

EXPONENTIAL THURSTON MAPS AND LIMITS OF QUADRATIC DIFFERENTIALS

JOHN HUBBARD, DIERK SCHLEICHER, AND MITSUHIRO SHISHIKURA

CONTENTS

1. Introduction	78
Organization of the paper	79
2. The characterization theorem	80
2.1. Definitions and statement of the main theorem	80
2.2. Classification of postsingularly finite exponential maps	80
3. Iteration in Teichmüller space	82
3.1. The Teichmüller space of a topological exponential map	83
3.2. The Teichmüller metric and its dual	84
3.3. Examples of quadratic differentials	85
3.4. Proof of the Main Theorem 2.4	86
4. The decomposition theorem	89
4.1. The thick-thin decomposition theorem for quadratic differentials	90
4.2. Quadratic differentials on annuli	91
4.3. The component of a quadratic differential adapted to a short closed geodesic	93
4.4. The component of a quadratic differential adapted to the thick part	96
5. Limits of quadratic differentials	98
5.1. Limit models	98
5.2. Thin parts	101
5.3. Thick parts	103
6. Push-forward of quadratic differentials	106
Appendix A. Some general results on Riemann surfaces	111
A.1. Collars on hyperbolic Riemann surfaces	112
A.2. Mass per modulus is convex	113
Acknowledgements	115
References	115

Received by the editors March 28, 2006.

2000 *Mathematics Subject Classification*. Primary 30F30; Secondary 30F60, 32G15, 37F20, 37F30.

Key words and phrases. Quadratic differential, decomposition, limit model, iteration, exponential map, classification.

©2008 American Mathematical Society
Reverts to public domain 28 years from publication

1. INTRODUCTION

In the theory of iterated rational maps, the easiest maps to understand are postcritically finite: maps whose critical orbits are all periodic or preperiodic. These maps are also the most important maps for understanding the combinatorial structure of parameter spaces of rational maps.

We know a lot about postcritically finite rational maps. The main result is a theorem of Thurston [DH] which gives a purely topological criterion for whether or not a given postcritically finite branched covering map $f : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ of the two-sphere to itself is equivalent (in a precise sense, given in Definition 2.2) to a rational map. Either a postcritically finite branched cover is equivalent to an essentially unique rational function or there is a “Thurston obstruction”. Such an obstruction is a collection of simple closed curves such that a certain associated matrix has leading eigenvalue at least 1.

Thurston’s theorem has two limitations. One is that the criterion is not easy to check, even though it is purely combinatorial-topological. More relevant to the present paper is the fact that the degree of the map enters in an essential way into the proof; the proof just does not go through for transcendental functions.

The simplest non-trivial transcendental maps are exponential maps $z \mapsto E_\lambda(z) = \lambda \exp(z)$ with $\lambda \in \mathbb{C}^* := \mathbb{C} \setminus \{0\}$. These have been investigated by many people; see for example [BR, EL, DGH, S1, RS1] and the references in these papers. Exponential maps have no critical values, but the unique singular value 0 plays an analogous role.

Postsingularly finite exponential maps are those for which the orbit of 0 is preperiodic. There are countably many such parameters. Bergweiler (unpublished) has used value distribution theory to estimate their density with respect to $|\lambda|$. There are no exponential maps with periodic singular orbits (but there are countably many hyperbolic components in exponential parameter space; these are completely classified in [S2]).

A *topological exponential map* is a covering map $f : \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{0, \infty\}$; this bears the same relation to exponentials as branched coverings $\mathbb{S}^2 \rightarrow \mathbb{S}^2$ bear to rational functions. Our Main Theorem 2.4 is the analog of Thurston’s characterization theorem: we show that a postsingularly finite topological exponential map is either equivalent to a holomorphic exponential map or it admits a (degenerate) Levy cycle. As with Thurston’s result, the complete classification of postsingularly finite maps is a separate step; we only state the result here and refer to [LSV] for details.

In the mid-1980’s, [DGH] gave a conjectural description for postsingularly finite exponential maps in analogy to and as a limit of results for polynomials $\lambda(1 + z/d)^d$ with a single finite critical point as $d \rightarrow \infty$. The theory of spiders [HS] was developed in the process. Our results confirm the conjecture in [DGH].

We use the same machinery for our proof as Thurston: given a postsingularly finite topological exponential map $f : \mathbb{S}^2 \rightarrow \mathbb{S}^2$, we set up a *Thurston map* $\sigma_f : \mathcal{T}_f \rightarrow \mathcal{T}_f$ in an appropriate Teichmüller space \mathcal{T}_f and show that either σ_f has a fixed point, in which case the topological exponential map is equivalent to a holomorphic exponential map, or the iteration of σ_f diverges in Teichmüller space, and there is a degenerate Levy cycle.

As mentioned above, the proof given in [DH] for Thurston’s result on rational maps depends in an essential way on the fact that rational maps have finite degree;

it does not work for exponentials. That paper shows that, depending on the initial point of the iteration of σ_f , there is a subset of Teichmüller space (with compact projection to moduli space) such that as soon as the iteration leaves this subset, the existence of a Thurston obstruction follows. A key ingredient in this proof is an estimate about how moduli of annuli on the finitely punctured Riemann sphere increase when erasing the points in $f^{-1}(P_f) \setminus P_f$; the cardinality of this set is bounded by $d|P_f|$, which diverges when $d \rightarrow \infty$ (see the beginning of Section 3 for details).

We use a different strategy to relate the failure of convergence to the existence of a Thurston obstruction. The cotangent space $T_\tau^* \mathcal{T}_f$ to Teichmüller space at $\tau \in \mathcal{T}_f$ is a certain space $Q^1(\tau)$ of integrable meromorphic quadratic differentials on \mathbb{C} , with at most as many poles as the length of the singular orbit (plus possibly a pole at ∞). The dual of the L^1 -norm on the $Q^1(\tau)$ space defines the infinitesimal Teichmüller metric on \mathcal{T}_f (see Subsection 3.2). The analytic map $\sigma_f : \mathcal{T}_f \rightarrow \mathcal{T}_f$ is weakly contracting for this metric in the sense that $\|d\sigma_f\| = \|(d\sigma_f)^*\| < 1$. This is not surprising: all analytic maps are non-expanding; but the norm may tend to 1 as we iterate σ_f .

More precisely, if the sequence

$$\tau_0, \tau_1 = \sigma_f(\tau_0), \dots, \tau_{n+1} = \sigma_f(\tau_n), \dots$$

does not converge in \mathcal{T}_f , then there must exist $q_n \in Q^1(\tau_n)$ with $\|q_n\|_{\mathbb{C}} = 1$ such that $\lim_{n \rightarrow \infty} \|(d\sigma_f)^*(q_n)\|_{\mathbb{C}} = 1$.

In this case, the q_n cannot converge in $L^1(\overline{\mathbb{C}})$. In fact, poles must coalesce, and very fat annuli in the complement of the poles of the q_n must develop; the core curves of these annuli will present us with the needed Levy cycle.

Proving this requires understanding how the mass of degenerating quadratic differentials is distributed. We prove in Section 4 a “thick-thin” decomposition theorem which describes this distribution in considerable detail.

Although in this paper we use this decomposition only for integrable meromorphic quadratic differentials on the Riemann sphere, it is proved for integrable quadratic differentials on an arbitrary Riemann surface; moreover, the constants that appear are independent of the topology. As such it may have many other applications: Veech curves (already in progress), compactifications of moduli spaces, complex dynamics, conformal field theory, and perhaps other subjects as well.

Exponential maps are of course rather special transcendental entire maps. However, we believe that our methods should help to prove a similar result for larger classes of transcendental maps.

Organization of the paper. In Section 2, we give the main definitions, state the main theorem (Theorem 2.4) and give the resulting classification of postsingularly finite exponential maps. In Section 3 we discuss Thurston’s iteration in Teichmüller space and in particular its contracting properties, and we prove the main theorem modulo a key proposition about contraction of quadratic differentials under repeated exponential push-forwards, which will be proved in Section 6. In Section 4 we state and prove our “thick-thin” theorem for quadratic differentials. This will be used in Section 5 to provide limit models for quadratic differentials when some annuli become infinitely fat; these limit models are either integrable or they are multiples of dz^2/z^2 on $\overline{\mathbb{C}}$. In Appendix A we provide some background information about the geometry of Riemann surfaces with short hyperbolic geodesics.

2. THE CHARACTERIZATION THEOREM

2.1. Definitions and statement of the main theorem. CONVENTIONS. Let \mathbb{S}^2 be an oriented topological 2-sphere with the two distinguished points 0 and ∞ . All homeomorphisms and coverings in this paper will be understood to preserve the orientation of \mathbb{S}^2 or \mathbb{C} . We write $f^{\circ n}$ for the n -th iterate $f \circ f \circ \cdots \circ f$ of f , and $E_\lambda(z) = \lambda e^z$ for $\lambda \in \mathbb{C}^* := \mathbb{C} \setminus \{0\}$.

Definition 2.1 (Topological exponential maps). A universal cover $g: (\mathbb{S}^2 \setminus \{\infty\}) \rightarrow (\mathbb{S}^2 \setminus \{\infty, 0\})$ will be called a *topological exponential map*. It is called *postsingularly finite* if the orbit of 0 is finite, hence preperiodic. The *postsingular set* is $P_g := \bigcup_{n \geq 0} g^{\circ n}(0) \cup \{\infty\}$.

Definition 2.2 (Thurston equivalence). Two postsingularly finite topological exponential maps f and g with postsingular sets P_f and P_g are called *Thurston equivalent* if there are two homeomorphisms $\varphi_1, \varphi_2: \mathbb{S}^2 \rightarrow \mathbb{S}^2$ with $\varphi_1|_{P_f} = \varphi_2|_{P_f}$, $P_g = \varphi_1(P_f) = \varphi_2(P_f)$ and $\varphi_1(\infty) = \varphi_2(\infty) = \infty$ such that the diagram

$$(1) \quad \begin{array}{ccc} \mathbb{S}^2 \setminus \{\infty\} & \xrightarrow{\varphi_2} & \mathbb{S}^2 \setminus \{\infty\} \\ \downarrow f & & \downarrow g \\ \mathbb{S}^2 \setminus \{\infty, 0\} & \xrightarrow{\varphi_1} & \mathbb{S}^2 \setminus \{\infty, 0\} \end{array}$$

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

Since 0 is the only omitted value of f and g , the relation $\varphi_1 \circ f = g \circ \varphi_2$ on $\mathbb{S}^2 \setminus \{\infty\}$ implies $\varphi_1(0) = 0$ and hence $\varphi_2(0) = 0$.

Definition 2.3 (Essential curves and Levy cycle). Let g be a postsingularly finite exponential map. A simple closed curve $\gamma \subset \mathbb{S}^2 \setminus P_g$ is called *essential* if both connected components of $\mathbb{S}^2 \setminus \gamma$ contain at least two points of P_g . A *Levy cycle* of g is a finite sequence of disjoint essential simple closed curves $\gamma_0, \gamma_1, \dots, \gamma_{m-1}, \gamma_m = \gamma_0$ such that for $i = 0, 1, \dots, m-1$, one component γ'_i of $g^{-1}(\gamma_{i+1})$ is homotopic to γ_i relative to P_g and $g: \gamma'_i \rightarrow \gamma_{i+1}$ is a homeomorphism.

Levy cycles are preserved under Thurston equivalences. Let U_i be the components of $\mathbb{S}^2 \setminus \gamma_i$ not containing ∞ ; if all restrictions $g: \overline{U}_i \rightarrow \overline{U}_{i+1}$ are homeomorphisms, then the Levy cycle is called *degenerate*. It is easy to see that in our case, every Levy cycle is degenerate. Degenerate Levy cycles have the convenient property that one can collapse all postsingular points surrounded by each simple closed curve in all Levy cycles and thus obtain another postsingularly finite topological exponential map without Levy cycle.

Main Theorem 2.4 (Characterization of exponential maps). *A postsingularly finite topological exponential map is Thurston equivalent to a (necessarily unique) postsingularly finite holomorphic exponential map if and only if it does not admit a degenerate Levy cycle.*

2.2. Classification of postsingularly finite exponential maps. The main theorem allows us to classify postsingularly finite exponential maps completely. This has been done in [LSV]; we briefly state the main result in Corollary 2.6 below. To do this, we need to introduce dynamic rays and to state one lemma.

Let f be an entire holomorphic map. A point $z \in \mathbb{C}$ is called an *escaping point* if $f^{on}(z) \rightarrow \infty$ as $n \rightarrow \infty$. A *dynamic ray* of f is an injective continuous map $\gamma: (0, \infty] \rightarrow \overline{\mathbb{C}}$ with $\gamma(\infty) = \infty$ such that $\gamma(t)$ is an escaping point for every $t \in \mathbb{R}^+$, subject to the condition that $\gamma((0, \infty])$ is maximal with respect to inclusion (in the sense that it is not a proper subset of the image of another curve consisting of escaping points). The dynamic ray γ *lands* at a point $a \in \mathbb{C}$ if $\lim_{t \rightarrow 0} \gamma(t)$ exists and equals a .

It is shown in [SZ1] that for every exponential map E_λ , every escaping point z is either on a dynamic ray or it is the landing point of a dynamic ray; in both cases the dynamic ray is unique. For every dynamic ray γ , the f -image $\hat{\gamma}: (0, \infty] \rightarrow \overline{\mathbb{C}}$ with $\hat{\gamma}(t) = f(\gamma(t))$ for $t \in \mathbb{R}^+$ and $\hat{\gamma}(\infty) = \infty$ is contained in another dynamic ray. A dynamic ray γ is *periodic* if $\gamma(\mathbb{R}^+) = f^{ok}(\gamma(\mathbb{R}^+))$ for some $k \in \mathbb{N}$ (no $\gamma(t)$ can be periodic, but the set $\gamma(\mathbb{R}^+)$ can be); the ray γ is *preperiodic* if $f^{ol}(\gamma)$ is periodic for some positive $l \in \mathbb{N}$. The following result is shown in [SZ2, Theorem 6.4].

Lemma 2.5 (Dynamic rays landing at singular value). *For every postsingularly finite exponential map, the singular value 0 is the landing point of at least one and at most finitely many dynamic rays.*

Let γ_1 be one of the rays landing at 0, and let $\gamma_n := f^{o(n-1)}(\gamma_1)$ for $n \geq 2$. By [SZ1], the limit $v_n := \lim_{t \rightarrow \infty} \text{Im}(\gamma_n(t))$ exists for every n , and $v_n + \text{Im} \log \lambda \in 2\pi\mathbb{Z}$ (the branch of $\log \lambda$ does not matter here as long as it remains fixed). We will associate an *external address* $\underline{s} = s_1 s_2 s_3 \dots$ of integers to γ_1 by setting $s_n := (v_n - v_1)/2\pi$.

Corollary 2.6 (Classification of exponential maps). *For every strictly preperiodic external address $\underline{s} \in \mathbb{Z}^{\mathbb{N}}$, starting with $s_1 = 0$, there is a unique postsingularly finite exponential map for which the dynamic ray γ_1 landing at 0 has external address \underline{s} .*

Different external addresses $\underline{s}, \underline{s}'$ may yield the same exponential map; this happens if and only if more than one dynamic ray lands at the singular value. There is a straightforward algorithm to tell when this happens; see [LSV]: a necessary condition is that \underline{s} and \underline{s}' have the same period and the same preperiod, and the precise answer involves either combinatorial itineraries or internal addresses [S3], [RS2, Appendix A]: \underline{s} and \underline{s}' describe the same exponential map if and only if they have the same angled internal address associated to them.

This result has useful implications on exponential parameter space (see [S1, S5, RS1]): this space is structured in terms of parameter rays associated to a precisely described set of external addresses in $\mathbb{Z}^{\mathbb{N}}$ [FS], and postsingularly finite exponential maps are exactly the landing points of parameter rays at preperiodic external addresses. More precisely, a postsingularly finite exponential map E_λ is the landing point of the parameter ray at the strictly preperiodic external address \underline{s} if and only if in the dynamical plane of E_λ , the dynamic ray at external address \underline{s} lands at the singular value 0 [LSV, Theorem 3.4]. Together with [FRS], this also allows us to answer a complex version of a question of Euler [E]: he determined for which $a > 0$ the sequence $a, a^a, a^{a^a}, a^{a^{a^a}}, \dots$ has a limit. The answer is that convergence to a fixed point happens for $a \in [e^{-e}, e^{1/e}]$; convergence to a 2-cycle happens for $a < e^{-e}$, and convergence to ∞ happens for $a \geq e^{1/e}$. An equivalent formulation of this question is to ask for which $\lambda \in \mathbb{R}$ (with $\lambda = \log a$) the sequence $z_0 := 0, z_{n+1} := \lambda e^{z_n}$ has a limit. In this form, this question makes sense for complex λ . It is easy to see that this sequence converges to a limit in \mathbb{C} ,

without eventually being constant, if and only if $\lambda = \mu e^{-\mu}$ for $|\mu| < 1$ or μ a root of unity (convergence to periodic cycles, without eventually being periodic, is classified in terms of hyperbolic components in [S2], as well as boundaries of hyperbolic components, as in [S5, RS1]). Eventually periodic dynamics happens exactly for postsingularly finite exponential maps, and the special case of eventually constant convergence happens for those postsingularly finite exponential maps where the singular orbit eventually falls onto a fixed point. Finally, convergence to ∞ leads to parameter rays as described above and is classified in [FRS].

3. ITERATION IN TEICHMÜLLER SPACE

In this section, we will describe the Teichmüller space setup that allows us to prove the Main Theorem 2.4. This is lifted almost verbatim from [DH]: for each postsingularly finite topological exponential map g , we will define an analytic map $\sigma_g : \mathcal{T}_g \rightarrow \mathcal{T}_g$, where \mathcal{T}_g is the Teichmüller space modeled on a sphere with punctures at P_g (this is not to be confused with the Teichmüller space modeled on a genus g surface, which is often denoted similarly). A fixed point of σ_g corresponds to a holomorphic exponential map which is Thurston equivalent to g .

We will see that σ_g is strictly contracting for the infinitesimal Teichmüller metric, i.e., that $\|d\sigma_g(\tau)\| < 1$ for every $\tau \in \mathcal{T}_g$. Since Teichmüller space is path connected and geodesically complete, it follows that if σ_g has a fixed point τ_0 , then this fixed point is unique, and every point $\tau \in \mathcal{T}_g$ is attracted to τ_0 under iteration of σ_g .

The problem is that the contraction of σ_g is not uniform: there is no constant $k < 1$ such that $\|d\sigma_g(\tau)\| \leq k$, so the existence of a fixed point does not follow from the Banach fixed point theorem. In fact, it is fairly easy to see that if g admits a Levy cycle, then σ_g has no fixed point. The main issue in the proof of Theorem 2.4 is the converse: to show that if for some $\tau \in \mathcal{T}_g$ the sequence $\sigma_g^m(\tau)$ does not converge in \mathcal{T}_g , then g admits a Levy cycle.

In [DH], finiteness of the degree of a rational map made it possible to describe a subset of Teichmüller space, depending only on the initial point of the iteration, such that the contraction on this subset is uniform, while the existence of a Thurston obstruction follows as soon as the iteration leaves this subset. An important ingredient in this argument is the following: let $f : \mathbb{P}^1 \rightarrow \mathbb{P}^1$ be a postcritically finite rational map of degree d with postcritical set P_f of finite cardinality $|P_f|$. Let $\tilde{P}_f = f^{-1}(P_f)$ be the set of pre-postcritical points; then clearly $|\tilde{P}_f| < d|P_f|$, and this bound depends on d . At most $|\tilde{P}_f|$ non-homotopic annuli in $\overline{\mathbb{C}} \setminus \tilde{P}_f$ may become homotopic in $\overline{\mathbb{C}} \setminus P_f$. Thus if an annulus in $\overline{\mathbb{C}} \setminus P_f$ is very fat, at least one of the annuli on $\overline{\mathbb{C}} \setminus \tilde{P}_f$ must have been fat too. If there is no bound on $|\tilde{P}_f|$, then very fat annuli in $\overline{\mathbb{C}} \setminus P_f$ may arise without any fat annuli in $\overline{\mathbb{C}} \setminus \tilde{P}_f$. This destroys the motor for the proof in [DH]. (The relevant parts in [DH] are Theorem 7.1, which describes how moduli of annuli, or equivalently inverses of lengths of simple geodesics, can increase when p punctures are removed, and Proposition 8.2, which identifies the number of removed points as d^m , where d is the degree and m is an integer which also depends on d .)

Since this strategy fails for transcendental maps, and we need a different argument: if there is no fixed point in Teichmüller space, the Thurston map σ_g cannot be uniformly contracting; this leads to a sequence of integrable meromorphic quadratic differentials with arbitrarily little loss of mass under the push-forward, and our Limit Theorem 5.2 provides control on the mass distribution of the quadratic

differentials. This will be used in Propositions 3.2 and 3.3 to conclude the existence of a degenerate Levy cycle.

We will now make this program precise.

3.1. The Teichmüller space of a topological exponential map. Let $g: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{\infty, 0\}$ be a postsingularly finite topological exponential map with singular orbit $p_j := g^{o(j-1)}(0)$ for $j \geq 1$. Let $k' \geq 1$ and $k'' \geq 1$ be the preperiod and period of the singular orbit so that $p_{k'} \neq p_{k'+k''}$ but $p_{k'+1} = p_{k'+k''+1}$. Set also $k := k' + k''$ and $p_0 := \infty$.

Since $g: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{S}^2 \setminus \{\infty, 0\}$ is a universal covering map, the group of deck transformations is canonically isomorphic to \mathbb{Z} ; denote by t the generator corresponding to a positively oriented simple loop around 0 in $\mathbb{S}^2 \setminus \{\infty, 0\}$. Let $v \in \mathbb{Z} \setminus \{0\}$ be the unique integer such that $t^v p_{k'} = p_{k'+k''}$. Equivalently, let $\gamma: [0, 1] \rightarrow \mathbb{S}^2 \setminus \{\infty\}$ be a path with $\gamma(0) = p_{k'}$ and $\gamma(1) = p_{k'+k''}$; then v is the winding number of the closed curve $g \circ \gamma$ around 0.

Let (\mathbb{S}^2, P_g) be a topological 2-sphere with the $k+1$ distinct points $P_g := \{p_0, p_1, \dots, p_{k'+k''}\}$ marked. The *Teichmüller space* \mathcal{T}_g of g (modeled on (\mathbb{S}^2, P_g)) is the space of homeomorphisms $\varphi: \mathbb{S}^2 \rightarrow \overline{\mathbb{C}}$ with $\varphi(\infty) = \infty$, $\varphi(0) = 0$ and $\varphi(p_{k'+k''}) - \varphi(p_{k'}) = 2\pi i v$, modulo the equivalence relation $\varphi \sim \varphi'$ if $\varphi|_{P_g} = \varphi'|_{P_g}$ and φ and φ' are isotopic relative to P_g . (The normalizations $\varphi(\infty) = \infty$, $\varphi(0) = 0$ and $\varphi(p_{k'+k''}) - \varphi(p_{k'}) = 2\pi i v$ allow us to avoid the usual quotient by conformal equivalences in the definition of Teichmüller space. Thus \mathcal{T}_g is isomorphic to the standard Teichmüller space modeled on \mathbb{S}^2 with $k+1$ marked points.)

The next step is to construct a map $\sigma_g: \mathcal{T}_g \rightarrow \mathcal{T}_g$ as follows: for a homeomorphism $\varphi: \mathbb{S}^2 \rightarrow \overline{\mathbb{C}}$ with $\varphi(0) = 0$ and $\varphi(\infty) = \infty$, the map $\varphi \circ g$ is a universal cover $\mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{C}^*$, which is analytic for a unique analytic structure on $\mathbb{S}^2 \setminus \{\infty\}$, and with this analytic structure $\mathbb{S}^2 \setminus \{\infty\}$ is isomorphic to \mathbb{C} (note that there is no universal covering from \mathbb{D} to \mathbb{C}^*). Let $\tilde{\varphi}: \mathbb{S}^2 \setminus \{\infty\} \rightarrow \mathbb{C}$ be the unique conformal isomorphism such that $\tilde{\varphi}(0) = 0$ and $\tilde{\varphi}(p_{k'+k''}) - \tilde{\varphi}(p_{k'}) = 2\pi i v$; the first requirement determines $\tilde{\varphi}$ up to a multiplicative factor, and the second determines the factor. If we set $\lambda = \varphi(p_2)$, we see that the diagram

$$(2) \quad \begin{array}{ccc} \mathbb{S}^2 \setminus \{\infty\} & \xrightarrow{\tilde{\varphi}} & \mathbb{C} \\ \downarrow g & & \downarrow E_\lambda \\ \mathbb{S}^2 \setminus \{\infty, 0\} & \xrightarrow{\varphi} & \mathbb{C}^* \end{array}$$

commutes: the map $\varphi \circ g \circ \tilde{\varphi}^{-1}$ is a holomorphic universal cover from \mathbb{C} to \mathbb{C}^* and sends $0 = p_1$ to $\lambda = \varphi(p_2)$, so it has the form $\lambda \exp(az)$ for some $a \in \mathbb{C}^*$. By construction, the points $p_{k'}$ and $p_{k''}$ differ by the v -th power of a generating deck transformation, so the same must be true for $\tilde{\varphi}(p_{k'+k''})$ and $\tilde{\varphi}(p_{k'})$ with $\tilde{\varphi}(p_{k'+k''}) - \tilde{\varphi}(p_{k'}) = 2\pi i v$; thus translation by $2\pi i$ is a generating deck transformation on the right hand side. The orientation of the generator t implies $a = 1$.

The homeomorphism $\tilde{\varphi}$ extends to a homeomorphism $\tilde{\varphi}: \mathbb{S}^2 \rightarrow \overline{\mathbb{C}}$ with $\tilde{\varphi}(\infty) = \infty$, since on both sides we have the one-point compactification. Therefore, $\tilde{\varphi}$ represents a point in \mathcal{T}_g .

Since an isotopy of φ lifts to an isotopy of $\tilde{\varphi}$, we see that the formula $\sigma_g(\langle \varphi \rangle) := \langle \tilde{\varphi} \rangle$, where $\langle \varphi \rangle$ denotes the point in Teichmüller space represented by φ , defines a mapping $\sigma_g: \mathcal{T}_g \rightarrow \mathcal{T}_g$. We will often write σ for σ_g .

The crucial observation is the following:

Theorem 3.1 (Fixed points and exponential maps). *A fixed point of σ gives a postsingularly finite holomorphic exponential map E_λ which is Thurston equivalent to g , and conversely any such holomorphic E_λ defines a fixed point of σ in \mathcal{T}_g .*

Proof. If $\langle \varphi \rangle$ and $\langle \tilde{\varphi} \rangle = \sigma(\langle \varphi \rangle)$ are the same point in Teichmüller space, then for $\lambda := \varphi(p_2)$, the relation $E_\lambda \circ \tilde{\varphi} = \varphi \circ g$ from (2) gives a Thurston equivalence between g and E_λ . Conversely, if $g = E_\lambda$ is holomorphic to begin with, then $\langle \text{id} \rangle$ is a fixed point of \mathcal{T}_g . \square

The case $|P_g| \leq 3$ is immediate: it implies $k' = k'' = 1$, $P_g = \{\infty, 0, p_2\}$ and $g(p_2) = p_2$, and in particular $|P_g| = 3$. There is no essential simple closed curve in $\overline{\mathbb{C}} \setminus P_g$ and hence no Levy cycle. On the other hand, we have $\varphi(p_2) = \varphi(0) + 2\pi iv = 2\pi iv$ and any two admissible maps $\varphi, \tilde{\varphi}: \mathbb{S}^2 \rightarrow \overline{\mathbb{C}}$ are isotopic to each other rel P_g . Thus Teichmüller space consists of a single point which is fixed under σ and g is equivalent to the exponential map $z \mapsto \lambda e^z$ with $\lambda = 2\pi iv$. We will from now on suppose that $k = k' + k'' \geq 3$, hence $|P_g| \geq 4$.

3.2. The Teichmüller metric and its dual. The Teichmüller metric on Teichmüller space is given as follows: let $\tau, \tau' \in \mathcal{T}_g$ be represented by two homeomorphisms $\varphi, \varphi': \mathbb{S}^2 \rightarrow \overline{\mathbb{C}}$ so that $\psi := \varphi' \circ \varphi^{-1}$ is quasiconformal. Then $d(\tau, \tau') = \inf \log K_\psi$, where ψ ranges over all quasiconformal homeomorphisms obtained from representatives φ and φ' of τ and τ' , and $K_\psi \geq 1$ is its maximal dilatation. Background on quasiconformal maps can be found in [A2, GL, H, IT, L]. We will need the facts that \mathcal{T}_g is complete for the Teichmüller metric and that the Teichmüller metric is a Finsler metric: there is a norm, called the Teichmüller norm, on each tangent space $T_\tau \mathcal{T}_g$ such that the distance between points τ, τ' is the infimum of lengths of curves joining τ to τ' : if $\tau' \neq \tau$, then this infimum is positive, and it is realized as the length of a curve in \mathcal{T}_g connecting τ to τ' .

The cotangent space to Teichmüller space is the space of integrable holomorphic quadratic differentials endowed with the L^1 -norm $\|q\|_{\mathbb{C}} = \int_{\mathbb{C}} |q(x + iy)| dx dy$. The infinitesimal metric inducing the Teichmüller metric is the dual norm to the norm on cotangent space.

For a finite set $Z \subset \overline{\mathbb{C}}$, let $Q^1(Z)$ be the finite-dimensional Banach space of integrable meromorphic quadratic differentials on $\overline{\mathbb{C}}$ with poles only on Z ; integrability implies that all poles are simple (note that for a Riemann surface X , the notation $Q^1(X)$ is often used for integrable quadratic differentials on X , also in our Section 4). We will make use of results from [DH] and [HS], which fit in with general results in Teichmüller theory as described in [A2, GL, H, IT, L].

The cotangent space to Teichmüller space at the point $\langle \varphi \rangle$ is canonically isomorphic to

$$T_{\langle \varphi \rangle}^* \mathcal{T}_g = Q^1(\varphi(P_g)) .$$

The L^1 -norm on $Q^1(\varphi(P_g))$ is dual to the infinitesimal Teichmüller norm on $T_{\langle \varphi \rangle} \mathcal{T}_g$. For a quadratic differential q on $\overline{\mathbb{C}}$, we denote its L^1 -norm on $\overline{\mathbb{C}}$ (or equivalently on \mathbb{C}) by $\|q\|_{\mathbb{C}}$. We use the words “mass” as a synonym for the norm of a quadratic differential.

The assignment $\langle \varphi \rangle \mapsto (\varphi(p_1), \dots, \varphi(p_k))$ defines a map $\mathcal{T}_g \rightarrow \mathbb{C}^k$; because of the two normalization conditions, this turns \mathcal{T}_g into a $k - 2$ -dimensional complex

manifold. The map σ is analytic. The crucial fact is that the co-derivative of σ at $\langle \tilde{\varphi} \rangle = \sigma(\langle \varphi \rangle)$ is the linear map

$$d\sigma_{\langle \tilde{\varphi} \rangle}^* : T_{\langle \tilde{\varphi} \rangle}^* \mathcal{T}_g \rightarrow T_{\langle \varphi \rangle}^* \mathcal{T}_g \quad \text{given by} \quad d\sigma_{\langle \tilde{\varphi} \rangle}^* q = (E_\lambda)_* q ,$$

for $q \in Q^1(\tilde{\varphi}(P_g))$ with $\lambda := \varphi(p_2)$. Here, $(E_\lambda)_* q$ denotes the push-forward of $q = q(z) dz^2$ under the map E_λ , i.e.

$$\begin{aligned} (E_\lambda)_* (q(z) dz^2) &= \sum_{z \in E_\lambda^{-1}(w)} \frac{q(z) dw^2}{(E'_\lambda(z))^2} = \frac{dw^2}{w^2} \sum_{z \in E_\lambda^{-1}(w)} q(z) \\ (3) \qquad \qquad \qquad &= \frac{dw^2}{w^2} \sum_{m \in \mathbb{Z}} q \left(\log \left(\frac{w}{\lambda} \right) + 2\pi i m \right) . \end{aligned}$$

This defines a meromorphic quadratic differential on \mathbb{C}^* with norm $\|(E_\lambda)_* q\|_{\mathbb{C}} \leq \|q\|_{\mathbb{C}}$. Since the isolated singularities at 0 and ∞ have finite mass, they must be at worst simple poles, so $(E_\lambda)_* q$ is a meromorphic quadratic differential with all poles simple, and these poles can only be at the E_λ -images of the poles of q , as well as at 0 and ∞ . Similar remarks apply to all push-forward maps induced by coverings.

For exponential maps E_λ , as for all transcendental entire maps, we have

$$\|(E_\lambda)_* q\|_{\mathbb{C}} < \|q\|_{\mathbb{C}} \quad \text{for } q \neq 0;$$

in case of equality, all preimages of a pole of $(E_\lambda)_* q$ (other than 0 or ∞) must also be poles of q , but q can have only finitely many poles.

Since $Q^1(\tilde{\varphi}(P_g))$ is finite dimensional, it follows that

$$\|d\sigma_{\langle \tilde{\varphi} \rangle}^*\| = \left\{ \sup_{q \in Q^1(\tilde{\varphi}(P_g)) \setminus 0} \frac{\|(E_\lambda)_* q\|_{\mathbb{C}}}{\|q\|_{\mathbb{C}}} \right\} < 1 .$$

By definition of the dual norm, we have $\|d\sigma_{\langle \varphi \rangle}\| = \|d\sigma_{\langle \tilde{\varphi} \rangle}^*\| < 1$.

This argument can be iterated: every $\tau_0 := \langle \varphi_0 \rangle \in \mathcal{T}_g$ defines a sequence $\tau_n := \langle \varphi_n \rangle$ in \mathcal{T}_g via $\tau_{n+1} = \sigma(\tau_n)$. For every $n \geq 0$ and $s \geq 1$, we have

$$(4) \qquad \qquad \qquad \|(d\sigma^{\circ k})_{\tau_n}^*\| = \left\{ \sup_{q \in Q^1(\varphi_{n+k}(P_g)) \setminus 0} \frac{\|(E^{(k)})_* q\|_{\mathbb{C}}}{\|q\|_{\mathbb{C}}} \right\}$$

where $E^{(k)} = E_{\lambda_k} \circ \dots \circ E_{\lambda_1}$ with $\lambda_s = \varphi_{n+k-s}(p_2)$.

The strategy of the proof of the main theorem is to show that the non-existence of a fixed point in Teichmüller space implies the existence of quadratic differentials with almost no contraction under the push-forward, and this will give us good enough control on the geometry to conclude that there is a Levy cycle. The key to this will be Proposition 3.2, which will be proved in Section 6. Before stating it, we will describe how to visualize certain quadratic differentials.

3.3. Examples of quadratic differentials. We have found the following construction useful when trying to visualize elements of $Q^1(Z)$, especially quadratic differentials which undergo little contraction under $(E_\lambda)_*$. Let $P \subset \mathbb{C}$ be a compact polygon with sides parallel to the real and imaginary axes, so the interior angles are all $\pi/2$ or $3\pi/2$. The double \tilde{P} of P , i.e., the surface obtained by gluing two copies P' and P'' of P along their common boundary, is homeomorphic to a sphere and carries a unique complex structure compatible with that of $P \subset \mathbb{C}$ on P' and compatible with the conjugate complex structure on P'' . Moreover \tilde{P} carries a

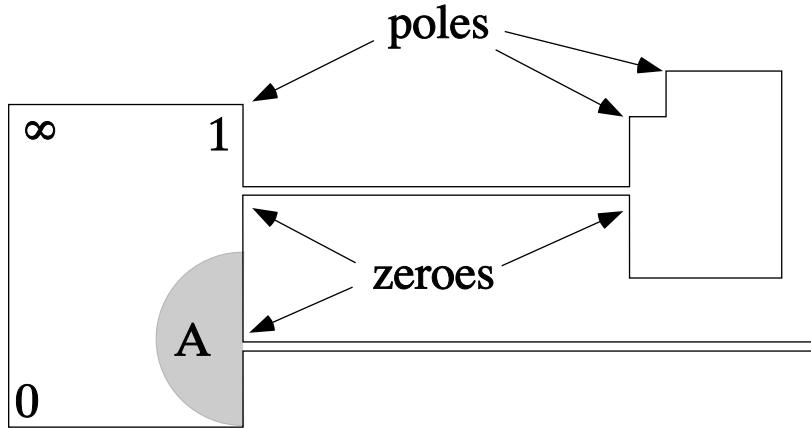


FIGURE 1. The quadratic differential corresponding to this polygon has 11 simple poles and 7 simple zeroes. The polygon P is composed of a “base rectangle” with three corners labeled $0, 1, \infty$, with decorations attached at the end of two “tubes”. When \tilde{P} is identified with \mathbb{P}^1 , the boundary of P becomes the real axis. The poles and zeroes, except for the four corners of the base rectangle, are concentrated in two small disks centered at points on the segment $[0, 1]$. As the tubes get longer (and perhaps thinner), more and more of the mass may “migrate” out to the decorations.

quadratic differential q which is simply dz^2 on P' and $d\bar{z}^2$ on P'' . The quadratic differential has simple poles at the vertices of P with angle $\pi/2$ and simple zeroes at those with angle $3\pi/2$. The measure $|q|$ is then simply the element of area, so that $\|q\|_{\tilde{P}}$ is equal to twice the area of P .

Since by the uniformization theorem \tilde{P} is isomorphic to \mathbb{P}^1 , by simply drawing P , we have completely specified an integrable meromorphic quadratic differential on \mathbb{P}^1 if we require three points of P to correspond to $0, 1, \infty$. Of course, actually writing it down might be difficult (already for a square the uniformization requires elliptic functions), but it is usually quite easy to grasp the qualitative aspects of q , and more particularly, it is easy to find the fat annuli in the complement of the poles and estimate their moduli.

Figure 1 illustrates this construction. Note that the annulus A has large modulus; since $0, 1, \infty$ are in the same component of $\mathbb{P}^1 \setminus A$, it follows that the bounded component of $\mathbb{P}^1 \setminus A$ has small diameter.

3.4. Proof of the Main Theorem 2.4. As mentioned above, the following result will be a key ingredient in the proof of our main theorem. In the remainder of this section, we will prove the main theorem using this result, and the rest of the paper will then give a proof of Proposition 3.2.

Proposition 3.2 (Contraction after several iterations). *For every number of poles N , for every number of iterates m and for every modulus $M > 0$, there is a positive ratio $r < 1$ with the following property: if q is an integrable meromorphic quadratic differential with at most N poles on $\overline{\mathbb{C}}$ and $\lambda_1, \dots, \lambda_m \in \mathbb{C}^*$ so that*

$$\|(E_{\lambda_m} \circ E_{\lambda_{m-1}} \circ \dots \circ E_{\lambda_1})_* q\| > r \|q\| ,$$

then there exist two concentric disks $\tilde{D} \subset D$ such that $E_{\lambda_m} \circ E_{\lambda_{m-1}} \circ \cdots \circ E_{\lambda_1}$ is injective on D , the disk \tilde{D} contains at least two poles of q , and the annulus $D \setminus \tilde{D}$ has modulus at least M .

We will need some properties of the geometry of Riemann surfaces which are collected in Subsection A.1. We will always use the hyperbolic metric on $\overline{\mathbb{C}} \setminus \varphi(P_g)$ with constant curvature -1 .

A simple closed curve or an annulus in $\overline{\mathbb{C}} \setminus \varphi(P_g)$ is called *essential* if both complementary components contain at least two points in $\varphi(P_g)$. We will call a closed geodesic γ “short” if it is a simple closed geodesic with length $\ell(\gamma) < \ell^* := \log(3 + 2\sqrt{2})$ (a closed geodesic is necessarily essential). Then any two short closed geodesics are either identical or disjoint and non-homotopic (see Corollary A.2). Since $|P_g| = k + 1$, it follows that the number of different short closed geodesics is at most $k - 2$.

Proposition 3.3 (Contraction or Levy cycle). *Let g be a postsingularly finite topological exponential map. Denote the preperiod of the singular orbit by k' and the period by k'' , and set $k := k'' + k'$. For every distance d_0 , there is a real number $r < 1$ with the following property: if $\tau \in \mathcal{T}_g$ satisfies*

$$d(\tau, \sigma(\tau)) < d_0 \quad \text{and} \quad \|d\sigma^{\circ k}(\tau)\| > r ,$$

then g has a degenerate Levy cycle.

Proof. Let $r < 1$ be as in Proposition 3.2, applied to $N = k + 1$ and

$$M := k \frac{\pi}{\ell^*} \cdot e^{kd_0} .$$

For this r , assume that $\tau \in \mathcal{T}_g$ satisfies the above conditions. Denote $\tau_s := \sigma^{k-s}(\tau)$ for $s = 0, \dots, k$. By the definition of σ , one can choose representatives φ_s of τ_s so that $E_{\lambda_s} \circ \varphi_{s-1} = \varphi_s \circ g$ for $s = k, \dots, 1$, where $\lambda_s = \varphi_s(p_2)$. It follows from (4) that the co-derivative of $\sigma^{\circ k}$ at τ_0 corresponds to the push-forward by $E^{(k)} = E_{\lambda_k} \circ \cdots \circ E_{\lambda_1}$. Hence there is a $q \in Q^1(\varphi_0(P_g))$ with $\|(E^{(k)})_* q\| > r \|q\|$.

By Proposition 3.2, there are two concentric disks $\tilde{D} \subset D$ such that $E^{(k)}|_D$ is injective, $A := D \setminus \tilde{D}$ is a round annulus with $\text{mod}(A) = M$, and q has at least two poles (which are automatically in $\varphi_0(P_g)$) in \tilde{D} .

The injective image $E^{(k)}(A)$ is an annulus which surrounds at least two points in $E^{(k)}(\varphi_0(P_g) \cap \mathbb{C}) \subset \varphi_k(P_g)$; it may contain some of the points $\varphi_k(P_g)$. Since $\varphi_k(P_g) \cap \mathbb{C}$ contains exactly k points, there is a parallel round subannulus $A_0 \subset A$ with $\text{mod}(A_0) > M/k$ such that $E^{(k)}(A_0)$ does not contain a point of $\varphi_k(P_g)$. Define A_s by $A_s = E_{\lambda_s}(A_{s-1})$ for $s = 1, \dots, k$. Then for $s = 0, \dots, k-1$, the map $E_{\lambda_{s+1}}|_{A_s} : A_s \rightarrow A_{s+1}$ is a conformal isomorphism, A_s does not contain a point of $\varphi_s(P_g)$ and A_s surrounds at least two points of $\varphi_s(P_g)$. These A_s are essential in $\overline{\mathbb{C}} \setminus \varphi_s(P_g)$, since $\varphi_s(p_{k'})$ and $\varphi_s(p_k) = \varphi_s(p_{k'}) + 2\pi iv$ cannot be surrounded at the same time because of the injectivity of $E_{\lambda_{s+1}}$ on the disk surrounded by A_s .

Let $\gamma_s \subset S^2 \setminus P_g$ be a simple closed curve such that $\varphi_s(\gamma_s)$ is the core curve of A_s . By the construction, g maps γ_s homeomorphically onto γ_{s+1} . Since $d(\tau_s, \tau_{s-1}) = d(\sigma^{\circ(n-s)}(\tau), \sigma^{\circ(n-s+1)}(\tau)) \leq d(\tau, \sigma(\tau)) < d_0$, we have $d(\tau_s, \tau_0) < kd_0$. There exists a quasiconformal map $\psi_s : \mathbb{C} \rightarrow \mathbb{C}$ isotopic to $\varphi_0 \circ \varphi_s^{-1}$ relative to $\varphi_s(P_g)$ such that its dilatation $K(\psi_s)$ is at most $e^{d(\tau_s, \tau_0)} < e^{kd_0}$. Then $\varphi_0(\gamma_s)$ is homotopic to $\psi_s(\varphi_s(\gamma_s))$ and the annulus $\hat{A}_s = \psi_s(A_s)$ has modulus at least $\frac{1}{K(\psi_s)} \text{mod}(A_s) > e^{-kd_0} \cdot \frac{M}{k} = \frac{\pi}{\ell^*}$.

(See Ahlfors [A1].) Therefore there is a closed geodesic $\hat{\gamma}_s$ in $\mathbb{C} \setminus \varphi_0(P_g)$ homotopic to $\varphi_0(\gamma_s)$ (and to the core curve of \hat{A}_s) and its length satisfies

$$\ell(\hat{\gamma}_s) \leq (\text{length of core curve of } \hat{A}_s) = \frac{\pi}{\text{mod}(\hat{A}_s)} < \ell^*.$$

There can be at most $k-2$ distinct closed geodesics of length less than ℓ^* . Therefore there are two indices $s_1 < s_2 \in \{0, 1, \dots, k-1\}$ so that $\hat{\gamma}_{s_1}$ and $\hat{\gamma}_{s_2}$ coincide. Hence corresponding γ_{s_1} and γ_{s_2} are homotopic in $\mathbb{S}^2 \setminus P_g$. Thus $\{\gamma_{s_1}, \dots, \gamma_{s_2}\}$ is a Levy cycle of g . It is automatically degenerate because the γ_s bound disks on which g is a homeomorphism. \square

Proof of Theorem 2.4. Let $g: \mathbb{C} \rightarrow \mathbb{C}^*$ be a postsingularly finite topological exponential map, and again let k be the length of the singular orbit. Choose a point $\tau_0 \in \mathcal{T}_g$ and run the Thurston iteration $\tau_n := \sigma^{\circ n}(\tau_0)$. Let $C_0: [0, 1] \rightarrow \mathcal{T}_g$ be a curve in Teichmüller space connecting τ_0 to τ_1 with finite length, say d_0 . Let $C_n: [0, 1] \rightarrow \mathcal{T}_g$ be the image curve with $C_n(t) := \sigma^{\circ n}(C_0(t))$, for $n \geq 1$. Then for $\tau' := C_n(t)$, we have

$$\begin{aligned} d(\tau', \sigma(\tau')) &\leq d(\tau', \tau_{n+1}) + d(\tau_{n+1}, \sigma(\tau')) \leq d(\tau', \tau_{n+1}) + d(\tau_n, \tau') \\ &\leq d(C_0(t), \tau_1) + d(\tau_0, C_0(t)) = d_0. \end{aligned}$$

Suppose that g has no degenerate Levy cycle. Then Proposition 3.3 yields an $r < 1$ such that for every $n \in \mathbb{N}$ and every $t \in [0, 1]$, $\|d\sigma^{\circ k}(C_n(t))\| \leq r$. Hence $\sigma^{\circ k}$ contracts the length of C_n by factor r and we have $d(\tau_{n+1}, \tau_n) \leq r^{(n-k)/k} d_0$. Therefore (τ_n) forms a Cauchy sequence which converges to a fixed point in \mathcal{T}_g ; this implies that g is Thurston equivalent to a holomorphic exponential map.

The uniqueness statement in the theorem follows because no contracting map can have more than one fixed point. It remains to show that g cannot simultaneously be Thurston equivalent to a holomorphic exponential map and have a Levy cycle. This follows just as for rational maps; we give an argument different from [DH]: since the existence of a Levy cycle is preserved under Thurston equivalences, it suffices to show that a postsingularly finite holomorphic exponential map g cannot have a Levy cycle. Let M be the maximal modulus of an essential annulus in $\overline{\mathbb{C}} \setminus P_g$ and let $\gamma \subset \overline{\mathbb{C}} \setminus P_g$ be a simple closed curve on a degenerate Levy cycle. Choose a point $\langle \varphi_0 \rangle \in \mathcal{T}_g$ in which there is an annulus with modulus $2M$ and with core curve homotopic to $\varphi_0(\gamma)$. Then every $\langle \varphi_n \rangle = \sigma^{\circ n}(\langle \varphi_0 \rangle)$ must have an annulus with modulus at least $2M$ in the homotopy class of γ , and the Thurston iteration cannot converge to the fixed point $\langle \text{id} \rangle \in \mathcal{T}_g$. \square

Remark. A different way to show that no holomorphic exponential map has a Levy cycle uses lengths of hyperbolic geodesics: let g be a holomorphic exponential map with finite postsingular set P_g , and set $\tilde{P}_g := g^{-1}(P_g) \supset P_g$. If a simple closed curve γ is part of a Levy cycle, then we may as well represent γ by the unique closed hyperbolic geodesic in its homotopy class; let $\ell(\gamma)$ be its length in $\overline{\mathbb{C}} \setminus P_g$. Since $g: (\overline{\mathbb{C}} \setminus \tilde{P}_g) \rightarrow \overline{\mathbb{C}} \setminus P_g$ is a covering map, hence a local hyperbolic isometry, $g^{-1}(\gamma)$ is a countable collection of simple closed curves in $\overline{\mathbb{C}} \setminus \tilde{P}_g$, each of length $\ell(\gamma)$ (or $g^{-1}(\gamma)$ is a single curve, homeomorphic to \mathbb{R} ; but then γ cannot be part of a Levy cycle). Since the inclusion $\iota: (\overline{\mathbb{C}} \setminus \tilde{P}_g) \rightarrow (\overline{\mathbb{C}} \setminus P_g)$ is a strict contraction with respect to the hyperbolic metrics, it follows that all bounded components of $g^{-1}(\gamma)$ are closed curves of length less than $\ell(\gamma)$; for those which are essential, the

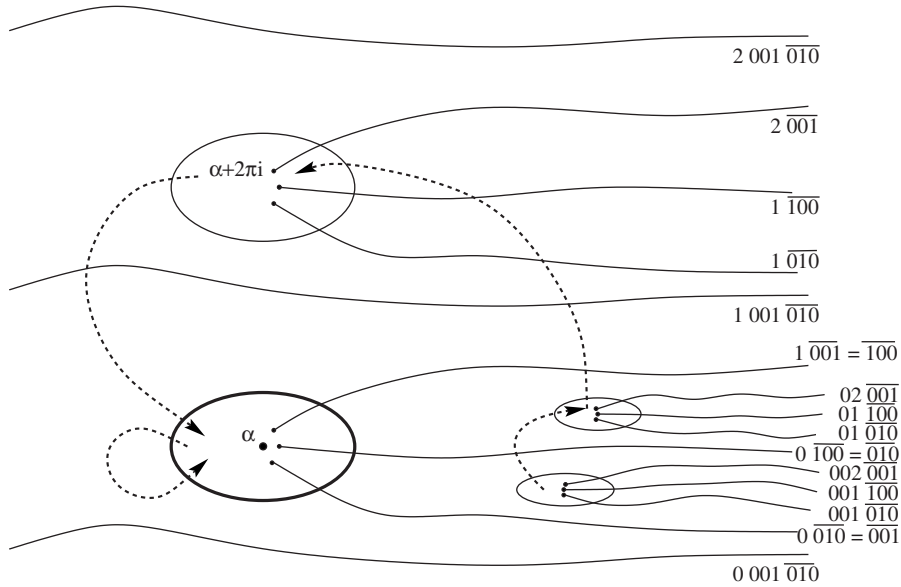


FIGURE 2. For the exponential map $z \mapsto \lambda e^z$ with $\lambda \approx 1.449 + 1.008i$, the singular orbit lands at a fixed point after three iterations, and this fixed point is the landing point of the three dynamic rays at external addresses $\overline{001}$, $\overline{010}$, and $\overline{100}$; these rays are permuted transitively by the dynamics. Consequently, the singular value is the landing point of the three preperiodic dynamic rays at external angles $001\overline{010}$, $001\overline{100}$, and $002\overline{001}$. If a topological exponential map is modeled after this example, except that the three dynamic rays are forced to land at three distinct fixed points, then the single curve surrounding the three fixed points is mapped to itself with degree one and forms a degenerate Levy cycle, hence a particularly simple Thurston obstruction. In the Thurston iteration, any annulus around this curve would acquire modulus tending to infinity, squeezing the three fixed points into a single point.

corresponding simple closed geodesics are yet shorter. But if γ were part of a Levy cycle, then a finite repetition of this argument would yield a simple closed geodesic of length less than $\ell(\gamma)$ in the same homotopy class as γ , and this is a contradiction.

An example of a topological exponential map with a Thurston obstruction is given in Figure 2. See also [HS, Sec. 5] for a discussion of Thurston obstructions of polynomials which applies also in our case.

4. THE DECOMPOSITION THEOREM

In this section we will give a “thick-thin” decomposition of Riemann surfaces of finite type with an integrable quadratic differential. This result is of interest in its own right, and it may help to prove an analog of Thurston’s theorem for entire functions other than exponentials.

4.1. The thick-thin decomposition theorem for quadratic differentials.

Let X be a hyperbolic Riemann surface of finite type. If $\gamma \subset X$ is a simple closed geodesic of length $\ell(\gamma) < \ell^* := \log(3 + 2\sqrt{2})$, then by the Collaring Theorem A.1, γ is surrounded by an annulus A_γ of modulus at least $\pi/\ell - 1$ (the *standard collar* around γ) so that γ is the core curve of A_γ ; moreover, for different simple closed geodesics, the standard collars are disjoint.

The standard collar A_γ is isomorphic to the cylinder $\{z \in \mathbb{C} : |\operatorname{Im} z| < h_\gamma\}/\mathbb{Z}$ with circumference 1 and height $2h_\gamma$, by a conformal isomorphism which identifies γ with \mathbb{R}/\mathbb{Z} . Set $h'_\gamma := h_\gamma - \sqrt{h_\gamma}$. We define the *central collar* $A_\gamma^0 \subset A_\gamma$ as the subannulus which corresponds to the region $\{z : |\operatorname{Im} z| < h'_\gamma\}$, so that $\operatorname{mod}(A_\gamma^0) = \operatorname{mod}(A_\gamma) - \sqrt{\operatorname{mod}(A_\gamma)}$. We will also denote by A_γ^\pm the two components of $A_\gamma \setminus A_\gamma^0$; in context we will need to be careful which is which.

Remark. The choice of the square root \sqrt{h} in this definition is not critical. It is sufficient to require that for each short curve γ , $h_\gamma - h'_\gamma$ is large, while h'_γ/h_γ is close to one: this means that $\operatorname{mod}(A_\gamma^\pm)$ (the ‘‘boundary padding’’) is large in absolute terms, but small compared to $\operatorname{mod}(A_\gamma)$. Under these assumptions, the estimate in Theorem 4.1 below still holds.

For any $\delta \in (0, \ell^*)$, the δ -*decomposition* of X consists of writing

$$X = \bigcup A_i^0 \cup \bigcup Y_j,$$

where the γ_i are the simple closed geodesics on X of length $\ell(\gamma_i) < \delta$, the A_i^0 are the central collars around the γ_i , and the Y_j are the connected components of $X \setminus \bigcup_i A_i^0$. We denote the standard collars around γ_i by A_i .

For any Riemann surface X , we denote by $Q^1(X)$ the space of integrable holomorphic quadratic differentials on X . If $Y \subset X$ is a measurable subset and $q \in Q^1(X)$, we will write

$$\|q\|_Y := \int_Y |q|.$$

The following theorem is the main result of this section.

Theorem 4.1 (Decomposition of quadratic differentials). *For every $\varepsilon > 0$, there exists a universal constant $\delta_0 > 0$ with the following property: if X is a hyperbolic Riemann surface of finite type and $0 < \delta < \delta_0$ and if*

$$X = \bigcup_{i \in I} A_i^0 \cup \bigcup_{j \in J} Y_j$$

is the δ -decomposition of X , then every $q \in Q^1(X)$ can be written as

$$q = \sum_{i \in I} q_{A_i} + \sum_{j \in J} q_{Y_j}$$

with all $q_{A_i}, q_{Y_j} \in Q^1(X)$ and adapted to the decomposition of X so that

$$\sum_i \left(\|q - q_{A_i}\|_{A_i^0} + \|q_{A_i}\|_{X \setminus A_i^0} \right) + \sum_j \left(\|q - q_{Y_j}\|_{Y_j} + \|q_{Y_j}\|_{X \setminus Y_j} \right) \leq \varepsilon \|q\|_{\bigcup A_i}.$$

Remark. The value of δ_0 depends only on ε , not on the topology of X , so the result might generalize to arbitrary hyperbolic Riemann surfaces.

4.2. Quadratic differentials on annuli. Choose $h > 1$ and let $A := \{z \in \mathbb{C} : |\operatorname{Im} z| < h\}/\mathbb{Z}$. Then any element $q \in Q^1(A)$ can be developed in a Fourier series

$$q = \left(\sum_{n=-\infty}^{\infty} a_n e^{2\pi i n z} \right) dz^2 .$$

Write

$$q^+ = \left(\sum_{n=-\infty}^{-1} a_n e^{2\pi i n z} \right) dz^2, \quad q^0 = a_0 dz^2, \quad q^- = \left(\sum_{n=1}^{\infty} a_n e^{2\pi i n z} \right) dz^2 .$$

Note that q^+ is actually defined in the semi-infinite cylinder $\{z \in \mathbb{C} : \operatorname{Im} z < h\}/\mathbb{Z}$ and q^- is defined in $\{z \in \mathbb{C} : \operatorname{Im} z > -h\}/\mathbb{Z}$.

Proposition 4.2 (Decomposition of differentials on annuli). *For $h > 1$, let $A := \{z \in \mathbb{C} : |\operatorname{Im} z| < h\}/\mathbb{Z}$ be an annulus of modulus $2h$, let $q \in Q^1(A)$ be an integrable quadratic differential and decompose it as $q = q^0 + q^+ + q^-$ as above. For every $\varepsilon > 0$, we have the following: suppose A_η^+ and A_η^- are parallel subannuli of A at the two ends of A , both of modulus $\eta > \frac{1}{2\pi} \log(2 + \frac{1}{\varepsilon})$, and $A^0 \subset A$ is another parallel subannulus of A with $A^0 \subset A \setminus (A^+ \cup A^-)$ and $\operatorname{mod}(A^0)/\operatorname{mod}(A) > 1 - \frac{\varepsilon}{1+2\varepsilon}$. Then we have*

$$(5) \quad \begin{aligned} \|q - q^0\|_{A^0} &\leq \varepsilon \|q\|_A, & \|q^0\|_{A_\eta^+ \cup A_\eta^-} &\leq \varepsilon \|q\|_A, \\ \|q - q^+\|_{A_\eta^+} &\leq \varepsilon \|q\|_A, & \|q^+\|_{\{\operatorname{Im} z < h-\eta\}/\mathbb{Z}} &\leq \varepsilon \|q\|_A, \\ \|q - q^-\|_{A_\eta^-} &\leq \varepsilon \|q\|_A, & \|q^-\|_{\{\operatorname{Im} z > -h+\eta\}/\mathbb{Z}} &\leq \varepsilon \|q\|_A, \\ & & \|q\|_{A \setminus (A_\eta^+ \cup A_\eta^- \cup A^0)} &\leq \varepsilon \|q\|_A . \end{aligned}$$

Proof. Let us consider

$$q^- = e^{2\pi i z} \left(\sum_{n=1}^{\infty} a_n e^{2\pi i (n-1)z} \right) dz^2 =: e^{2\pi i z} f(z) dz^2 .$$

Set $F(y) = \int_0^1 |f(x + iy)| dx$ as in Proposition A.4. For $\alpha < \beta \in \mathbb{R} \cup \{\pm\infty\}$, let $B_{\alpha,\beta} := \{z \in \mathbb{C} : \alpha < \operatorname{Im}(z) < \beta\}$.

If $\beta > \alpha \geq -h$, then

$$\|q^-\|_{B_{\alpha,\beta}} := \int_\alpha^\beta \left(\int_{\mathbb{R}/\mathbb{Z}} |q^-(x + iy)| dx \right) dy = \int_\alpha^\beta e^{-2\pi y} F(y) dy .$$

Moreover, by Proposition A.4(b), the function F is decreasing, so we have

$$\|q^-\|_{B_{-h,-h+\eta}} \geq F(-h') \int_{-h}^{-h+\eta} e^{-2\pi y} dy = F(-h') \cdot \frac{e^{2\pi h} - e^{2\pi(h-\eta)}}{2\pi} ,$$

whereas

$$\|q^-\|_{B_{-h+\eta,\infty}} \leq F(-h') \int_{-h+\eta}^{\infty} e^{-2\pi y} dy = F(-h') \cdot \frac{e^{2\pi(h-\eta)}}{2\pi}$$

and hence

$$(6) \quad \frac{\|q^-\|_{A \setminus A_\eta^-}}{\|q^-\|_{A_\eta^-}} \leq \frac{\|q^-\|_{B_{-h+\eta,\infty}}}{\|q^-\|_{B_{-h,-h+\eta}}} \leq \frac{e^{2\pi(h-\eta)}}{e^{2\pi h} - e^{2\pi(h-\eta)}} = \frac{1}{e^{2\pi\eta} - 1} < \frac{\varepsilon}{1 + \varepsilon} =: \varepsilon' .$$

The case of q^+ is similar, and the case of q^0 is easier: since $q_0 = a_0 dz^2$, we have

$$(7) \quad \|q^0\|_{A \setminus A^0} < \frac{\varepsilon}{1 + \varepsilon} \|q^0\|_{A^0} = \varepsilon' \|q^0\|_{A^0} .$$

The argument from here is just bookkeeping: we have

$$(8) \quad \begin{aligned} \|q\|_A &\geq \|q\|_{A_\eta^+} + \|q\|_{A^0} + \|q\|_{A_\eta^-} \\ &\geq \left(\|q^+\|_{A_\eta^+} - \|q^0\|_{A_\eta^+} - \|q^-\|_{A_\eta^+} \right) \\ &\quad + \left(\|q^0\|_{A^0} - \|q^+\|_{A^0} - \|q^-\|_{A^0} \right) \\ &\quad + \left(\|q^-\|_{A_\eta^-} - \|q^+\|_{A_\eta^-} - \|q^0\|_{A_\eta^-} \right) \\ &\geq \left(\|q^+\|_{A_\eta^+} + \|q^0\|_{A^0} + \|q^-\|_{A_\eta^-} \right) \\ &\quad - \varepsilon' \left(\|q^+\|_{A_\eta^+} + \|q^0\|_{A^0} + \|q^-\|_{A_\eta^-} \right) \\ &= (1 - \varepsilon') \left(\|q^+\|_{A_\eta^+} + \|q^0\|_{A^0} + \|q^-\|_{A_\eta^-} \right) . \end{aligned}$$

Now inequality (6) leads to

$$\|q^-\|_{B_{-h+\eta, \infty}} \leq \varepsilon' \|q^-\|_{A_\eta^-} \leq \frac{\varepsilon'}{1 - \varepsilon'} \|q\|_A = \varepsilon \|q\|_A ;$$

the analogous inequality for q^+ is proved the same way, and the analogous formula for q^0 follows in the same way from inequality (7). This proves half of Proposition 4.2.

The other half follows also: for instance, using (8),

$$\begin{aligned} \|q - q^0\|_{A^0} &\leq \|q^-\|_{A^0} + \|q^+\|_{A^0} \leq \|q^-\|_{A \setminus A_\eta^-} + \|q^+\|_{A \setminus A_\eta^+} \\ &\leq \varepsilon' \left(\|q^-\|_{A_\eta^-} + \|q^+\|_{A_\eta^+} \right) \leq \frac{\varepsilon'}{1 - \varepsilon'} \|q\|_A = \varepsilon \|q\|_A . \end{aligned}$$

□

Corollary 4.3 (Little mass near end). *For $h > 1$, again let $A := \{z \in \mathbb{C} : |\operatorname{Im} z| < h\} / \mathbb{Z}$ be an annulus of modulus $2h$, let $q \in Q^1(A)$ be an integrable quadratic differential and let A_η^+ and A_η^- be the two parallel subannuli at the two ends of modulus $\eta > 0$. If $\eta > \frac{1}{2\pi} \log(2 + \frac{1}{\varepsilon})$ and $\eta/h < \frac{\varepsilon}{2+4\varepsilon}$, then $\|q\|_{A_\eta^+} + \|q\|_{A_\eta^-} \leq \varepsilon \|q\|_A$ implies that there is a $c \in \mathbb{C}$ with*

$$\frac{\|q - c dz^2\|_A}{\|q\|_A} < 5\varepsilon .$$

Proof. Using Proposition 4.2 and its notation, write $q = q^+ + q^- + q^0$ with $q^0 = a_0 dz^2$ and estimate

$$\|q^+\|_A \leq \|q\|_{A_\eta^+} + \|q^+ - q\|_{A_\eta^+} + \|q^+\|_{A \setminus A_\eta^+} \leq \|q\|_{A_\eta^+} + (\varepsilon + \varepsilon) \|q\|_A .$$

The estimate for q^- is the same. The claim follows. □

4.3. The component of a quadratic differential adapted to a short closed geodesic. We now start to discuss an arbitrary hyperbolic Riemann surface X , using ideas from McMullen [M3]. If $\gamma \subset X$ is a simple closed geodesic of length $\ell(\gamma)$, we denote by $\pi_\gamma: \tilde{X}_\gamma \rightarrow X$ the covering space in which there is a unique simple closed geodesic $\tilde{\gamma}$, and $\tilde{\gamma}$ projects under π_γ to γ by an isometry. We will make the identification

$$\tilde{X}_\gamma = \left\{ z \in \mathbb{C} : |\operatorname{Im} z| < \frac{\pi}{2\ell(\gamma)} \right\} / \mathbb{Z},$$

so that $\tilde{\gamma} = \mathbb{R}/\mathbb{Z}$. Then $dz^2 \in Q^1(\tilde{X}_\gamma)$ with $\|dz^2\|_{X_\gamma} = \pi/\ell(\gamma)$, and the quadratic differential

$$q_\gamma = (\pi_\gamma)_* dz^2$$

is an element of $Q^1(X)$ with $\|q_\gamma\|_X \leq \frac{\pi}{\ell(\gamma)}$.

Again let A_γ be the standard collar around γ (compare Subsection A.1), and set $\operatorname{mod}(A_\gamma) =: M_\gamma = 2h_\gamma$. As a covering map, $\pi_\gamma: \tilde{X}_\gamma \rightarrow X$ is a local hyperbolic isometry. Let \tilde{A}_γ be the component of $(\pi_\gamma)^{-1}(A_\gamma)$ containing $\tilde{\gamma}$; hence

$$\tilde{A}_\gamma = \{z \in \mathbb{C} : |\operatorname{Im} z| < h_\gamma\} / \mathbb{Z}.$$

Set $h'_\gamma = h_\gamma - \sqrt{h_\gamma}$, and as before decompose the cylinder \tilde{A}_γ as $\tilde{A}_\gamma = \tilde{A}_\gamma^0 \cup \tilde{A}_\gamma^+ \cup \tilde{A}_\gamma^-$ corresponding, respectively, to points $z \in \tilde{A}_\gamma$ with

$$|\operatorname{Im} z| < h'_\gamma, \quad h'_\gamma \leq \operatorname{Im} z < h_\gamma, \quad \text{and} \quad -h_\gamma < \operatorname{Im} z \leq -h'_\gamma.$$

These project to annuli $A_\gamma^0, A_\gamma^+, A_\gamma^-$ in X by conformal isomorphisms; then A_γ^0 is the central collar around γ . Let γ^+ and γ^- be the components of the boundary ∂A_γ with $\gamma^+ \subset \partial A_\gamma^+$, $\gamma^- \subset \partial A_\gamma^-$, and let $\tilde{\gamma}^+$ and $\tilde{\gamma}^-$ be the corresponding boundary curves of \tilde{A}_γ .

In the proposition below, we speak of $q \in Q^1(A_\gamma)$. In practice, such a differential will be obtained by taking a differential in $Q^1(X)$, restricting it to A_γ , and extending it by 0 elsewhere. In particular, it will be quite discontinuous on the boundary of A_γ .

Proposition 4.4 (Decomposition of differentials on annuli in X). *For every $\varepsilon > 0$, there exists $\delta > 0$ (independent of X) such that if $\ell(\gamma) < \delta$, then for any $q \in Q^1(A_\gamma)$, there exist quadratic differentials q_γ^0, q_γ^+ and q_γ^- subject to the following conditions:*

- $q_\gamma^0 \in Q^1(X)$, $q_\gamma^+ \in Q^1(X \setminus \gamma^+)$, $q_\gamma^- \in Q^1(X \setminus \gamma^-)$;
- $q = q_\gamma^0 + q_\gamma^+ + q_\gamma^-$ on X (where we use the extension $q \equiv 0$ on $X \setminus A_\gamma$);
- the differential $q - q_\gamma^+$ has an analytic extension to a neighborhood of γ^+ , and $q - q_\gamma^-$ has an analytic extension to a neighborhood of γ^- ;
- these differentials satisfy the inequalities

$$(9) \quad \begin{aligned} \|q - q_\gamma^0\|_{A_\gamma^0} &\leq \varepsilon \|q\|_{A_\gamma}, & \|q_\gamma^0\|_{X \setminus A_\gamma^0} &\leq \varepsilon \|q\|_{A_\gamma}, \\ \|q - q_\gamma^+\|_{A_\gamma^+} &\leq \varepsilon \|q\|_{A_\gamma}, & \|q_\gamma^+\|_{X \setminus A_\gamma^+} &\leq \varepsilon \|q\|_{A_\gamma}, \\ \|q - q_\gamma^-\|_{A_\gamma^-} &\leq \varepsilon \|q\|_{A_\gamma}, & \|q_\gamma^-\|_{X \setminus A_\gamma^-} &\leq \varepsilon \|q\|_{A_\gamma}. \end{aligned}$$

Proof. Use the restricted projection $\pi_\gamma: \tilde{A}_\gamma \rightarrow A_\gamma$ to define a pull-back quadratic differential $\tilde{q} := \pi_\gamma^* q$ on \tilde{A}_γ and extend it by 0 to $\tilde{X}_\gamma \setminus \tilde{A}_\gamma$. We then have $\|q\|_{A_\gamma} = \|\tilde{q}\|_{\tilde{A}_\gamma}$. Expand \tilde{q} on \tilde{A}_γ into a Fourier series as in the beginning of Subsection 4.2,




FIGURE 3. On the left we see \tilde{X}_γ , identified with $\{z \in \mathbb{C} : |\operatorname{Im} z| < \pi/(2\ell(\gamma))\}$. On the right we see the Riemann surface X , together with the various subannuli. The key issue is that only a tiny part of \tilde{X}_γ , the little sliver at the top and the bottom, can project to $X \setminus A_\gamma$ (and part of the sliver also projects to A_γ). Thus the quadratic differential dz^2 , corresponding to the element of area on the left, has very little mass to spare for the remainder of X , or to compensate for the mass of \tilde{A}_γ^0 .

and decompose $\tilde{q} =: \tilde{q}^+ + \tilde{q}^0 + \tilde{q}^-$ corresponding to the positive terms of the Fourier series, the constant term $a_0 dz^2$, and the negative terms.

Recall that the series for \tilde{q}^+ converges if $\operatorname{Im} z < h_\gamma$ and that for \tilde{q}^0 it converges if $\operatorname{Im} z > -h_\gamma$; we extend them by 0 to $\{z : \operatorname{Im} z \geq h_\gamma\}$ and $\{z : \operatorname{Im} z \leq -h_\gamma\}$, respectively, so that they give quadratic differentials on \tilde{X}_γ holomorphic on the complement of $\tilde{\gamma}^+$ and $\tilde{\gamma}^-$, respectively. The differential \tilde{q}^0 is holomorphic on \tilde{X}_γ anyway, with $\tilde{q}^0 \in Q^1(\tilde{X}_\gamma)$.

Since \tilde{A}_γ is a parallel subannulus of \tilde{X}_γ with $\operatorname{mod}(\tilde{A}_\gamma) \geq \operatorname{mod}(\tilde{X}_\gamma) - 1$, the complement $\tilde{X}_\gamma \setminus \tilde{A}_\gamma$ consists of two parallel subannuli E^+ and E^- , each of modulus at most 1 and so that E^+ is adjacent to \tilde{A}_γ^+ and E^- is adjacent to \tilde{A}_γ^- .

Define quadratic differentials on \tilde{X}_γ as

$$\tilde{q}_{\text{rest}}^+ := (\tilde{q}^- + \tilde{q}^0)|_{E^+} \quad \text{and} \quad \tilde{q}_{\text{rest}}^- := (\tilde{q}^+ + \tilde{q}^0)|_{E^-}$$

(the restrictions to E^\pm , extended by zero elsewhere). This gives an exact equality $\tilde{q} + \tilde{q}_{\text{rest}}^+ + \tilde{q}_{\text{rest}}^- = \tilde{q}^+ + \tilde{q}^0 + \tilde{q}^-$ on all of $\tilde{X} = E^+ \cup \tilde{A} \cup E^-$. We can thus write

$$\tilde{q} = (\tilde{q}^- - \tilde{q}_{\text{rest}}^-) + \tilde{q}^0 + (\tilde{q}^+ - \tilde{q}_{\text{rest}}^+).$$

While this is valid on \tilde{X} , the left hand side has its support on \tilde{A} ; on the right hand side, the first term is discontinuous only at $\tilde{\gamma}^-$, the last term is discontinuous only at $\tilde{\gamma}^+$, and the middle term is continuous everywhere.

Thus we can define

$$q^0 := (\pi_\gamma)_* \tilde{q}^0, \quad q^+ := (\pi_\gamma)_*(\tilde{q}^+ - \tilde{q}_{\text{rest}}^+) \quad \text{and} \quad q^- := (\pi_\gamma)_*(\tilde{q}^- - \tilde{q}_{\text{rest}}^-).$$

This yields holomorphic quadratic differentials $q^0 \in Q^1(X)$, $q^+ \in Q^1(X \setminus \gamma^+)$ and $q^- \in Q^1(X \setminus \gamma^-)$. Moreover, we have $q^+ + q^0 + q^- = q$ on all of X (where we have extended q by zero to $X \setminus A_\gamma$). The construction assures that $q - q^+ = q^0 + q^-$ has a holomorphic extension to a neighborhood of γ^+ and $q - q^- = q^0 + q^+$ has a holomorphic extension to a neighborhood of γ^- .

Now subdivide $A_\gamma = A_\gamma^+ \cup A_\gamma^0 \cup A_\gamma^-$ as above into disjoint parallel subannuli, and similarly for \tilde{A}_γ . We will show that q^- , q^0 and q^+ satisfy the inequalities as claimed in the proposition.

On A_γ^+ , we have the decomposition

$$\begin{aligned} (q - q^+)|_{A_\gamma^+} &= (\pi_\gamma)_* \left((\tilde{q} - \tilde{q}^+)|_{\tilde{A}_\gamma^+} \right) \\ &\quad - (\pi_\gamma)_* \left(\tilde{q}^+|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^+} \right) \\ &\quad + (\pi_\gamma)_* \left((\tilde{q}^0 + \tilde{q}^-)|_{E^+} \right), \end{aligned}$$

because $q|_{A_\gamma} = (\pi_\gamma)_*(\tilde{q})$, the differential \tilde{q} has its support in \tilde{A}_γ , and $\pi_\gamma^{-1}(A_\gamma^+) \cap \tilde{A}_\gamma = \tilde{A}_\gamma^+$.

Since the push-forward $(\pi_\gamma)_*$ does not increase the L^1 -norm of quadratic differentials, we have

$$(10) \quad \|q - q^+\|_{A_\gamma^+} \leq \|\tilde{q} - \tilde{q}^+\|_{\tilde{A}_\gamma^+} + \|\tilde{q}^+\|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^+} + \|\tilde{q}^0\|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^0} + \|\tilde{q}^-\|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^-}.$$

Now suppose that δ is sufficiently small so that $\ell(\gamma) < \delta$ implies

$$\text{mod}(A_\gamma^\pm) = \sqrt{h_\gamma} > \frac{1}{2\pi} \log \left(2 + \frac{1}{\varepsilon/4} \right)$$

and

$$\frac{\text{mod}(A_\gamma^0)}{\text{mod}(A_\gamma)} = 1 - \frac{1}{\sqrt{h_\gamma}} > 1 - \frac{\varepsilon/4}{1 + 2\varepsilon/4}.$$

Then we can apply Proposition 4.2 with an error term $\varepsilon/4$, and (10) becomes

$$\|q - q^+\|_{A_\gamma^+} \leq 4 \frac{\varepsilon}{4} \|\tilde{q}\|_{\tilde{A}_\gamma} = \varepsilon \|q\|_{A_\gamma}.$$

Similarly, from the decompositions

$$\begin{aligned} q^+|_{X \setminus A_\gamma^+} &= (\pi_\gamma)_* \left(\tilde{q}^+|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^+} \right) - (\pi_\gamma)_* \left((\tilde{q}^0 + \tilde{q}^-)|_{E^+} \right), \\ (q - q^0)|_{A_\gamma^0} &= (\pi_\gamma)_* \left((\tilde{q} - \tilde{q}^0)|_{\tilde{A}_\gamma^0} \right) - (\pi_\gamma)_* \left(\tilde{q}^0|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^0} \right), \\ q^0|_{X \setminus A_\gamma^0} &= (\pi_\gamma)_* \left(\tilde{q}^0|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^0} \right), \end{aligned}$$

we obtain

$$\begin{aligned} \|q^+\|_{X \setminus A_\gamma^+} &\leq \|\tilde{q}^+\|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^+} + \|\tilde{q}^0\|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^0} + \|\tilde{q}^-\|_{\tilde{X} \setminus \tilde{A}_\gamma^-} \leq 3\frac{\varepsilon}{4} \|\tilde{q}\|_{\tilde{A}_\gamma} \leq \varepsilon \|q\|_{A_\gamma}, \\ \|q - q^0\|_{A_\gamma^0} &\leq \|\tilde{q} - \tilde{q}^0\|_{\tilde{A}_\gamma^0} + \|\tilde{q}^0\|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^0} \leq 2\frac{\varepsilon}{4} \|\tilde{q}\|_{\tilde{A}_\gamma} \leq \varepsilon \|q\|_{A_\gamma}, \\ \|q^0\|_{X \setminus A_\gamma^0} &\leq \|\tilde{q}^0\|_{\tilde{X}_\gamma \setminus \tilde{A}_\gamma^0} \leq \frac{\varepsilon}{4} \|\tilde{q}\|_{\tilde{A}_\gamma} \leq \varepsilon \|q\|_{A_\gamma}. \end{aligned}$$

The case of q^- is similar. \square

4.4. The component of a quadratic differential adapted to the thick part.

Given a finite type hyperbolic Riemann surface X and a length $\delta \in (0, \ell^*)$, there is a bounded finite number of simple closed geodesics γ_i of length less than δ . Each of them has an associated standard collar A_i with core curve γ_i and a parallel central collar $A_i^0 \subset A_i$.

Now consider a quadratic differential $q \in Q^1(X)$ and a particular standard collar A_i . Applying Proposition 4.4 to the restriction $q|_{A_i}$, we associate to A_i a quadratic differential $q_{A_i} := q_{\gamma_i}^0 \in Q^1(X)$, so that for all i , we have the inequalities $\|q - q_{A_i}\|_{A_i^0} \leq (\varepsilon/4) \|q\|_{A_i}$ and $\|q_{A_i}\|_{X \setminus A_i^0} \leq (\varepsilon/4) \|q\|_{A_i}$, provided δ is sufficiently small depending only on ε . We also have two additional quadratic differentials $q_{A_i}^+$ and $q_{A_i}^-$ which approximate q near the two ends of A_i .

This automatically defines the domains Y_j , as the components of $X \setminus \bigcup_i A_i^0$. Every component Y_j contains one component of $X \setminus \bigcup_i A_i$ (the actual thick part) plus finitely many adjacent annuli of diverging moduli (the ‘‘boundary padding’’). We still need to define the q_{Y_j} associated to the thick parts Y_j and show that they satisfy the corresponding inequalities of Theorem 4.1.

For each Y_j , consider its adjacent annuli A_{i_1}, \dots, A_{i_n} . To lighten notation, we will write Y for Y_j and denote the adjacent annuli by A_1, \dots, A_n . These will be oriented so that A_i^+ is on the same side of γ_i as Y (suppose for now that no A_i is adjacent to Y on both ends). Write

$$Z := Y \setminus \bigcup_{i=1}^n A_i$$

and consider the (non-analytic) quadratic differential $q|_Z$ obtained by restricting q to Z and extending by 0 to $X \setminus Z$ (see Figure 4).

We can now define q_Y to be

$$(11) \quad q_Y = q|_Z + \sum_{i=1}^n q_{A_i}^+.$$

In the case that an annulus A_{γ_i} is adjacent to Y on both ends, it will be counted twice, and both differentials $q_{A_i}^-$ and $q_{A_i}^+$ are included in the sum for q_Y .

Proof of Theorem 4.1. We start by observing that q_Y is holomorphic on X : it is by definition holomorphic on $X \setminus \bigcup_i \gamma_i^+$, and by Proposition 4.4, $q|_{A_j} - q_{A_j}^+$ is holomorphic near γ_j^+ for every j . Now in a neighborhood of γ_j^+ ,

$$q_Y = (q|_Z + q|_{A_j}) - (q|_{A_j} - q_{A_j}^+) + \sum_{i \neq j} q_{A_i}^+,$$

so q_Y is holomorphic in a neighborhood of γ_j^+ also.




FIGURE 4. One thick piece Y with its adjacent annuli A_1, \dots, A_n . (Some annuli A_i might be adjacent to Y at both ends.)

We claim that $q = \sum_Y q_Y + \sum_A q_A$, where \sum_Y and \sum_A run over the thick and thin parts of the δ -decomposition of X . Indeed,

$$q = \sum_Z q|_Z + \sum_A q|_A = \sum_Z q|_Z + \sum_A (q_{A^-} + q_A + q_{A^+}) = \sum_Y q_Y + \sum_A q_A$$

because every A^- and every A^+ is adjacent to exactly one component Y .

Now we start to establish the inequalities. We have

$$\|q - q_Y\|_Z \leq \sum_i \|q_{A_i^+}\|_Z$$

and

$$\|q - q_Y\|_{A_j^+} \leq \sum_{i \neq j} \|q_{A_i^+}\|_{A_j^+},$$

hence

$$\|q - q_Y\|_Y \leq \sum_i \|q_{A_i^+}\|_{X \setminus A_i^+} \leq \varepsilon' \sum_i \|q\|_{A_i}$$

where the summation is over the annuli A_i adjacent to Y (counting an annulus twice if both ends are adjacent to Y). Similarly,

$$\|q_Y\|_{X \setminus Y} \leq \sum_i \|q_{A_i^+}\|_{X \setminus A_i^+} \leq (\varepsilon/4) \sum_i \|q\|_{A_i}.$$

Combining this with the inequalities on q_{A_i} from the beginning of Subsection 4.4, we obtain

$$\sum_Y \|q - q_Y\|_Y + \sum_Y \|q_Y\|_{X \setminus Y} + \sum_A \|q - q_A\|_{A^0} + \sum_A \|q_A\|_{X \setminus A^0} \leq \varepsilon \sum_A \|q\|_{A}.$$

The sum \sum_Y runs over all thick components, and together they count all annuli exactly twice (once from each end); the sum \sum_A runs over all thin annuli. \square

5. LIMITS OF QUADRATIC DIFFERENTIALS

In Theorem 4.1, we have described how a meromorphic integrable quadratic differential on a Riemann surface of finite type can be decomposed according to a decomposition of the Riemann surface into annuli around short simple closed geodesics and the complementary components. In this section, we describe the possible limit models when the lengths of the simple closed geodesics tend to zero. This will be used in Section 6 to prove our key Proposition 3.2.

5.1. Limit models. We continue to investigate sequences of quadratic differentials (such as those arising as cotangent vectors in the iteration procedure in Teichmüller space, especially when the contraction under the push-forward tends to zero). In particular, we want to describe possible limits (for each of the components in the thick-thin decomposition as described in Section 4) that can arise in such sequences: we show that up to choosing a subsequence, we can describe simple models for the limiting mass distribution of holomorphic quadratic differentials on the complex plane. For our purposes it is sufficient to investigate quadratic differentials up to automorphisms of their domain \mathbb{C} , so we replace a sequence of quadratic differentials q_n by a rescaled version $M_n^*q_n$, where the M_n are automorphisms of \mathbb{C} . (The second of our two possibilities involves holomorphic quadratic differentials on \mathbb{C}^* .)

Definition 5.1 (Limit model for quadratic differential). Let $(q_n)_{n \in \mathbb{N}}$ be a sequence of measurable quadratic differentials on \mathbb{C} with $0 < \|q_n\|_{\mathbb{C}} < \infty$ for each n , such that $\lim_{n \rightarrow \infty} \|q_n\|_{\mathbb{C}}$ exists and is non-zero.

- If q is a meromorphic quadratic differential with $0 < \|q\|_{\mathbb{C}} < \infty$, then we say that (q_n) has *limit model* q if there exist conformal automorphisms $M_n(z) = a_n z + b_n$ of \mathbb{C} (with $a_n, b_n \in \mathbb{C}$) satisfying

$$\lim_{n \rightarrow \infty} \|M_n^*q_n - q\|_{\mathbb{C}} = 0.$$

In this case, we call M_n the *scaling* and a_n the *scaling factor*.

- We say that the sequence (q_n) has *limit model* dz^2/z^2 on annuli

$$A_n = \{z \in \mathbb{C} : r_n < |z - b_n| < R_n\} \quad (\text{with } 0 < r_n < R_n \text{ and } b_n \in \mathbb{C})$$

if there exist complex numbers $c_n \neq 0$ such that $R_n/r_n \rightarrow \infty$ and the affine maps $M_n(z) = z + b_n$ satisfy

$$\lim_{n \rightarrow \infty} \left\| M_n^*q_n - c_n \frac{dz^2}{z^2} \Big|_{\{r_n < |z| < R_n\}} \right\|_{\mathbb{C}} = 0.$$

In this case, we call r_n and R_n the *inner and outer radii* of A_n and b_n the *center* of A_n . (Note that r_n , R_n and b_n are not unique.)

In the first case, it is easy to see that $\|q_n\|_{\mathbb{C}} \rightarrow \|q\|_{\mathbb{C}}$. In the second case, the differential $c_n \frac{dz^2}{z^2} \Big|_{\{r_n < |z| < R_n\}}$ is a finite-mass restriction of $c_n \frac{dz^2}{z^2}$ to the annulus between radii r_n and R_n , so in particular the claim implies that

$$\|M_n^*q_n\|_{\{|z| < r_n\} \cup \{|z| > R_n\}} \rightarrow 0.$$

Note that the q_n in this definition are not global meromorphic differentials on $\overline{\mathbb{C}}$ (unlike differentials q_n used in Theorem 5.2): our quadratic differentials are not

required to be continuous; they will be the restrictions of meromorphic quadratic differentials to measurable subsets of \mathbb{C} .

Figure 5 shows a sequence of quadratic differentials in which both kinds of limit models occur.




FIGURE 5. The shaded figure is a polygon P_n representing a sequence of quadratic differentials q_n on $\overline{\mathbb{C}}$. The symbols $0, 1, \infty$ written inside P_n indicate that the double \tilde{P}_n should be normalized so that the three labeled vertices go to $0, 1,$ and ∞ , respectively. The numbers written around P_n (and the $1/n$ written inside it) are intended to be the actual Euclidean lengths of the sides, so P_n has area $13 + 1/n$, and since q_n is a quadratic differential on the double \tilde{P}_n of P_n , we have $\|q_n\| = 26 + 2/n$ which does not tend to 0 or ∞ . The symbols $0, 1, \infty$ written inside indicate that those points of the double of the polygon correspond to $0, 1, \infty$ on the Riemann sphere. Around it are the three limit models of the thick parts. The actual limit of the quadratic differentials “is” the model on the right, because $0, 1, \infty$ belong to distinct components of the complement, or to the region itself. This convergence is actually uniform on any compact subset of $\overline{\mathbb{C}}$ omitting 0 and ∞ . Small neighborhoods of these points need to be blown up to “see” the other limit models. There is only one thin part contributing a region to Theorem 5.2, the one labeled A . Although the part labeled B is an annulus of modulus n , it carries mass $2/n$ which disappears in the limit. The region labeled C (shaded dark) is also a large annulus, of modulus $\sim (\log n)/(2\pi)$, and it also contributes a definite share of the mass, since its area is $\sim \pi$ and does not tend to 0 . But the quadratic differential in that region does not behave like dz^2/z^2 (it surrounds as many zeroes as poles of q).

Now we show that every sequence of integrable meromorphic quadratic differentials with bounded numbers of poles and constant mass has a subsequence so that the entire mass is recovered in finitely many limit models.

Theorem 5.2 (Decomposition of mass). *Let (q_n) be a sequence of meromorphic quadratic differentials on $\overline{\mathbb{C}}$ with at most N poles and $\|q_n\|_{\mathbb{C}} = 1$ for all n . Then there are a subsequence (also denoted (q_n)), a number $l \in \mathbb{N}$ and disjoint regions $V_n^{[1]}, \dots, V_n^{[l]} \subset \overline{\mathbb{C}}$ for each n with the following properties:*

(a) *for each j , the non-zero limit $\lim_{n \rightarrow \infty} \|q_n\|_{V_n^{[j]}}$ exists in $(0, 1]$, while*

$$\lim_{n \rightarrow \infty} \|q_n\|_{\overline{\mathbb{C}} \setminus \bigcup_{j=1}^l V_n^{[j]}} = 0;$$

(b) *for each $j = 1, 2, \dots, l$, the restriction $q_n|_{V_n^{[j]}}$ has limit model $q_{\infty}^{[j]}$, where either*

(b1) *$q_{\infty}^{[j]}$ is a meromorphic quadratic differential on $\overline{\mathbb{C}}$ with at most N poles and $0 < \|q_{\infty}^{[j]}\|_{\mathbb{C}} \leq 1$; in this case, for each pole w_{∞} of $q_{\infty}^{[j]}$, there exists a pole w_n for q_n such that $(M_n)^{-1}(w_n) \rightarrow w_{\infty}$, where M_n is the scaling associated to q_n ; or*

(b2) *$q_{\infty}^{[j]} = \frac{dz^2}{z^2}$, and $V_n^{[j]}$ is contained in the annulus A_n associated to q_n in Definition 5.1; moreover, the bounded component of $\mathbb{C} \setminus V_n^{[j]}$ contains at least two poles of q_n .*

Proof. Let $Z^0(q_n)$ be the set of poles of q_n , and define $Z(q_n) := Z^0(q_n) \cup \{\infty\}$ and $X_n = \overline{\mathbb{C}} \setminus Z(q_n)$: each X_n is a Riemann surface on which we have a holomorphic quadratic differential $q_n \in Q^1(X_n)$, to which we will be able to apply Theorem 4.1. By choosing subsequences repeatedly, we may assume:

- The number of poles of q_n on $\overline{\mathbb{C}}$ and on \mathbb{C} is constant, so that all X_n have the same topology. Moreover, the poles are labeled so that the cross-ratios of all quadruples of distinct poles or of any triple of distinct poles and ∞ converge in the Riemann sphere as $n \rightarrow \infty$.
- There is an integer s , a length $\delta^* < \log(3 + 2\sqrt{2})$ and a sequence $\delta_n \rightarrow 0$ such that on every X_n , there are precisely s simple closed geodesics $\gamma_n^{[1]}, \dots, \gamma_n^{[s]}$ of length $\ell(\gamma_n^{[i]}) < \delta^*$, and these lengths satisfy $\ell(\gamma_n^{[i]}) \leq \delta_n$. Let $I := \{1, \dots, s\}$. (To accomplish this, note first that by Corollary A.2, the number of simple closed geodesics of length less than $\log(3 + \sqrt{2})$ is bounded by $N - 3$; by restricting to a subsequence, we may assume that all their lengths converge. Relabel so that $\ell(\gamma_n^{[1]}) \leq \ell(\gamma_n^{[2]}) \leq \dots$, and let s be the largest index with $\lim \ell(\gamma_n^{[s]}) = 0$. Then δ^* is half the smallest positive limit length, or $\delta^* = 1$ if all limit lengths are zero. Finally, let $\delta_n = \max\{\ell(\gamma_n^{[1]}), \dots, \ell(\gamma_n^{[s]})\}$.)
- The components of $X_n \setminus \bigcup_i \gamma_n^{[i]}$ can be labeled by a set J so that components with a given label have the same boundary curves and contain the same number of poles, independently of n .

Around each curve $\gamma_n^{[i]}$ put the standard collar $A_n^{[i]}$ and the central collar $B_n^{[i]} \subset A_n^{[i]}$ (compare Subsection 4.1; we write $B_n^{[i]}$ instead of $(A_n^{[i]})^0$), and let $Y_n^{[j]}$ be the components of

$$X_n \setminus \bigcup_{i=1}^s B_n^{[i]}.$$

By passing to a further subsequence, we may assume that the masses

$$b_n^{[i]} := \|q_n\|_{B_n^{[i]}} \quad \text{and} \quad c_n^{[j]} := \|q_n\|_{Y_n^{[j]}}$$

converge to $b^{[i]}$ and $c^{[j]}$, respectively. Note that $\sum_i b^{[i]} + \sum_j c^{[j]} = 1$.

The decomposition of X_n into the components $B_n^{[i]}$ and $Y_n^{[j]}$ is related to the standard thick-thin decomposition of X_n for the constant δ^* , but there are two significant differences. First, the punctures are all in the thick part; by our construction none of the poles are in the closure in $\overline{\mathbb{C}}$ of the $A_n^{[i]}$, much less in the closure of the $B_n^{[i]}$. Second, since the padding $A_n^{[i]} \setminus B_n^{[i]}$ has been added to the thick parts, but itself has modulus tending to infinity with n , the boundary curves of the $Y_n^{[j]}$ have hyperbolic lengths tending to 0 with n .

In our limit $n \rightarrow \infty$, we have $\delta_n \rightarrow 0$ in our δ_n -decomposition of q_n on X_n . Therefore, there exists a sequence $\varepsilon_n \rightarrow 0$ such that $\delta_n < \delta_0(\varepsilon_n)$, where $\delta_0(\varepsilon)$ is as in Theorem 4.1. We restrict to such n for which $\delta_n < \delta^*$. We will use Theorem 4.1 to decompose the quadratic differentials q_n on X_n with error size ε_n as follows: for each region $B_n^{[i]} \subset A_n^{[i]}$ and $Y_n^{[j]}$, we have a quadratic differential $q_{A_n^{[i]}} \in Q^1(\mathbb{C})$ or $q_{Y_n^{[j]}} \in Q^1(\mathbb{C})$ which ‘‘represents’’ q_n in its respective region $B_n^{[i]}$ or $Y_n^{[j]}$, so that $\|q_n|_{B_n^{[i]}} - q_{A_n^{[i]}}\|_{\mathbb{C}} \leq 2\varepsilon_n \|q_n\|_{\mathbb{C}} \rightarrow 0$ and $\|q_n|_{Y_n^{[j]}} - q_{Y_n^{[j]}}\|_{\mathbb{C}} \leq 2\varepsilon_n \|q_n\|_{\mathbb{C}} \rightarrow 0$. Hence, as $n \rightarrow \infty$, we have $\|q_{A_n^{[i]}}\|_{\mathbb{C}} \rightarrow b^{[i]}$ and $\|q_{Y_n^{[j]}}\|_{\mathbb{C}} \rightarrow c^{[j]}$. Let $V_n^{[1]}, \dots, V_n^{[l]}$ be the collection of $B_n^{[i]}$ with $b^{[i]} > 0$ and $Y_n^{[j]}$ with $c^{[j]} > 0$. For such regions it suffices, in view of these estimates, to show that the $q_{A_n^{[i]}}$ and $q_{Y_n^{[j]}}$ have limit models. This will be done in Lemmas 5.3 and 5.4 below. (For the remaining regions with $b^{[i]} = 0$ or $c^{[j]} = 0$, we could say that the limit model is zero.) It also follows that $\lim_{n \rightarrow \infty} \|q_n\|_{\overline{\mathbb{C}} \setminus \bigcup_{j=1}^l V_n^{[j]}} = 0$.

5.2. Thin parts.

Lemma 5.3 (Limit model on thin parts). *For every index i with $b^{[i]} > 0$, $q_{A_n^{[i]}}$ has limit model dz^2/z^2 ; moreover, $B_n^{[i]}$ surrounds at least two poles of q_n .*

Proof. The annulus $B_n^{[i]}$ surrounds at least two poles of q_n because its core curve is essential by construction. Let w_n be one of them; after an affine change of coordinates, we may assume that $w_n = 0$. Let \hat{A}_n be the largest round annulus centered at w_n with $\hat{A}_n \subset A_n^{[i]}$, let $\hat{A}_n^0 \subset \hat{A}_n$ be the smallest parallel subannulus containing $B_n^{[i]}$, and let \hat{A}_n^+ and \hat{A}_n^- be the outer and inner components of $\hat{A}_n \setminus \hat{A}_n^0$ (see Figure 6). The annuli exist and are non-empty by the Round Annuli Lemma A.3, and there is a universal constant μ^* with

$$\text{mod}(A_n^{[i]}) \geq \text{mod}(\hat{A}_n) \geq \text{mod}(A_n^{[i]}) - \mu^* ;$$

since the two components of $A_n^{[i]} \setminus B_n^{[i]}$ contain \hat{A}_n^+ and \hat{A}_n^- , we get

$$\sqrt{\text{mod}(A_n^{[i]})} \geq \text{mod}(\hat{A}_n^\pm) \geq \sqrt{\text{mod}(A_n^{[i]})} - \mu^* \longrightarrow \infty .$$

In particular,

$$\text{mod}(\hat{A}_n^\pm)/\text{mod}(\hat{A}_n) \longrightarrow 0 .$$




FIGURE 6. Since the padding $A_n^{[i]} \setminus B_n^{[i]}$ is itself the union of two annuli of large modulus (even though much smaller than the modulus of $A_n^{[i]}$), each contains a large round subannulus, of approximately the same modulus; we may take these annuli \hat{A}_n^\pm to be centered at the same point w_n . By definition \hat{A}_n is the smallest annulus containing both, and $\hat{A}_n^0 = \hat{A}_n \setminus \bigcup \hat{A}_n^\pm$.

Since $\hat{A}_n^\pm \subset X_n \setminus B_n^{[i]}$ and $\hat{A}_n \supset B_n^{[i]}$, Theorem 4.1 implies that $\|q_{A_n}\|_{B_n^{[i]}} \rightarrow b^{[i]} > 0$ and $\|q_{A_n}\|_{X_n \setminus B_n^{[i]}} \rightarrow 0$, so we have

$$\frac{\|q_{A_n}\|_{\hat{A}_n^\pm}}{\|q_{A_n}\|_{\hat{A}_n}} \leq \frac{\|q_{A_n}\|_{X_n \setminus B_n^{[i]}}}{\|q_{A_n}\|_{B_n^{[i]}}} \rightarrow 0.$$

Let $h_n := \text{mod}(\hat{A}_n)/2$, so that \hat{A}_n is conformally equivalent to

$$\hat{B}_n = \{z \in \mathbb{C}: |\text{Im}(z)| \leq h_n\}/\mathbb{Z};$$

conformal isomorphisms $\psi_n: \hat{B}_n \rightarrow \hat{A}_n$ have the form $\psi_n(z) = \alpha_n(\exp(2\pi iz))$ with $\alpha_n \in \mathbb{C}$. Define $Q_n := \psi_n^* q_{A_n}$. Then Corollary 4.3 applies to the restricted differentials $Q_n|_{\hat{B}_n}$; since $\|Q_n\|_{\hat{B}_n} = \|q_{A_n}\|_{\hat{A}_n} \rightarrow b^{[i]}$, this yields

$$\left\| Q_n|_{\hat{B}_n} - c_n dz^2 \right\|_{\hat{B}_n} \rightarrow 0$$

for an appropriate sequence of numbers $c_n \in \mathbb{C}$. But since $(\psi_n)_*(dz^2) = -dz^2/(2\pi z)^2$, we obtain

$$\left\| q_{A_n} - \frac{-c_n}{(2\pi)^2} \frac{dz^2}{z^2} \right\|_{\hat{A}_n} \longrightarrow 0 .$$

Since $\|q_{A_n}\|_{\mathbb{C} \setminus \hat{A}_n} \rightarrow 0$, the claim follows. \square

5.3. Thick parts. From now on, consider an index j describing thick regions $Y_n^{[j]}$ with $c^{[j]} = \lim_{n \rightarrow \infty} c_n^{[j]} > 0$.

Lemma 5.4 (Limit model on thick parts). *There is a limiting integrable meromorphic quadratic differential $q_\infty^{[j]}$ on $\overline{\mathbb{C}}$ with $\|q_\infty^{[j]}\|_{\mathbb{C}} = c^{[j]}$ and a sequence of Möbius transforms $M_n(z) = a_n z + b_n$ so that $\|M_n^* q_{Y_n^{[j]}} - q_\infty^{[j]}\|_{Y_n^{[j]}} \rightarrow 0$.*

Proof. Choose two poles $w_n^0, w_n^1 \in \mathbb{C}$ of q_n with the same labels for all n , not belonging to the same component of $\mathbb{C} \setminus Y_n^{[j]}$, and not belonging to the unbounded component of $\mathbb{C} \setminus Y_n^{[j]}$ (if any); it is perfectly acceptable to have w_n^0 and/or w_n^1 in $\overline{Y}_n^{[j]}$. Such poles exist: each component of $\mathbb{C} \setminus \overline{Y}_n^{[j]}$ contains two elements of $\overline{\mathbb{C}} \setminus X_n$ (or the boundary annuli of $Y_n^{[j]}$ would not be essential); each element is either a pole of q_n or the point ∞ . If $X_n \setminus \overline{Y}_n^{[j]}$ has at least three components, then let w_n^0 and w_n^1 be two poles from two different bounded components. If $X_n \setminus \overline{Y}_n^{[j]}$ has only two components, then at least one of them is bounded, and additionally $\overline{Y}_n^{[j]}$ must contain one pole (or the two annuli in the boundary padding would be homotopic). If $X_n \setminus \overline{Y}_n^{[j]}$ has only one component, then $\overline{Y}_n^{[j]}$ must contain two poles itself (or the adjacent annulus $B_n^{[j]}$ would not be essential), and if $X_n \setminus \overline{Y}_n^{[j]}$ is empty, then $\overline{Y}_n^{[j]}$ must contain at least three poles, at most one of which is at ∞ .

Let $M_n \in \text{Aut}(\mathbb{C})$ be given by $M_n(z) = a_n z + b_n := (w_n^1 - w_n^0)z + w_n^0$, so that $M_n(p) = w_n^p$ for $p = 0, 1$. By construction, all poles of $M_n^* q_n$ converge in $\overline{\mathbb{C}}$; let $\{w_0 = 0, w_1 = 1, \dots, w_s = \infty\}$ be the set of their limits, including ∞ even if it is not a limit of poles. Set $X^{[j]} := \overline{\mathbb{C}} \setminus \{w_0, \dots, w_s\}$.

We argue in two steps: first we show that $Y'_n := M_n^{-1} Y_n^{[j]}$ fills $X^{[j]}$ as $n \rightarrow \infty$ (in the sense that every compact $K \subset X^{[j]}$ is contained in almost all Y'_n). Then we show that, after choosing a further subsequence, there is a $q_\infty^{[j]} \in Q^1(\mathbb{C})$ with the required properties.

For the first step, recall that $Y_n^{[j]}$ is surrounded by annuli with diverging moduli. Let D_n be an open component of $\overline{\mathbb{C}} \setminus Y'_n$; it contains at least two poles of $M_n^* q_n$ (or one pole and ∞), and the positions of the poles converge in $\overline{\mathbb{C}}$. There is an annulus $C_n \subset Y'_n$ which separates D_n from all poles in Y'_n and all components of $\overline{\mathbb{C}} \setminus Y'_n$ other than D_n (this annulus is part of the boundary padding in Y'_n), and $\text{mod}(C_n) \rightarrow \infty$. By construction, C_n separates D_n from at least two of the three points $0, 1, \infty$. This implies that the spherical diameter of D_n tends to zero: for every four distinct points $a_1, a_2, b_1, b_2 \in \overline{\mathbb{C}}$, the modulus of annuli separating $\{a_1, a_2\}$ from $\{b_1, b_2\}$ is bounded above.

Since $\overline{\mathbb{C}} \setminus Y'_n$ consists of a bounded number of open components, and each component either converges to ∞ or has diameter tending to 0, it follows that Y'_n fills $X^{[j]}$ as claimed. Distinct sequences of poles in \overline{Y}'_n or components of $\overline{\mathbb{C}} \setminus \overline{Y}'_n$ converge to

distinct points in $\overline{\mathbb{C}} \setminus X^{[j]}$: otherwise there would be a short curve which separated the poles of q_n in a way the previously chosen short curves γ_i never did.

Now for convergence of the $\hat{q}_n := M_n^* q_{Y_n^{[j]}}$, which are meromorphic quadratic differentials in $Q^1(M_n^{-1}(Z(q_n)))$. For one thing, there is no problem in choosing a subsequence that converges algebraically in the following sense: if we write $\hat{q}_n = \frac{P_n(z)}{Q_n(z)} dz^2$ with polynomials $P_n(z)$ and $Q_n(z)$, after adjusting by multiplicative constants and extracting a subsequence, we may suppose that P_n and Q_n converge to limit polynomials P_∞ and Q_∞ , so that at least one of them is not identically 0. Moreover we have $Q_\infty \not\equiv 0$; otherwise in a neighborhood of a point z_0 where $P_\infty(z_0) \neq 0$, the mass of \hat{q}_n must tend to ∞ , and this contradicts the fact that $\|\hat{q}_n\|_{\mathbb{C}}$ is uniformly bounded. Therefore $\hat{q}_n \rightarrow q_\infty^{[j]} := \frac{P_\infty(z)}{Q_\infty(z)} dz^2$ uniformly on compact subsets of $\overline{\mathbb{C}} \setminus X^{[j]}$. (At this point, we cannot exclude the possibility that $P_\infty \equiv 0$.)

The differential $q_\infty^{[j]}$ cannot have a multiple pole: otherwise, there would be a compact set near a multiple pole of $q_\infty^{[j]}$ in which the mass of $q_\infty^{[j]}$ exceeded the available mass of the \hat{q}_n . It follows that $q_\infty^{[j]} \in Q^1(X^{[j]})$ and the poles of $q_\infty^{[j]}$ are the limit of poles of \hat{q}_n . (The same conclusion can be reached using the Cauchy formula and the uniform bound on the L^1 -norm, together with the Arzela-Ascoli theorem and Fatou's lemma.)

It remains to show that the sequence \hat{q}_n converges in L^1 -norm, since this will imply $\|\hat{q}_n - q_\infty^{[j]}\|_{\mathbb{C}} \rightarrow 0$ and $\|q_\infty^{[j]}\|_{\mathbb{C}} = c^{[j]}$. Since we had shown above that \hat{q}_n converges in $\overline{\mathbb{C}} \setminus X^{[j]}$, it suffices to show the following:

Claim 5.5 (Small mass in small neighborhood). For every $w \in \overline{\mathbb{C}} \setminus X^{[j]}$ and $\varepsilon > 0$, there exists $\delta > 0$ such that $\|\hat{q}_n\|_{N_\delta(w)} < \varepsilon$ for sufficiently large n , where $N_\delta(w)$ is the δ -neighborhood of w in the spherical metric.

If $w \in \overline{\mathbb{C}} \setminus X^{[j]}$ is the limit of simple poles of \hat{q}_n in $\overline{Y_n^{[j]}}$, then the claim follows from the following fact whose proof is left to the reader (it can be derived from Proposition 4.2 via $z \mapsto e^{2\pi iz}$).

Claim 5.6 (Mass near simple pole). Suppose $h(z)$ is a meromorphic function in $\mathbb{D} = D_1(0)$ with at most one simple pole in $D_{1/2}(0)$ and no poles in the annulus $\mathbb{D} \setminus D_{1/2}(0)$. For every $\varepsilon > 0$, there exists a $\delta > 0$ (depending only on ε , not on h) such that $\|h\|_{D_\delta(0)} \leq \varepsilon \|h\|_{D_1(0)}$.

The remaining case is when $w \in \overline{\mathbb{C}} \setminus X^{[j]}$ is the limit of components of $\overline{\mathbb{C}} \setminus \overline{Y_n^{[j]}}$. Even though we already know that $\|\hat{q}_n\|_{\overline{\mathbb{C}} \setminus \overline{Y_n^{[j]}}} \rightarrow 0$ by Theorem 4.1, it is conceivable that some or all of the mass of the \hat{q}_n could drift into the punctures (such as w) that the complement of $Y_n^{[j]}$ becomes. While $\|\hat{q}_n\|_{Y_n^{[j]}} \rightarrow c^{[j]} > 0$, the $Y_n^{[j]}$ contain annuli in the boundary padding with diverging moduli, and parts of these annuli tend into the punctures. We need to know that almost all the mass of \hat{q}_n is not near the punctures.

Given w as above, for each n , there exist a component D_n of $\overline{\mathbb{C}} \setminus \overline{Y_n^{[j]}}$ and a pole (or ∞) w_n of \hat{q}_n within D_n such that D_n and w_n tend to w . Let $\hat{\gamma}_n$ be the closed geodesic in $M_n^{-1}(X_n)$ homotopic to ∂D_n , and let \hat{A}_n be its standard collar. Then $\gamma_n^{[i]} = M_n(\hat{\gamma}_n)$ and $A_n^{[i]} = M_n(\hat{A}_n)$ for some i . Let \hat{A}_n^+ be the boundary padding of

\hat{A}_n attached to Y'_n , i.e. $\hat{A}_n^+ = Y'_n \cap \hat{A}_n$. For $\eta > 0$, let $\hat{A}_{n,\eta}^+$ be the parallel subannulus of \hat{A}_n^+ of modulus η , sharing the boundary with \hat{A}_n (as in Proposition 4.2).

By Proposition 4.2 applied to \hat{q}_n on \hat{A}_n (or, more directly, applied to $q_{A_i}^+$ that appeared in Subsection 4.4), we obtain the following: given $\varepsilon > 0$, there exists $\eta > 0$ such that $\|\hat{q}_n\|_{\hat{A}_n^+ \setminus \hat{A}_{n,\eta}^+} < \frac{\varepsilon}{2}$ for sufficiently large n . Since $\|\hat{q}_n\|_{\overline{\mathbb{C}} \setminus \overline{Y'_n}} < \frac{\varepsilon}{2}$ for large n , in order to prove Claim 5.5, it suffices to show that, with η fixed as above, for some $\delta > 0$ and large n ,

$$(12) \quad N_\delta(w_n) \subset (\overline{\mathbb{C}} \setminus \overline{Y'_n}) \cup (\hat{A}_n^+ \setminus \hat{A}_{n,\eta}^+).$$

Let C_n be the maximal round annulus within $\hat{A}_n \setminus \hat{A}_{n,\eta}^+$ with center w_n , and let C'_n be the maximal round annulus within $M_n^{-1}(X_n)$ with center w_n and containing C_n . Then we have

$$(13) \quad \begin{aligned} \text{mod}(C_n) &\geq \text{mod}(\hat{A}_n \setminus \hat{A}_{n,\eta}^+) - \mu^* = \text{mod}(\hat{A}_n) - \eta - \mu^* \\ &\geq \frac{\pi}{\ell(\hat{\gamma}_n)} - 1 - \eta - \mu^* \geq \text{mod}(C'_n) - 1 - \eta - \mu^* ; \end{aligned}$$

indeed, the first inequality is the Round Annuli Lemma A.3, the second is the Collaring Theorem A.1, and finally $\text{mod}(C'_n) \leq \pi/\ell(\hat{\gamma}_n)$ is the fact that $\hat{\gamma}_n$ has shorter length in $M_n^{-1}(X_n)$ than the core curve of C'_n , where the latter has length $\pi/\text{mod}(C'_n)$ in C'_n .

Now suppose that $w \neq \infty$. Then the outer boundary of C'_n must pass through a point in $\overline{\mathbb{C}} \setminus M_n(X_n)$ which does not converge to w or ∞ ; otherwise it contradicts the choice of scaling M_n at the beginning of the proof. Therefore the outer radius of C'_n converges to a positive finite limit r' . By (13) and $C_n \subset C'_n$, the outer radius of C_n must be at least $r = r' \exp(-2\pi(1 + \eta + \mu^*))$. This implies (12) with $N_\delta(w_n) = D_r(w_n)$.

The same argument works for $w = \infty$, by considering the inner radius of C_n , and thus Claim 5.5 has been proved for all cases. This concludes the proof of Lemma 5.4. \square

Now we can finish the proof of Theorem 5.2. For a thick region $Y_n^{[j]}$ with $c^{[j]} > 0$, we have found a limiting quadratic differential $q_\infty^{[j]} \in Q^1(\mathbb{C})$ with $\|q_\infty^{[j]}\|_{\mathbb{C}} = c^{[j]}$. Its number of poles is bounded by N , and each pole is a limit of poles of $M_n^* q_n$. Moreover,

$$\begin{aligned} \|M_n^*(q_n|_{Y_n^{[j]}}) - q_\infty^{[j]}\|_{\mathbb{C}} &\leq \|M_n^*(q_n|_{Y_n^{[j]}}) - M_n^*q_{Y_n^{[j]}}\|_{\mathbb{C}} + \|M_n^*q_{Y_n^{[j]}} - q_\infty^{[j]}\|_{\mathbb{C}} \\ &\leq \|q_n - q_{Y_n^{[j]}}\|_{Y_n^{[j]}} + \|q_{Y_n^{[j]}}\|_{\mathbb{C} \setminus Y_n^{[j]}} + \|M_n^*q_{Y_n^{[j]}} - q_\infty^{[j]}\|_{\mathbb{C}} \\ &\longrightarrow 0. \end{aligned}$$

This shows that the restrictions of the q_n to the $Y_n^{[j]}$ have limit model $q_\infty^{[j]}$.

As for an annulus $B_n^{[i]}$ with $b^{[i]} > 0$, by a similar estimate and Lemma 5.3, we conclude that the restrictions of the q_n to the annuli $B_n^{[i]}$ have limit model dz^2/z^2 .

Since the mass of q_n on the remaining domains must tend to 0, this finishes the proof of Theorem 5.2. \square

6. PUSH-FORWARD OF QUADRATIC DIFFERENTIALS

In this section, we finally conclude the proof of our main theorem, by proving Proposition 3.2: if the Thurston map σ from Subsection 3.1 has no fixed point, then there exists a sequence of quadratic differentials with almost no contraction under the push-forward under any finite number of exponential maps. Now that we have established good control on the geometry of quadratic differentials, it will be relatively easy to conclude that this implies the existence of a degenerate Levy cycle.

We would like to point out that the results in this section are the only places in this paper where we really need the precise properties of the exponential function. These are the statements which need to be generalized in order for a generalization of the main theorem to hold.

We call a sequence (q_n) of measurable quadratic differentials *exp-efficient* if $\|q_n\|_{\mathbb{C}} \neq 0$ for all n and $\lim_{n \rightarrow \infty} \frac{\|\exp_* q_n\|_{\mathbb{C}}}{\|q_n\|_{\mathbb{C}}} = 1$.

Proposition 6.1 (Exp-efficient integrable push-forwards). *Let q_n be a sequence of measurable quadratic differentials on \mathbb{C} which has limit model q with associated scalings $M_n(z) = a_n z + b_n$, where q is a meromorphic quadratic differential on \mathbb{C} with $0 < \|q\|_{\mathbb{C}} < \infty$.*

If the sequence (q_n) is exp-efficient, then the scaling factors $a_n \rightarrow 0$, and $(\exp_ q_n)$ also has limit model q with scaling $z \mapsto \hat{M}_n(z) := a_n \exp(b_n)z + \exp(b_n)$.*

Remark. The scaling \hat{M}_n is chosen so that if $a_n \rightarrow 0$, then

$$S_n(z) := (\hat{M}_n)^{-1} \circ \exp \circ M_n(z) = \frac{e^{a_n z} - 1}{a_n} \rightarrow z \quad \text{as } n \rightarrow \infty$$

uniformly on compact sets on \mathbb{C} .

Proof. No non-zero meromorphic integrable quadratic differential has absolutely efficient push-forward under the exponential map (in the sense that $\|\exp_* q\|_{\mathbb{C}} = \|q\|_{\mathbb{C}}$ because that would require that $q(z + 2\pi i) = \beta(z)q(z)$ where $\beta(z)$ is holomorphic and positive real-valued (see equation (3)); hence β is constant, and this implies that $\|q\|_{\mathbb{C}} = \infty$. Therefore, there is an $\alpha \in (0, 1)$ depending on q such that $\|\exp_* q\|_{\mathbb{C}} \leq \alpha \|q\|_{\mathbb{C}}$.

We first discuss the special case that q_n has the limit model q with trivial scalings, i.e. $M_n(z) = z$ for all n . Then for every $\varepsilon > 0$ we have $\|q_n - q\|_{\mathbb{C}} < \varepsilon \|q\|_{\mathbb{C}}$ for sufficiently large n ; hence $\|q\|_{\mathbb{C}} \leq \|q_n\|_{\mathbb{C}} / (1 - \varepsilon)$. For such n we have

$$\begin{aligned} \|\exp_* q_n\|_{\mathbb{C}} &\leq \|\exp_* q\|_{\mathbb{C}} + \|\exp_*(q_n - q)\|_{\mathbb{C}} \leq \alpha \|q\|_{\mathbb{C}} + \|q_n - q\|_{\mathbb{C}} \\ &\leq (\alpha + \varepsilon) \|q\|_{\mathbb{C}} \leq \frac{\alpha + \varepsilon}{1 - \varepsilon} \|q_n\|_{\mathbb{C}}. \end{aligned}$$

If $\varepsilon > 0$ is chosen small enough so that $(\alpha + \varepsilon)/(1 - \varepsilon) < 1$, this implies that q_n cannot be exp-efficient, a contradiction.

Now consider the case of general scalings $M_n(z) = a_n z + b_n$. It is clear that the translations b_n have no effect on the efficiency under the push-forward, so we may as well translate the differentials q_n so that all $b_n = 0$. If the sequence a_n is bounded in \mathbb{C}^* , this leads to the same contradiction as before: for bounded scaling factors a_n , there is a uniform contraction rate $\alpha < 1$. If the sequence $a_n \rightarrow \infty$, then we write $a_n = 2^{k_n} a_n^\sharp$ with $k_n \in \mathbb{N}$ and a sequence a_n^\sharp with $1 \leq |a_n^\sharp| \leq 2$; define an

auxiliary scaling $M_n^\sharp(z) := a_n^\sharp z$ and set $Q(z) := z^2$. Now the functional equation $\exp(2z) = Q(\exp(z))$ implies

$$\|\exp_*((M_n)_*q)\|_{\mathbb{C}} = \|(Q^{\circ k_n})_* \exp_*(M_n^\sharp)_*q\|_{\mathbb{C}} \leq \|\exp_*(M_n^\sharp)_*q\|_{\mathbb{C}},$$

and this is at least as bad as before. Therefore, the sequence a_n can have neither a bounded subsequence nor a subsequence tending to ∞ , so $a_n \rightarrow 0$ as claimed. Since $b_n = 0$, it remains to show that $(\exp_* q_n)$ has limit model q with scalings $\hat{M}_n(z) = a_n z + 1$.

Recall that $S_n = (\hat{M}_n)^{-1} \circ \exp \circ M_n$ converges to the identity uniformly on compact sets of \mathbb{C} . Since the push-forward by S_n cannot increase the L^1 -norm, we have

$$\begin{aligned} \left\| (\hat{M}_n)^*(\exp_* q_n) - q \right\|_{\mathbb{C}} &= \left\| (\hat{M}_n^{-1})_* \circ \exp_* \circ (M_n)_* \circ (M_n)^*(q_n) - q \right\|_{\mathbb{C}} \\ &= \|(S_n)_*(M_n^* q_n) - q\|_{\mathbb{C}} \\ &\leq \|(S_n)_*(M_n^* q_n) - (S_n)_*q\|_{\mathbb{C}} + \|(S_n)_*q - q\|_{\mathbb{C}} \\ &\leq \|(M_n)^* q_n - q\|_{\mathbb{C}} + \|(S_n)_*q - q\|_{\mathbb{C}}. \end{aligned}$$

Given $\varepsilon > 0$, it follows from Claim 5.6 that there exists a neighborhood W of $Z(q) \cup \{0, \infty\}$ such that for sufficiently large n , we have $\|(S_n)_*q\|_W < \varepsilon/3$ and $\|q\|_W < \varepsilon/3$. Here uniformity with respect to n follows from uniformity in Claim 5.6 and the convergence of poles of $(S_n)_*q$ to $Z(q) \cup \{0, \infty\}$. Since $S_n \rightarrow \text{id}$ uniformly in $\overline{\mathbb{C}} \setminus W$ as $n \rightarrow \infty$, we have $\|(S_n)_*q - q\|_{\overline{\mathbb{C}} \setminus W} < \varepsilon/3$ for sufficiently large n . Thus we have

$$\|(S_n)_*q - q\|_{\mathbb{C}} \leq \|(S_n)_*q - q\|_{\overline{\mathbb{C}} \setminus W} + \|(S_n)_*q\|_W + \|q\|_W < \varepsilon.$$

This implies $\left\| (\hat{M}_n)^*(\exp_* q_n) - q \right\|_{\mathbb{C}} \rightarrow 0$ as claimed. \square

Proposition 6.2 (Exp-efficient annular push-forwards). *Let q_n be a sequence of measurable quadratic differentials on \mathbb{C} . Suppose that the sequence q_n has the limit model dz^2/z^2 on annuli $A_n = \{z \in \mathbb{C}: r_n < |z - z_n| < R_n\}$.*

If the sequence (q_n) is exp-efficient, then there is a sequence $R_n^ \in (r_n, R_n)$ with $R_n^* \rightarrow 0$ such that*

$$\|q_n\|_{A_n^*} / \|q_n\|_{A_n} \rightarrow 1$$

where $A_n^* := \{z \in \mathbb{C}: r_n < |z - z_n| < R_n^*\}$.

Moreover, the quadratic differentials $\exp_ q_n$ also have the limit model dz^2/z^2 on annuli $\{z \in \mathbb{C}: r'_n < |z - z'_n| < R'_n\}$, where $z'_n = \exp(z_n)$, $r'_n = |\exp(z_n)|r_n$ and $R'_n = |\exp(z_n)|R_n^*$.*

Proof. The general idea of the proof is similar to Proposition 6.1; the main difference lies in the fact that there is no integrable limit model for our quadratic differentials, and this requires a preparatory argument. We will use the notation $A(r, R) := \{z \in \mathbb{C}: r < |z| < R\}$ for the centered annulus between radii $r < R$. Let $q := dz^2/z^2$ on \mathbb{C} .

As a toy example, consider the annulus $A := A(\pi, 2\pi)$ and the restricted differential $q|_A$. Then there is a constant $\alpha < 1$ such that $\|\exp_*(q|_A)\|_{\mathbb{C}} \leq \alpha \|q|_A\|_{\mathbb{C}}$, simply because the contributions near the points $1 - i\pi$ and $1 + i\pi$ are not aligned.

Similarly as in Proposition 6.1, the functional equation $\exp(2z) = (\exp(z))^2$ implies that for any $r \geq \pi$ and any integer $i \geq 0$, we have

$$(14) \quad \frac{\|\exp_*(q|_{A(2^i r, 2^{i+1} r)})\|_{\mathbb{C}}}{\|q|_{A(2^i r, 2^{i+1} r)}\|_{\mathbb{C}}} \leq \frac{\|\exp_*(q|_{A(r, 2r)})\|_{\mathbb{C}}}{\|q|_{A(r, 2r)}\|_{\mathbb{C}}}.$$

But this yields for any $r \geq \pi$ and any integer $k \geq 1$

$$\frac{\|\exp_*(q|_{A(r, 2^k r)})\|_{\mathbb{C}}}{\|q|_{A(r, 2^k r)}\|_{\mathbb{C}}} \leq \frac{\|\exp_*(q|_{A(r, 2r)})\|_{\mathbb{C}}}{\|q|_{A(r, 2r)}\|_{\mathbb{C}}}$$

because on every subannulus $A(2^i r, 2^{i+1} r)$, the relative loss of mass is at least the same as on $A(r, 2r)$, and the interference between these various subannuli can only increase the amount of cancellation. This implies that there is a universal constant $\alpha' < 1$ such that for every pair of radii R, r with $\pi \leq r \leq R/2$, we have

$$\frac{\|\exp_*(q|_{A(r, R)})\|_{\mathbb{C}}}{\|q|_{A(r, R)}\|_{\mathbb{C}}} \leq \alpha'.$$

Now suppose that q_n is an exp-efficient sequence of measurable quadratic differentials, and suppose there are annuli $A_n = \{z \in \mathbb{C} : r_n < |z - z_n| < R_n\}$ as in the claim of the proposition such that q_n has limit model dz^2 on A_n . Since nothing changes under translations, we may as well assume that all $z_n = 0$.

Write $A'_n := \{z \in A_n : |z - z_n| < \pi\}$ and $A''_n := \{z \in A_n : |z - z_n| > \pi\}$. Under the push-forward, a definite fraction of the mass on A''_n cancels, so if the sequence q_n on the annuli A_n is exp-efficient, then $\|q_n\|_{A''_n} / \|q_n\|_{A_n} \rightarrow 0$.

Now fix any small $R^* \in (0, \pi)$ and restrict to n with $r_n < R^*$. Then the annulus $A(R^*, \pi)$ has modulus $\frac{1}{2\pi} \log(\pi/R^*)$. Since the mass of dz^2/z^2 on concentric subannuli of A_n is proportional to the moduli of these subannuli and $\text{mod}(A_n) \rightarrow \infty$, it follows that $\|q_n\|_{A(R^*, \pi)} \rightarrow 0$ as $n \rightarrow \infty$. Combining this, it follows that

$$\frac{\|q_n\|_{A(r_n, R^*)}}{\|q_n\|_{A_n}} \rightarrow 1.$$

The rest of the argument is analogous to Proposition 6.1, using the scalings $M_n(z) = a_n z + b_n$ with $a_n = R_n^*$. \square

Lemma 6.3 (Sequence of exp-efficient push-forwards). *Suppose that q_n is a sequence of measurable quadratic differentials with $0 < \|q_n\|_{\mathbb{C}} < \infty$ for all n such that*

$$\frac{\|(E_{\lambda_{n,m}} \circ E_{\lambda_{n,m-1}} \circ \cdots \circ E_{\lambda_{n,1}})_* q_n\|_{\mathbb{C}}}{\|q_n\|_{\mathbb{C}}} \rightarrow 1,$$

where $\lambda_{n,i} \in \mathbb{C}^*$ are arbitrary. Let $V_n \subset \mathbb{C}$ be domains such that $\|q_n\|_{V_n} / \|q_n\|_{\mathbb{C}} \geq \alpha > 0$ for some fixed $\alpha > 0$ and all n . Then for every $s = 0, 1, \dots, m-1$, the sequence of quadratic differentials

$$(E_{\lambda_{n,s}} \circ E_{\lambda_{n,s-1}} \circ \cdots \circ E_{\lambda_{n,1}})_*(q_n|_{V_n})$$

is exp-efficient.

Proof. Set $\hat{q}_n := q_n|_{V_n}$. We obviously have

$$\begin{aligned} 0 &\leq \|\hat{q}_n\|_{\mathbb{C}} - \|(E_{\lambda_{n,m}} \circ E_{\lambda_{n,m-1}} \circ \cdots \circ E_{\lambda_{n,1}})_* \hat{q}_n\|_{\mathbb{C}} \\ &\leq \|q_n\|_{\mathbb{C}} - \|(E_{\lambda_{n,m}} \circ E_{\lambda_{n,m-1}} \circ \cdots \circ E_{\lambda_{n,1}})_* q_n\|_{\mathbb{C}} \end{aligned}$$

because all the cancellation which occurs for \hat{q}_n also occurs for q_n . But this implies

$$\begin{aligned} & \frac{\|\hat{q}_n\|_{\mathbb{C}} - \|(E_{\lambda_{n,m}} \circ E_{\lambda_{n,m-1}} \circ \cdots \circ E_{\lambda_{n,1}})_* \hat{q}_n\|_{\mathbb{C}}}{\|\hat{q}_n\|_{\mathbb{C}}} \\ & \leq \frac{\|q_n\|_{\mathbb{C}} - \|(E_{\lambda_{n,m}} \circ E_{\lambda_{n,m-1}} \circ \cdots \circ E_{\lambda_{n,1}})_* q_n\|_{\mathbb{C}}}{\alpha \|q_n\|_{\mathbb{C}}} \longrightarrow 0 ; \end{aligned}$$

hence

$$\frac{\|(E_{\lambda_{n,m}} \circ E_{\lambda_{n,m-1}} \circ \cdots \circ E_{\lambda_{n,1}})_* \hat{q}_n\|_{\mathbb{C}}}{\|\hat{q}_n\|_{\mathbb{C}}} \longrightarrow 1 .$$

Since all push-forwards are weakly contracting, we have

$$\|\hat{q}_n\|_{\mathbb{C}} \geq \|(E_{\lambda_{n,1}})_* \hat{q}_n\|_{\mathbb{C}} \geq \cdots \geq \|(E_{\lambda_{n,m}} \circ E_{\lambda_{n,m-1}} \circ \cdots \circ E_{\lambda_{n,1}})_* \hat{q}_n\|_{\mathbb{C}} ,$$

and this implies for every $s < m$

$$\frac{\|(E_{\lambda_{n,s+1}} \circ E_{\lambda_{n,s}} \circ \cdots \circ E_{\lambda_{n,1}})_* \hat{q}_n\|_{\mathbb{C}}}{\|(E_{\lambda_{n,s}} \circ \cdots \circ E_{\lambda_{n,1}})_* \hat{q}_n\|_{\mathbb{C}}} \longrightarrow 1$$

as claimed. \square

Proof of Proposition 3.2. Suppose that the statement is false. Then there exist N, m, M such that for all $r < 1$, there exists $E^{(m)} =: E_{\lambda_m} \circ \cdots \circ E_{\lambda_1}$ and a quadratic differential q with at most N poles such that $\|(E^{(m)})_* q\|_{\mathbb{C}} \geq r \|q\|_{\mathbb{C}}$ and for every nested pair of disks $\tilde{D} \subset D$ such that \tilde{D} is an essential annulus in the complement of the poles of q and with $\text{mod}(D \setminus \tilde{D}) \geq M$, the map $E^{(m)}$ is not injective on D .

Thus choose N, m, M , a sequence r_n tending to 1, a sequence

$$E_n^{(m)} = E_{\lambda_{n,m}} \circ \cdots \circ E_{\lambda_{n,1}}$$

and a corresponding sequence of quadratic differentials q_n such that

$$\|(E_n^{(m)})_* q_n\|_{\mathbb{C}} \geq r_n \|q_n\|_{\mathbb{C}} .$$

If we can show that there exists a subsequence for which there are disks $\tilde{D}_n \subset D_n$ which are essential in the complement of the poles of q_n , with $\text{mod}(D_n \setminus \tilde{D}_n) > M$ and $E_n^{(m)}$ injective on D_n , we will have derived the contradiction needed to prove Proposition 3.2.

Extract a subsequence as in Theorem 5.2 (and omit double indices for simplicity of notation). Pick an index j and a sequence of regions $V_n^{[j]}$ such that $q_n|_{V_n^{[j]}}$ has a limit model q .

Define

$$q_n^{(0)} := q_n \quad \text{and} \quad q_n^{(i)} := (E_{\lambda_{n,i}})_* (q_n^{(i-1)}) \quad \text{for } i = 1, 2, \dots, m.$$

Similarly, set

$$\hat{q}_n^{(0)} := q_n|_{V_n^{[j]}} \quad \text{and} \quad \hat{q}_n^{(i)} := (E_{\lambda_{n,i}})_* (\hat{q}_n^{(i-1)}) \quad \text{for } i = 1, 2, \dots, m.$$

By Lemma 6.3, the sequences $\hat{q}_n^{(0)}, \dots, \hat{q}_n^{(m-1)}$ are also exp-efficient. We will discuss the two cases of Theorem 5.2 separately.

Case (b1): The limit model q of $\hat{q}_n^{(0)}$ is an integrable meromorphic quadratic differential on $\overline{\mathbb{C}}$. In this case, there exist scalings $M_n^{(0)}(z) = a_n^{(0)} z + b_n^{(0)}$ with $a_n^{(0)} \rightarrow 0$ as $n \rightarrow \infty$. We will determine the scalings $M_n^{(i)}$ for $\hat{q}_n^{(i)}$ inductively.




FIGURE 7. The quadratic differentials $(E_n^{(s)})_*q_n$ look more and more like q in $M_n^{(i)}(D_R(0))$. Since $A_n^{[i]}$ tends to 0, these regions become smaller and smaller (much more drastically than suggested by the picture) and correspondingly are surrounded by bigger and bigger annuli containing no poles.

Suppose that, for $i \in \{0, 1, \dots, m-1\}$, the differentials $\hat{q}_n^{(i)}$ have limit model q with scalings $M_n^{(i)}(z) = a_n^{(i)}z + b_n^{(i)}$. Then by Proposition 6.1 and the remark thereafter, $a_n^{(i)} \rightarrow 0$, and $\exp_* \hat{q}_n^{(i)}$ has limit model q with scalings $\hat{M}_n^{(i)}$.

Let $M_n^{(i+1)}(z) := \lambda_{n,i+1} \hat{M}_n^{(i)}(z)$. Then $\hat{q}_n^{(i+1)}$ has limit model q with scalings $\hat{M}_n^{(i+1)}$ because $(M_n^{(i+1)})_* \hat{q}_n^{(i+1)} = (\hat{M}_n^{(i)})_*(\exp_* \hat{q}_n^{(i)})$. We also have

$$S_n^{(i+1)} := (M_n^{(i+1)})^{-1} \circ E_{\lambda_{n,i+1}} \circ M_n^{(i)} = (\hat{M}_n^{(i)})^{-1} \circ \exp \circ M_n^{(i)} \longrightarrow \text{id}$$

uniformly on compact sets in \mathbb{C} . Hence for any $R > 0$, the composition $E_{\lambda_{n,m}} \circ E_{\lambda_{n,m-1}} \circ \dots \circ E_{\lambda_{n,1}}$ is injective on $D_n := M_n^{(0)}(D_R(0))$.

Choose \tilde{R} so that $D_{\tilde{R}-1}(0)$ contains all poles of q in \mathbb{C} (there must be at least three of them). Then $\tilde{D}_n := M_n^{(0)}(D_{\tilde{R}}(0))$ must contain at least two poles of q_n for large n ; in fact, the poles of the limit model q in \mathbb{C} are the limits of poles of $(M_n)_*q_n$ by Theorem 5.2. Taking $R := \tilde{R}e^{2\pi M}$ and large n , the conclusion follows.

Case (b2): The limit model is $q = dz^2/z^2$. In this case, we can also introduce the scaling $M_n^{(i)}(z) = a_n^{(i)}z + b_n^{(i)}$ so that we can compare $(M_n^{(i)})^*\hat{q}_n^{(i)}$ with $c_n(dz^2/z^2)|_{\{r_n < |z| < 1\}}$. It follows from Proposition 6.2 that, by successively cutting out the outer radii of the annuli, we can achieve the following for $i = 0, 1, \dots, n$:

$$\left\| (M_n^{(i)})^*\hat{q}_n^{(i)} - c_n \frac{dz^2}{z^2} \Big|_{\{r_n < |z| < 1\}} \right\|_{\mathbb{C}} \longrightarrow 0$$

with $a_n^{(i)} \rightarrow 0$ and

$$(M_n^{(i+1)})^{-1} \circ E_{\lambda_{n,i+1}} \circ M_n^{(i)} \rightarrow \text{id}$$

uniformly on compacts in \mathbb{C} . The rest is similar to Case (b1). □




FIGURE 8. The polygon labeled q_n represents a quadratic differential on $\overline{\mathbb{C}}$. Note that the small rectangle at the lower left is most of $\overline{\mathbb{C}}$, since it contains 0, 1 and ∞ , and it is separated from the remainder by an annulus with large modulus. But most of the mass of q_n is elsewhere; in particular, much of the mass is in the shaded thick part $V_n^{[j]}$, which, if we imagine squeezing the two necks as $n \rightarrow \infty$, will tend to the copy of the Riemann sphere labeled q , with the corresponding quadratic differential (with five poles and one zero). Note that the position of ∞ on that copy of $\overline{\mathbb{C}}$ is imposed, but the position of 0 is arbitrary; the plane carrying q is only defined up to affine transformations. The shaded part of that plane is supposed to represent a disk around the chosen 0 containing all the poles of q except ∞ . The affine map M_n pushes q forward to a quadratic differential which closely approximates q_n in D_n^i and has very little mass elsewhere.

APPENDIX A. SOME GENERAL RESULTS ON RIEMANN SURFACES

In this appendix we review some important facts about Riemann surfaces and quadratic differentials.




FIGURE 9. Left: It is often convenient to use the band model of the hyperbolic plane; this is the region $|\operatorname{Im} z| < \pi/2$, with the metric $|dz|/\cos y$; on the real axis, the Euclidean metric coincides with the hyperbolic metric. On the left we have drawn the configuration defining $\eta(\ell)$. Right: If you divide out by translation by ℓ and rescale the resulting Euclidean cylinder so that the circumference is 1, then the (Euclidean) heights are as indicated. Note that when ℓ becomes small, so that h becomes large, the part of the cylinder not in the standard collar becomes negligible.

A.1. Collars on hyperbolic Riemann surfaces. We will consider Riemann surfaces which have a hyperbolic metric of constant negative curvature. We will always normalize the metric so that the curvature is -1 everywhere.

The Collaring Theorem A.1 says that every sufficiently short simple closed geodesic is surrounded by an annulus of definite modulus. We make this precise as follows.

Let $\eta : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be the *collaring function* which is defined as follows (compare Figure 9): in a simply connected hyperbolic Riemann surface, let L be a hyperbolic geodesic and let $I \subset L$ be a segment of length ℓ . Draw perpendiculars to L through the endpoints of I , both on the same side of L , and extend them until they meet the circle at infinity. Join these points at infinity by a line M . Then $\eta(\ell)$ is the distance between L and M .

Theorem A.1 (The collaring theorem and standard collars). *Let X be a hyperbolic Riemann surface. For every simple closed geodesic γ on X of length $\ell(\gamma) < \ell^* := \log(3 + 2\sqrt{2})$, the set of points at distance less than $\eta(\ell(\gamma))$ from γ is an embedded annulus A_γ (called the standard collar around γ) whose modulus M_γ satisfies*

$$\frac{\pi}{\ell(\gamma)} - 1 < M_\gamma < \frac{\pi}{\ell(\gamma)}.$$

For the hyperbolic metric of the annulus A_γ , the curve γ is the unique simple closed geodesic (the core curve of A_γ).

Moreover, if γ_1 and γ_2 are disjoint simple closed geodesics, then the standard collars A_{γ_1} and A_{γ_2} are disjoint.

It will be convenient to write $M_\gamma = 2h_\gamma$ (h for “height”) so that A_γ is isomorphic to the standard cylinder of circumference 1 and height h_γ :

$$A_\gamma \simeq \{z \in \mathbb{C} : |\operatorname{Im} z| < h_\gamma\} / \mathbb{Z} .$$

Note that if $\ell(\gamma)$ is small (the only case of interest to us), the estimate of Theorem A.1 on M_γ is extremely precise (it is also the best possible).

We also need the *central collar* B_γ around γ : this is the unique parallel subannulus of A_γ of modulus $M_\gamma - \sqrt{2M_\gamma}$ with core curve γ .

Corollary A.2 (Short geodesics disjoint). *(a) If two simple closed curves γ_1 and γ_2 on X satisfy $\ell(\gamma_2) < 2\eta(\ell(\gamma_1))$, then they are disjoint or they coincide.*

(b) If two simple closed geodesics γ_1 and γ_2 both have length less than $\log(3 + 2\sqrt{2})$, then they are disjoint or they coincide. In particular, on a Riemann surface of genus g with N punctures, there are at most $3g - 3 + N$ such curves.

Proofs for these results can be found in [DH, Section 6]¹.

The Collaring Theorem A.1 guarantees large annuli around short curves on hyperbolic Riemann surfaces X . When X is a subset of \mathbb{C} , we can do better: a short curve on such an X is surrounded by a large *round* annulus; this is an annulus whose two boundaries are concentric Euclidean circles. We say that a round annulus $A \subset \mathbb{C}$ is *centered at* $a \in \mathbb{C}$ if there are radii $R > r > 0$ so that $A = \{z \in \mathbb{C} : r < |z - a| < R\}$.

Lemma A.3 (Round annuli). *There are universal constants $\mu^*, \mu^{**} > 0$ with the following property: if $A \subset \mathbb{C}$ is an annulus with modulus $\mu > \mu^{**}$ and a is a point in the bounded component of $\mathbb{C} \setminus A$, then there is a round annulus $B \subset A$ centered around a and with modulus $\operatorname{mod}(B) \geq \operatorname{mod}(A) - \mu^*$.*

For a proof, see [M2, Theorem 2.1].

A.2. Mass per modulus is convex. If q is a quadratic differential on an annulus $\{z \in \mathbb{C} : a < \operatorname{Im}(z) < b\} / \mathbb{Z}$, we define *mass per modulus* as the function $F : (a, b) \rightarrow \mathbb{R}_0^+$ with $F(y) := \int_{\mathbb{R}/\mathbb{Z}} |f(x + iy)| dx$. We will need the following elementary result, reminiscent of Jensen’s formula.

Proposition A.4 (Mass per modulus on annulus). *(a) Let*

$$f(z) = \sum_{n=-\infty}^{\infty} a_n e^{2\pi i n z}$$

and let the band $\{z \in \mathbb{C} : A < \operatorname{Im} z < B\}$ be the maximal open set on which the

¹ Douady and Hubbard [DH] use the hyperbolic metric with constant curvature -4 , so their value ℓ^* is half as large.




FIGURE 10. Left: The graph of η_ρ . It is normalized to have integral 1, and as $\rho \rightarrow 0$, it is a good approximation to the Dirac δ -measure. Right: A convex function with a limit at infinity is decreasing.

series converges. Then the function

$$F(y) = \int_{\mathbb{R}/\mathbb{Z}} |f(x + iy)| dx$$

is convex on (A, B) .

(b) If $a_n = 0$ for $n < 0$, then $B = \infty$, $\lim_{y \rightarrow \infty} F(y) = a_0$, and $F(y)$ is decreasing on (A, ∞) .

Similarly, if $a_n = 0$ for $n > 0$, we have $A = -\infty$, $\lim_{y \rightarrow -\infty} F(y) = a_0$ and F is increasing on $(-\infty, B)$.

Proof. We will require the following characterization of convexity. Define the function

$$\eta_\rho(x) = \begin{cases} \frac{2}{\pi\rho^2} \sqrt{\rho^2 - x^2} & \text{if } |x| \leq \rho, \\ 0 & \text{if } |x| > \rho. \end{cases}$$

An easy computation (or an easier argument using symmetry) shows that if $h(x) = ax + b$ is a linear function, then $h(x) = (h * \eta_\rho)(x)$. \square

Lemma A.5 (Convexity on interval). *Let $I \subset \mathbb{R}$ be an open interval and let $f : I \rightarrow \mathbb{R}$ be continuous. Then f is convex if and only if $f(x) \leq (f * \eta_\rho)(x)$ for all sufficiently small ρ .*

Proof. The direction \implies is clear. For the converse, note first that a function is convex if for all a, b the function $g(x) = f(x) + ax + b$ has no local strict maximum. This is true in our case: at a strict local maximum x we would have

$$g(x) > (g * \eta_\rho)(x) = (f * \eta_\rho)(x) + ax + b \geq f(x) + ax + b = g(x)$$

for sufficiently small ρ , a contradiction. \square

Proof of Proposition A.4. Choose $y \in (A, B)$ and $\rho > 0$ such that $(y - \rho, y + \rho) \subset (A, B)$. Let $D_z \subset \mathbb{C}$ be the disk of radius ρ centered at z . Then by the mean-value property of analytic functions we have

$$f(z) = \frac{1}{\pi\rho^2} \int_{D_z} f(u + iv) du dv.$$

Thus

$$\begin{aligned}
 F(y) &= \int_{\mathbb{R}/\mathbb{Z}} |f(x + iy)| dx = \int_{\mathbb{R}/\mathbb{Z}} \left| \frac{1}{\pi\rho^2} \int_{D_{x+iy}} f(u + iv) du dv \right| dx \\
 &\leq \int_{\mathbb{R}/\mathbb{Z}} \int_{-\rho}^{\rho} \left(\frac{1}{\pi\rho^2} \int_{x-\sqrt{\rho^2-y^2}}^{x+\sqrt{\rho^2-y^2}} |f(u + i(y + v'))| du \right) dv' dx \\
 (15) \quad &= \frac{2}{\pi\rho^2} \int_{-\rho}^{\rho} \int_{\mathbb{R}/\mathbb{Z}} |f(x + i(y + v'))| \sqrt{\rho^2 - (v')^2} dx dv' = (F * \eta_{\rho})(y).
 \end{aligned}$$

This proves part (a). For part (b), in the case where $a_n = 0$ for $n < 0$, the function $x \mapsto f(x + iy)$ converges uniformly to the constant function a_0 as $y \rightarrow \infty$. Figure 10, right, illustrates what a convex function on (A, ∞) with a finite limit at ∞ must look like. The other case works the same way. \square

ACKNOWLEDGEMENTS

This project has resulted from common stays at the Mathematical Sciences Research Institute in Berkeley, the Institute for Mathematical Sciences in Stony Brook, the Institut Henri Poincaré in Paris, and the Fields Institute in Toronto. We would like to thank these institutes and the organizers of the special programs, including Bill Thurston, Curt McMullen, John Milnor, Misha Lyubich, and Adrien Douady. Discussions with them, as well as with a number of participants at these programs, have been very helpful, in particular with Adam Epstein, Jeremy Kahn, Kevin Pilgrim, Mary Rees, and Nikita Selinger; we wish to thank them all. We also wish to thank Ben Bielefeld and Yuval Fisher for their early work on the spider algorithm.

REFERENCES

- [A1] Lars Ahlfors: *Conformal invariants*. McGraw-Hill (1973). MR0357743 (50:10211)
- [A2] Lars Ahlfors: *Lectures on quasiconformal mappings*. Van Nostrand (1966). MR0200442 (34:336)
- [BDGHHR] Clara Bodelon, Robert Devaney, Lisa Goldberg, Martin Hayes, John Hubbard, Gareth Roberts: *Hairs for the complex exponential family*. Internat. J. Bif. Chaos **9** No. 8 (1999), 1517–1534. MR1721835 (2001a:37055)
- [BFH] Ben Bielefeld, Yuval Fisher, John Hubbard: *The classification of critically preperiodic polynomials as dynamical systems*. Jour. AMS **5** 4 (1992). MR1149891 (93h:58128)
- [BKS] Henk Bruin, Alexandra Kaffl, Dierk Schleicher: *Symbolic dynamics of quadratic polynomials*. Manuscript, in preparation. Earlier version circulated as [BS].
- [BR] Noel Baker, Phil Rippon: *Iteration of exponential functions*. Ann. Acad. Sci. Fenn., Series A.I. Math. **9** (1984), 49–77. MR752391 (86d:58065)
- [BS] Henk Bruin, Dierk Schleicher: *Symbolic dynamics of quadratic polynomials*. Report, Institute Mittag-Leffler, Djursholm **7** (2001/02). (128 pp.)
- [DGH] Robert Devaney, Lisa Goldberg, John Hubbard: *A dynamical approximation to the exponential map of polynomials*. Preprint, MSRI (1986).
- [DH] Adrien Douady, John Hubbard: *A proof of Thurston's topological characterization of rational functions*. Acta Math. **171** (1993), 263–297. MR1251582 (94j:58143)
- [E] Leonhardt Euler: *De formulis exponentialibus replicatis*. Acta Acad. Petropolitanae **1** (1777), 38–60.
- [EL] Alexandre Eremenko, Mikhail Lyubich: *Dynamical properties of some classes of entire functions*. Ann. Sci. Inst. Fourier, Grenoble **42** 4 (1992), 989–1020. MR1196102 (93k:30034)
- [FRS] Markus Förster, Lasse Rempe, Dierk Schleicher: *Classification of Escaping Exponential Maps*. Proc. Am. Math. Soc. **136** 2 (2008), 651–663. MR2358507

- [FS] Markus Förster, Dierk Schleicher: *Parameter rays for the complex exponential family*. Ergod. Thy and Dynam. Syst., to appear. ArXiv math.DS/050597.
- [GL] Frederick Gardiner, Nikola Lakić: *Quasiconformal Teichmüller theory*. Mathematical Surveys and Monographs, 76. AMS (2000) MR1730906 (2001d:32016)
- [H] John Hubbard: *Teichmüller theory and applications to geometry, topology, and dynamics*, Volume I: Teichmüller theory. Matrix editions, Ithaca/NY (2006). MR2245223
- [HS] John Hubbard, Dierk Schleicher: *The spider algorithm*. In: Complex dynamical systems, Robert Devaney (ed.), AMS (1994), 155–180. MR1315537
- [IT] Yoichi Iwayoshi, Masahiko Taniguchi: *An introduction to Teichmüller spaces*. Springer-Verlag, Tokyo (1992). MR1215481 (94b:32031)
- [K] Jan Kiwi: *Rational laminations of complex polynomials*. In: *Laminations and foliations in dynamics, geometry and topology*, Mikhail Lyubich, John Milnor, Yair Minsky (eds), Contemporary Mathematics **269**, AMS (2000), 111–154. MR1810538 (2002e:37063)
- [L] Olli Lehto: *Univalent functions and Teichmüller spaces*. Graduate Texts in Mathematics **109**, Springer-Verlag, New York (1987). MR867407 (88f:30073)
- [Le] Silvio Levy: *Critically finite rational maps*. Ph.D. thesis, Princeton (1985).
- [LS] Eike Lau, Dierk Schleicher: *Internal addresses in the Mandelbrot set and irreducibility of polynomials*. Preprint, Inst. Math. Sci., Stony Brook **19** (1994). ArXiv math.DS/9411238.
- [LSV] Bastian Laubner, Dierk Schleicher and Vlad Vicol: *A dynamical classification of postsingularly finite exponential maps*. Discr. and Cont. Dyn. Sys., to appear. ArXiv math.DS/0602602.
- [M1] Curtis McMullen: *Amenability, Poincaré series and quasiconformal maps*. Invent. Math. **97** (1989), 95–127. MR999314 (90e:30048)
- [M2] Curtis McMullen: *Complex dynamics and renormalization*. Ann. Math. Studies **135**, Princeton Univ. Press (1994). MR1312365 (96b:58097)
- [M3] Curtis McMullen: *The moduli space of Riemann surfaces is Kähler hyperbolic*. Annals of Mathematics **151** (2000), 327–357. MR1745010 (2001m:32032)
- [P] Alfredo Poirier: *On postcritically finite polynomials I/II*. Preprint, Inst. Math. Sci., Stony Brook **5/7** (1993).
- [RS1] Lasse Rempe, Dierk Schleicher: *Bifurcations in the space of exponential maps*. Submitted. Preprint, Institute for Mathematical Sciences at Stony Brook **3** (2004). ArXiv math.DS/0311480.
- [RS2] Lasse Rempe, Dierk Schleicher: *Combinatorics of bifurcations in exponential parameter space*. In: P. Rippon, G. Stallard (eds), *Transcendental dynamics and complex analysis*, volume in honour of Professor I. N. Baker, LMS Lecture Note Series. ArXiv math.DS/0408011.
- [RS3] Lasse Rempe, Dierk Schleicher: *Bifurcation Loci of Exponential Maps and Quadratic Polynomials: Local Connectivity, Triviality of Fibers, and Density of Hyperbolicity*, with Lasse Rempe. In: M. Lyubich, M. Yampolsky (eds), *Holomorphic Dynamics and Renormalization*, in Honour of John Milnor’s 75th birthday. Fields Institute Communications **53** (2008), to appear.
- [S1] Dierk Schleicher: *On the dynamics of iterated exponential maps*. Habilitationsschrift, TU München (1999).
- [S2] Dierk Schleicher: *Attracting dynamics of exponential maps*. Annales Academiae Scientiarum Fennicae **28** 1 (2003), 3–34. MR1976827 (2004k:37091)
- [S3] Dierk Schleicher: *Internal addresses in the Mandelbrot set and Galois groups of polynomials*. Manuscript (2007), submitted. Earlier version circulated as [LS].
- [S4] Dierk Schleicher: *The dynamical fine structure of iterated cosine maps and a dimension paradox*. Duke Mathematics Journal **136** 2 (2007), 343–356. MR2286634 (2008d:37078)
- [S5] Dierk Schleicher: *Hyperbolic Components in Exponential Parameter Space*. Comptes Rendus — Mathématiques **339** 3 (2004), 223–228. ArXiv math.DS/0406256. MR2078079 (2005e:37105)

- [SZ1] Dierk Schleicher and Johannes Zimmer: *Escaping points of exponential maps*. Journal Lond. Math. Soc. (2) **67** (2003), 380–400. MR1956142 (2003k:37067)
- [SZ2] Dierk Schleicher and Johannes Zimmer: *Periodic points and dynamic rays of exponential maps*. Ann. Acad. Scient. Fenn. **28** 2 (2003), 327–354. MR1996442 (2004e:37068)

DEPARTMENT OF MATHEMATICS, MALOTT HALL, CORNELL UNIVERSITY, ITHACA, NEW YORK 14853, AND CENTRE DE MATHÉMATIQUES ET D'INFORMATIQUE, UNIVERSITÉ DE PROVENCE, 39 RUE FRÉDÉRIC JOLIO-CURIE, 13453 MARSEILLE CEDEX 13, FRANCE

E-mail address: jhh8@cornell.edu

SCHOOL OF ENGINEERING AND SCIENCE, JACOBS UNIVERSITY BREMEN, POSTFACH 750 561, D-28725 BREMEN, GERMANY

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

DEPARTMENT OF MATHEMATICS, FACULTY OF SCIENCES, KYOTO UNIVERSITY, KYOTO 606-8502, JAPAN

E-mail address: mitsu@math.kyoto-u.ac.jp