $a \cdot b \notin \mathbb{R}$) [71].

Let us focus on the latter family. It was shown by McMullen that the set of escaping points in the cosine family has positive Lebesgue measure. In fact, it was recently shown by Schubert [78] that for the map sin, the measure of the set of nonescaping points in every period strip is finite. (This question appears in [56, Page 64].) The set of escaping points for cosine maps was described by Rottenfußer [71] in the same way as for exponential maps.

However, there is a major difference to the exponential family, caused by the absence of asymptotic values in the cosine family. Recall that the existence of an asymptotic value in the Julia set leads to rigidity phenomena on the set of escaping points and nonlanding external rays in the exponential family. For the cosine family, however, these proofs are no longer valid. In fact, it is possible to show — using a hyperbolic contraction argument — that every external ray of a critically preperiodic cosine map lands. (This is unpublished work by Schleicher [72].) This leads to the following curious phenomenon: The Julia set of such a map is the entire plane, and every point of the Julia set is either on a ray or the landing point of an external ray. However, the set of rays still has Hausdorff dimension one.

In view of these facts, it seems reasonable to ask whether all maps in the cosine family for which both critical values do not escape are conjugate on their sets of escaping points, or whether there is some other obstruction. Note that the obstruction which prevents exponential maps on the same parameter ray to be conjugate is still present in the cosine family.

Apart from further investigations into the exponential and cosine families, we believe that the time is ripe also to depart from the considerations of specialized families and attempt to achieve results which are valid for larger classes of transcendental functions.

Chapter 7

Open Questions

In this chapter, we will discuss some of the questions about exponential dynamics which are left open. This is by no means an exhaustive list; rather we will concentrate largely on those questions which arise naturally from the considerations in this thesis.

7.1 Behavior of External Rays

We have seen that for many parameters with $\kappa \in J(E_\kappa)$, there exist external rays which do not land, and even accumulate on themselves. It seems reasonable to believe that this is always true in this case.

7.1.1 Question (Nonlanding Rays)
If $\kappa \in J(E_\kappa)$, does there always exist an exponentially bounded address $\underline{s}$ such that $g_{\underline{s}}^\kappa$ does not land? Can $\underline{s}$ always be chosen so that $g_{\underline{s}}$ accumulates on itself? Are there always uncountably many such addresses?

In fact, one can ask whether the singular value is *always* contained in the accumulation set of an external ray:

7.1.2 Question (Rays at the Singular Value)
Suppose that $\kappa \in J(E_\kappa)$. Does there always exist an $\underline{s}$ such that $\kappa \in \overline{g_{\underline{s}}}$?

Among the problem that may occur when trying to answer such question is that the possible limiting behavior of external rays is not well-understood. Even for quadratic polynomials, the possibility that an external ray can accumulate on the entire Julia set has not been excluded. In our setting, we can thus ask:

7.1.3 Question (Rays Accumulating on the Plane)
Can the accumulation set of an external ray be the entire plane?

We can ask an even stronger question: can a sequence rays whose addresses converge to an address which is not exponentially bounded have a finite accumulation point? We know that this does not happen in parameter space; however, the answer in the dynamical plane does not seem clear to us.

7.1.4 Question (Accumulating Rays)
If $\underline{s}^n$ is a sequence of addresses with $\underline{s}^n \to \infty$, is it true that

$$z_n \to \infty$$

whenever $z_n \in g_{\underline{s}^n}$ for all n?

More generally, is this true whenever $(\underline{s}^n)$ converges to an address which is not exponentially bounded?

Finally, let us depart from questions concerning single rays and consider the escaping sets of exponential maps in their entirety. We have already formulated the conjecture that two exponential maps whose singular value lies in the Julia set are never conjugate on their escaping sets by an order-preserving conjugacy. We can ask whether the map is already determined by the topology of this set. We say that a homeomorphism between $I(E_{\kappa_1})$ and $I(E_{\kappa_2})$ is *natural* if it preserves the addresses of external rays.

7.1.5 Question (Natural Homeomorphisms)
If $I(E_{\kappa_1})$ and $I(E_{\kappa_2})$ are naturally homeomorphic, are E_{κ_1} and E_{κ_2} conjugate on their sets of escaping points?

The answer to this question can easily seen to be "yes" when κ_1 and κ_2 are Misiurewicz-parameters, using the construction of nonlanding external rays 3.8.4. With some more care one can also do this when κ_1 and κ_2 are escaping with $\operatorname{addr}(\kappa_1) \neq \operatorname{addr}(\kappa_2)$. It seems thus interesting to investigate this question in the case where $\operatorname{addr}(\kappa_1) = \operatorname{addr}(\kappa_2)$; for example if $\kappa_1, \kappa_2 \in (-1, \infty)$.

We have given a model for the topological dynamics of attracting and parabolic exponential maps on their sets of escaping points. As we have seen, the situation becomes much more complicated when the singular value moves into the Julia set. Nevertheless, Misiurewicz parameters are uniquely determined by their combinatorics. One would thus hope that their topological dynamics can also be completely understood in terms of their combinatorics.

7.1.6 Question (Topological Dynamics of Misiurewicz Maps)
Let κ be a Misiurewicz parameter. Can one construct a model for the topological dynamics of $E_\kappa|_{I(E_\kappa)}$ in terms of $\operatorname{addr}(\kappa)$*?*

7.2 Parameter Space

The main open question about exponential parameter is, of course, the density of hyperbolicity, or equivalently the nonexistence of queer components. However, there are several simpler questions one can ask.

7.2.1 Conjecture (Bifurcation Locus Connected)
The bifurcation locus $\mathcal{B}$ is connected.

Theorem 1.8 shows that the nonhyperbolic locus — which is conjecturally equal to $\mathcal{B}$ — is connected. Thus the question remains whether a queer component can disconnect the plane. In fact, it seems reasonable to hope that one can prove the following.

7.2.2 Conjecture (Queer Components Bounded)
Every queer component is bounded.

We have already seen that ∞ cannot be accessible from a queer component. A promising way to prove Conjecture 7.2.2 would be to show the following strengthening of the Squeezing Lemma. Let ϕ be a curve which does not intersect closures of hyperbolic components and which accumulates at ∞. Define $\mathrm{ADDR}(\phi)$ in analogy to Section 3.7. Our proof of the Squeezing Lemma shows that ADDR can consist only of exponentially bounded external addresses.

7.2.3 Conjecture (Strong Squeezing Lemma)
Let $\phi : [0, \infty) \to \mathbb{C}$ be a curve which does not contain attracting or indifferent parameters, and suppose that $\limsup |\phi(t)| = \infty$. If $\underline{s} \in \mathrm{ADDR}(\phi)$, then ϕ contains an end piece of the parameter ray $\mathcal{G}_{\underline{s}}$.

Finally, recall from Section 5.14 that the question whether repelling periodic points are always landing points can be phrased as a conjecture about parameter space.

7.2.4 Conjecture (Wakes Consist of Subwakes)
Let W be a hyperbolic component. Then any nonescaping parameter in $\mathcal{W}(W) \setminus \overline{W}$ lies in a subwake of W.

In particular, this would imply the following.

7.2.5 Conjecture (Indifferent Parameters are Landing Points)
Every indifferent parameter is the landing point of a parameter ray.

In view of the construction of Section 5.4, the proof of Conjecture 7.2.4 seems to require some sort of "Yoccoz Inequality" which bounds the size of, say, an initial segment of the central internal ray of bifurcating hyperbolic components. Similarly, a proof of the Strong Squeezing Lemma would probably need to control the behavior of internal rays of hyperbolic components. It seems that proofs of these facts might also yield a quantitative proof of Theorem 1.8.

Finally, Theorem 5.12.2 suggests a "pinched Cantor Bouquet" model for the bifurcation locus of exponential maps.

7.2.6 Question (Pinched Cantor Bouquet)
Is the map $\mathcal{G} : X \to \{\kappa : \kappa \in I(E_\kappa)\}$ a homeomorphism? Does $\mathcal{G}$ extend to a continuous map

$$\mathcal{G} : \overline{X} \to \mathcal{B}?$$

It seems that a positive answer to this question would imply density of hyperbolicity.

7.3 Other Questions

As mentioned in Section 6.1, it is not known whether Bryuno's condition is necessary for linearizability. This is true for unicritical polynomials, and it seems likely that the same holds in the exponential case.

7.3.1 Conjecture (Necessary Condition for Linearizability)
An exponential map with an irrationally indifferent periodic orbit has a Siegel disk if and only if its rotation number satisfies Bryuno's condition.

Another question about Siegel disks concerns the (fixed) Siegel disk whose rotation number is the golden mean. In the case of quadratic polynomials, it is known [61] that every Siegel disk whose rotation number is of bounded type has locally connected Julia set. Several further results about Siegel disks of this type, including self-similarity near the critical value, were shown by McMullen [55]. In fact, Buff and Henriksen [14] were able to show that the Siegel disk contains some triangle based at the critical value. Computer pictures suggest that — in the parametrization $z \mapsto \lambda z(1-z)$ — the critical value $\frac{1}{2}$ is the closest point to 0 on the Siegel disk boundary. This also seems to be true in the exponential case (cf Figure 6.1). The known results, however, are even less satisfying. Indeed, the following question, asked by Baker and Dominguez [4], is still open.

7.3.2 Conjecture (Accessibility of Infinity)
∞ is accessible in the Siegel disk of $E_{\kappa(\theta)}$, where θ is the golden mean.

(Baker and Dominguez actually ask whether *any* such θ exists; however the golden mean is the most likely candidate.)

As a final question, recall from Theorem 4.3.1 that there is no topological renormalization for exponential maps. However, there seem to be similarity features in the parameter space of exponential maps as witnessed e.g. by the combinatorial tuning formula. So it is natural to ask whether some — different — notion of renormalization does not exist also in the exponential family. Let us ask this (vaguely) in a special case.

7.3.3 Question (Renormalization)
Let W be a hyperbolic component, and let $\Phi : \mathbb{H} \to \mathcal{W}$ be its preferred parametrization. Let $W' := W_\infty$ be the unique period 1 component, and let $\Psi : \mathbb{H} \to W'$ be its parametrization, as given by Lemma 5.2.6. Then

$$\mathcal{R} := \Psi \circ \Phi^{-1} : \overline{W} \to \overline{W'}$$

maps each parameter of $\overline{W}$ to a period 1 parameter with the same multiplier. Is there some analytic way to construct the parameter $\mathcal{R}(\kappa)$ from the dynamics of κ in such a way that dynamical features such as linearizability etc. are preserved?

Bibliography

[1] Jan M. Aarts and Lex G. Oversteegen, *The geometry of Julia sets*, Trans. Amer. Math. Soc. **338** (1993), no. 2, 897–918.

[2] Lars V. Ahlfors, *Conformal invariants: topics in geometric function theory*, McGraw-Hill Book Co., New York, 1973, McGraw-Hill Series in Higher Mathematics.

[3] I. N. Baker, *The domains of normality of an entire function*, Ann. Acad. Sci. Fenn. Ser. A I Math. **1** (1975), no. 2, 277–283.

[4] I. N. Baker and P. Domínguez, *Boundaries of unbounded Fatou components of entire functions*, Ann. Acad. Sci. Fenn. Math. **24** (1999), no. 2, 437–464.

[5] I. N. Baker and P. J. Rippon, *Iteration of exponential functions*, Ann. Acad. Sci. Fenn. Ser. A I Math. **9** (1984), 49–77.

[6] I. N. Baker and J. Weinreich, *Boundaries which arise in the dynamics of entire functions*, Rev. Roumaine Math. Pures Appl. **36** (1991), no. 7-8, 413–420, Analyse complexe (Bucharest, 1989).

[7] Alan F. Beardon, *Iteration of rational functions*, Graduate Texts in Mathematics, vol. 132, Springer-Verlag, New York, 1991. MR 92j:30026

[8] Walter Bergweiler, *Iteration of meromorphic functions*, Bull. Amer. Math. Soc. (N.S.) **29** (1993), no. 2, 151–188.

[9] Ranjit Bhattacharjee and Robert L. Devaney, *Tying hairs for structurally stable exponentials*, Ergodic Theory Dynam. Systems **20** (2000), no. 6, 1603–1617.

[10] Ranjit Bhattacharjee, Robert L. Devaney, R. E. Lee Deville, Krešimir Josić, and Monica Moreno-Rocha, *Accessible points in the Julia sets of stable exponentials*, Discrete Contin. Dyn. Syst. Ser. B **1** (2001), no. 3, 299–318.

[11] Clara Bodelón, Robert L. Devaney, Michael Hayes, Gareth Roberts, Lisa R. Goldberg, and John H. Hubbard, *Hairs for the complex exponential family*, Internat. J. Bifur. Chaos Appl. Sci. Engrg. **9** (1999), no. 8, 1517–1534.

[12] ———, *Dynamical convergence of polynomials to the exponential*, J. Differ. Equations Appl. **6** (2000), no. 3, 275–307.

[13] D. A. Brannan and W. K. Hayman, *Research problems in complex analysis*, Bull. London Math. Soc. **21** (1989), no. 1, 1–35.

[14] Xavier Buff and Christian Henriksen, *Scaling ratios and triangles in Siegel disks*, Math. Res. Lett. **6** (1999), no. 3-4, 293–305.

[15] Lennart Carleson and Theodore W. Gamelin, *Complex dynamics*, Universitext: Tracts in Mathematics, Springer-Verlag, New York, 1993.

[16] Stephen B. Curry, *One-dimensional nonseparating plane continua with disjoint ϵ-dense subcontinua*, Topology Appl. **39** (1991), no. 2, 145–151.

[17] Robert L. Devaney, *Structural instability of* $\exp(z)$, Proc. Amer. Math. Soc. **94** (1985), no. 3, 545–548.

[18] ———, *Knaster-like continua and complex dynamics*, Ergodic Theory Dynam. Systems **13** (1993), no. 4, 627–634.

[19] Robert L. Devaney, Núria Fagella, and Xavier Jarque, *Hyperbolic components of the complex exponential family*, Fund. Math. **174** (2002), no. 3, 193–215.

[20] Robert L. Devaney and Lisa R. Goldberg, *Uniformization of attracting basins for exponential maps*, Duke Math. J. **55** (1987), no. 2, 253–266.

[21] Robert L. Devaney, Lisa R. Goldberg, and John H. Hubbard, *A dynamical approximation to the exponential map by polynomials*, Preprint, 1986.

[22] Robert L. Devaney and Xavier Jarque, *Indecomposable continua in exponential dynamics*, Conform. Geom. Dyn. **6** (2002), 1–12.

[23] Robert L. Devaney and Michał Krych, *Dynamics of* $\exp(z)$, Ergodic Theory Dynam. Systems **4** (1984), no. 1, 35–52.

[24] Adrien Douady, *Disques de Siegel et anneaux de Herman*, Astérisque (1987), no. 152-153, 4, 151–172 (1988), Séminaire Bourbaki, Vol. 1986/87.

[25] ———, *Descriptions of compact sets in* $\mathbf{C}$, Topological methods in modern mathematics (Stony Brook, NY, 1991), Publish or Perish, Houston, TX, 1993, pp. 429–465.

[26] Adrien Douady and Lisa R. Goldberg, *The nonconjugacy of certain exponential functions*, Holomorphic functions and moduli, Vol. I (Berkeley, CA, 1986), Math. Sci. Res. Inst. Publ., vol. 10, Springer, New York, 1988, pp. 1–7.

[27] Adrien Douady and John Hubbard, *Etude dynamique des polynômes complexes*, Prépublications mathémathiques d'Orsay (1984 / 1985), no. 2/4.

[28] A. È. Erëmenko, *On the iteration of entire functions*, Dynamical systems and ergodic theory (Warsaw, 1986), Banach Center Publ., vol. 23, PWN, Warsaw, 1989, pp. 339–345.

[29] A. È. Erëmenko and G. M. Levin, *Periodic points of polynomials*, Ukrain. Mat. Zh. **41** (1989), no. 11, 1467–1471, 1581.

[30] A. È. Erëmenko and M. Yu. Lyubich, *Структурная устоичивость в некоторых семейтвах целых функции (structural stability in some families of entire functions)*, Preprint, Physico-Technical Institute of Low-Temperature Physics Kharkov, 1984.

[31] ______, *Dynamical properties of some classes of entire functions*, Ann. Inst. Fourier (Grenoble) **42** (1992), no. 4, 989–1020.

[32] Leonhardt Euler, *De formulis exponentialibus replicatis*, Acta Acad. Petropolitanae **1** (1777), 38–60.

[33] Núria Fagella, *Limiting dynamics for the complex standard family*, Internat. J. Bifur. Chaos Appl. Sci. Engrg. **5** (1995), no. 3, 673–699.

[34] Pierre Fatou, *Sur l'itération des fonctions transcendantes entières*, Acta Math. **47** (1926), 337–370.

[35] Markus Förster, *Parameter rays for the exponential family*, Diplomarbeit, Techn. Univ. München, 2003.

[36] Otto Forster, *Riemannsche Flächen*, Springer-Verlag, Berlin, 1977, Heidelberger Taschenbücher, Band 184.

[37] Lukas Geyer, *Siegel discs, Herman rings and the Arnold family*, Trans. Amer. Math. Soc. **353** (2001), no. 9, 3661–3683 (electronic).

[38] Lisa R. Goldberg and John Milnor, *Fixed points of polynomial maps. II. Fixed point portraits*, Ann. Sci. École Norm. Sup. (4) **26** (1993), no. 1, 51–98.

[39] Peter Haïssinsky, *Chirurgie parabolique*, C. R. Acad. Sci. Paris Sér. I Math. **327** (1998), no. 2, 195–198.

[40] Martin Hemke, *Verhalten von typischen Orbits der Exponentialfamilie*, Diplomarbeit, CAU Kiel, 2002.

[41] ______, *Measurable dynamics of transcendental entire functions on their Julia sets*, Manuscript, 2003.

[42] Michael-R. Herman, *Are there critical points on the boundaries of singular domains?*, Comm. Math. Phys. **99** (1985), no. 4, 593–612.

[43] ———, *Conjugaison quasi symètrique des difféomorphismes du cercle à des rotations et applications aux disques singuliers des Siegel*, Manuscript, ca. 1986, http://www.kusm.kyoto-u.ac.jp/~mitsu/Herman/.

[44] Bogusława Karpińska, *Area and Hausdorff dimension of the set of accessible points of the Julia sets of λe^z and $\lambda \sin z$*, Fund. Math. **159** (1999), no. 3, 269–287.

[45] Boguslawa Karpińska, *Hausdorff dimension of the hairs without endpoints for $\lambda \exp z$*, C. R. Acad. Sci. Paris Sér. I Math. **328** (1999), no. 11, 1039–1044.

[46] Steven G. Krantz, *Function theory of several complex variables*, John Wiley & Sons Inc., New York, 1982.

[47] Eike Lau and Dierk Schleicher, *Internal addresses in the mandelbrot set and irreducibility of polynomials*, Preprint #1994/14, Institute for Mathematical Sciences, SUNY Stony Brook, 1994.

[48] M. Yu. Lyubich, *Generic behavior of trajectories of the exponential function*, Uspekhi Mat. Nauk **41** (1986), no. 2(248), 199–200.

[49] ———, *Measurable dynamics of an exponential*, Dokl. Akad. Nauk SSSR **292** (1987), no. 6, 1301–1304.

[50] ———, *The measurable dynamics of the exponential*, Siberian Math. J. **28** (1987), no. 5, 780–797.

[51] R. Mañé, P. Sad, and D. Sullivan, *On the dynamics of rational maps*, Ann. Sci. École Norm. Sup. (4) **16** (1983), no. 2, 193–217.

[52] John C. Mayer, *An explosion point for the set of endpoints of the Julia set of $\lambda \exp(z)$*, Ergodic Theory Dynam. Systems **10** (1990), no. 1, 177–183.

[53] Curt McMullen, *Area and Hausdorff dimension of Julia sets of entire functions*, Trans. Amer. Math. Soc. **300** (1987), no. 1, 329–342.

[54] Curtis T. McMullen, *Complex dynamics and renormalization*, Annals of Mathematics Studies, vol. 135, Princeton University Press, Princeton, NJ, 1994.

[55] ———, *Self-similarity of Siegel disks and Hausdorff dimension of Julia sets*, Acta Math. **180** (1998), no. 2, 247–292.

[56] John Milnor, *Dynamics in one complex variable*, Friedr. Vieweg & Sohn, Braunschweig, 1999.

[57] ———, *On rational maps with two critical points*, Experiment. Math. **9** (2000), no. 4, 481–522.

[58] ———, *Periodic orbits, externals rays and the Mandelbrot set: an expository account*, Astérisque (2000), no. 261, xiii, 277–333, Géométrie complexe et systèmes dynamiques (Orsay, 1995).

[59] Michał Misiurewicz, *On iterates of e^z*, Ergodic Theory Dynamical Systems **1** (1981), no. 1, 103–106.

[60] Ricardo Pérez-Marco, *Fixed points and circle maps*, Acta Math. **179** (1997), no. 2, 243–294.

[61] Carsten Lunde Petersen, *Local connectivity of some Julia sets containing a circle with an irrational rotation*, Acta Math. **177** (1996), no. 2, 163–224.

[62] Weiyuan Qiu, *Hausdorff dimension of the M-set of $\lambda \exp(z)$*, Acta Math. Sinica (N.S.) **10** (1994), no. 4, 362–368, A Chinese summary appears in Acta Math. Sinica **38** (1995), no. 5, 719.

[63] Mary Rees, *The exponential map is not recurrent*, Math. Z. **191** (1986), no. 4, 593–598.

[64] Lasse Rempe, *An answer to a question of Herman, Baker and Rippon*, Berichtsreihe des Math. Sem. der CAU Kiel 02/7, 2002.

[65] ———, *A landing theorem for periodic rays of exponential maps*, Manuscript, 2002.

[66] ———, *On the landing of periodic rays for exponential maps*, Berichtsreihe des Math. Sem. der CAU Kiel 02/8, 2002.

[67] ———, *Topology of the set of escaping points for exponential maps*, Manuscript, 2002.

[68] Lasse Rempe and Dierk Schleicher, *Bifurcations in the space of exponential maps*, Manuscript, 2001.

[69] Philip J. Rippon, *On the boundaries of certain Siegel discs*, C. R. Acad. Sci. Paris Sér. I Math. **319** (1994), no. 8, 821–826.

[70] Mónica Moreno Rocha, *Existence of indecomposable continua for unjstable exponentials*, Manuscript, 2002.

[71] Günter Rottenfußer, *Entkommende Punkte einer Familie ganzer transzendenter Funktionen*, Diplomarbeit, TU München, 2002.

[72] Dierk Schleicher, *Dynamics of cosine maps*, Talk at the conference "Holomorphic Iteration, Non-Uniform Hyperbolicity"at the Banach Center, Warsaw, May 2002.

[73] Dierk Schleicher, *On the dynamics of iterated exponential maps*, Habilitation thesis, TU München, 1999.

[74] ———, *Attracting dynamics of exponential maps*, Ann. Acad. Sci. Fenn **28** (2003), 3–34.

[75] Dierk Schleicher and Mitsuhiro Shishikura, *Topology and combinatorics of postsingularly finite exponential maps*, Manuscript, 1999.

[76] Dierk Schleicher and Johannes Zimmer, *Periodic points and dynamic rays of exponential maps*, Ann. Acad. Sci. Fenn, to appear.

[77] ———, *Escaping points of exponential maps*, J. London Math. Soc. (2) **67** (2003), no. 2, 380–400.

[78] Hendrik Schubert, *Das Maß nichtentkommender Punkte der Sinusfunktion*, Diplomarbeit, CAU Kiel, 2003.

[79] Mitsuhiro Shishikura, *The boundary of the Mandelbrot set has Hausdorff dimension two*, Astérisque (1994), no. 222, 7, 389–405, Complex analytic methods in dynamical systems (Rio de Janeiro, 1992).

[80] Norbert Steinmetz, *Rational iteration*, de Gruyter Studies in Mathematics, vol. 16, Walter de Gruyter & Co., Berlin, 1993.

[81] Dennis Sullivan, *Quasiconformal homeomorphisms and dynamics. I. Solution of the Fatou-Julia problem on wandering domains*, Ann. of Math. (2) **122** (1985), no. 3, 401–418.

[82] Mariusz Urbanski and Anna Zdunik, *The finer geometry and dynamics of the hyperbolic exponential family*, Preprint, 2001, To appear in Michigan Math. J.

[83] ———, *Real analyticity of Hausdorff dimension of finer Julia sets of exponential family*, Preprint, 2002, To appear in Ergodic Theory Dynamical Systems.

[84] M. Viana da Silva, *The differentiability of the hairs of* $\exp(Z)$, Proc. Amer. Math. Soc. **103** (1988), no. 4, 1179–1184.

[85] Zhuan Ye, *Structural instability of exponential functions*, Trans. Amer. Math. Soc. **344** (1994), no. 1, 379–389.

[86] Anna Zdunik, *Conformal measures and Hausdorff measures for the exponential family*, Talk at the Holomorphic Dynamics Meeting at the University of Warwick, March 2003.

Referent/in: ______________________________

Koreferent/in: ______________________________

Tag der mündlichen Prüfung: ______________________

Zum Druck genehmigt: Kiel, ______________________

Der Dekan