

Models for the Speiser class

Christopher J. Bishop

ABSTRACT

The Eremenko–Lyubich class \mathcal{B} consists of transcendental entire functions with bounded singular set and the Speiser class $\mathcal{S} \subset \mathcal{B}$ is made up of functions with a finite singular set. In an earlier work (*J. Lond. Math. Soc.* 92 (2015) 202–221), I gave a method for constructing Eremenko–Lyubich functions that approximate certain simpler functions called models. In this paper, I show that all models can be approximated in a weaker sense by Speiser class functions, and that the stronger approximation of the earlier work (*J. Lond. Math. Soc.* 92 (2015) 202–221) can fail for the Speiser class. In particular, I give geometric restrictions on the geometry of a Speiser class function that need not be satisfied by general Eremenko–Lyubich functions.

1. Introduction

If f is an entire function, we say f is transcendental if it is not a polynomial. The singular set of an entire function f is the closure of its finite critical values and finite asymptotic values, and will be denoted by $S(f)$. The Eremenko–Lyubich class \mathcal{B} consists of transcendental entire functions such that $S(f)$ is a bounded set. The Speiser class $\mathcal{S} \subset \mathcal{B}$ consists of functions for which $S(f)$ is a finite set. We let $\mathcal{S}_{n,k} \subset \mathcal{S}$ denote the subcollection of functions with at most n finite critical values and k finite asymptotic values. In this paper, we will be particularly concerned with $\mathcal{S}_{2,0}$.

The Eremenko–Lyubich and Speiser classes are important in the study of transcendental dynamics and it is known that the dynamical behavior in the Speiser class is more restricted than in the Eremenko–Lyubich class. For example, a Speiser class function cannot have a wandering domain (proved by Eremenko and Lyubich in [9], and Goldberg and Keen in [11]), whereas an Eremenko–Lyubich function can have a wandering domain [2]. On the other hand, various types of pathological behavior, such as a Julia set with no nontrivial path components can be constructed in either class (see [2, 17]).

In this paper, we prove an approximation theorem involving the Speiser class that is analogous to a result proven for the Eremenko–Lyubich class in [4]. However, the function we construct here fails to satisfy some of the side conditions that could be imposed in [4]. Comparing the two results helps illustrate the differences between the two classes of functions. To state our results precisely, we need to introduce some notation.

Suppose $\Omega = \bigcup_j \Omega_j$ is a disjoint union of unbounded simply connected domains so that sequences of components of Ω accumulate only at infinity. See Figure 1. Also suppose that there exists a map $\tau: \Omega \rightarrow \mathbb{H}_r + \rho_0 = \{x + iy : x > \rho_0\}$ that is holomorphic and such that

- (1) the restriction of τ to each Ω_j is a conformal map $\tau_j: \Omega_j \rightarrow \mathbb{H}_r + \rho_0$, and
- (2) if $\{z_n\} \subset \Omega$ and $\tau(z_n) \rightarrow \infty$ then $z_n \rightarrow \infty$.

An open set Ω as above will be called a model domain and $F = e^\tau$ will be called a model function. Note that $F: \Omega \rightarrow \{z : |z| > e^{\rho_0}\}$ is a covering map. A choice of both a model domain




FIGURE 1. A normalized model consists of an open set Ω with possibly several tracts, each of which is mapped conformally by τ to \mathbb{H}_r and then by e^z to $\{|z| > 1\}$, giving the model function F on Ω . The points $F^{-1}(1)$ partition each boundary component into arcs. In the Eremenko–Lyubich class, τ can be rescaled independently on different tracts, so that the partitions on different tract boundaries are unrelated, but we will prove that for the Speiser class, the partitions for different tracts satisfy certain geometric relations.

Ω and a model function F on Ω will be called a model. If $\rho_0 = 0$, we say the model is normalized; this is the main case we will consider.

We call the connected components, $\{\Omega_j\}$, of a model domain Ω the tracts of Ω . In many cases of interest, the tracts will be Jordan domains on the Riemann sphere with the point ∞ on the boundary. The number of tracts can be either finite or infinite. (Usually a domain refers to an open connected set, so using ‘model domain’ for regions that may have several connected components might be confusing. We are using the phrase to abbreviate ‘the domain of definition of the model function’ rather than invent a new term for this — terrain, territory, archipelago, etc. Except for this usage, the term domain will retain its usual meaning).

Given a normalized model (Ω, F) , we let

$$\Omega(\rho) = \{z \in \Omega : |F(z)| > e^\rho\} = \tau^{-1}(\{x + iy : x > \rho\}),$$

and

$$\Omega(\rho, \delta) = \{z \in \Omega : e^\rho < |F(z)| < e^\delta\} = \tau^{-1}(\{x + iy : \rho < x < \delta\}).$$

Given a tract Ω_j of Ω , we let $\Omega_j(\rho) = \Omega(\rho) \cap \Omega_j$ and similarly for $\Omega_j(\rho, \delta)$.

Suppose Ω is a normalized model domain and $\rho > 0$. The boundary of $\Omega_j(\rho)$ has a natural partition into subarcs with endpoints that satisfy $\tau_j(z) \in \rho + \pi i\mathbb{Z}$. We call this a τ -partition or conformal partition of $\partial\Omega_j(\rho)$. It is easy to see from the distortion theorems for conformal maps (for example, see Section 2 of this paper or [10, Theorem I.4.5]) that these subarcs of $\partial\Omega_j(\rho)$ are smooth with bounds depending only on ρ , and adjacent arcs have comparable lengths (again with a constant depending only on ρ).

Suppose f is a transcendental entire function and that $S(f) \subset \mathbb{D}_R = \{z : |z| < R\}$ (when $R = 1$, we write $\mathbb{D} = \mathbb{D}_1$). In [9], Eremenko and Lyubich observed that $\Omega = f^{-1}(\{z : |z| > R\})$ is a disjoint union of analytic, unbounded simply connected domains and f acts a covering map $f : \Omega_j \rightarrow \{|z| > R\}$ on each tract Ω_j of Ω . Thus, each function f in the Eremenko–Lyubich class that satisfies $S(f) \subset \mathbb{D}$ gives rise to a normalized model domain $\Omega = \{z : |f(z)| > 1\}$ and a model function $F = f|_\Omega$ (hence, $\tau(z)$ is a branch of $\log f(z)$). The components of Ω are called

the tracts of f . We call a model arising in this way an Eremenko–Lyubich model. If f is in the Speiser class, we call it a Speiser model.

The purpose of this paper is to quantify the differences between \mathcal{B} and \mathcal{S} in terms of models. In [4], I showed that Eremenko–Lyubich functions can essentially behave like arbitrary models near ∞ ; the tracts can have any shape and the choice of τ on each tract is independent of the choice in other tracts. In this paper, I show that in Speiser models, the choice of τ in different tracts must satisfy certain geometric constraints (for example, Theorems 1.4 and 10.1); however, given any model Ω , it is always possible to add extra tracts and define τ on these new tracts so that the geometric conditions are satisfied. Thus, informally we say ‘every model is an Eremenko–Lyubich model’ and ‘every model is a submodel of a Speiser model’. More precisely, the following theorem is proved in [4].

THEOREM 1.1 (All models occur in \mathcal{B}). *Suppose (Ω, F) is a normalized model and $\rho > 0$. Then there is a $f \in \mathcal{B}$ and a quasiconformal homeomorphism $\varphi: \mathbb{C} \rightarrow \mathbb{C}$ so that $F = f \circ \varphi$ on $\Omega(\rho)$. In addition,*

- (1) *we have $|f \circ \varphi| \leq e^\rho$ off $\Omega(\rho)$ (that is, f is bounded off $\varphi(\Omega)$);*
- (2) *the singular set satisfies $S(f) \subset D(0, e^\rho)$;*
- (3) *the maximal dilatation K of φ depends only on ρ ;*
- (4) *the map φ is conformal except on $\Omega(\rho, 2\rho)$.*

We will review the definition and basic properties of quasiconformal mappings in Section 2. One of the main goals of this paper is to prove the following analog of Theorem 1.1 for the Speiser class.

THEOREM 1.2 (All models occur as submodels in \mathcal{S}). *Suppose (Ω, F) is a normalized model and $\rho > 0$. Then there is a $f \in \mathcal{S}$ and a quasiconformal homeomorphism $\varphi: \mathbb{C} \rightarrow \mathbb{C}$ so that $F = f \circ \varphi$ on $\Omega(\rho)$. In addition,*

- (1) *the function f has no finite asymptotic values and two critical values, $\pm e^\rho$;*
- (2) *every critical point of f has degree at most 12;*
- (3) *the maximal dilatation K of φ depends only on ρ ;*
- (4) *the map φ is conformal on $\Omega(2\rho)$.*

The maximal dilatation bound for φ remains bounded as $\rho \rightarrow \infty$, but blows up as $\rho \rightarrow 0$ (we will not be explicit about the dependence of K on ρ , but estimates could be derived from a careful reading of [2]). It is not true that the maximal dilatation K tends to 1 as $\rho \rightarrow \infty$, at least for the construction given here, since the use of the folding maps from [2] introduces a fixed amount of distortion, independent of ρ .

The degree of a critical point z of a holomorphic map f is taken to be the local valence of f near z , for example, $f(z) = z^3$ has a critical point of degree 3 at 0. The bound in (2) follows immediately from the proof of the folding theorem in [2]. However, by making some simple changes to the construction in [2], the 12 can be improved to 4. This will be discussed in more detail at the end of Section 3.

The crucial difference between Theorems 1.1 and 1.2 is that the latter omits the conclusion ‘ $|f \circ \varphi| \leq e^\rho$ off $\Omega(\rho)$ ’. The function $f \in \mathcal{B}$ constructed in Theorem 1.1 is only large where the model is large (inside Ω), so it has the same number of tracts as the model has. However, the function $f \in \mathcal{S}$ in Theorem 1.2 might also be large outside Ω , and so it can have ‘extra’ tracts. This is the sense in which approximation by Speiser functions is weaker than approximation by Eremenko–Lyubich functions. In fact, our proof will always introduce extra tracts; we will first give a construction that creates an infinite number of extra tracts, and then give a more intricate construction that shows

THEOREM 1.3. *The function f in Theorem 1.2 may be chosen so that the number of tracts of f is at most twice the number of tracts of the model (Ω, F) .*

Simple examples show that some models with n tracts require the approximating Speiser class function to have $2n$ tracts, so the bound in Theorem 1.3 is sharp. Roughly speaking, if Ω has n tracts, then the domain $W = \mathbb{C} \setminus \Omega(\rho)$ has n distinct ‘ends’ at infinity. If these ends are each ‘large’ compared to the tracts of the model, then each end must contain at least one extra tract of the approximating Speiser class function. A very concrete example is

THEOREM 1.4. *The half-strip $S = \{x + iy : x > 0, |y| < 1\}$ cannot be mapped to any Speiser class model domain by any quasiconformal homeomorphism of the plane.*

In other words, there is no Speiser class function with a single tract, so that this tract is the image of a half-strip under a quasiconformal map of the plane. However, there are Speiser class functions with two tracts, one of which can be sent to a half-strip by a quasiconformal map of the plane; moreover, this tract can approximate the half-strip in the Hausdorff metric on the plane as closely as we wish (see Figure 17 and the remarks in Section 13). On the other hand, Theorem 1.1 implies that there are Eremenko–Lyubich functions with single tracts that approximate the half-strip as closely as we wish in the Hausdorff metric.

The referee of this paper asked if Theorem 1.4 also holds for any tract that is contained in the half-strip S . While our proof of Theorem 1.4 extends to cover many cases of this type, and it is not hard to see that no subdomain of S can itself be a Speiser class model domain, there might be such a subdomain that can be mapped to a Speiser class model domain by some quasiconformal map of the plane; deciding this would be an interesting problem. It would also be very interesting to have a geometric characterization (even up to quasiconformal maps) of the tracts of Speiser class functions that have a single tract.

Another difference between Theorems 1.1 and 1.2 concerns the proofs. The proof of Theorem 1.1 given in [4] is mostly self-contained and depends on constructing a Blaschke product in the disk that approximates a certain inner function arising from the model. On the other hand, the proof of Theorem 1.2 in this paper depends on the more difficult quasiconformal folding construction of Speiser class functions in [2]. The precise statement we use will be reviewed in Section 3.

Finally, we mention an application of Theorem 1.2 to dynamics. We call a model (Ω, F) disjoint type, if it is normalized and $\bar{\Omega} \cap \bar{\mathbb{D}} = \emptyset$. An entire function is usually called disjoint type if (1) it is hyperbolic (the singular set is bounded and every point in it tends to an attracting periodic cycle of f under iteration) and (2) the Fatou set is connected (the Fatou set is the largest open set on which the iterates of f form a normal family; its complement is called the Julia set of f). Alternatively, [16, Proposition 2.1] states that a transcendental entire function is disjoint type if and only if there is a Jordan domain D so that $S(f) \subset D$ and $f(\bar{D}) \subset D$. This implies that if (Ω, F) is a disjoint-type Eremenko–Lyubich model, then $F = f|_{\Omega}$, where f is an Eremenko–Lyubich entire function that is disjoint type in the sense above (just take $D = \mathbb{D}$).

We can iterate a model function F as long as the iterates keep landing in Ω , and we define the Julia set of a model as

$$\mathcal{J}(F) = \bigcap_{n \geq 0} \{z \in \Omega : F^n(z) \in \Omega\}.$$

If F is a disjoint-type Eremenko–Lyubich model, then this is the same as the usual Julia set of the extension of F . Lasse Rempe-Gillen has pointed out that Theorem 1.1 implies that any disjoint-type model function is conjugate on its domain to a disjoint-type $f \in \mathcal{B}$, in particular, the Julia set and the escaping set for the model function F are homeomorphic

via a quasiconformal mapping of the whole plane to the corresponding sets for f . Thus, various pathological examples in \mathcal{B} can be constructed simply by exhibiting a model with the desired property, for example, see [16].

For the Speiser class, the approximating function f may have extra tracts that do not correspond to tracts of the model. In this case, Rempe-Gillen’s argument implies the model function F restricted to its Julia set can be conjugated to a Speiser class function f restricted to a certain closed subset $A \subset \mathcal{J}(f)$. More precisely,

THEOREM 1.5. *Suppose that (Ω, F) is any normalized, disjoint-type model, f is a Speiser class function, and φ is a quasiconformal mapping of the plane so that $f = F \circ \varphi$ on $U = \varphi^{-1}(\Omega)$ (this is a subcollection of tracts of f). Assume that $(U, f|_U)$ is also a normalized, disjoint-type model. Then there is a quasiconformal map $\Phi: \mathbb{C} \rightarrow \mathbb{C}$ so that $\Phi \circ f = F \circ \Phi$ on U .*

In other words, the Julia set of the model function F is quasiconformally conjugate to a closed subset A of the Julia set of the Speiser class function f . The set A consists of those points whose orbits stay within U forever, where U is the subcollection of the tracts of f corresponding to the tracts of F via Φ . This result is a straightforward application of [4, Theorem 9.1] (which itself is simply a summary of an argument of Rempe-Gillen from [15]).

The use of quasiconformal techniques to build and understand entire functions with finite singular sets has a long history with its roots in the work of Grötzsch, Speiser, Teichmüller, Ahlfors, Nevanlinna, Lavrentieff, and many others. The earlier work was often phrased in terms of Riemann surfaces and deciding if a simply connected surface built by branching over a finite singular set was conformally equivalent to the plane or to the disk (the type problem; in the first case, the uniformizing map gives a Speiser class function). Such constructions play an important role in value distribution theory; see [7] for an excellent survey of these methods and a very useful guide to this literature, also see [12, Chapter VII]. More recent work (including this paper) is motivated by applications to dynamics, where the Speiser class provides an interesting mix of structure (like polynomials, the quasiconformal equivalence classes are finite dimensional [9]) and flexibility (as indicated by the results of [2, 3, 16], and the current paper).

In this paper, the notation $A \lesssim B$ means that $A \leq CB$, where A, B are quantities that depend on some parameter and $C < \infty$ is a constant that is independent of the parameter. The notation means the same as $A = O(B)$. Similarly, $A \gtrsim B$ is equivalent to $B \lesssim A$ or $B = O(A)$. If $A \lesssim B$ and $A \gtrsim B$, then we say $A \simeq B$, that is, A and B are comparable, independent of the parameter.

2. Modulus and quasiconformal maps

Many of our arguments involve the modulus of path families, conformal maps, and quasiconformal maps, so we briefly review the basic facts here for the convenience of the reader. Everything in this section can be found (in greater detail and with proofs) in standard references such as [1] or [10].

An orientation preserving homeomorphism φ of the plane to itself is quasiconformal if it is absolutely continuous on all lines and $|\varphi_{\bar{z}}| \leq k|\varphi_z|$ almost everywhere (with respect to area measure) for some $k < 1$. At points of differentiability, this means that the tangent map of φ sends circles to ellipses of eccentricity at most $K = (1 + k)/(1 - k) \geq 1$. The smallest K that works for φ at almost every point is called the maximal dilatation of φ ; such a map is also called K -quasiconformal. A K -quasiconformal map φ satisfies a Beltrami equation $f_{\bar{z}} = \mu f_z$ almost everywhere for some bounded measurable function μ called the dilatation of f and $\|\mu\|_\infty \leq k = (K - 1)/(K + 1)$. A 1-quasiconformal map is conformal. The family of K -quasiconformal maps of the plane to itself that fix two finite points (usually taken to be 0,1) is compact.

The measurable Riemann mapping theorem (for example, see [1]) says that given any measurable μ with $\|\mu\|_\infty = k < 1$, there is a K -quasiconformal map with dilatation μ almost everywhere. An important consequence of this is that if f is entire and φ is quasiconformal, then there exists a quasiconformal ψ so that $g = \varphi \circ f \circ \psi$ is entire. Two entire functions f and g that are related in this way are called quasiconformally equivalent. Eremenko and Lyubich proved that if f has q singular values, then the collection of entire functions that are quasiconformally equivalent to f forms a $(q + 2)$ -dimensional complex manifold (see [9, Section 3]).

Suppose Ω is a planar open set. A nonnegative Borel function ρ on Ω is called a metric on Ω . Suppose Γ is a collection of locally rectifiable curves in Ω . We say a metric ρ is an admissible metric for Γ if

$$\inf_{\gamma \in \Gamma} \int_{\gamma} \rho ds \geq 1,$$

and we define the modulus of Γ as

$$M(\Gamma) = \inf_{\rho} \int_{\Omega} \rho^2 dx dy,$$

where the infimum is over all admissible metrics for Γ . The reciprocal of $M(\Gamma)$ is called the extremal length of Γ and is denoted by $\lambda(\Gamma)$. The most important facts that we will need are as follows.

Conformal invariance: If $f: \Omega \rightarrow \Omega'$ is conformal, Γ is a path family in Ω and $\Gamma' = f(\Gamma)$, then $M(\Gamma') = M(\Gamma)$.

Quasi-invariance: If $f: \Omega \rightarrow \Omega'$ is K -quasiconformal, Γ is a path family in Ω and $\Gamma' = f(\Gamma)$, then $M(\Gamma)/K \leq M(\Gamma') \leq K \cdot M(\Gamma)$.

Extension: If Γ, Γ' are path families such that each path in Γ' contains a subpath in Γ then $M(\Gamma') \leq M(\Gamma)$. In particular, if $\Gamma' \subset \Gamma$, then $M(\Gamma') \leq M(\Gamma)$.

Parallel rule: If $\Gamma_1, \dots, \Gamma_n$ are defined on disjoint open sets, and every $\gamma \in \cup_j \Gamma_j$ contains some curve in Γ then $M(\Gamma) \geq \sum_j M(\Gamma_j)$.

Round annuli: the modulus of the path family separating the two boundary components of the round annulus $A(r, R) = \{z : r < |z| < R\}$ is $(\log R/r)/2\pi$. We call this the modulus of the annulus. Every topological annulus $\Omega \subset \mathbb{C}$ is conformally equivalent to a round annulus, and its modulus is equal to the modulus of the corresponding round annulus.

Topological annuli: There is a $M_0 < \infty$ so that if Ω is a topological annulus with modulus $M \geq M_0$, then Ω contains a round annulus of modulus $M' > 1$ and M' tends to ∞ as M tends to ∞ .

Reciprocity: The modulus of the path family separating the two boundary components of a topological annulus Ω is the reciprocal of the modulus of the path family in Ω that connects the two boundary components.

Rectangles: The modulus of the path family connecting the sides of length a in a $a \times b$ rectangle is a/b .

Another fact that we will use repeatedly is as follows.

LEMMA 2.1. *Suppose $e, f \subset \mathbb{C}$ are disjoint Jordan arcs and let Γ be the family of closed curves in $\mathbb{C} \setminus (e \cup f)$ that separates them. Let M be the modulus of Γ . Then*

$$\text{dist}(e, f) \geq \epsilon \cdot \min(\text{diam}(e), \text{diam}(f)), \quad (2.1)$$

where $\epsilon > 0$ depends only on a lower bound for M . Conversely, if (2.1) holds, then M is bounded away from zero with an estimate depending only on ϵ . Moreover, ϵ tends to infinity if and only if M tends to infinity.

Proof. This is fairly standard. Let $r = \min(\text{diam}(e), \text{diam}(f))$. If there are points $x \in e$ and $y \in f$ with $|x - y| \leq \epsilon$, then we define a metric ρ on $\{z : \epsilon r < |x - z| < r/2\}$ by

setting $\rho(z) = (|z - x| \log \frac{2}{\epsilon})^{-1}$. It is a standard exercise to show that ρ is admissible and integrating ρ^2 gives $M \leq (\log \frac{2}{\epsilon})^{-1}$, which tends to 0 with ϵ . This proves the first claim. For the other direction, suppose $\text{dist}(e, f) \geq \epsilon r$. Then setting $\rho(z) = (\epsilon r)^{-1}$ on an ϵr -neighborhood of e (if $\text{diam}(e) = r$) or f (otherwise) gives an admissible metric for the path family connecting e to f . Since this neighborhood has area at most $\pi(\epsilon r + r)^2$, computing the integral of ρ^2 shows this family has modulus at most

$$(\epsilon r)^{-2} \pi(\epsilon r + r)^2 \leq \pi(1 + \epsilon^{-2}).$$

Since this modulus is the reciprocal of the modulus of the path family separating e and f , we get a lower bound for the latter modulus in terms of ϵ . If M is large, then by the topological annuli property, there is a large round annulus separating e and f , and hence ϵ is large. Conversely, if ϵ is large, then there is clearly a large round annulus separating the curves and so the modulus M is large. \square

We will use the following in Section 11.

LEMMA 2.2. *If I, J are disjoint intervals on \mathbb{R} , let $M(I, J)$ be the modulus of the path family in $\mathbb{H}_u = \{x + iy : y > 0\}$ (the upper half-plane) separating I and J . If I, J have unit length and are distance $r \geq 2$ apart, then $M(I, J) \simeq \log r$.*

Proof. There are several ways to estimate this modulus, but we will use a conformal map. Without loss of generality, assume $I = [-1, 0], J = [r, r + 1]$. The Schwarz-Christoffel formula (for example, see [8] and its references) that says that \mathbb{H}_u is conformally mapped to an $a \times b$ rectangle with I, J going to the sides of length a by the map

$$f(z) = \int^z \frac{dw}{(w + 1)^{1/2} w^{1/2} (w - r)^{1/2} (w - r - 1)^{1/2}}.$$

Moreover,

$$a = \int_{-1}^0 \frac{dx}{|x + 1|^{1/2} |x|^{1/2} |x - r|^{1/2} |x - r - 1|^{1/2}} \simeq \frac{1}{r} \int_{-1}^0 \frac{dx}{|x + 1|^{1/2} |x|^{1/2}} \simeq \frac{1}{r},$$

and similarly

$$\begin{aligned} b &= \int_0^r \frac{dx}{|x + 1|^{1/2} |x|^{1/2} |x - r|^{1/2} |x - r - 1|^{1/2}} \\ &\simeq \frac{1}{r} \int_0^{r/2} \frac{dx}{|x + 1|^{1/2} |x|^{1/2}} \\ &\simeq \frac{1}{r} + \frac{1}{r} \int_1^{r/2} \frac{dx}{x} \\ &\simeq \frac{1}{r} (1 + \log r). \end{aligned}$$

Therefore, by conformal invariance and the rectangle rule, $M(I, J) = b/a \simeq 1 + \log r$ (and $1 + \log r \simeq \log r$ since $r \geq 2$). \square

Other proofs of the lemma are possible. For example, one can use a Möbius transformation to map I to $[-1, 1]$, map J to the complement of $[-y, y]$ for some $y \simeq r$, and then estimate the modulus of the planar complement of these using explicit metrics.

Several times in this paper, we will use Koebe's $\frac{1}{4}$ -theorem and its consequences. Koebe's theorem says that if $f: \mathbb{D} \rightarrow \Omega$ is conformal (holomorphic and 1-to-1), then

$$\frac{1}{4}|f'(z)|(1 - |z|^2) \leq \text{dist}(f(z), \partial\Omega) \leq |f'(z)|(1 - |z|^2).$$

See [10, Theorem I.4.3]. A consequence of this is that if f is conformal on a region W and $E \subset W$ is compact, then $|f'|$ is comparable at any two points of E with a constant that depends only on E and W (in fact, it only depends on the diameter of E in the hyperbolic metric for W).

The image γ of a line under a quasiconformal mapping of the plane to itself is called a quasiline. Such curves γ are exactly characterized by the three-point condition: there is a $M < \infty$ so that given any three points $x, y, z \in \gamma$ with x, y in different connected components of $\gamma \setminus \{z\}$, we have $|x - z| \leq M|x - y|$. Equivalently, the subarc of γ connecting x and y has diameter $O(|x - y|)$. We will use this in the following way.

A quasidisk is the image of \mathbb{D} under a quasiconformal map of the plane. Abusing notation slightly, we will say Ω is an unbounded quasidisk if it is the image of a half-plane under a quasiconformal map of the plane (this sounds better than 'quasi-half-plane', and would be technically correct if we simply considered quasiconformal maps of the Riemann sphere to itself, rather than just maps that fix ∞).

LEMMA 2.3. *Suppose Ω an unbounded quasidisk. Then there is a $C < \infty$ so that given any $x \in \partial\Omega$, there is a curve γ in Ω that connects x to ∞ and satisfies*

$$\text{dist}(z, \partial\Omega) \geq |z - x|/C,$$

for every $z \in \gamma$. If $|z| \geq 2|x|$, then $\text{dist}(z, \partial\Omega) \geq |z|/(2C)$, for every $z \in \gamma$.

Proof. Suppose $\Omega = f(\mathbb{H}_r)$. Without loss of generality, we may assume $x = f(0)$. The right half-plane can easily be quasiconformally mapped to a quarter-plane by a quasiconformal map g of the plane (leave radii fixed and contract angles by a factor of two in one half-plane and expand them by a factor of $3/2$ in the remaining half-plane; we leave the details to the reader). Thus, if $\Omega_1 \subset \Omega$ is the image of the first quadrant under f , then it is also an unbounded quasidisk, and hence $\partial\Omega_1$ satisfies the three-point condition with some constant C .

Suppose that there was a point on $\gamma = f(\mathbb{R}^+)$ that was 'too close' to $f(i\mathbb{R}^+) \subset \partial\Omega$, that is, suppose there were $s, t > 0$ so that

$$|f(s) - f(it)| < \epsilon|f(s) - f(0)|.$$

Then the arc of $\partial\Omega_1$ connecting $f(s)$ and $f(t)$ must have diameter $\leq C\epsilon|f(s) - f(0)|$ by the three-point condition, but it contains both $x = f(0)$ and $f(s)$ so it has diameter at least $|f(s) - f(0)|$. Thus, $C\epsilon \geq 1$. The same argument applies to the image of the fourth quadrant and the negative imaginary axis, and this proves the first part of the lemma. The final claim follows easily. \square

Mori's theorem states that K -quasiconformal maps of the plane are bi-Hölder, that is,

$$\frac{1}{C|z - w|^\alpha} \leq |f(z) - f(w)| \leq C|z - w|^\alpha,$$

where α depends only on K . Quasiconformal maps of the plane are also quasimetric: there is a homeomorphism η from $[0, \infty)$ to itself such that $|x - y| \leq t|a - b|$ implies $|f(x) - f(y)| \leq \eta(t)|f(x) - f(y)|$ (see [13] and its references).

LEMMA 2.4. *Given a Jordan arc $\gamma \subset \mathbb{C}$ define*

$$\gamma(r) = \{z \in \mathbb{C} : \text{dist}(z, \gamma) \leq r \cdot \text{diam}(\gamma)\}.$$

If f is a K -quasiconformal map of the plane to itself, then there are $0 < s < t < \infty$ depending only on r and K so that if $\sigma = f(\gamma)$ then

$$\sigma(s) \subset f(\gamma(r)) \subset \sigma(t).$$

Proof. Without loss of generality, we may assume $\text{diam}(\gamma) = \text{diam}(\sigma) = 1$. Taking the metric $\rho = 1/r$ on $\gamma(r)$, we see that the modulus of the path family connecting γ to $\partial\gamma(r)$ is bounded above by

$$\frac{\text{area}(\gamma(r))}{r^2} \leq \frac{\pi(1+r)^2}{r^2} = \pi \left(1 + \frac{1}{r^2} \right).$$

Hence, the modulus of the path family separating γ and $\partial\gamma(r)$ is bounded below by the reciprocal of this upper bound. Thus, the f -image of this family therefore also has modulus bounded below (by quasi-invariance), and thus the distance between σ and $f(\partial\gamma(r))$ is bounded below by a constant s times $\text{diam}(\sigma)$. This gives the left-hand inclusion of the lemma.

The other inclusion is easier. If t is large, then the modulus of the path family surrounding σ in $\sigma(t)$ is also large and hence its preimage under f is also large. This means that the preimage contains a large round annulus that surrounds γ , and hence contains $\gamma(r)$ if t is large enough compared to r (depending on K). □

We will also use the following well-known result of Teichmüller, Wittich, Belinskiĭ, and Lehto (for example, [6; 12, Theorem 7.3.1; 14, 18]).

THEOREM 2.5. *Suppose $\varphi: \mathbb{C} \rightarrow \mathbb{C}$ is K -quasiconformal with dilatation μ and*

$$\iint_{|z|>R} |\mu(z)| \frac{dx dy}{|z|^2} < \infty,$$

for some $R < \infty$. Then there is a nonzero, finite complex constant A so that $\varphi(z)/Az \rightarrow 1$ as $|z| \rightarrow \infty$.

3. Quasiconformal folding

In this section, we review notation and results from [2]. Recall that $\mathcal{S}_{2,0} \subset \mathcal{S}$ is the subcollection of Speiser class functions that have two critical values and no finite asymptotic values. We will start by describing how an element of $\mathcal{S}_{2,0}$ gives rise to a locally finite, infinite planar tree; we then describe how to start with such a tree (satisfying some geometric regularity conditions) and obtain an element of $\mathcal{S}_{2,0}$. This construction is the main result of [2], and contains much of the work needed to prove Theorem 1.2.

Suppose $f \in \mathcal{S}_{2,0}$ and that the critical values of f are exactly $\{-1, 1\}$. Let $T = f^{-1}([-1, 1])$. Let $U = \mathbb{C} \setminus [-1, 1]$ and let $\Omega = f^{-1}(U)$. Then each component of Ω is simply connected and f acts as a covering map from each component of Ω to U . The boundary of Ω is an infinite tree, where the vertices are the preimages of $\{-1, 1\}$. For each connected component of Ω , there is a conformal map τ to \mathbb{H}_r so that $f = \cosh \circ \tau$. The edges of $\partial\Omega$ are mapped to intervals of length π on $\partial\mathbb{H}_r$ (see Figure 2).

Given $r > 0$ and an edge e on $\partial\Omega$, we define a neighborhood

$$e(r) = \{z : \text{dist}(z, e) < r \cdot \text{diam}(e)\},$$




FIGURE 2. A function with two critical values at $\{-1, 1\}$ and no finite asymptotic values. $T = f^{-1}([-1, 1])$ is a tree with vertices mapping to ± 1 (shown as black and white dots). τ is a conformal map from each complementary component of T to the right half-plane and $f = \cosh \circ \tau$.




FIGURE 3. The neighborhood $T(r)$ of a tree (a finite tree is shown, but the definition also makes sense for infinite trees and graphs). The dashed regions show $e(r)$ for two edges.

and define a neighborhood of $T = \partial\Omega$ by taking the union over all edges. This neighborhood will be denoted by $T(r)$ (see Figure 3).

Now suppose that we start with an infinite planar tree T and a holomorphic map $\tau : \Omega \rightarrow \mathbb{H}_r$ where $\Omega = \mathbb{C} \setminus T$, and τ is conformal from each connected component Ω_j of Ω to \mathbb{H}_r . We want to construct an $f \in \mathcal{S}_{2,0}$ so that T approximates $f^{-1}([-1, 1])$ and f approximates e^z away from T . There are two basic conditions that we impose.

(1) *Bounded geometry*: This holds for T if

- (1) every edge is C^2 with uniform bounds on the derivatives;
- (2) edges meet at angles uniformly bounded away from zero;
- (3) any two adjacent edges have uniformly comparable lengths and their union is uniformly quasiconvex;
- (4) nonadjacent edges e, f satisfy $\text{dist}(e, f)/\text{diam}(e) > \epsilon$ with a uniform $\epsilon > 0$.

Here ‘quasiconvex’ means that the arc-length distance between two points, x, y , on the curve is $O(|x - y|)$. Note that condition (2) implies that the vertex degrees of T are uniformly bounded. A very useful alternate version of (4) comes from Lemma 2.1: (4) holds if and only if any two nonadjacent edges of T are separated by a path family with modulus bounded uniformly away from zero. Because of the conformal invariance of modulus, this allows us to

easily verify that under certain conditions, conformal images of bounded geometry trees still have bounded geometry (see Section 4).

Later, we will also consider a bounded geometry ‘forest’ G that is a disjoint union of bounded geometry trees, where (1)–(3) hold for all edges in the forest and (4) holds for all pairs of nonadjacent edges in G (either from the same or from different components of G).

Each edge e of the planar tree T has two sides and each side may be considered as a boundary arc of one of the complementary components Ω_j of T (possibly both sides belong to the same component). Conversely, the boundary of each component Ω_j is partitioned into arcs by the sides of the tree T . We say that two sides of T are adjacent if they are sides of adjacent edges of T that are on the boundary of the same complementary component Ω_j and the two sides correspond to adjacent intervals after conformally mapping Ω_j to \mathbb{H}_r . Two sides of T can also be adjacent if they are opposite sides of a single edge of T that has an endpoint of degree 1.

When Ω_j is mapped to \mathbb{H}_r by τ_j , the sides of T map to intervals on $\partial\mathbb{H}_r$. The Euclidean length of the image interval is called the τ -length of the corresponding side of T . The collection of resulting intervals on $\partial\mathbb{H}_r$ form a partition, denoted by \mathcal{P}_j , of this line.

For us, a partition of a line is a locally finite collection of disjoint open intervals whose closures cover the whole line. The endpoints of the partition intervals form a countable, discrete set that accumulates only at ∞ . We say that a partition has *bounded geometry* if adjacent elements (that is, partition intervals that share an endpoint) have comparable lengths with a constant that is independent of the intervals. The bounded geometry constant of the partition is the supremum $|I|/|J|$ over all adjacent pairs of intervals. Occasionally, we will also consider bounded geometry partitions of bounded open segments or arcs that are defined in the same way (adjacent partition intervals have comparable lengths and accumulate only at the endpoints).

If the infinite tree T has bounded geometry then the partitions \mathcal{P}_j of $\partial\mathbb{H}_r$, corresponding to each complementary component Ω_j , also have bounded geometry, with a constant depending only on the bounds in the definition of bounded geometry (see [2, Lemma 4.1]). In other words, if T has bounded geometry, then adjacent sides of T have comparable τ -length. This fact is the main way that we utilize the bounded geometry assumption.

(II) *The τ -length lower bound:* The second condition we require is that every side of T has τ -length $\geq \pi$ (but no upper bound is assumed).

An apparently weaker form of this is to simply require that for each complementary component Ω_j of T , there is a strictly positive lower bound $\epsilon_j > 0$ for the length of every interval in the partition \mathcal{P}_j . If this weaker condition holds, then on each component Ω_j of the model domain Ω , we can replace τ_j by a positive multiple of itself, namely $(\pi/\epsilon_j) \cdot \tau_j$. This is still a conformal map of Ω_j to the right half-plane but now every partition arc on $\partial\mathbb{H}_r$ has length at least π . Thus if each tract has a positive τ -length lower bound, we can easily choose a new model function for which it satisfies the stronger at least π bound. Therefore, in most cases, we only need to check the weaker condition. Note that having a positive τ -length, lower bound is a geometric property of each tract in Ω ; if each tract has such a lower bound, then having a positive lower bound that works simultaneously for all the tract depends on the particular choice of model function $F = e^\tau$.

The following is the main result from [2].

THEOREM 3.1. *Suppose (T, τ) has bounded geometry and every side of T has τ -length at least π . Then there is a $f \in \mathcal{S}_{2,0}$, $r > 0$ and a quasiconformal $\varphi: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ so that $f \circ \varphi = \cosh \circ \tau$ off $T(r)$. In addition,*

- (1) *the map φ is conformal off $T(r)$;*
- (2) *the function f has only critical values ± 1 , and no finite asymptotic values;*
- (3) *the number r and the maximal dilatation K of φ only depend on the bounded geometry constants of T ;*

- (4) *the degree of any critical point of f is bounded by $4D$, where D is the maximum degree of vertices in T .*

This result is essentially [2, Theorem 1.1]. The statement there includes conclusions (1), (2), and (3) explicitly. Conclusion (4) follows from the proof given in [2]; by construction, the degree of any critical point of f equals the graph degree of a corresponding vertex v of a tree T' that is obtained by adding finite trees to the vertices of T . The bounded geometry condition implies that the vertices of T have uniformly bounded degree, D . Each of the added trees has maximum vertex degree 4 and has degree 3 at the vertex that is attached to $v \in T$. At most $\deg(v)$ such trees are added at v , so the new tree has maximum degree at most $4D$. This gives (4). Since the tree T that we construct in the proof of Theorem 1.2 will have maximum vertex degree 3, this gives the ‘12’ in part (2) of Theorem 1.2.

Conclusion (4) (and hence the estimate in Theorem 1.2) can be improved by making some alterations to the folding construction in [2], but we only sketch the possibilities here; details will appear elsewhere. It is fairly easy to modify the construction so that the finite trees we add to T have degree 1 at the vertex that is attached to T ; this requires redrawing the diagrams in [2, Figure 8] with a new vertical segment at endpoints of the horizontal segment and changing some corresponding bookkeeping in the proof. This change lowers the bound in (4) above from $4D$ to $2D$. If we replace τ by a positive multiple of itself, we can do even better. Subdivide each edge of T into three subedges, so as to give a bounded geometry tree T' and rescale τ so that each new edge has τ -length at least π . We then further modify the construction in [2] by adding two extra vertices to the horizontal edges in [2, Figure 8] and make corresponding changes to the bookkeeping. This results in a finite tree being added to every third vertex of T' viewed from a complementary component, and we can easily arrange all these to be the ‘new’ vertices of T' , and these all have degree 2. We can add at most two trees to any such vertex, making the degree at most 4. This makes the bound in (4) equal to $\max(D, 4)$.

As noted above, in the proof of Theorem 1.2, we apply Theorem 3.1 to a tree T with maximal degree 3. Some of the complementary components of T correspond to components of $\Omega(\rho)$, and on these components all the τ -lengths equal π . This implies that the folding construction does not add any trees inside these components. The other complementary components of T are ‘new’ components and we are free to choose any positive multiple of τ we want on these components. Therefore, we may assume the τ -lengths in these new components are all large, and hence the sketched argument above applies: finite trees are adjoined to vertices of degree 2 in T , and the added trees have maximum degree 4 and degree 1 at the vertex attached to T (so the degree 2 vertex becomes degree at most 4). Thus, making the appropriate changes to the folding construction in [2] will give the upper bound 4 in part (2) of Theorem 1.2.

One could improve the bound $\max(D, 4)$ to $\max(D, 3)$ if the finite trees we add have maximum vertex degree 3. Figure 12 of [2] shows how degree 4 vertices arise when adjacent trees are attached to each other. It seems plausible that this can be avoided, but requires more extensive modifications and needs to be verified. We would also need to add at most one tree to any vertex of T' (the tree obtained from T by splitting each edge into three edges). However, this is easy to arrange by always attaching trees one vertex to the left of a vertex of T , when viewed from the corresponding complementary component; this will attach at most one tree to each of the two vertices of T' that lie on any edge of T . Together, these improvements would give the bound 3 in part (2) of Theorem 1.2.

4. *Glueing trees using conformal maps*

In this section, we describe a way to combine two or more bounded geometry trees to obtain a new bounded geometry tree.

It is convenient to introduce a stronger version of bounded geometry. We say that a Jordan arc γ is ϵ -analytic if there is conformal map on

$$\gamma(\epsilon) = \{z : \text{dist}(z, \gamma) < \epsilon \cdot \text{diam}(\gamma)\}$$

that maps γ to a line segment. We call a bounded geometry tree T *uniformly analytic* if there is an $\epsilon > 0$ so that every edge of T is ϵ -analytic. We say a vertex v of T is ϵ -analytic if it has degree 2 and the union of the two edges meeting at v form a single ϵ -analytic Jordan arc. Note that vertices of a uniformly analytic tree need not be analytic (the edges may meet at various angles), but if we add vertices to the edges of a uniformly analytic tree T_1 to form a new bounded geometry tree T_2 , then all the new vertices are analytic with the same constant as T_1 . A bounded geometry forest in which all the edges are uniformly analytic will be called a *uniformly analytic forest*. An important example of such a forest is $\partial\Omega(\rho)$, where Ω is a model domain and the vertices are the usual ones, $\tau^{-1}(\rho + i\pi\mathbb{Z})$. In this case, all the vertices are analytic as well (with a uniformly bounded constant).

LEMMA 4.1. *Suppose T is a uniformly analytic forest and Ω is a connected component of $\mathbb{C} \setminus T$. Suppose W is either \mathbb{D} or \mathbb{H}_r and $\tau: \Omega \rightarrow W$ is conformal. Suppose $T_0 \subset \overline{W}$ is a uniformly analytic tree (all open edges of T_0 are in W , but some vertices may lie on the boundary of W). We assume that $\tau^{-1}(T_0)$ is locally finite in \mathbb{C} . Suppose there is a $M < \infty$ so that for every edge e of T_0 either*

- (1) *the edge e has hyperbolic diameter at most M (we call these the internal edges of T_0), or*
- (2) *the edge e has one endpoint $x = \tau(v) \in \partial W$, where v is an analytic vertex of T and*

$$\frac{1}{M} \text{diam}(\tau(I \cup J)) \leq \text{diam}(e) \leq M \cdot \text{diam}(\tau(I \cup J)),$$

where we take Euclidean diameters, and I and J are the two edges of T adjacent to v . We call such an edge e a boundary edge of T_0 .

Then $T' = T \cup \tau^{-1}(T_0)$ is a uniformly analytic forest. The constants for T' depend only on the bounded geometry and uniform analyticity constants for T and T_0 . If a component of $W \setminus T_0$ satisfies a positive τ -length lower bound, the same bound is satisfied by the image of this component under τ^{-1} .

Proof. The Koebe distortion theorem easily implies that conditions (1)–(4) in the definition of bounded geometry are transferred from T_0 to $\tau^{-1}(T_0)$ for all pairs of internal edges in T_0 . Similarly, the images of all internal edges are clearly uniformly analytic. Moreover, because of our assumption on the hyperbolic diameters, each internal edge e of T_0 is separated from ∂W by a path family in W with modulus bounded uniformly away from zero. By conformal invariance of modulus, this also holds for $\tau^{-1}(e)$ and $\partial\Omega$ and hence (4) holds whenever one edge corresponds to an internal edge of T_0 and the other is an edge of T . A similar argument works for an internal edge of T_0 and a nonadjacent boundary edge of T_0 .

Each boundary edge e of T_0 has an endpoint x corresponding to an analytic vertex v of T and by Schwarz reflection the map τ^{-1} extends analytically a uniform neighborhood of $S = \tau(I \cup J)$, where I, J are the edges of T adjacent to v . Since $\text{diam}(S) \gtrsim \text{diam}(I \cup J)$, this implies $\tau^{-1}(e)$ is uniformly analytic. Moreover, there is an $\epsilon \cdot \text{diam}(e)$ neighborhood of e , where τ^{-1} extends to be conformal and whose image hits I and J , but no other edges of T . This (and the Koebe distortion theorem) implies the separation property (4) holds for $\tau^{-1}(e)$. The final statement holds simply because having a positive lower bound for τ -lengths is conformally invariant by definition. □




FIGURE 4. The idea of the proof of Theorem 1.2 is to reduce it to Theorem 3.1 by joining the components of $\Gamma = \partial\Omega(\rho)$ to form a bounded geometry tree T_1 and then add extra edges to the ‘new’ components make a tree T_2 so that the τ -length lower bound holds.

5. Proof of Theorem 1.2: part 1, bounded geometry

It is stated in [2] that Theorem 3.1 reduces constructing functions in $\mathcal{S}_{2,0}$ to ‘drawing a picture’ of the correct tree. One then has to verify that the tree has bounded geometry and the complementary components each satisfy a positive τ -length lower bound. This is exactly what we will do to prove Theorem 1.2. In this section, we connect the various components of $\Gamma = \partial\Omega(\rho)$ to form a bounded geometry tree T_1 . If $\partial\Omega(\rho)$ has $N < \infty$ components, then the tree T_1 will have $2N$ complementary components; N of these are the original components of $\Omega(\rho)$ and the other N are subdomains of $W = \mathbb{C} \setminus \overline{\Omega}$. When N is finite, it is easy to make the connections if we are willing to allow the bounded geometry constant to grow. However, we will give a more intricate construction that can also deal with infinitely many components and gives uniformly bounded geometry.

The new components might not satisfy a τ -length lower bound condition, but we will fix this in the next section by a simple trick that subdivides each of these new components into infinitely many components, each with a positive τ -length lower bound. Later, in Section 9, we will show how to replace each component by a single subdomain that has the desired τ -length lower bound (see Figure 4).

We should note that the construction of T_1 and T_2 given in this paper was chosen for its generality, but it might not be the most elegant or efficient choice for a particular Ω that arises in some application. Very likely, the geometry of the model domain will suggest a natural way of connecting the different components of $\partial\Omega(\rho)$, while satisfying the bounded geometry and τ -length conditions. Theorems 1.2 and 1.3 simply ensure that there is always at least one way to accomplish this.

Now we start the construction of T_1 . Let $\Gamma = \partial\Omega(\rho)$. This is a union of unbounded, analytic Jordan curves and each curve comes with a set of marked points (or vertices) defined by $\text{Im}(\tau(z)) \in \pi i\mathbb{Z}$ (recall that this is called a conformal partition of $\partial\Omega(\rho)$). Γ is a uniformly analytic forest and every vertex is analytic with a uniform constant.

Let $W = \mathbb{C} \setminus \overline{\Omega(\rho)}$. This is a proper simply connected domain in the plane, so by the Riemann mapping theorem, there is a conformal map $\Psi: W \rightarrow \mathbb{D}$. Each curve $\Gamma_j = \partial\Omega_j(\rho)$ maps to an open arc $I_j \subset \mathbb{T}$ under Ψ . We let $E = \mathbb{T} \setminus \cup_j I_j$; this compact set corresponds to ∞ under Ψ^{-1} , hence it has zero Lebesgue length (even stronger, it has zero logarithmic capacity, but we will not need this). The partition of $\Gamma_j = \partial\Omega_j(\rho)$ with endpoints $\tau_j^{-1}(i\pi\mathbb{Z})$ corresponds via Ψ to a partition of I_j . Because $\partial\Omega(\rho)$ is a bounded geometry forest (with constant depending only on ρ), adjacent intervals in the partition of I_j have comparable lengths with a fixed constant, depending only on ρ . In particular, we can choose a point $v_j \in I_j$ so that the distances from v_j to each endpoint of I_j are comparable to each other (just take an endpoint of a partition interval that contains the actual center of I_j). We call v_j the ‘approximate center’ of I_j . The main objective of this section is to prove




FIGURE 5. The Whitney decomposition of the disk.

LEMMA 5.1. Suppose that the notation is as above, that is, T_0 is the tree consisting of $\Gamma = \partial\Omega(\rho)$ with vertices given by the conformal partition on each component of $\partial\Omega(\rho)$. There is a bounded geometry tree T_1 that contains T_0 , so that

- (1) all new edges are in $W = \mathbb{C} \setminus \overline{\Omega(\rho)}$;
- (2) the vertices of T_1 on $\partial\Omega(\rho)$ are exactly the vertices of T_0 (no new vertices are added on $\partial\Omega(\rho)$);
- (3) if $\Omega(\rho)$ has $N < \infty$ components then $\mathbb{C} \setminus T_1$ has $2N$ connected components, N of which are the connected components of $\Omega(\rho)$ and N are subdomains of W .

Using Lemma 4.1, the proof of Lemma 5.1 reduces to the following construction on the disk.

LEMMA 5.2. Suppose $E \subset \mathbb{T}$ is closed and has length zero, $\mathbb{T} \setminus E = \cup_j I_j$ and $v_j \in I_j$ are the approximate centers, as above. Then there is a tree T in \mathbb{D} that satisfies the following conditions.

- (1) The tree T has bounded geometry. In fact, T is uniformly analytic and every edge is either a line segment or a circular arc. The maximum vertex degree is 3.
- (2) For each v_j , there is a boundary edge of T that has v_j as a common endpoint. The length of this edge is comparable to the lengths of the two partition arcs of I_j that have v_j as an endpoint. This edge makes an angle with \mathbb{T} that is bounded uniformly away from zero.
- (3) Every other arc of T has uniformly bounded hyperbolic diameter.
- (4) The closure of every component of $\mathbb{D} \setminus T$ meets E in exactly one point. In particular, if E is a finite set with N elements, then $\mathbb{D} \setminus T$ has N components.

Proof. Consider a Whitney decomposition of the disk, as illustrated in Figure 6. The innermost part of the decomposition is a central disk of radius $1/4$. Outside the central disk, the annulus $A_1 = \{1/4 < |z| < 1/2\}$ is divided into eight equal sectors, the annulus $A_2 = \{1/2 < |z| < 3/4\}$ into sixteen sectors, and so on, as shown in Figure 5. These sectors are called Whitney boxes.

Each Whitney box has two radial sides and two circular arc sides concentric with the origin. The circular arc closer to the origin is called the top of the box and the arc further from the origin is called the bottom. Each bottom arc is divided into two pieces by the tops of the Whitney boxes below it ('below' means between the given box and the unit circle). We call




FIGURE 6. The paths from the boundary to the central disk described in the text. Any boundary point can be joined to the central disk by a path moving along edges of Whitney boxes: move radially toward the origin whenever possible, and move counterclockwise (right in the figure) otherwise.

these the left and right sides of the bottom arc (left is the one further clockwise). We will call the sides and bottoms of Whitney boxes as the Whitney edges, their endpoints as Whitney vertices. The union of these edges and vertices forms an infinite graph in \mathbb{D} , which we call the Whitney graph. The radial projection of a closed Whitney box B onto the unit circle, \mathbb{T} , is a closed arc that we denote B^* (this is sometimes called the ‘shadow’ of B , thinking of a light source at the origin). The union of a closed Whitney box B and all the closed Whitney boxes B' so that $(B')^* \subset B^*$ is called the Carleson square with base $I = B^*$.

Each point on the unit circle can be connected to the central disk by a path in the Whitney graph that moves toward the origin whenever possible and moves counterclockwise otherwise (see Figure 6).

Note that such a path never contains the ‘left-half’ of the bottom of any Whitney box (otherwise the path would have moved up the left radial side of the box). For each arc $I_j \subset \mathbb{T}$, we connect the approximate center v_j of I_j to the central disk by such a path. The union of all such paths, together with the boundary of the central disk, is a closed set and divides the disk into countably many simply connected subdomains $\{U_j\}$. By removing one of the eight arcs that bounds the central disk, we join the central disk to one of the domains U_j . This makes every subdomain U_j an infinite union of Whitney boxes; a finite union would contain a box closest to the unit circle and the bottom of this box would be on a path, which is impossible since the left side of the bottom cannot be on any path.

Thus, every subdomain U_j has a boundary that hits \mathbb{T} , and $J_j = \partial U_j \cap \mathbb{T}$ must be a closed interval; if J_j is not connected, then there is a component of $\mathbb{D} \setminus U_j$ that is separated from the central disk by U_j , but this is impossible by construction (points on the boundary of this component are on a path that continues all the way to the central disk).

The closed interval J_j must hit $E = \mathbb{T} \setminus \cup_j I_j$, otherwise two paths were generated in the same component I_j of $\mathbb{T} \setminus E$, contrary to the construction. Also J_j must hit E in a single point, x_j , otherwise U_j separates some component I_k of $\mathbb{T} \setminus E$ from the central disk, contradicting the fact that the approximate center of I_k is connected to the central disk.

We would like to turn the curve $\partial U_j \setminus E$ into a tree by using the partition vertices on $\mathbb{T} \cap \partial U_j$ and using the Whitney vertices on ∂U_j , but there are infinitely many Whitney vertices on ∂U_j that accumulate at each approximate center v_j . To fix this, recall that v_j is the endpoint of two partition intervals of comparable length. Suppose r is the length of the shorter of these




FIGURE 7. The paths connecting v_j to the central disk approach v_j through a nontangential cone near the boundary. Thus, if the path to $v_j \in I_j$ is truncated at an appropriate scale and replaced by a line segment, this segment makes an angle with \mathbb{T} that is uniformly bounded away from zero.

two and let $w_j \in \partial U_j \cap \mathbb{D}$ be a Whitney vertex on the path starting at v with distance from \mathbb{T} between r and $r/2$. We call this the truncation point associated to v_j . The path terminating at v_j lies in a cone in \mathbb{D} with radial axis, fixed angle, and vertex at v_j and there are no other paths that hit the disk $D(v_j, |v_j - w_j|)$, so we can replace the part of this path between these points by the line segment $[v_j, w_j]$. The length of this segment is comparable to the distance between v_j and its neighboring partition points and is also comparable to the adjacent segment in the path ∂U_j . Moreover, the angle between this segment and the unit circle is uniformly bounded away from zero, so we obtain a bounded geometry tree (even uniformly analytic), as desired (see Figure 7). \square

6. Proof of Theorem 1.2: part 2, the τ -length bound

In the previous section, we showed how to connect the components of $\Gamma = \partial\Omega(\rho)$ into a single, connected, bounded geometry (even uniformly analytic) tree T_1 . In this section, we modify the construction further to give a positive τ -length lower bound on each complementary component; as noted earlier, it is then easy to modify τ by multiplying it by a positive constant on each component of $\Omega(\rho)$ to make the lower bound π , as required in Theorem 3.1. We only have to prove a lower bound on the ‘new’ complementary components that we create; the sides of the components of $\Omega(\rho)$ have τ -length equal to π by definition.

Let $T \subset \mathbb{D}$ be the tree constructed in the Section 5 and let T' be the tree we obtain by adding a vertex at the midpoint of each edge of T (all the edges are segments or circular arcs so the midpoint is well defined). Note that all the ‘new’ vertices are analytic vertices with a uniform constant. These vertices all have degree 2, and later we will attach single edges to them, giving vertices of degree 3; the resulting trees will have maximum degree 3.

Let $\{U_j\} = \mathbb{D} \setminus T$ be the complementary components of T . Let $x_j = \partial U_j \cap E$ be as defined in the previous section, and let $\Phi_j: U_j \rightarrow \mathbb{H}_r$ be conformal with $\Phi_j(x_j) = \infty$. The vertices of T' on ∂U_j map to points on $\partial\mathbb{H}_r$; let \mathcal{P}_j be the bounded geometry partition of $\partial\mathbb{H}_r$ induced by these points. The ‘new’ vertices of T' induce a bounded geometry partition \mathcal{Q}_j whose endpoints are alternating endpoints of \mathcal{P}_j .

LEMMA 6.1. *With notation as above, fix j and consider the union of horizontal rays in \mathbb{H}_r that start at each endpoint for the partition \mathcal{Q}_j . Along each ray, add vertices that are equally spaced, with a spacing that is equal to the distance between that ray and the closer of the two adjacent rays (see Figure 8). This is a uniformly analytic forest that we denote by G_j . Then*




FIGURE 8. Given a partition of $\partial\mathbb{H}_r$ where adjacent intervals have comparable lengths, add a horizontal ray in \mathbb{H}_r at each partition point in $\partial\mathbb{H}_r$ and place equally spaced vertices on each ray, where the spacing equals the smaller width of the two adjacent half-strips. It is easy to see that this gives a bounded geometry tree that satisfies a positive τ -length lower bound. Note that the vertices of the tree have maximum degree 3.

$T_2 = T_1 \cup \cup_j \Phi_j^{-1}(G_j)$ is a bounded geometry, uniformly analytic tree that satisfies a positive τ -length lower bound.

Proof. It is obvious that G_j has bounded geometry and is uniformly analytic and that it satisfies all the other hypotheses of Lemma 4.1, so that $T_2 = T_1 \cup \cup_j \Phi_j^{-1}(G_j)$ is indeed a bounded geometry, uniformly analytic tree. To prove that each connected component of T_2 satisfies a τ -length lower bound, we simply note the conformal map of a half-strip to a half-plane has exponential growth (we can check this via an explicit formula involving $\sinh(z)$), so if the partition segments on the boundary of the half-strip have Euclidean lengths bounded below by a constant times the width of the strip, then the τ -lengths grow exponentially with a uniform bound. In particular, the τ -lengths are uniformly bounded away from zero. \square

Verifying the τ -length condition for a half-strip above is simple because there is an explicit formula for the conformal map to \mathbb{H}_r . One could also use the more geometrical and more general Lemma 8.1, which will be stated and proved later.

7. Proof of Theorem 1.2: part 3, final details

The construction in the previous section created a forest in \mathbb{H}_r with infinitely many complementary components, each of which satisfies a positive τ -length lower bound. We now apply this construction to the partitions $\{\mathcal{Q}_j\}$ corresponding to the analytic vertices of the tree T' constructed at the beginning of Section 6. Using Lemma 4.1, we can attach a conformal image of the forest created in Lemma 6.1 to T' to give a bounded geometry tree. A positive τ -length lower bound holds automatically by the conformal invariance of this condition. As before, we can multiply τ by positive constants on each component, so that the τ -length lower bound is π . We then apply Theorem 3.1 and multiply the resulting function by e^ρ to get

THEOREM 7.1. *Suppose Ω is as in Theorem 1.1. Then there is a $f \in \mathcal{S}_{2,0}$ and a K -quasiconformal map ϕ of the plane so that $f \circ \phi = \cosh \circ \tau$ on Ω and ϕ is conformal on $\Omega \setminus T(r)$. The constants $K, r < \infty$ depend on ρ but are otherwise independent of Ω and τ . The function f has no finite asymptotic values, exactly two critical values, $\pm e^\rho$, and every critical point has degree at most 12.*

There are a few slight differences between this and Theorem 1.2, but it is easy to deduce Theorem 1.2 from Theorem 7.1 as follows.

First, Theorem 7.1 uses \cosh instead of \exp . However, these functions are almost the same in \mathbb{H}_r away from the boundary. Consider the map $z \rightarrow \frac{1}{2}(z + \frac{1}{z})$; this is a conformal homeomorphism of $\{|z| > 1\}$ to $U = \mathbb{C} \setminus [-1, 1]$ and maps the circle $C = \{|z| = e^\rho\}$ to some ellipse E . Define a quasiconformal map ψ that equals the inverse of this map outside E and extends it diffeomorphically to the interior. Since $\cosh = \frac{1}{2}(e^z + e^{-z})$, we get $\exp(z) = \psi(\cosh(z))$ when $|\exp(z)| > e^\rho$. Therefore, if we use the measurable Riemann mapping theorem to find a quasiconformal φ so that $F = \psi \circ f \circ \varphi$, this function satisfies Theorem 7.1 with \cosh replaced by \exp .

Second, Theorem 7.1 only claims that φ is conformal off $T(r)$, whereas Theorem 1.2 says that it is conformal on all of $\Omega(2\rho)$. The first step in verifying this stronger condition is to prove

LEMMA 7.2. *With notation as above, there is a $A < \infty$, depending only on ρ and r so that $T(r) \cap \Omega(A \cdot \rho) = \emptyset$.*

Proof. The proof is a modulus argument. The modulus of the path family in $\mathbb{H}_r + \rho$ separating a segment of length π on $\{x = \rho\}$ from the vertical line $\{x = A\rho\}$ is easily seen to increase to infinity as A increases to infinity. Thus, by conformal invariance of modulus and Lemma 2.1

$$\frac{\text{dist}(I, \partial\Omega(A\rho))}{\min(\text{diam}(I), \text{diam}(\partial\Omega(A\rho)))} = \frac{\text{dist}(I, \partial\Omega(A\rho))}{\text{diam}(I)} \rightarrow \infty$$

as $A \rightarrow \infty$. This proves the intersection is empty if A is large enough. □

We can easily choose a quasiconformal map $H: \mathbb{H}_r \rightarrow \mathbb{H}_r$ so that H

- (1) is the identity on $\{0 < x < \rho\}$;
- (2) is of the form $(x, y) \rightarrow (ax + b, y)$ mapping $\{\rho < x < 2\rho\}$ to $\{\rho < x < A\rho\}$;
- (3) is a horizontal translation from $\{x > 2\rho\}$ to $\{x > A\rho\}$.

Defining $G = \tau_j^{-1} \circ H \circ \tau_j$ on each Ω_j and letting G be the identity elsewhere gives a quasiconformal map of the plane to itself so that $\tilde{\varphi} = \varphi \circ G$ is conformal off $\Omega(2\rho)$ and satisfies all the other conclusions of Theorem 1.2.

Finally, the tree we have explicitly constructed has maximal vertex degree 3. Hence by the folding theorem (Theorem 3.1), the corresponding entire function will have critical points of degree at most 12. (As noted at the end of Section 3, this bound can be improved to 4 by some modifications to the construction in [2].)

8. Two estimates on τ -length

In this section, we give two explicit estimates for τ -lengths that we will use during the proof of Theorem 1.3 in the next section.

Given two disjoint intervals K, J on the real line, let $M(K, J)$ be the modulus of the path family in the upper half-plane, \mathbb{H}_u , that separates K from J (these are the paths in \mathbb{H}_u with

two endpoints in $\mathbb{R} \setminus (K \cup J)$, exactly one of which separates K and J). This is the reciprocal of the modulus of the path family that joins K and J (paths in \mathbb{H}_u that have one endpoint in each of K and J).

The first result is helpful for domains that look like ‘tubes’ built by attaching quadrilaterals of bounded modulus end-to-end (for example, as the half-strip is a union of squares joined end-to-end).

LEMMA 8.1. *Suppose $K = (-\infty, -1]$ and $\{J_j\}$ is a sequence of disjoint intervals in $[1, \infty)$ such that $M = \sup_j M(J_j, K) < \infty$. Also assume that the $\{J_j\}$ are in increasing order (that is, J_{j+1} is to the right of J_j). Then the lengths of J_j grow exponentially in j ; in particular, these lengths are uniformly bounded below by a constant depending only on M .*

Proof. Fix some $J_j = [a_j, b_j]$ with $1 < a_j < b_j$. If $b_j - a_j \leq \epsilon a_j$ for some $0 < \epsilon < 1$, then J_j is separated from K by the annulus

$$A = \{z \in \mathbb{H}_u : \epsilon a_j < |z - a_j| < a_j\}.$$

Any path connecting different components of $\partial A \cap \mathbb{R}$ also separates J and K so the modulus of the first family is a lower bound for the modulus of the second. However, this modulus is $\frac{1}{2\pi} \log \frac{1}{\epsilon}$, so $\epsilon \geq \epsilon_M = \exp(-2\pi M)$. Hence, $b_j > (1 + \epsilon_M)a_j$. By induction $|J_j| = b_j - a_j \geq \epsilon_M(1 + \epsilon_M)^{j-1}$, as desired. \square

The next lemma is helpful when we build a domain by taking a ‘tube domain’ and attach ‘rooms’ along the sides of the tube.

LEMMA 8.2. *Suppose $K = [s, t]$, $I = [x, y]$, $J = [u, v]$ are intervals on the real line so that $t \leq x < u$. If $M(J, K) \leq M(I, K)$, then $|I| \leq |J|$.*

Proof. We prove the contrapositive. Suppose $|I| > |J|$. After translating (if necessary) we may assume that $t = 0$. Then dilate by $\lambda = u/x > 1$. Note that $K \subset \lambda K$ and $J \subset \lambda I$ (strictly), so using the monotonicity and conformal invariance of modulus, we deduce $M(J, K) > M(\lambda I, \lambda K) = M(I, K)$. \square

9. Proof of Theorem 1.3

In this section, we improve Lemma 6.1 by showing that instead of creating infinitely many complementary components, we can accomplish the same result using a single complementary component.

LEMMA 9.1. *Suppose \mathcal{Q} is a bounded geometry partition of $\partial\mathbb{H}_r$ with constant C (that is, adjacent intervals have length ratio at most C). Then there is a bounded geometry, uniformly analytic forest $T'' \subset \mathbb{H}_r$ which satisfies the hypotheses of Lemma 4.1 and so that $W'' = \mathbb{H}_r \setminus T''$ consists of a single component that satisfies a positive τ -length lower bound. The constants associated to T'' depend only on C . Moreover, the boundary edges of T meet $\partial\mathbb{H}_r$ exactly at the partition points of \mathcal{Q} .*

Given the lemma, we can complete the proof of Theorem 1.3 just as we finished the proof of Theorem 1.2 in Section 7. Briefly, we had constructed a tree T' that contained the analytic arcs $\partial\Omega(\rho)$ as well as arcs that connected the various components of $\partial\Omega(\rho)$. The complementary components of T' consist of the N components of $\Omega(\rho)$ (which already satisfy the τ -length condition because all their sides have τ -length π by definition) and N other components $\{U'_j\}$ (which might not satisfy the τ -length condition; this is what we want to fix). The tree T' is




FIGURE 9. This figure encapsulates the proof of Lemma 9.1 by showing the relevant tree. On the left is the central region, in the center is the central tube, and on the right is the tree T'' . We can easily add vertices to make this a bounded geometry tree, and we will use extremal length estimates to show W'' satisfies a positive τ -length lower bound.

uniformly analytic, and alternate vertices are analytic (T' was obtained from an analytic tree T by adding midpoints of edges). We then map each of the U'_j conformally to \mathbb{H}_r , take \mathcal{Q} to be the image of the new analytic vertices and apply Lemma 9.1 to this partition. The resulting forest T'' is then mapped conformally back to U'_j and attached to T' . Using Lemma 4.1, we see that the resulting tree has bounded geometry and satisfies a lower τ -length condition. The rest of the proof of Theorem 1.3 then exactly follows the proof of Theorem 1.2. Thus to prove Theorem 1.3, it suffices to establish Lemma 9.1.

We now start the proof of Lemma 9.1. Fix a partition \mathcal{Q} of $\partial\mathbb{H}_r$. Choose a base interval I_0 in the partition. Without loss of generality, we may assume $I_0 = [-i, i]$ and label the partition endpoints $\{z_j\} = \{ix_j\} \subset \partial\mathbb{H}_r$ so that

$$\cdots x_{-3} < x_{-2} < x_{-1} = -1 < x_1 = 1 < x_2 < x_3 \cdots$$

Note that the elements of \mathcal{Q} are labeled by \mathbb{Z} , but the endpoints are labeled by $\mathbb{Z}^* = \mathbb{Z} \setminus \{0\}$. It will be convenient to define $k^* = k + 1$ if $k > 0$ and $k^* = k - 1$ if $k < 0$; for $k \in \mathbb{Z}^*$, k^* is the integer adjacent to k , but farther from 0. With this notation, we can write $I_k = (x_k, x_{k^*})$ without needing to have special cases for $k > 0$ and $k < 0$ (although we accept that an interval can be written as either (a, b) or (b, a)).

Define the central region in \mathbb{H}_r as the union of the rectangle $[0, 1] \times [-1, 1]$ and the sector $\{x + iy \in \mathbb{H}_r : |y| < x\}$. This region is illustrated in the left part of Figure 9. The boundary of the central region consists of the segment $I_0 \subset \partial\mathbb{H}_r$ and two infinite paths in \mathbb{H}_r that we will call the upper and lower boundaries.

It would be very convenient for us if the partition \mathcal{Q} was symmetric with respect to the origin, that is, $J_{-k} = -J_k$. Since this need not be the case, we will build a new partition \mathcal{Q}' that is symmetric, has bounded geometry, and is ‘finer’ than \mathcal{Q} in a certain sense. More precisely,

LEMMA 9.2. *Given a bounded geometry partition \mathcal{Q} of $\partial\mathbb{H}_r$ normalized as above, and a real number $1 < M < \infty$, there is a bounded geometry partition $\mathcal{Q}' = \{I'_j\}_{\mathbb{Z}}$ of $\partial\mathbb{H}_r$ so that*

- (1) the interval $I'_0 = I_0$ is an element of \mathcal{Q}' ;
- (2) the partition is symmetric, that is, $I'_{-j} = -I'_j$;
- (3) the length of I'_j is a nonincreasing function of $|j|$;
- (4) the length of every I_j is an integer power of 2;
- (5) if any interval $I' \in \mathcal{Q}' \setminus \{I_0\}$ intersects an interval $I \in \mathcal{Q}$, then $|I'| < |I|/M$;
- (6) the bounded geometry constant of \mathcal{Q}' is bounded above depending only on M and the bounded geometry constant C of \mathcal{Q} .

Proof. Cover I_1 by a collection \mathcal{D}_1 of closed dyadic intervals that all hit I_1 and have lengths strictly less than $|I_1|/4M$ and greater or equal to $|I_1|/8M$ (since there is exactly one power of two in this range, all the chosen intervals have the same length, call it ℓ_1). In general, suppose we have already covered $I_1 \cup \dots \cup I_{j-1}$ by a collection \mathcal{D}_{j-1} of closed dyadic intervals such that

- (1) the interiors are disjoint;
- (2) every $J \in \mathcal{D}_{j-1}$ hits some I_k , $1 \leq k < j$;
- (3) the lengths are nonincreasing;
- (4) if $J \in \mathcal{D}_{j-1}$ hits I_k , $1 \leq k < j$ then $|J| < |I_k|/M$;
- (5) adjacent dyadic intervals have comparable lengths.

Let ℓ_{j-1} be the length of the last (rightmost) dyadic interval in \mathcal{D}_{j-1} . Cover I_j by a collection \mathcal{C}_j of dyadic intervals, all with the same length ℓ_j , where ℓ_j is the integer power of 2 satisfying $|I_j|/8M \leq \ell_j < |I_j|/4M$.

First, suppose $\ell_j \leq \ell_{j-1}$. Remove the last interval in \mathcal{D}_{j-1} and replace it by its dyadic subintervals of length ℓ_j that do not hit I_j . Also add the dyadic intervals in \mathcal{C}_j to \mathcal{D}_{j-1} to get the collection \mathcal{D}_j . Clearly, (1)–(4) all hold. Moreover,

$$\ell_{j-1} \geq \ell_j \geq |I_j|/8M \geq |I_{j-1}|/8MC \geq \ell_{j-1}/2C,$$

where C is the bounded geometry constant of \mathcal{Q} . Thus, (5) also holds.

Next, suppose $\ell_j > \ell_{j-1}$. Then subdivide each dyadic interval in the cover \mathcal{C}_j of I_j into dyadic subintervals of length ℓ_{j-1} and redefine $\ell_j = \ell_{j-1}$. Add these intervals to \mathcal{D}_{j-1} to give \mathcal{D}_j (except possibly the first interval, if it is already in the collection). Then (1)–(5) are all obvious.

Next, do the analogous construction for $j < 0$ and reflect the resulting dyadic cover of $(-\infty, -1]$ across zero to get a dyadic covering of $[1, \infty)$. By taking the shortest interval covering each point, we get a dyadic covering of $[1, \infty)$, which satisfies all the desired conditions. \square

Fix $M \geq 8C$ (recall C is the bounded geometry constant of \mathcal{Q}) and apply Lemma 9.2 to get a symmetric partition $\mathcal{Q}' = \{I_j\}_{j \in \mathbb{Z}^*}$. We will use the partition \mathcal{Q}' to fill the central region with a meandering tube. Let $\{a_j\} \subset [1, \infty)$ be the positive endpoints of \mathcal{Q}' ($I_j = (ia_j, ia_{j+1})$). Let $\eta_j = a_{j+1} - a_j = |I_j|$. Now add the vertical segments

$$V_j = \left\{ x + iy : x = a_j, \quad -a_j + \eta_j \frac{1 + (-1)^j}{2} \leq y \leq a_j - \eta_j \frac{1 + (-1)^{j+1}}{2} \right\}$$

inside the central region; more geometrically, we are adding segments on the vertical lines $\{x = a_n\}$ that lie inside the central region so that one endpoint lies on the boundary of the central region and the other is distance η_j below or above the boundary. Alternate segments alternately touch the ‘top’ and ‘bottom’ sides of the central region. This defines a simply connected subregion of the central region that we call the central tube. The boundary of the tube consists of I'_0 and two connected components that we call the upper and lower components. The central tube is illustrated in the center of Figure 9.




FIGURE 10. We build approximate trapezoids by joining partition points on $\partial\mathbb{H}_r$ by almost horizontal lines to partition vertices on the boundary of the central region. A small segment is then removed from the boundary of the central region so that the interior of the trapezoid is joined to the central tube. In this figure, $k < 0$, so $k^* = k - 1$.

We make the boundary of the central tube into a tree by adding vertices on the vertical segments at their endpoints and at points spaced η_j apart on V_j . On the upper and lower boundaries of the central region, we add a vertex at all the points $\{a_j \pm ia_j\}_{j \geq 1}$. It is easy to see that this makes the boundary of the central tube into a bounded geometry tree (since $\eta_j \simeq \eta_{j+1}$).

As before, let $\{z_j\} = \{ix_j\}$ be the endpoints of the partition $\mathcal{Q} = \{I_j\}$. For each $k \in \mathbb{Z}^*$, choose a $y_k = a_{n(k)}$ so that $|x_k - y_k| < |I_k|/M$. This is possible by condition (5) of Lemma 9.2. Define the segment $H_k = [ix_k, w_k]$, where $w_k = |y_k| + iy_k$. This segment connects ix_k to a vertex on the boundary of the central tube and is close to horizontal (the absolute value of its slope is $\leq 1/M$; by abusing notation, we will refer to these segments as ‘horizontal’. Doing this for every k divides the complement of the central region into quadrilaterals that look like trapezoids, and which we will call trapezoids by another abuse of notation.

We can also make the choice above so $n(k)$ is even if $k < 0$ and is odd if $k > 0$. This means that the right-hand vertex w_k of the segment H_k is a degree 2 vertex of the central tube, and hence forms a degree 3 vertex when the segment H_k is added. This is important because we want the final tree to have maximum degree 3 (otherwise we would end up with the upper bound 16 instead of 12 in part (2) of Theorem 1.2).

For $k \in \mathbb{Z}^*$, the k th trapezoid has left side $I_k \subset \partial\mathbb{H}_r$, two ‘horizontal’ sides H_k and H_{k^*} , and a right-hand side of slope ± 1 along the boundary of the central region (see Figure 10). Vertices are added to the ‘horizontal’ sides H_k that break H_k into segments that start at the left with length comparable to the length of I_k and end on the right with lengths comparable to (but larger than) $\eta_n = a_{n+1} - a_n$, where $n = n(k)$. Thus, the subsegment of H_k that meets the boundary of the central tube has length comparable to the edges of the central tube at the meeting point. Hence, we have a bounded geometry forest in \mathbb{H}_r . However, this forest cuts the plane into infinitely many components (the trapezoids and the central tube). We want to form a single component by removing some segments.

For each $k \in \mathbb{Z}^*$, choose $n = n(k)$ so $w_{k^*} = a_n \pm ia_n$ then remove the open segment $(a_{n-1} \pm ia_{n-1}, w_{k^*})$ from the boundary of the central region; this is the segment on the boundary of the central tube that is also on the boundary of the k th trapezoid and has w_{k^*} as one endpoint (see Figure 10). Removing this segment connects the k th trapezoid to the central tube. When we have removed all such segments, we have a bounded geometry forest T'' with a single complementary component W'' in \mathbb{H}_r (see the right side of Figure 9).

All that remains is to prove that W'' satisfies a positive τ -length lower bound. First, consider sides of W'' that are also sides of the central tube. If such a side lies on the upper boundary, it can clearly be separated from the lower boundary by a path family with uniformly bounded




FIGURE 11. The modulus of the shaded tube in the top panel gives an upper bound for $M(J, K)$. The interval J is on the side of the trapezoid, and K is on the boundary of the central tube. The modulus of the shaded rectangle in the bottom panel gives a lower bound for $M(I_0, K)$.

modulus. Thus, Lemma 8.1 implies that the τ -lengths of such sides grow exponentially and hence are bounded away from zero. A similar argument applies to sides on the lower boundary of the central tube.

Next, we have to consider sides of W'' that are sides of the k th trapezoid. We want to use Lemma 8.2 with $I = I_0$, J a side of k th trapezoid, and K a side of the central tube, chosen as shown in Figure 11. To prove $M(J, K) \leq M(I, K)$, we will first give an upper bound for $M(J, K)$ and then give a lower bound for $M(I, K)$ that is larger than this bound.

Replacing J by a subinterval only increases $M(J, K)$, and every side of the k th trapezoid has length at least η_{n^*} , where $n^* = n(k^*)$ is defined by the relation $y_{k^*} = a_{n^*}$. So, we assume that J is any interval of length η_{n^*} on the side of the k th trapezoid. Any such J is one side of a generalized quadrilateral $Q \subset W''$ whose opposite side is K and so that the two remaining sides are at least distance η_{n^*} apart (see Figure 11).

Moreover, we can choose Q so its area is at most $(|I_k| + y_{k^*})\eta_{n^*}$. Therefore, the modulus of the path family separating J from K in the quadrilateral is at most $(|I_k| + y_{k^*})/\eta_{n^*}$ (just take the constant metric $\rho = 1/\eta_{n^*}$ in the definition of modulus; any admissible metric gives an upper bound for the modulus.) Any path separating J and K in W'' contains a subpath that separates them in Q , so by the extension principle

$$M(J, K) \leq (|I_k| + y_{k^*})/\eta_{n^*}.$$

Now, we give a lower bound for $M(I, K)$. Note that there is $\eta_{n^*} \times (2y_{k^*} - 2\eta_{n^*})$ rectangle that separates K from I_0 ; it is contained in the vertical section of the central tube just above the opening to the k th trapezoid and the lower portion of the rectangle is shown as a shaded region in the bottom panel of Figure 11 (the rectangle extends upwards almost to the upper boundary of the central region). The horizontal segments that cross this rectangle separate K from I_0 , so by the extension principle again, we see that the modulus of these horizontal

segments is a lower bound for $M(I, K)$, hence (since the modulus of a rectangle is the ratio of its sides),

$$M(I, K) \geq \frac{2y_{k^*}}{\eta_{n^*}} - 2. \tag{9.1}$$

Now we have to compare our lower bound for $M(I, K)$ to our upper bound for $M(J, K)$. By the definition of the bounded geometry constant C for \mathcal{Q} , if $|k| = 1$ then $I_k \in \mathcal{Q}$ is distance 1 from the origin and $|I_k| \leq C|I_0| \leq 2C$, so $1 \geq |I_k|/2C$. If $|k| > 1$, then $I_k \in \mathcal{Q}$ is separated from the origin by another partition interval of length at least $|I_k|/C$. So in either case,

$$y_{k^*} \geq x_{k^*} \geq |I_k| + |I_k|/2C,$$

so that $|I_k| \leq (1 - \frac{1}{4C})y_{k^*}$ (here we use that $(1 + \epsilon)^{-1} < 1 - \epsilon/2$ if $0 < \epsilon < 1$). Hence,

$$\begin{aligned} M(J, K) &\leq \frac{|I_k| + y_{k^*}}{\eta_{n^*}} \leq \left(2 - \frac{1}{4C}\right) \frac{y_{k^*}}{\eta_{n^*}} \\ &\leq \frac{2y_{k^*}}{\eta_{n^*}} - \frac{y_{k^*}}{4C\eta_{n^*}} \leq \frac{2y_{k^*}}{\eta_{n^*}} - \frac{|I_k|}{4C\eta_{n^*}} \\ &\leq \frac{2y_{k^*}}{\eta_{n^*}} - \frac{M}{4C}. \end{aligned}$$

Since we assumed $M \geq 8C$, and using (9.1), we have

$$M(J, K) \leq \frac{2y_{k^*}}{\eta_{n^*}} - 2 \leq M(I, K).$$

Thus by Lemma 8.2, the τ -length of J is greater than that of $I = I_0$ and hence is bounded uniformly from below. As before, it is easy to check that the constructed tree has maximum vertex degree 3. This completes the proof of Lemma 9.1 and hence the proof of Theorem 1.3 (the number of tracts in the approximation is at most double the number of tracts in the model).

10. Geometric restrictions on Speiser models

So far, this paper has dealt with methods for building Speiser class functions. The remainder of the paper is devoted to placing limits on what can be accomplished in this direction. In this section, we show that the choice of τ on different tracts of a Speiser class function must satisfy certain constraints; no such restriction needs hold for the Eremenko–Lyubich class by the results of [4]. This is a clear difference between the two classes.

Suppose $f \in \mathcal{S}$ and $S(f) \subset \mathbb{D}$. Since $S(f)$ is finite, there is an $\epsilon > 0$, so that

$$\text{dist}(S(f), \partial\mathbb{D}) > 4\epsilon$$

and

$$\min\{|a - b| : a, b \in S(f), a \neq b\} > 4\epsilon.$$

For $a \in S(f)$, let $D_a = D(a, \epsilon)$ and let $2D_a = D(a, 2\epsilon)$. The open disks D_a are evidently pairwise disjoint and all lie inside \mathbb{D} .

For each $a \in S(f)$, the set $W(a, \epsilon) = f^{-1}(D_a)$ only has simply connected components, and on each such component U the map $f : U \rightarrow D_j$ acts either as

- (1) a 1-to-1 map onto D_j ;
- (2) a finite-to-1 branched cover of D_j with a single critical value at a ; or
- (3) an ∞ -to-1 cover of $D_a \setminus \{a\}$.




FIGURE 12. Each $D(a, \epsilon)$ can be connected to $\{|z| = 1\}$ by a path γ of length $\leq \pi$ that stays at least distance ϵ from every singular point. Such a path can be covered by $O(\epsilon^{-1})$ disks, each of which has a double that misses the singular set. Therefore, the derivative of any branch of f^{-1} on γ has comparable sizes at any two points of γ (with a constant depending only on ϵ).

In the first two cases U is bounded, and in the third case it is unbounded and contains a path to ∞ along which f has asymptotic value a . Moreover, the preimages $f^{-1}(a + \epsilon)$ partition $\partial W(a, \epsilon)$ into arcs. Let $X = \mathbb{D} \setminus \bigcup_{a \in S(f)} D(a, \epsilon)$.

Recall that given $r > 0$ and an arc I , we define a neighborhood of I by

$$I(r) = \{z : \text{dist}(z, I) < r \cdot \text{diam}(I)\}.$$

THEOREM 10.1. *There is an $r < \infty$, depending only on ϵ , so that for each partition arc I of $\partial W(a, \epsilon)$ there is an edge J of $\partial\Omega$ with $I \subset J(r)$ and $J \subset I(r)$. Moreover, $\text{diam}(I) \simeq \text{diam}(J) \simeq \text{dist}(I, J)$ and the lengths of I and J are comparable to their diameters.*

Proof. Fix $a \in S(f)$. Then ∂D_a can be covered by a uniformly bounded number of disks whose doubles do not hit $S(f)$ and so f^{-1} is conformal from each such disk to any of its preimages under f . If I is a partition arc of $\partial W(a, \epsilon)$, this fact and Koebe’s distortion theorem imply that I has bounded geometry and its diameter is comparable to its length with constants that are uniform (the constants only depend on the number of disks covering ∂D_a , which is uniformly bounded).

Similarly, $\mathbb{T} = \partial\mathbb{D}$ is covered by $O(\epsilon^{-1})$ disks of radius ϵ whose doubles do not intersect $S(f)$, so the same argument shows that a partition arc J of $\partial\Omega$ has bounded geometry, but now with constants that depend on ϵ .

Finally, ∂D_a and \mathbb{T} can be joined by a curve γ that is never closer than ϵ to any point of $S(f)$ (use a straight line segment and replace its intersection with any D_a by the shorter arc of ∂D_a connecting the same two points; see Figure 12). This arc can also be covered by $O(\epsilon^{-1})$ distinct ϵ -disks whose doubles miss $S(f)$.

Therefore, all the lifts of γ via f have bounded geometry with constants depending only on ϵ . Thus if we take I as above and a lift of γ with one endpoint on I , then the other endpoint of the lifted curve γ' is on a partition arc J of $\partial\Omega$. By Koebe’s theorem, $|f'|$ has comparable values at all points of $I \cup \gamma' \cup J$ with constant that depends only on ϵ . Since ∂D_a , γ and \mathbb{T} all have comparable lengths (within a factor of $O(\epsilon^{-1})$), so do I , γ' , and J . This proves the lemma (see Figure 13). □




FIGURE 13. As noted earlier, each $D(a, \epsilon)$ can be connected to $\{|z| = 1\}$ by a path γ that stays at least distance ϵ away from the singular set. A lift of this path via f is a path that connects partition arcs of $\partial\Omega$ and $\partial W(a, \epsilon)$. As explained in the text, the lengths and diameters of the lifted curve and the arcs it connects are all comparable with a constant depending only on ϵ .

COROLLARY 10.2. *With notation as above, there is a $r < \infty$, depending only on ϵ , so that $f^{-1}(X) \subset T_\Omega(r)$.*

Proof. Note that $\partial W(a, \epsilon) \subset T(r)$ and this implies $f^{-1}(X) \subset T(r)$, as claimed. □

Theorem 10.1 immediately implies.

COROLLARY 10.3. *Suppose f is in the Speiser class and $S(f) \subset \mathbb{D}$. For every $\epsilon > 0$, there is a $r < \infty$ so that each connected component of $\mathbb{C} \setminus (\Omega \cup T_\Gamma(r))$ (where $\Gamma = \partial\Omega$) maps under f into some disk D_a for some $a \in S(f)$. If the connected component is unbounded, then a must be an asymptotic value of f .*

If the critical points of f have uniformly bounded degree D , then the components of $W(a, \epsilon)$ containing critical points have boundaries with at most D partition arcs, each with diameter comparable to the whole component (the constant depending only on D). Since one of these arcs is contained in some $J(r)$ the whole component will be contained in $J(Cr)$ for some C depending only of D . Thus,

COROLLARY 10.4. *If $f \in \mathcal{S}$, $S(f) \subset \mathbb{D}$, and f have no finite asymptotic values and every critical point has uniformly bounded degree D , then there is a $r > 0$ so that $\mathbb{C} = \Omega \cup T(r)$ (as before, $\Omega = \{|f| > 1\}$, and r depends only on D and ϵ , where ϵ is the minimal separation between different points of $S(f)$ and between $S(f)$ and \mathbb{T}).*

For the half-strip $S = \{x + iy : x > 0, |y| < 1\}$, the partition elements decay exponentially fast and it is obvious that $\Omega \cup T_\Omega(r)$ is not the whole plane for any finite r . Thus, this model cannot be approximated in the sense of Theorem 1.2 without using extra tracts. In the remainder of the paper, we will show that something even stronger is true: the approximation of a half-strip by a Speiser model domain with a single tract is not possible even if we allow

- (1) finite asymptotic values;
- (2) high degree critical points; and
- (3) φ to be nonconformal everywhere in the plane.

Very briefly, the problem with the half-strip is that the τ -lengths for it and its complement behave so differently, that the comparability implied by Theorem 10.1 cannot hold (see Figure 14).




FIGURE 14. Suppose T is the boundary of a half-strip with unit spacing of the vertices, as shown above. The conformal map of the half-strip (left) to a half-plane is $\sinh(z)$, and it expands exponentially, so the unit segments shown in each contain exponentially many conformal partition elements for the half-strip. However, the conformal map of the exterior domain (right) to \mathbb{H}_r behaves like $z^{1/2}$ near ∞ , so the unit segments are much smaller than conformal partitions elements should be. This ‘imbalance’ of τ -sizes is what prevents the half-strip (or any quasiconformal image of it) from being a Speiser model domain. The following sections will make this precise.




FIGURE 15. The family of all separating curves for two unit intervals distance n apart on $\partial\mathbb{H}_r$ has modulus $\simeq \log n$. Thus, in Ω any separating family must have modulus $\gtrsim \log n$. Applying this observation to the two path families defined in the text gives the desired size and separation estimates for the image intervals.

11. A polynomial lower bound for thick tracts

Suppose Ω is a simply connected planar domain bounded by a Jordan curve on the sphere that passes through ∞ . Suppose $\tau: \Omega \rightarrow \mathbb{H}_r$ is conformal, maps ∞ to ∞ , and \mathcal{J} is the partition of $\partial\Omega$ that corresponds via τ to the partition of $\partial\mathbb{H}_r$ with endpoints $i\pi\mathbb{Z}$ (recall that this is called a conformal partition of $\partial\Omega$). In this section, we want to prove that if a tract Ω is ‘large’ in a certain sense, then the size of elements in \mathcal{J} cannot tend to zero too quickly. By ‘large’ we will mean that Ω contains an unbounded quasidisk. For example, a sector $W_\theta = \{z : |\arg(z)| < \theta\}$, $0 < \theta < \pi$ is an example of an unbounded quasidisk, so if Ω contains a sector, its partition elements cannot have diameters that decrease exponentially quickly.

LEMMA 11.1. *Suppose Ω is bounded by a Jordan curve through ∞ and $\{\mathcal{J}\}$ is a conformal partition of $\partial\Omega$. Suppose Ω contains an unbounded quasidisk W . Then there is a $R_0 < \infty$ so that any partition arc J that hits a circle $\{|z| = R\}$ with $R > R_0$ satisfies $\text{diam}(J) \geq CR^{-\sigma}$ for some $C > 0$, $\sigma < \infty$, independent of J .*

Proof. Since Ω contains an unbounded quasidisk, Lemma 2.3 says that there is a curve $\gamma \subset \Omega$ that connects some point $z_0 \in \Omega$ to ∞ and has the property that there is a $C_1 < \infty$ so that $\text{dist}(z, \partial\Omega) \geq C_1|z|$, for all $z \in \Gamma$ (see Figure 15).

Choose R_0 large enough so that $|z_0| < R_0$ and so that $\{|z| < R_0\}$ contains some partition element $I' = \tau^{-1}(I) \in \partial\Omega \cap \mathcal{J}$. Fix $R \gg 4R_0$ and choose a partition element J so that

$J' = \tau^{-1}(J) \subset \partial\Omega$ hits $\{|z| = R\}$. If J' has diameter greater than 1, there is nothing to do, so we may assume $\epsilon = \text{diam}(J') \leq 1$ and hence $J' \subset \{z : R - 1 \leq |z| \leq R + 1\}$ (see Figure 15).

Suppose I and J (which are equal length intervals on $\partial\mathbb{H}_r$) are separated by n other partition elements. By Lemma 2.2, $M(I, J) \simeq \log n$ (recall that $M(I, J)$ is the modulus of the path family that separates I and J in \mathbb{H}_r).

We consider two other path families in Ω . First, let Γ_1 be the family of circular arcs in $\Omega \cap \{z : R_0 < |z| < R/2\}$ concentric with 0 that connect the two components of $\partial\Omega \setminus I'$ (see Figure 15). Let $z_{J'}$ be any point of J' and let Γ_2 be the path family consisting of circular arcs in $\{z : \text{diam}(J') < |z - z_{J'}| < 1\}$ that are concentric with $z_{J'}$ and connect different components of $\partial\Omega \setminus J'$ (see Figure 15).

Each path in Γ_1 and Γ_2 separates I' from J' in Ω , so by conformal invariance and the parallel rule for modulus, we have

$$\log n \simeq M(\Gamma_0) \geq M(\Gamma_1) + M(\Gamma_2),$$

and thus both terms on the right are bounded by $C_2 \log n$ for some $C_2 < \infty$.

Because γ crosses each element of Γ_1 and each crossing point z is at least distance $C_1|z|$ from $\partial\Omega$, we deduce that there is a quadrilateral region

$$Q = \{w : |\arg(w) - \arg(z)| \leq A, 1 \leq |w/z| \leq B\}$$

contained in Ω for some $A > 0, B > 1$ depending only on the constant C_1 . The path family in Q connecting the two radial sides of Q has fixed modulus M_Q and this modulus is a lower bound for the modulus of the path family in $\Omega \cap \{|z| < w < B|z|\}$ connecting different components of $\partial\Omega \setminus I'$. Thus, by the parallel rule, the modulus of Γ_1 is bounded below by $M_Q \cdot \lfloor \log_B(R/R_0) \rfloor$. In other words, $\log n \gtrsim M(\Gamma_1) \gtrsim \log R$, for R large. Hence, when R is large, we have $R \leq n^\alpha$ for some α that only depends on the constant C_1 .

On the other hand, the usual estimate of the modulus of an annulus says that

$$\log \frac{1}{\text{diam}(J')} \lesssim \log n,$$

so $\text{diam}(J') \gtrsim n^{-\beta}$, for some $\beta > 0$. Thus

$$\text{diam}(J') \gtrsim (R^{1/\alpha})^{-\beta} \gtrsim R^{-\beta/\alpha} = R^{-\sigma},$$

as desired. □

12. An exponential upper bound for thin tracts

We now want to do the opposite of the previous section: show that if Ω is ‘thin’, then the diameters of partition elements decay faster than any polynomial.

LEMMA 12.1. *Suppose Ω is the image of the half-strip $S = \{x + iy : x > 0, |y| < \frac{1}{2}\}$ under a K -quasiconformal map ϕ of the plane fixing 0 and 1 and \mathcal{J} is a conformal partition of $\partial\Omega$. Then all the partition elements satisfy*

$$\text{diam}(J) \leq C_1 \exp(-C_2 \text{dist}(J, 0)^\alpha),$$

for some finite, positive constants C_1, C_2, α depending on K .

Proof. Using the measurable Riemann mapping theorem, we can write ϕ as a composition of two K -quasiconformal maps $\phi = g_2 \circ g_1$, where both maps also fix both 0 and 1, g_1 is conformal outside S , and g_2 is conformal on $W = g_1(S)$.

Consider the square S_n inside S between the vertical lines $\{x = n\}$ and $\{x = n + 1\}$. Then $W_n = g_1(S_n)$ is a quasidisk and its image in \mathbb{H}_r under the conformal map $\tau: \mathbb{H}_r \rightarrow W$ is

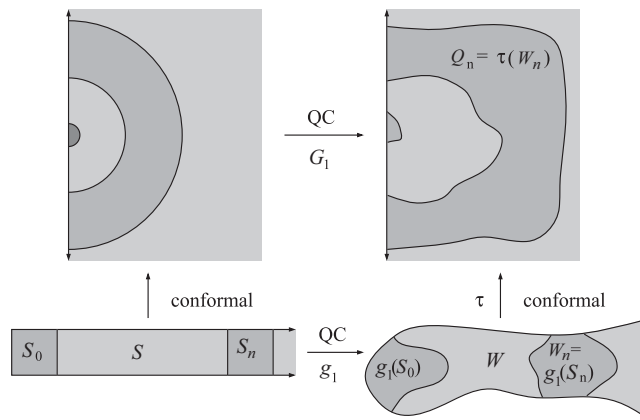


FIGURE 16. The modulus separating S_0 and S_n in the half-strip is comparable to n , so by quasi-invariance the same is true for Q_0 and Q_n in \mathbb{H}_r . This implies that the Euclidean diameter of Q_n grows exponentially, hence $\partial Q_n \cap \partial \mathbb{H}_r$ contains exponentially many partition points. The map from Q_n to W_n is conformal, so the same is true of W_n and partition arcs for $W = g_1(S)$.

generalized quadrilateral Q_n with two sides on $\partial \mathbb{H}_r$ and modulus bounded above and below, depending only on K (see Figure 16).

Since the extremal length between S_0 and S_n in S is $\simeq n$, the same is true for the extremal distance between Q_0 and Q_n in \mathbb{H}_r in \mathbb{H}_r (with a constant depending on the maximal dilatation g_1). This implies that the diameters of Q_n must grow exponentially, as do the component intervals of $\partial Q_n \cap \partial \mathbb{H}_r$ (this is the same argument as in the proof of Lemma 8.1). Thus, ∂Q_n hits $\geq ce^{cn}$ partition intervals on $\partial \mathbb{H}_r$ for some fixed $c > 0$ (depending only on K). Hence, W_n hits the same number of partition arcs on ∂W . Because W_n is the image of U_n under a bi-Hölder map (since it is a quasidisk), each of these partition arcs has diameter bounded by $C \text{diam}(W_n) \exp(-an)$ for another constant a depending only on K .

Since g_1 has dilatation supported in the half-strip S and

$$\int_S \frac{dx dy}{1 + x^2 + y^2} < \infty,$$

Theorem 2.5 implies that $|g_1(z)/z|$ has a limit as $z \rightarrow \infty$. Thus, if $R_n = \text{dist}(W_n, 0)$ we have $R_n \simeq n$ and $\text{diam}(W_n) \lesssim n$. Thus, all the partition elements hitting W_n have diameters less than cne^{-an} , where a, c are positive constants that depend only on K .

Since g_2 is conformal on W , the partition for Ω is just the image of the partition for W under g_2 , and since g_2 is bi-Hölder (with exponent depending only on K), the estimate in the lemma follows. \square

The proof can be applied to other tracts that look like thin tubes, for example,

$$\Omega = \{x + iy : x > 0, |y| \leq \eta(x)\},$$

if $\eta(x), \eta'(x) \rightarrow 0$ as $x \rightarrow \infty$. However, the proof does not work for all subdomains of the half-strip, since adding ‘rooms’ to the sides of a half-strip can create partition arcs whose diameters do not tend to zero (in fact, we used a similar construction in the proof of Lemma 9.1.)

13. The half-strip is not the quasiconformal-image of a Speiser model domain

Before starting the proof of Theorem 1.4, we record the following result that is immediate from Lemma 2.4.




FIGURE 17. Although the half-strip cannot be a Speiser class model domain (or even quasiconformally mapped to a Speiser class model domain), it can be approximated by a tract of such a model. These figures show some ways this can occur using two, three or infinitely many tracts.

COROLLARY 13.1. Suppose Ω is a model domain. If ϕ is a K -quasiconformal map of the plane that is conformal on Ω , then

$$T_{\phi(\Omega)}(t) \subset \phi(T_{\Omega}(r)) \subset T_{\phi(\Omega)}(s),$$

where t, s depend only on r and K .

We can now prove Theorem 1.4: the half-strip $S = \{x + iy : x > 0, |y| < 1\}$ cannot be mapped to any Speiser class model domain by any quasiconformal homeomorphism of the plane.

Proof. Suppose there were a K -quasiconformal map ϕ of the plane taking S to the tract $\Omega = \{z : |f(z)| > R\}$ of some $f \in \mathcal{S}$. Choose ϵ as in Theorem 10.1 and let r be as given by that theorem. Let s be as given by Corollary 13.1.

As in the proof of Lemma 12.1, write $\phi = g_2 \circ g_1$, where g_1 is conformal off S , g_1 is conformal on $W = g_1(S)$. Using Theorem 2.5 again implies that we can choose g_1 so that

$$W \cup T_W(s) \subset V = \{z : |z| < R\} \cup \{z : |\arg(z)| < \pi/4\},$$

if R is large enough (depending on s). Note that V^c is a quasidisk and hence $V' = g_2(V^c)$ is a quasidisk as well. By Lemma 13.1, this domain is contained in the complement of $\Omega \cup T_{\Omega}(t)$. Therefore, V' is contained inside some component U of $W(a, \epsilon)$ for $a \in S(f)$.

Lemma 11.1 applies to U and Lemma 12.1 applies to Ω , giving estimates that contradict the conclusion of Theorem 10.1 (partition elements for ∂U are contained in r -neighborhoods of partition elements for $\partial\Omega$). This proves that Ω could not have been the tract of any $f \in \mathcal{S}$. \square

Although the half-strip cannot be approximated by Speiser class model domains with a single tract, Figure 17 shows how it can be approximated by functions in $\mathcal{S}_{2,0}$ with two, three, or infinitely many tracts. To apply Theorem 3.1, we must add vertices so that the bounded geometry and τ -length conditions are satisfied. Bounded geometry is trivial in these figures using vertices from the obvious lattice and Lemma 8.1 easily gives a strictly positive τ -length lower bound for each tract (as usual, we can then multiply τ on each component by a positive constant to attain a τ -length lower bound of π).

Lemma 8.1 also implies that the conformal partition of each of the drawn tracts has elements whose sizes decay exponentially quickly. This implies that $T(r)$ (the set where the quasiconformal correction map φ in Theorem 3.1 has its dilatation supported) has finite Lebesgue area. If we replace τ by a large positive integer multiple of itself, say $N \cdot \tau$, each

edge of the conformal partition is divided into N segments. This implies that the area of $T(r)$ decreases to zero as $N \nearrow \infty$. However, the maximal dilatation of φ remains bounded, independent of N . A standard argument then shows that the corresponding map φ tends uniformly on compact sets (or uniformly on the Riemann sphere) to the identity map as N increases. A more careful argument shows that we can normalize the correction maps so that they tend to the identity uniformly with respect to the Euclidean metric on the whole plane (for example, see [5, Theorem 1.1]; our examples satisfy the (ϵ, φ) -thin hypothesis of that result). Thus, the tracts of the resulting Speiser functions can approximate the tracts in Figure 17 as closely as we wish in the Hausdorff metric on the plane.

In general, it seems that the shapes of individual tracts of Speiser model domains and Eremenko–Lyubich model domains do not differ significantly. However, Speiser models only allow disjoint tracts to be combined in certain ways depending on the choice of τ_j in each tract, whereas Eremenko–Lyubich models allow disjoint tracts to be combined arbitrarily, and τ_j can be chosen on each tract independently of the choices in other tracts.

Roughly speaking, each tract of a Speiser class function can ‘see’ other nearby tracts in the sense that it can be connected to such tracts by path families that avoid the singular set and come with nice geometric estimates (see the proof of Theorem 10.1). However, the singularities of an Eremenko–Lyubich function can effectively ‘block the view’ between different tracts. It would be reasonable to expect that if the singular set of an Eremenko–Lyubich functions was infinite but ‘sparse’ in an appropriate sense, then nearby tracts would be forced to satisfy compatibility conditions similar to Speiser class functions. What can be said about entire functions where the singular set is restricted to lie in a given closed set E , or satisfies some bound on its area, dimension or capacity, or satisfies some other natural geometric restriction?

Acknowledgement. Many thanks to Simon Albrecht, Adam Epstein, Alex Eremenko, and Lasse Rempe-Gillen for numerous helpful discussions about the content of this paper and about the quasiconformal folding construction and its applications. The introduction of the paper and the formulation of the main result in terms of models was inspired by a lecture of Lasse Rempe-Gillen at an ICMS conference on transcendental dynamics in Edinburgh, May 2013. The results of both this paper and [4] originally appeared in a single 2013 preprint titled ‘The geometry of bounded type entire functions’. Based partly on a referee’s report, I decided to split that manuscript in order to improve the exposition and separate the self-contained arguments for the Eremenko–Lyubich class (now contained in [4]) from the proofs for the Speiser class that depend crucially on the quasiconformal folding techniques in [2]. Theorem 1.3 is new and did not appear in the earlier manuscript. Malik Younsi read the revised manuscript and I greatly appreciate his comments and suggestions. The referee of the current paper produced two detailed and thoughtful reports that contained numerous comments and suggestions that improved the exposition, and I am most thankful for the great deal of time and effort that went into these reports. Finally, I am indebted to Aimo Hinkkanen for a great deal of encouragement and constructive advice that substantially improved both this paper and its companion [4].

References

1. L. V. AHLFORS, *Lectures on quasiconformal mappings*, 2nd edn, University Lecture Series 38 (American Mathematical Society, Providence, RI, 2006) With supplemental chapters by C. J. Earle, I. Kra, M. Shishikura and J. H. Hubbard.
2. C. J. BISHOP, ‘Constructing entire functions by quasiconformal folding’, *Acta Math.* 214 (2015) 1–60.
3. C. J. BISHOP, ‘The order conjecture fails in \mathcal{S} ’, *J. Anal. Math.* 127 (2015) 283–302.
4. C. J. BISHOP, ‘Models for the Eremenko–Lyubich class’, *J. Lond. Math. Soc.* 92 (2015) 202–221.
5. C. J. BISHOP, ‘Quasiconformal maps with dilatations of small support’, Preprint, 2016.
6. D. DRASIN, ‘On the Teichmüller–Wittich–Belinskii theorem’, *Results Math.* 10 (1986) 54–65.
7. D. DRASIN, A. A. GOLDBERG and P. POGGI-CORRADINI, ‘Quasiconformal mappings in value-distribution theory’, *Handbook of complex analysis: geometric function theory*, vol. 2 (Elsevier Science, Amsterdam, 2005) 755–808.

8. T. A. DRISCOLL and L. N. TREFETHEN, *Schwarz-Christoffel mapping*, Cambridge Monographs on Applied and Computational Mathematics 8 (Cambridge University Press, Cambridge, 2002).
9. A. E. EREMEENKO and M. YU. LYUBICH, 'Dynamical properties of some classes of entire functions', *Ann. Inst. Fourier (Grenoble)* 42 (1992) 989–1020.
10. J. B. GARNETT and D. E. MARSHALL, *Harmonic measure*, New Mathematical Monographs 2 (Cambridge University Press, Cambridge, 2005).
11. L. R. GOLDBERG and L. KEEN, 'A finiteness theorem for a dynamical class of entire functions', *Ergodic Theory Dynam. Systems* 6 (1986) 183–192.
12. A. A. GOLDBERG and I. V. OSTROVSKII, *Value distribution of meromorphic functions*, Translations of Mathematical Monographs 236 (American Mathematical Society, Providence, RI, 2008) Translated from the 1970 Russian original by Mikhail Ostrovskii, With an appendix by A. Eremenko and J. K. Langley.
13. J. HEINONEN, *Lectures on analysis on metric spaces*, Universitext (Springer, New York, 2001).
14. O. LEHTO, 'On the differentiability of quasiconformal mappings with prescribed complex dilatation', *Ann. Acad. Sci. Fenn. Ser. A I* 275 (1960) 1–28.
15. L. REMPE, 'Rigidity of escaping dynamics for transcendental entire functions', *Acta Math.* 203 (2009) 235–267.
16. L. REMPE-GILLEN, 'Arc-like continua, Julia sets of entire functions and Eremenko's conjecture', Preprint, 2014, arXiv:1610.06278 [math.DS].
17. G. ROTTENFUSSER, J. RÜCKERT, L. REMPE and D. SCHLEICHER, 'Dynamic rays of bounded-type entire functions', *Ann. of Math. (2)* 173 (2011) 77–125.
18. O. TEICHMÜLLER, 'Eine Umkehrung des Zweiten Hauptsatzes der Wertverteilungslehre', *Deutsche Math.* 2 (1937) 96–107.

Christopher J. Bishop
Mathematics Department
Stony Brook University
Stony Brook, NY 11794-3651
USA

bishop@math.sunysb.edu