

ROOTS, SCHOTTKY SEMIGROUPS, AND A PROOF OF BANDT'S CONJECTURE

DANNY CALEGARI, SARAH KOCH, AND ALDEN WALKER

ABSTRACT. In 1985, Barnsley and Harrington defined a “Mandelbrot Set” \mathcal{M} for pairs of similarities — this is the set of complex numbers z with $0 < |z| < 1$ for which the limit set of the semigroup generated by the similarities

$$x \mapsto zx \text{ and } x \mapsto z(x - 1) + 1$$

is connected. Equivalently, \mathcal{M} is the closure of the set of roots of polynomials with coefficients in $\{-1, 0, 1\}$. Barnsley and Harrington already noted the (numerically apparent) existence of infinitely many small “holes” in \mathcal{M} , and conjectured that these holes were genuine. These holes are very interesting, since they are “exotic” components of the space of (2 generator) Schottky semigroups. The existence of at least one hole was rigorously confirmed by Bandt in 2002, and he conjectured that the interior points are dense away from the real axis. We introduce the technique of *traps* to construct and certify interior points of \mathcal{M} , and use them to prove Bandt's Conjecture. Furthermore, our techniques let us certify the existence of infinitely many holes in \mathcal{M} .

CONTENTS

1. Introduction	1
2. Semigroups of similarities	4
3. Elementary estimates	11
4. Roots, polynomials, and power series with regular coefficients	13
5. Topology and geometry of the limit set	17
6. Limit sets of differences	23
7. Interior points in \mathcal{M}	25
8. Holes in \mathcal{M}	32
9. Infinitely many holes in \mathcal{M} and renormalization	39
10. Whiskers	53
11. Holes in \mathcal{M}_0	60
References	64

1. INTRODUCTION

Consider the similarity transformations $f, g : \mathbb{C} \rightarrow \mathbb{C}$ given by

$$f : x \mapsto zx \quad \text{and} \quad g : x \mapsto z(x - 1) + 1,$$

where $z \in \mathbb{D}^* := \{z \in \mathbb{C} \mid 0 < |z| < 1\}$. Because these maps are contractions, there is a nonempty compact attractor $\Lambda_z \subseteq \mathbb{C}$ associated with the *iterated function*

Date: October 30, 2014.

system (or IFS) given by the pair $\{f, g\}$. The attractor Λ_z coincides with the set of accumulation points of the forward orbit of any $x \in \mathbb{C}$ under the semigroup $G_z := \langle f, g \rangle$.

In this article, we study the topology of certain subsets of the parameter space \mathbb{D}^* for G_z . The first set we consider is the *connectedness locus*, denoted by \mathcal{M} ; that is, the set of parameters z for which Λ_z is connected. Standard IFS arguments prove that the limit set Λ_z is either connected, or it is a Cantor set (for details, see Lemma 5.2.1).

The second subset of the parameter space we examine is related to the geometry of Λ_z . For all values of the parameter $z \in \mathbb{D}^*$, the map f fixes 0, and the map g fixes 1. As both of these maps are contracting by the same factor (in fact, by a factor of $|z|$) around their respective fixed points, the limit set Λ_z has a center of symmetry about the point $1/2$ in the dynamical plane. The set \mathcal{M}_0 is defined to be the set of parameters z for which Λ_z contains the point $1/2$.

The sets \mathcal{M} and \mathcal{M}_0 have been studied by various mathematicians over the past 30 years: Barnsley-Harrington [2], Bousch [3, 4], Bandt [1], Solomyak [11, 12], Shmerkin-Solomyak [10], and Solomyak-Xu [13], to name a few.

There is a profound and unexpected connection between the sets \mathcal{M} and \mathcal{M}_0 and the set of roots of power series with prescribed coefficients (see Section 4). In particular, \mathcal{M} can be identified with the closure of the set of roots of polynomials with coefficients in $\{-1, 0, 1\}$ (which are in \mathbb{D}^*), and \mathcal{M}_0 can be identified with the closure of the set of roots of polynomials with coefficients in $\{-1, 1\}$ (which are in \mathbb{D}^*). Via this formulation, the set \mathcal{M}_0 is related to roots of the minimal polynomials associated to the *core entropy* of real quadratic polynomials as defined by Thurston [14], and established by Tiozzo [15]. We further elaborate on the history of \mathcal{M} and \mathcal{M}_0 in Section 2.6.

In [3] and [4], Bousch proved that the sets \mathcal{M} and \mathcal{M}_0 are connected and locally connected. However, the complement of \mathcal{M} and the complement of \mathcal{M}_0 are *disconnected*. The complement of \mathcal{M} and the complement of \mathcal{M}_0 both contain a prominent central component (see Figure 2 and Figure 3). In 1985, Barnsley and Harrington numerically observed other connected components of the complement, or “holes” in \mathcal{M} , and they conjectured that these holes are genuine. In 2002, Bandt rigorously established the existence of one hole in \mathcal{M} . In Theorem 9.1.1, we prove that there are *infinitely many* holes in \mathcal{M} .

These “exotic holes” in \mathcal{M} are quite interesting and somewhat mysterious; they appear to be very well-organized in parameter space, suggesting that there may be a combinatorial classification of them. We currently have found no such classification.

1.1. Statement of results. We prove that all of the connected components of $\mathbb{D}^* \setminus \mathcal{M}$ are *Schottky*, in the sense that if z in $\mathbb{D}^* \setminus \mathcal{M}$, there is a topological disk D containing Λ_z , so that $f(D) \cap g(D) = \emptyset$, and $f(D)$ and $g(D)$ are contained in the interior of D .

Theorem 5.2.3 (Disconnected is Schottky). The semigroup G_z has disconnected Λ_z if and only if G_z is Schottky.

To prove that these exotic components in the complement of \mathcal{M} exist, we introduce the method of *traps* (see Section 7.1), which allows us to numerically certify that a parameter $z \in \mathbb{D}^*$. This technique is different from Bandt’s proof of the existence of these exotic holes. In fact, the existence of a trap is an open condition,

so if there is a trap for the parameter $z \in \mathbb{D}^*$, then necessarily $z \in \text{int}(\mathcal{M})$. Traps therefore allow us to access the interior points of \mathcal{M} . In [1], Bandt conjectured that the interior of \mathcal{M} is dense away from $\mathcal{M} \cap \mathbb{R}$ (see Conjecture 2.6.3). In Theorem 7.2.7, we prove Bandt's conjecture using traps.

Theorem 7.2.7 (Interior is almost dense). The interior of \mathcal{M} is dense away from the real axis; that is,

$$\mathcal{M} = \overline{\text{int}(\mathcal{M})} \cup (\mathcal{M} \cap \mathbb{R}).$$

Interestingly, the proof of Theorem 7.2.7 requires a complete characterization of the set of $z \in \mathcal{M}$ for which the limit set Λ_z is convex. This is established in Lemma 7.2.3.

In Section 9, we examine families of exotic holes in \mathcal{M} which appear to spiral down and limit on a distinguished point $z \in \partial\mathcal{M}$ (see Figure 20).

Theorem 9.1.1 (Limit of holes). Let $\omega \sim 0.371859 + 0.519411i$ be the root of the polynomial $1 - 2z + 2z^2 - 2z^5 + 2z^8$ with the given approximate value. Then

- (1) ω is in \mathcal{M} , and \mathcal{M}_0 ; in fact, the intersection of $f\Lambda_\omega$ and $g\Lambda_\omega$ is exactly the point $1/2$;
- (2) there are points in the complement of \mathcal{M} arbitrarily close to ω ; and
- (3) there are infinitely many rings of concentric loops in the interior of \mathcal{M} which nest down to the point ω .

Thus, \mathcal{M} contains infinitely many holes which accumulate at the point ω .

We continue Section 9 by generalizing the methods of Theorem 9.1.1. We define the notion of *renormalization* and *limiting traps* to show that at certain renormalization points $z \in \mathcal{M}$, the set \mathcal{M} is asymptotically similar to Γ_z , where Γ_z is the limit set of the 3 generator IFS

$$x \mapsto z(x+1) - 1 \quad x \mapsto zx \quad x \mapsto z(x-1) + 1.$$

Previous results of Solomyak established this asymptotic similarity at certain ‘landmark points’ in $\partial\mathcal{M}$. We reprove his results with a more algorithmic approach using traps, and as a consequence, we obtain “asymptotic interior.”

Theorem 9.2.2 (Renormalizable traps). Suppose that ω is a renormalization point. There are constants A and B , depending only on ω , such that

- (1) If $C \in (A + B\Gamma_\omega)$, then for all $\epsilon > 0$, there is a C' such that $|C - C'| < \epsilon$ and for all sufficiently large n , there is a trap for $\omega + C'\omega^{bn}$.
- (2) If $f\Lambda_z \cap g\Lambda_z$ is a single point, then there is $\delta > 0$ such that for all $C \notin (A + B\Gamma_\omega)$ with $|C| < \delta$, the limit set for the parameter $\omega + C\omega^{bn}$ is disconnected for all sufficiently large n .

In Section 11, we prove that the complement of \mathcal{M}_0 is also disconnected by numerically certifying a loop in \mathcal{M}_0 which bounds a component of the complement.

Theorem 11.3 (Hole in \mathcal{M}_0). There is a hole in \mathcal{M}_0 .

1.2. Outline. In Section 2, we establish key definitions and survey some previous results about \mathcal{M} and \mathcal{M}_0 . In Section 3, we collect a few elementary estimates about the geometry of Λ_z . In Section 4, we explore the connection the sets \mathcal{M} and \mathcal{M}_0 have with roots of power series with prescribed coefficients in a more general context involving regular languages. In Section 5, we establish some important

results about the topology and geometry of the limit set, and we prove Theorem 5.2.3. We also present an algorithm (similar to an algorithm used by Bandt in [1]) to certify that the limit set Λ_z is disconnected.

In Section 6, we examine the set of differences between points in Λ_z . This set of differences is actually the limit set Γ_z of the 3 generator IFS

$$x \mapsto z(x+1) - 1 \quad x \mapsto zx \quad x \mapsto z(x-1) + 1.$$

In Section 7, we introduce the notion of traps, and characterize the set of $z \in \mathcal{M}$ for which Λ_z is convex in Lemma 7.2.3. In Theorem 7.2.7, we prove that the interior of \mathcal{M} is dense away from the real axis, establishing Bandt's Conjecture 2.6.3.

In Section 8, we describe our trap-finding algorithm and prove the estimates required to certify that \mathcal{M} has holes. In Section 9, we introduce the notions of renormalization and limiting traps, and we prove Theorem 9.1.1 and Theorem 9.2.2. In Section 10, we discuss the “real whiskers” of \mathcal{M} , and we use a 2-dimensional real IFS for this analysis. And lastly, in Section 11, we prove that there is a hole in \mathcal{M}_0 ; that is, we prove that the complement of \mathcal{M}_0 is disconnected.

1.3. Acknowledgements. We would like to thank Christoph Bandt, Emmanuel Breuillard, Giulio Tiozzo and especially Boris Solomyak for comments, corrections, pointers to references, and enthusiasm and interest in this project. Danny Calegari was supported by NSF grant DMS 1405466. Sarah Koch was supported by NSF grant DMS 1300315 and a Sloan research fellowship. Alden Walker was supported by NSF grant DMS 1203888.

2. SEMIGROUPS OF SIMILARITIES

2.1. Definitions.

Definition 2.1.1. A *contracting similarity* (or just a *similarity*) with center $c \in \mathbb{C}$ and dilation $z \in \mathbb{C}$ with $0 < |z| < 1$ is the complex affine map $\mathbb{C} \rightarrow \mathbb{C}$ given by

$$x \mapsto z(x - c) + c.$$

The composition of any positive number of similarities is again a similarity. The set of all similarities is topologized as $\mathbb{C} \times \mathbb{D}^*$. We are concerned in the sequel with semigroups generated by finitely many similarities.

Definition 2.1.2. Let G be a finitely generated semigroup of contracting similarities. The *limit set* Λ (also called the *attractor*) is the closure of the set of fixed points of elements of G .

The limit set of G is the unique compact, nonempty invariant subset of \mathbb{C} for the action of G . In particular the action of G on Λ is minimal (every orbit is dense).

Example 2.1.3 (Middle third Cantor set). The semigroup $f : x \mapsto \frac{1}{3}x$, $g : x \mapsto \frac{1}{3}(x-1)+1$ has the middle third Cantor set as limit set.

Example 2.1.4 (Sierpinski carpet). The semigroup $f : x \mapsto \frac{1}{2}x$, $g : x \mapsto \frac{1}{2}(x-1)+1$, $h : x \mapsto \frac{1}{2}(x-\omega)+\omega$ for $\omega = e^{i\pi/3}$ has the Sierpinski triangle as limit set.

Definition 2.1.5 (Schottky semigroup). Let S be a finite set of contracting similarities, and let G be the semigroup they generate. We say that G is a *Schottky semigroup* if there is an embedded loop $\gamma \subseteq \mathbb{C}$ bounding a closed (topological) disk D , so that the elements of S take D to disjoint disks contained in the interior of D .

A loop γ with this property, and the disk D it bounds is said to be *good* for G .

Lemma 2.1.6. *The Schottky semigroup G is free (on S) and discrete as a subset of $\mathbb{C} \times \mathbb{D}^*$.*

Proof. Actually, every finitely generated semigroup which is strictly contracting is discrete, since the set of dilations accumulates only at 0; so the point is to prove freeness. This follows from Klein's ping-pong argument applied to a good disk D and its translates. \square

Note that if S generates a Schottky semigroup, the centers of generators are distinct. Indeed, a good disk D must contain all of the centers, and since the generators map D to disjoint disks, the centers must be distinct. Thus for a Schottky semigroup G , the limit set is a Cantor set, which is the intersection of the images of a good disk D under elements of G , and which can be identified (topologically) with the set of right-infinite words in the generators. Thus, any two Schottky semigroups with the same number of generators have topologically conjugate actions on their limit sets. In fact, we can say more:

Lemma 2.1.7. *Any two isomorphic Schottky semigroups G, G' are topologically conjugate on their restriction to good disks D, D' .*

Proof. If S and S' are the generators of G and G' , then choose any homeomorphism $h : D - S(D) \rightarrow D' - S'(D')$ which extends a conjugacy on their boundaries, and extend to $h : D - \Lambda \rightarrow D' - \Lambda'$ using $D - \Lambda = G(D - S(D))$ and $D' - \Lambda' = G'(D' - S'(D'))$. Then extend to $h : D \rightarrow D'$ by the canonical (abstract) isomorphism $h : \Lambda \rightarrow \Lambda'$ coming from the identification of these limit sets with the right-infinite words in the generators. \square

Remark 2.1.8. Note that Schottky semigroups G, G' are very rarely topologically conjugate on all of \mathbb{C} ; for, they are invertible on \mathbb{C} , and therefore any conjugacy would extend to a conjugacy between the *groups* they generate. But these are indiscrete, and indiscrete subgroups of $\mathrm{PSL}(2, \mathbb{C})$ are rarely topologically conjugate.

2.2. Pairs of similarities. For the remainder of the paper we focus almost entirely on semigroups generated by a pair of similarities with the same dilation z . After conjugation by a similarity of \mathbb{C} we may assume that the two centers of the generators are at 0 and 1 respectively. Thus the space of conjugacy classes of such semigroups is parameterized by $z \in \mathbb{D}^*$.

Notation 2.2.1. For $z \in \mathbb{D}^*$, let G_z denote the semigroup with generators

$$f : x \mapsto zx, \quad g : x \mapsto z(x - 1) + 1,$$

and let Λ_z denote the limit set of G_z . We omit the subscript z from f and g to lighten notation.

Other normalizations have some nice features. Barnsley and Harrington [2], Bousch [3] and others use the normalization

$$f : x \mapsto zx + 1, \quad g : x \mapsto zx - 1,$$

and Solomyak [12] uses

$$f : x \mapsto zx, \quad g : x \mapsto zx + 1.$$

Our normalization has the convenient property that 0 and 1 are always in Λ as the centers of the two generators, independent of z .

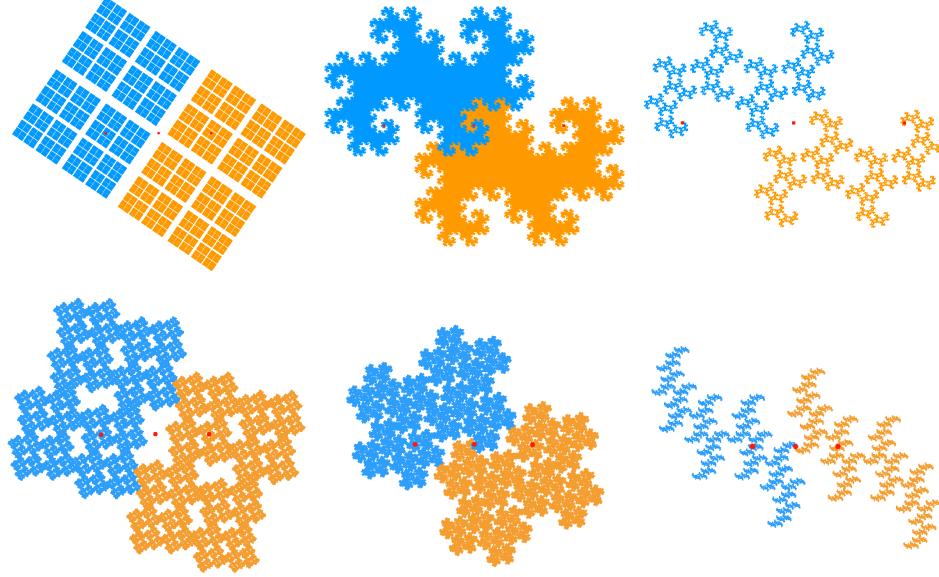


FIGURE 1. Some limit sets Λ_z for various parameters. In each case, we show the decomposition of Λ_z as the union of $f\Lambda_z$ (blue) and $g\Lambda_z$ (orange). The points 0, $1/2$ and 1 are marked in red. Along the bottom from left to right, the parameters lie in $\mathcal{M} - \mathcal{M}_0$, $\mathcal{M}_0 - \mathcal{M}_1$, and \mathcal{M}_1 , respectively.

2.3. Basic symmetries. Complex conjugation “conjugates” G_z to $G_{\bar{z}}$. Thus Λ_z and $\Lambda_{\bar{z}}$ are mirror images of each other. In particular, they are homeomorphic, and are therefore connected, simply connected etc. for the same values of z .

The semigroup G_z has another basic symmetry: rotation through π about the point $1/2$ interchanges the two generators f and g . Thus the limit set Λ_z is invariant under this symmetry: $\Lambda_z = 1 - \Lambda_z$. On the other hand, by definition

$$\Lambda_z = (z\Lambda_z) \cup (z\Lambda_z + (1 - z)).$$

Using the relation $\Lambda_z = 1 - \Lambda_z$ we obtain the identity

$$\Lambda_z = (z\Lambda_z) \cup (-z\Lambda_z + 1)$$

which is the limit set of the semigroup H_z with generators

$$f : x \mapsto zx, \quad g : x \mapsto 1 - zx.$$

Thus, although G_z and H_z are *not conjugate* (not even topologically, and in general not even when restricted to Λ_z), they have the same limit set. Now, from the definition, the limit sets of H_z and H_{-z} are similar. It follows that the same is true for G_z and G_{-z} .

We record this observation as a lemma:

Lemma 2.3.1 (Similar limit sets). *The limit sets $\Lambda_z, \Lambda_{-z}, \Lambda_{\bar{z}}$ and $\Lambda_{-\bar{z}}$ are similar or mirror images of each other.*

2.4. Three sets. We now define three subsets in parameter space \mathbb{D}^* of our semigroups G_z . These sets are the basic objects of interest in this paper.

- (1) \mathcal{M} is the set of z such that Λ_z is connected;
- (2) \mathcal{M}_0 is the set of z such that Λ_z contains $1/2$; and
- (3) \mathcal{M}_1 is the set of z such that Λ_z is connected and full.

Recall that a set is *full* if its complement is connected. These sets are all closed.

As far as we know, the set \mathcal{M} was first introduced by Barnsley-Harrington [2], and the set \mathcal{M}_0 was first introduced by Bousch [3]. We are not aware of any previous explicit mention of \mathcal{M}_1 , although Bandt [1], Solomyak [12] and others have studied the (closely related) set of z for which Λ_z is a dendrite. Figure 2 is a picture of \mathcal{M} , and Figure 3 is a picture of \mathcal{M}_0 . The set \mathcal{M}_1 is much less substantial, and it is harder to draw a good picture.

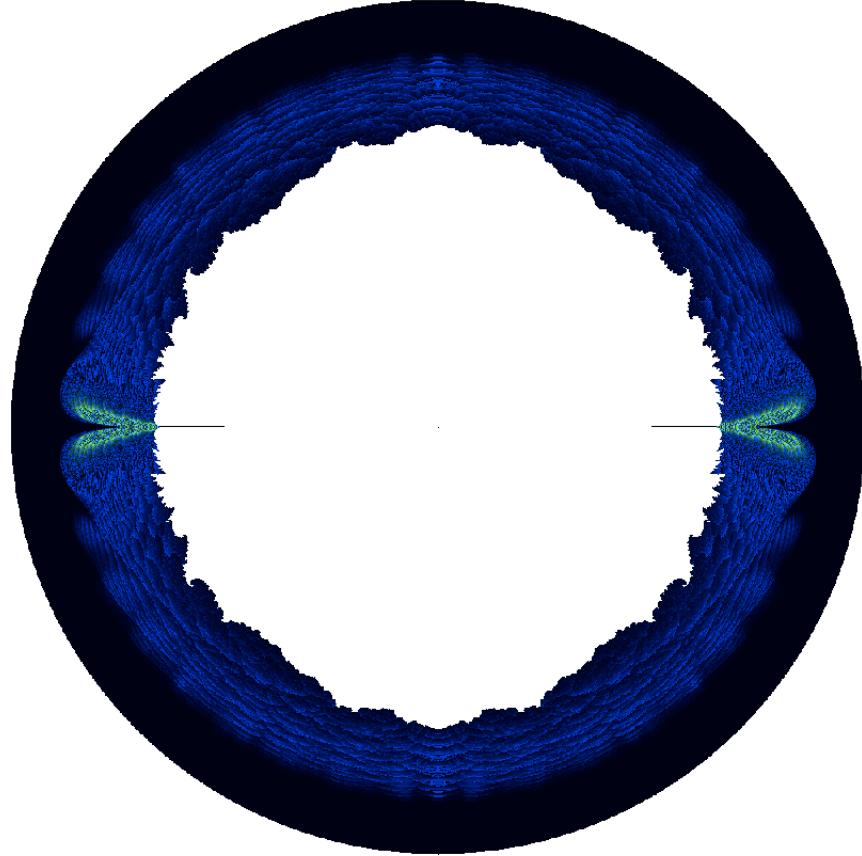
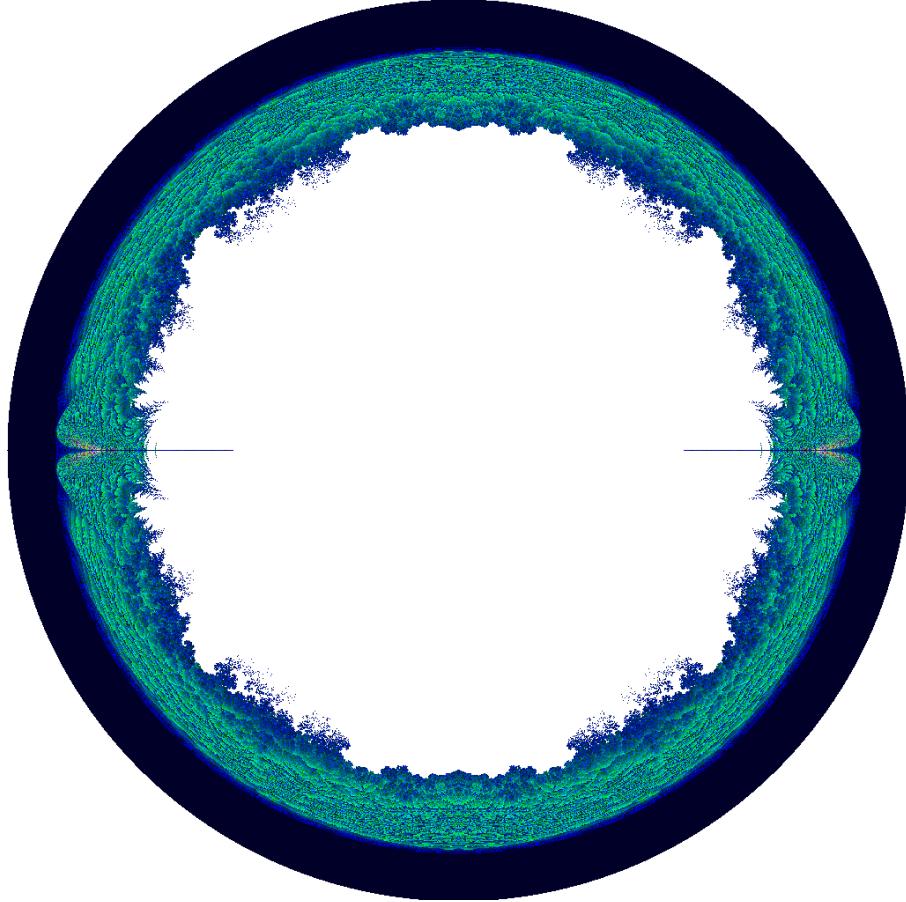


FIGURE 2. \mathcal{M} drawn in \mathbb{D}^* .

Proposition 2.4.1. *We have $\mathcal{M}_1 \subsetneq \mathcal{M}_0 \subsetneq \mathcal{M}$.*

Proof. It is straightforward to show (see Lemma 5.2.1) that $z \in \mathcal{M}$ — i.e. the limit set Λ_z is connected — if and only if $f\Lambda_z := f(\Lambda_z)$ intersects $g\Lambda_z := g(\Lambda_z)$. Since Λ_z is rotationally symmetric about the point $1/2$, it follows that \mathcal{M}_0 is contained in \mathcal{M} . Likewise, if Λ_z is connected and simply-connected, then because it is rotationally

FIGURE 3. \mathcal{M}_0 drawn in \mathbb{D}^* .

symmetric about $1/2$, it follows that Λ_z contains $1/2$. No two of these sets are equal; see Figure 1. \square

We will focus on the sets \mathcal{M} and \mathcal{M}_0 for the remainder of the paper.

2.5. Holes. We will show (see Theorem 5.2.3) that z is in the complement of \mathcal{M} if and only if G_z is Schottky. We have already observed that all Schottky semigroups are topologically conjugate when restricted to good disks. The set of z for which G_z is Schottky is evidently open. However, an examination of Figure 2 with a microscope reveals the apparent existence of tiny “holes” in \mathcal{M} , corresponding to “exotic” components of Schottky space.

One hole in \mathcal{M} is clearly visible in Figure 2; it is shaped approximately like a round disk except for two “whiskers” of \mathcal{M} along the real axis. But it turns out that there are also much smaller holes in \mathcal{M} , which can be thought of as exotic components of Schottky space. This is in stark contrast to the situation of Kleinian groups, where the (Teichmüller) spaces of (quasifuchsian) representations of a surface of fixed topological type are connected, as can be proved by means of the measurable Riemann mapping theorem.

Figure 4 shows a collection of holes in \mathcal{M} centered near the point $0.372368 + 0.517839i$, which we refer to colloquially as *hexaholes*. The diameter of the picture is approximately 0.0005, so these holes are much too small to see in Figure 2. It is one of the main goals of this paper to prove rigorously that infinitely many holes such as these really do exist in \mathcal{M} .

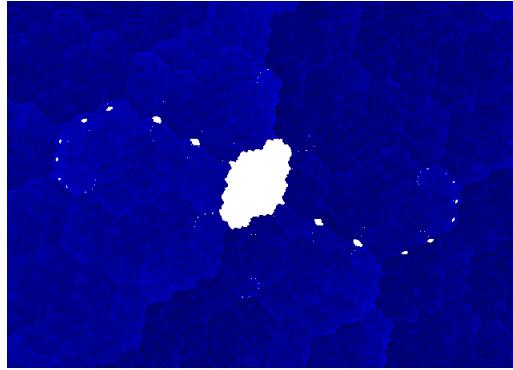


FIGURE 4. Apparent holes in \mathcal{M} near the point $z = 0.372368 + 0.517839i$. The width of the figure is about 0.0005.

2.6. Some history. The sets \mathcal{M} and \mathcal{M}_0 have a long history, and these sets (and some close relatives) were discovered independently several times by people working in quite different areas of mathematics. In fact, we ourselves did not learn of the work of Bandt and Solomyak until an advanced stage of our investigations. Therefore we believe it would be useful to briefly mention some of the important papers on this subject that have appeared over the last 30 years, and say something about their contents.

- In 1985, Barnsley and Harrington [2] initiated a (mainly numerical) study of \mathcal{M} . They discovered much structure evident in this set, most significantly the presence of apparent holes, whose rigorous existence they conjectured. Another phenomenon they discovered was the real whiskers in \mathcal{M} , and they proved rigorously that \mathcal{M} is entirely real in some definite neighborhood of the endpoints ± 0.5 of these whiskers:

Theorem 2.6.1 (Barnsley-Harrington, whiskers). *There is a neighborhood of the points ± 0.5 in which \mathcal{M} is contained in \mathbb{R} .*

Let α be the supremum of the real numbers t for which \mathcal{M} intersects some neighborhood of $[0.5, t]$ only in real points. Barnsley-Harrington obtained a rigorous estimate $\alpha > 0.53$ but observed that this estimate is far from sharp.

- In 1988 Thierry Bousch began a systematic study of \mathcal{M} and \mathcal{M}_0 in his unpublished papers [3] and [4]. Bousch proved many remarkable theorems about \mathcal{M} and \mathcal{M}_0 , including the following:

Theorem 2.6.2 (Bousch, connectivity). *\mathcal{M} and \mathcal{M}_0 are both connected and locally connected.*

Bousch interpreted both sets as the zeros of power series with coefficients of a particular form; we will return to this perspective in Section 4.

- In 1993, Odlyzko and Poonen [9] studied zeroes of polynomials with 0, 1 coefficients (a set closely related to \mathcal{M}_0) and showed the closure of this set is path connected; their techniques are similar to those of Bousch. They also noted the presence of apparent holes, and conjectured that they really existed.
- In 2002 Bandt [1] developed some fast algorithms to draw accurate pictures of \mathcal{M} , and managed to rigorously prove the existence of a hole in \mathcal{M} , thus positively answering the conjecture of Barnsley-Harrington. Bandt first realized the importance of understanding the set of interior points in \mathcal{M} , and made the following conjecture:

Conjecture 2.6.3 (Bandt, interior almost dense). *The interior of \mathcal{M} is dense away from the real axis.*

which has been at the center of much subsequent work. Note that the necessity to exclude the real axis from this conjecture is already implied by Theorem 2.6.1.

Bandt's algorithm explicitly related $z \in \mathcal{M}$ to the dynamics of a 3-generator semigroup $f : x \mapsto zx - 1$, $g : x \mapsto zx$, $h : x \mapsto zx + 1$ which we denote H_z , and remarked on the apparent similarity of \mathcal{M} and the limit set Γ_z of H_z at certain algebraic points on $\partial\mathcal{M}$ that he called *landmark points*.

- In 2003 Solomyak [11] and Solomyak-Xu [13] made partial progress on Bandt's conjecture, finding some interior points in \mathcal{M} with $|z| < 2^{-1/2}$, and showing that interior points are dense in \mathcal{M} in some definite neighborhood of the imaginary axis. They also obtained strong results on the structure of the natural invariant measures on Λ_z , relating this to the classical study of Bernoulli convolutions, and were able to compute the Hausdorff dimension and measure of the limit set for almost all z .
- In 2005 Solomyak [12] proved the asymptotic similarity of \mathcal{M} and Γ_z at certain points z which satisfy the condition that z is a root of a rational function of a particular form. Following Solomyak, we refer to these points as *landmark points*. Then Solomyak shows

Theorem 2.6.4 (Solomyak [12]). *If $z \in \mathcal{M} - \mathbb{R}$ is a landmark point then \mathcal{M} is asymptotically similar at z to the set Γ_z at a certain specific point, and both of these sets are asymptotically self-similar at these points.*

Here asymptotic similarity of two sets X and Y at 0 (for simplicity) means that the Hausdorff distance between $t^{-1}(X)$ and $t^{-1}(Y)$ restricted to balls of fixed radius (and ignoring the boundary) goes to zero as $t \rightarrow 0$; and asymptotic self-similarity means that there is a complex z with $|z| < 1$ so that the sets $z^n X$ converge on compact subsets in the Hausdorff topology to a limit.

- In 2011, Thurston [14] studied the set of Galois conjugates of algebraic numbers e^λ where λ is the core entropy of a postcritically finite interval map $x \mapsto x^2 + c$, for which the parameter c is taken from the main “limb” of the Mandelbrot set (the intersection of the Mandelbrot set with \mathbb{R}). He asserted that the closure of this set of roots (in \mathbb{C}) is connected and path connected. In \mathbb{D}^* the closure of this set agrees with \mathcal{M}_0 , and therefore the assertion generalizes Theorem 2.6.2. For $|z| \geq 1$ this assertion was verified in an elegant paper by Giulio Tiozzo [15], who also went on to plot Galois conjugates associated with core entropies of postcritically finite maps $x \mapsto x^2 + c$, where c comes from other limbs of the Mandelbrot set; these sets display a “family resemblance” to \mathcal{M}_0 .

These papers describe some remarkable connections related to the theory of post-critically finite interval maps, Perron numbers, Galois theory and so on. The richness and mathematical depth of these various sets has barely begun to be plumbed. We emphasize that the survey above is not exhaustive, and the papers cited contain a substantial amount beyond the part we summarize here.

3. ELEMENTARY ESTIMATES

In this section we collect a few elementary estimates about the geometry of Λ_z .

3.1. Geometry of Λ_z . Recall our notation G_z for the semigroup generated by $f : x \mapsto zx$ and $g : x \mapsto z(x - 1) + 1$. The map f fixes 0 and the map g fixes 1. Any element $e \in G_z$ of length n acts as a similarity on \mathbb{C} with dilation $|z^n|$ and center some point of Λ_z . We make some *a priori* estimates on the geometry of Λ_z .

Lemma 3.1.1 (Diameter bound). *The limit set Λ_z is contained in the ball of radius $|z - 1|/2(1 - |z|)$ centered at 1/2.*

Proof. Let D denote the ball of radius R centered at 1/2. Then $fD := f(D)$ and $gD := g(D)$ are the balls of radius $|z|R$ centered at $z/2$ and $1 - z/2$ respectively. So providing $R \geq |z - 1|/2(1 - |z|)$ we have $fD, gD \subseteq D$. But this means $\Lambda_z \subseteq D$. \square

Lemma 3.1.2. *Let e, e' be words with a common prefix of length n . Let x be contained in D , the ball of radius $|z - 1|/2(1 - |z|)$ centered at 1/2. Then*

$$d(ex, e'x) \leq \frac{|z|^n|z - 1|}{1 - |z|}.$$

Proof. Write $e = uv$ and $e' = uv'$. Then $vx, v'x \in D$ so $d(vx, v'x) \leq |z - 1|/(1 - |z|)$ by Lemma 3.1.1. But the dilation of u is $|z|^n$, so the estimate follows. \square

Definition 3.1.3 (Compactification). Let Σ denote the set of finite words in the alphabet $\{f, g\}$, and let $\bar{\Sigma}$ denote all right-infinite words in this alphabet, such that if a word contains $*$, all successive letters are also $*$. Metrize $\bar{\Sigma}$ with the metric $d(e, e') = 2^{-n}$ where n is the length of the biggest common prefix of e and e' .

The set $\bar{\Sigma}$ decomposes naturally into the subset $\partial\Sigma$ of words not containing the symbol $*$, and words that do contain the symbol $*$ which are in natural bijection with Σ , under the map that takes a finite word in f, g to the infinite word obtained by padding with infinitely many $*$ symbols.

Lemma 3.1.4. *The space Σ is compact. The subspace $\partial\Sigma$ is homeomorphic to a Cantor set, and Σ is homeomorphic to a discrete set, whose accumulation points are precisely $\partial\Sigma$.*

Proof. This is immediate from the definition. \square

There is a natural symmetry of $\bar{\Sigma}$ interchanging the symbols f and g and fixing the symbol $*$.

Note that $\bar{\Sigma}$ is formally distinct from G_z , which is the semi-group generated by compositions of the affine maps f and g . It's important to make the distinction between Σ and G_z as we are interested in how the semigroup changes as the parameter z varies.

Definition 3.1.5. There is an obvious map

$$\sigma_z : \Sigma \rightarrow G_z$$

such that $\sigma_z(u) \in G_z$ is the appropriate composition of the maps

$$f : x \mapsto zx, \quad \text{and} \quad g : x \mapsto z(x-1) + 1.$$

Definition 3.1.6. Let $u \in \Sigma$ be a word of length n , and (by abusing notation), define the map $u : \mathbb{D}^* \times \mathbb{C} \rightarrow \mathbb{C}$ given by

$$u : (z, x) \mapsto \sigma_z(u)(x).$$

We will also use the notation $u(z)(x) := u(z, x)$, and we often consider the map $\mathbb{D}^* \rightarrow \mathbb{C}$ given by $z \mapsto u(z, x)$. The map u is continuous in both z and x , which is evident in Section 4.

Definition 3.1.7. The map $\pi : \partial\Sigma \times \mathbb{D}^* \rightarrow \mathbb{C}$ is defined by

$$\pi(u, z) = \lim_{n \rightarrow \infty} u_n(z, x)$$

where u_n is the prefix of u of length n , and $x \in \mathbb{C}$ is any point. By Lemma 3.1.2, this limit is well-defined, independent of the point $x \in \mathbb{C}$.

Lemma 3.1.8. For $u, v \in \Sigma$ and $x \in \mathbb{C}$, we have $uv(z, x) = u(z, v(z, x))$. That is, $uv(z) = u(z) \circ v(z)$. For $u \in \Sigma$ and $v \in \partial\Sigma$, we have $\pi(uv, z) = u(z, \pi(v, z))$.

Proof. Obvious from the definitions. □

Lemma 3.1.9 (Hölder continuous). *The map $\pi(\cdot, z) : \partial\Sigma \rightarrow \mathbb{C}$ is Hölder continuous with exponent $\log|z|/\log(0.5)$, and the image is Λ_z .*

Proof. Evidently if e is a periodic word $e := vvvv \dots$ then $\pi(e, z)$ is the center (i.e. the fixed point) of v ; since $\partial\Sigma$ is compact, if π is continuous, then the image is closed and is therefore equal to Λ_z . So it suffices to show π is Hölder, and estimate the exponent.

From the definition, if e, e' have a common maximal prefix of length n then $d_G(e, e') = 2^{-n}$. On the other hand, by Lemma 3.1.2 we obtain

$$d(\pi(e, z), \pi(e', z)) \leq \frac{|z|^n |z-1|}{(1-|z|)} = \frac{(0.5^n)^\alpha |z-1|}{(1-|z|)}$$

for $\alpha = \log|z|/\log(0.5)$. □

3.2. Geometry of \mathcal{M} . The following result is proved in [3]; we include a proof for completeness.

Lemma 3.2.1 (inner and outer annuli). *\mathcal{M} (the set of z for which the semigroup G_z has connected Λ_z) contains the region $|z| \geq 1/\sqrt{2} = 0.7071067 \dots$ and is contained in the region $|z| \geq 1/2$.*

Proof. We shall see (Lemma 5.2.1) that the limit set Λ_z of the semigroup G_z is disconnected if and only if $f\Lambda_z \cap g\Lambda_z$ is empty, in which case Λ_z is a Cantor set. In this case, the Hausdorff dimension of Λ_z can be computed from Moran's Theorem (see [6], Ch. 2), as the unique d for which

$$2|z|^d = 1.$$

In fact, this is easy to see directly: for a subset of Euclidean space, the d -dimensional Hausdorff measure transforms by λ^d when the set is scaled linearly by the factor

λ . When Λ_z is disconnected, it is the disjoint union of $f\Lambda_z$ and $g\Lambda_z$, which are obtained (up to translation) by scaling Λ_z by z ; the formula follows.

If $|z| > 1/\sqrt{2}$ then $d > 2$ which is absurd, since Λ_z is a subset of \mathbb{C} . Thus $|z| > 1/\sqrt{2}$ is in \mathcal{M} , and since this set is closed, so is $|z| \geq 1/\sqrt{2}$.

Conversely, if $|z| < 1/2$ then the round disk D of radius 1 centered at $1/2$ is a good disk, so G_z is Schottky, and Λ_z is disconnected. \square

Example 3.2.2. The estimates in Lemma 3.2.1 are sharp. Taking $z = 1/2$, we see that $f(1) = g(0) = 1/2$, so $1/2 \in f\Lambda_{1/2} \cap g\Lambda_{1/2}$; in fact, $\Lambda_z = [0, 1]$ in this case (and for all $z \in [1/2, 1)$).

Likewise, taking $z = i/\sqrt{2}$ the rectangle R with corners $\{-1, i/\sqrt{2}, 2, 2 - i/\sqrt{2}\}$ satisfies $R = fR \cup gR$, so that $R = \Lambda_{i/\sqrt{2}}$, whereas for $z = it$ with $t < 1/\sqrt{2}$ the rectangle with corners $\{-1, it, 2, 2 - it\}$ is good and G_z is Schottky; see Figure 1, left.

Solomyak–Xu [13] Thm. 2.8 show that the set of z with $|z| < 2^{-1/2}$ for which the Hausdorff dimension of Λ_z is different from $d := -\log 2 / \log |z|$ itself has Hausdorff dimension less than 2. Since (as we shall show) interior points are dense in \mathcal{M} away from \mathbb{R} , this implies that the simple formula for the Hausdorff dimension of Λ_z is valid on a dense subset of \mathcal{M} . Finer results about the “exceptions” are known, but we do not pursue that here.

4. ROOTS, POLYNOMIALS, AND POWER SERIES WITH REGULAR COEFFICIENTS

The most interesting mathematical objects are those that can be defined in many different — and apparently unrelated — ways. The sets \mathcal{M} and \mathcal{M}_0 can be defined in a way which is (at first glance) entirely unconnected to dynamics, namely as the closures of the set of *roots* of certain classes of polynomials. This connection is quite sensitive to the choice of normalization for our semigroups, and in fact the freedom to choose several different normalizations is itself of some theoretical interest.

4.1. The Barnsley-Harrington and Bousch normalization. Recall the normalization $f : x \mapsto zx + 1$, $g : x \mapsto zx - 1$. If w is a word of length n in f and g , we can express wx as a polynomial of a particularly simple form, namely

$$w(z, x) = \sum_{j=0}^{n-1} a_j z^j + x z^n$$

where the $a_j \in \{-1, 1\}$ are equal to 1 or -1 according to whether each successive letter of w is equal to f or g . In particular, the limit set Λ_z is precisely equal to the set of values of power series in z with coefficients in $\{-1, 1\}$. In this normalization, the center of symmetry is 0 (rather than $1/2$ as in our normalization), so we obtain the following characterization of \mathcal{M}_0 :

Proposition 4.1.1 (Power series with $\{-1, 1\}$ coefficients). *The set \mathcal{M}_0 is the set of $z \in \mathbb{D}^*$ which are zeros of power series with coefficients in $\{-1, 1\}$.*

Similarly, the subsets $f\Lambda_z$ and $g\Lambda_z$ are the sets of *values* of power series with $\{-1, 1\}$ coefficients which start with 1 and -1 respectively. Thus $z \in \mathcal{M}$ if and only if $f\Lambda_z \cap g\Lambda_z$ is nonempty, which happens if and only if z is a root of a power series with coefficients in $\{-2, 0, 2\}$ starting with ± 2 . Equivalently, after dividing such a

power series by 2, we see that $z \in \mathcal{M}$ if and only if z is a root of a power series with coefficients in $\{-1, 0, 1\}$ starting with ± 1 :

Proposition 4.1.2 (Power series with $\{-1, 0, 1\}$ coefficients). *The set \mathcal{M} is the set of $z \in \mathbb{D}^*$ which are zeros of power series with coefficients in $\{-1, 0, 1\}$ with constant term $= \pm 1$.*

In either case, zeros of power series with prescribed coefficients can be approximated by zeros of *polynomials* with the same constraints on the coefficients. This suggests defining an “extended” \mathcal{M} (resp. \mathcal{M}_0) to be the closures of the set of *all* roots z (not just those in \mathbb{D}^*) of polynomials with coefficients in $\{-1, 0, 1\}$ (resp. $\{-1, 1\}$). Reversing the order of the coefficients replaces a root by its reciprocal, so these extended sets are exactly the sets obtained by taking the union of \mathcal{M} (resp. \mathcal{M}_0) together with its image under inversion in the unit circle.

Using this interpretation of Λ_z as the set of values of power series with $\{-1, 1\}$ coefficients, Bousch noted an interesting relationship between \mathcal{M} and \mathcal{M}_0 . We give the proof here for several reasons. Firstly, Bousch’s paper is unpublished, and this argument is not easy to extract from the paper. Secondly, it is short and illuminating. Thirdly, it depends on a geometric fact which we use later in the proof of Proposition 6.1.3.

Proposition 4.1.3 (Bousch [3], Prop. 2). *If $z^2 \in \mathcal{M}$ then $z \in \mathcal{M}_0$. Consequently \mathcal{M}_0 contains the annulus $2^{-1/4} \leq |z| < 1$.*

Proof. Let \mathcal{P} denote the set of power series with coefficients in $\{-1, 1\}$. Then for any $p \in \mathcal{P}$ we can write $p(z) = p_e(z^2) + zp_o(z^2)$ for unique $p_e, p_o \in \mathcal{P}$. But this means that $\Lambda_z = \Lambda_{z^2} + z\Lambda_{z^2}$.

Now, in this normalization, limit sets all have rotational symmetry around 0. So if Λ_{z^2} is connected but doesn’t contain 0, there is some symmetric innermost loop γ around 0. If Λ_z doesn’t contain 0, then (since $\Lambda_{z^2} = -\Lambda_{z^2}$) it must be that Λ_{z^2} and $z\Lambda_{z^2}$ are disjoint, so that $z\Lambda_{z^2}$ is contained in the disk bounded by γ , and similarly $z^2\Lambda_{z^2}$ is contained in the disk bounded by $z\gamma$, and therefore Λ_{z^2} is disjoint from $z^2\Lambda_{z^2}$. But Bousch shows this is absurd in the following way.

Write $L := z^2\Lambda_{z^2}$ so that $\Lambda_{z^2} = (L + 1) \cup (L - 1)$. By hypothesis, both L and $(L + 1) \cup (L - 1)$ are compact and connected, so that $(L - 1)$ intersects $(L + 1)$. But then L must intersect $(L + 1) \cup (L - 1)$, since if it is disjoint from them both, the union of L with vertical rays from its top-most and bottom-most point to infinity separates $(L - 1)$ from $(L + 1)$. \square

4.2. Our normalization. Now let’s return to our normalization $f : x \mapsto zx$, $g : x \mapsto z(x - 1) + 1 = zx + (1 - z)$. If we fix $e \in \partial\Sigma$ and vary z , note that $z \mapsto \pi(e, z)$ is a function of z . For any fixed e , we can express $\pi(e, z)$ as a very simple power series in z . In fact, the set of power series that can be obtained are precisely those whose coefficients, listed in order, are (right-infinite) words in an explicit regular language (for an introduction to the theory of regular languages, see e.g. [7]).

Proposition 4.2.1 (Power series). *For any fixed $e \in \Sigma$ of length m there is a formula*

$$e(z, x) = xz^m + \sum_{j=0}^m a_j z^j$$

where each $a_j \in \{-1, 0, 1\}$, and $a_m = 0$ if e ends with f and $a_m = -1$ if e ends with g .

Furthermore, the string of digits a_j for $j < m$ can be recursively obtained as follows. Read the letters of e from left to right, and express this as a walk on the edges of the directed labeled graph in Figure 5, starting at the vertex labeled $*$.

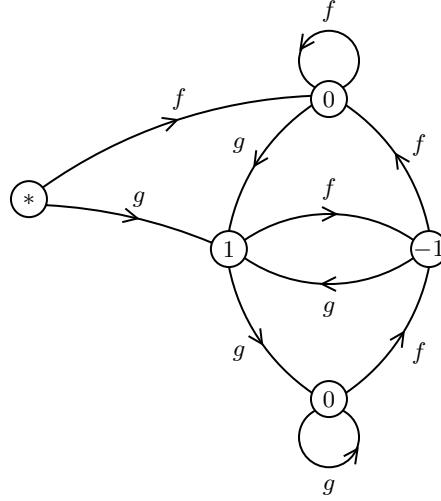


FIGURE 5. The coefficients a_j are the vertex labels visited in order on the walk associated to a word.

The coefficients a_j for $j < m$ (in order) are the labels on the vertices visited in this path, after the initial vertex. Thus, the sequences that occur are precisely the sequences in which the nonzero coefficients alternate between 1 and -1 , starting with 1.

Similarly, for any fixed $e \in \partial\Sigma$ we can write

$$\pi(e, z) = \sum_{j=0}^{\infty} a_j z^j$$

where each $a_j \in \{-1, 0, 1\}$, and the a_j are obtained as the labels on the vertices associated to the right-infinite walk on the graph as above.

Proof. This is immediate by induction. \square

Example 4.2.2. From Proposition 4.2.1 we can quickly generate the formula for $e(z, x)$ for any finite word e . For example, taking $e = gfgfffgg$, and writing -1 as $\bar{1}$, we compute the sequence as follows:

$$\emptyset \xrightarrow{g} 1 \xrightarrow{f} 1\bar{1} \xrightarrow{g} 1\bar{1}1 \xrightarrow{f} 1\bar{1}1\bar{1} \xrightarrow{f} 1\bar{1}1\bar{1}0 \xrightarrow{f} 1\bar{1}1\bar{1}00 \xrightarrow{g} 1\bar{1}1\bar{1}001 \xrightarrow{g} 1\bar{1}1\bar{1}0010$$

so that we obtain the formula

$$e(z, x) = 1 - z + z^2 - z^3 + z^6 + (x - 1)z^8$$

4.3. Regular coefficients. We now show that in general families of semigroups of similarities with centers depending polynomially on the (common) dilation give rise to limit sets which are the values of power series with “regular” coefficients. Very similar, but somewhat complementary observations were made by Mercat [8].

Definition 4.3.1. Fix some finite alphabet S of complex numbers, and fix a prefix-closed regular language $L \subseteq S^*$. Let \bar{L} denote the set of right-infinite words in S whose finite prefixes are in L . Call a power series

$$e(z) := a_0 + a_1 z + a_2 z^2 + \dots$$

L-regular if the sequence $(a_0, a_1, \dots) \in \bar{L}$.

Proposition 4.3.2 (Coefficient language). *Let p_i for $1 \leq i \leq m$ be a finite set of complex polynomials, and define K_z to be the semigroup generated by contractions*

$$f_i : x \mapsto zx + p_i(z)$$

Then there is a regular language L in a finite alphabet of complex numbers so that a power series of the form

$$e(z) := a_0 + a_1 z + a_2 z^2 + \dots$$

is L -regular if and only if $e \in \partial K_z$; that is e is an infinite composition of the generators f_i , thought of as a function in z .

Proof. The effect of f_i on some element of ∂K_z is to shift the sequence by one (the $x \mapsto zx$ part) and to add the coefficients of p_i to the first $d_i + 1$ coefficients, where d_i is the degree of p_i . Introduce the notation

$$p_i(z) = b_{i,0} + b_{i,1}z + \dots + b_{i,d_i}z^{d_i}$$

and pad coefficients up to $b_{i,d}$ where $d = \max_i d_i$ by defining $b_{i,j} = 0$ if $d_i < j \leq d$. Then if we let $e = f_{s_1} f_{s_2} f_{s_3} \dots$ be an arbitrary element of ∂K_z , the n th coefficient a_n of the power series expansion of $e(z)$ is given by the formula

$$a_n = b_{s_n,0} + b_{s_{n-1},1} + \dots + b_{s_{n-d},d}$$

provided $n \geq d$, and for $n < d$ we simply omit the terms $b_{s_{n-i},d}$ for $n - i < 0$ and $i \leq d$. This coefficient depends only on the last $d + 1$ letters visited in order, and a finite state automaton can store this information as a vertex.

Explicitly, we build a finite graph with $(m^{d+1} - 1)/(m - 1)$ vertices in bijection with words of length at most d in the f_i , and with an edge from each vertex corresponding to the word u to the vertex corresponding to v , with the edge labeled f_j if uf_j has v as a suffix of length $\min(d, |uf_j|)$.

Now at each vertex associated to a word u of length $d' \leq d$ of the form $u = s_1 s_2 \dots s'_d$, put the coefficient

$$a(u) := b_{s_{d'},0} + b_{s_{d'-1},1} + \dots + b_{s_{d'-d},d}$$

We now relabel the edges in such a way that the new label on each edge is equal to the coefficient at the vertex it points to. The resulting directed graph is a *nondeterministic* finite state automaton in a finite alphabet (the alphabet of possible coefficients), and the set of possible edge paths is some language L . It is a standard theorem in the theory of automata due to Kleene–Rabin–Scott (see [7], Thm. 1.2.7) that there is a *deterministic* finite state automaton in the same alphabet recognizing L ; this means (by definition) that L is regular. Moreover by construction, \bar{L} is precisely the language of coefficient sequences. \square

We refer to the language of coefficient sequences as the *coefficient language* of the parameterized family K_z , and denote it $L(K_z)$. In the special case that K_z is generated by two elements p_1, p_2 , then at least on the subset where $p_1(z) \neq p_2(z)$, the semigroup K_z is conjugate to the semigroup generated by $f : x \mapsto zz$, $g : x \mapsto z(x-1)+1$ that we have been studying up to now.

Question 4.3.3. *Which regular languages in a finite alphabet arise as $L(K)$ for some K ?*

Example 4.3.4 (Differences). Let K_z be a holomorphic family of semigroups of similarities, parameterized by z , whose IFS has coefficient language $L(K_z)$. The set of differences $DL(K_z) := \{a - b \text{ such that } a, b \in L(K_z)\}$ is of the form $L(DK_z)$ for a suitable holomorphic family DK_z .

In Section 6 we illustrate this difference operation in the context of our 2-generator IFSs, obtaining a sequence of “iterated Mandelbrot sets”, of which \mathcal{M}_0 and \mathcal{M} are the first two terms.

5. TOPOLOGY AND GEOMETRY OF THE LIMIT SET

In this section we establish basic facts about the geometry and topology of Λ_z , establishing quantitative versions of the fundamental dichotomy that either Λ_z is (path) connected, or Λ_z is a Cantor set and G_z is Schottky. These facts lead to an explicit algorithm (essentially due to Bandt) to (numerically) certify that a particular G_z is Schottky. It is important to describe this algorithm and its justification in some detail for several reasons. Firstly, this algorithm powers our program `schottky`, which provided numerical certificates for many of the assertions we make in this paper (and produced most of the pictures!). Secondly, understanding the *theoretical* behaviour of this algorithm, were it run on an ideal computer for infinite time, leads to some of the key theoretical insights that underpin our main theorems.

5.1. Constructing Λ_z . Recall that Σ is the set of all finite words in f and g . For each $n \in \mathbb{N}$ define Σ_n to be the set of words of length n .

Because Λ_z is minimal, for any $p \in \Lambda_z$,

$$\Lambda_z = \overline{\bigcup_n \Sigma_n(z, p)}.$$

Furthermore, the limit set is well-approximated by $\Sigma_n(z, p)$ for any p which is close to Λ_z :

Lemma 5.1.1. *Let $p \in \mathbb{C}$. Then $\Lambda_z \subseteq N_\delta(\Sigma_n(z, p))$ where*

$$\delta = |z|^n \left(\frac{|z-1|}{1-|z|} + d(p, \Lambda_z) \right)$$

Proof. Let $x \in \Lambda_z$ be such that $d(p, x) = d(p, \Lambda_z)$. We can write $x = \pi(u, z)$ for $u \in \partial\Sigma$. Now let $y \in \Lambda_z$ be given, and write $y = \pi(v, z)$ for $v \in \partial\Sigma$. Let $v_n \in \Sigma_n$ be the prefix of y of length n . Consider $w = \pi(v_n u, z) \in \Lambda_z$. Note $w = v_n(z, x)$, so $d(v_n(z, p), w) = d(v_n(z, p), v_n(z, x)) = |z|^n d(p, x)$. By Lemma 3.1.2, $d(w, y) \leq |z|^n |z-1| / (1-|z|)$, so by the triangle inequality,

$$d(v_n(z, p), y) \leq d(v_n(z, p), w) + d(w, y) \leq \delta$$

□

Let D_z be any compact set containing Λ_z (for example a disk of radius $|z - 1|/2(1 - |z|)$ centered at $1/2$). Let $D_n = \Sigma_n(z, D_z)$; this is a union of 2^n copies of D_z scaled by the factor $|z|^n$. We can construct Λ_z as a descending intersection:

Lemma 5.1.2. *We have*

$$\Lambda_z = \bigcap_n D_n.$$

Proof. Observe that $\bigcap_n D_n$ is a compact, nonempty invariant set. Since Λ_z is the unique such set, they must be equal. \square

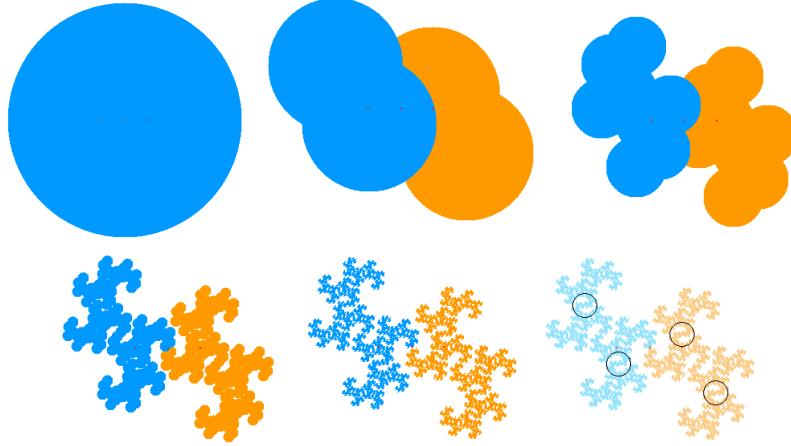


FIGURE 6. Constructing Λ_z by intersecting the unions of disks D_n . The bottom right picture indicates how Λ_z decomposes as a union of 4 copies of itself centered at the indicated circles.

5.2. Connectivity.

Lemma 5.2.1. *The following are equivalent:*

- (1) Λ_z is disconnected;
- (2) Λ_z is a Cantor set; or
- (3) $f\Lambda_z \cap g\Lambda_z$ is empty.

Moreover, any of these conditions is implied by G_z Schottky.

Proof. The implications (3) \rightarrow (2) \rightarrow (1) are obvious, and (1) \rightarrow (3) is a standard result in the theory of IFS. For a proof (in exactly this context) see [3], p.2 (alternately, it follows from the estimates in Lemma 5.2.2).

G_z Schottky immediately proves (3), since if D is a good disk for G_z , then D contains Λ , but then fD and gD contain $f\Lambda_z$ and $g\Lambda_z$ and are disjoint by the definition of a good disk. \square

Lemma 5.2.1 implies that Schottky semigroups have disconnected limit sets. The next Lemma, although elementary, is the key to proving the converse. See Figure 7 for an illustration of the paths produced by Lemma 5.2.2.

Lemma 5.2.2 (Short Hop Lemma). *Suppose that $f\Lambda_z$ and $g\Lambda_z$ contain points at distance δ apart. Then the $\delta/2$ neighborhood of Λ_z is path connected.*

Proof. Since $|z| < 1$ there is some n so that for any two $e, e' \in \partial\Sigma$ with a common prefix of length at least n we have $d(\pi(e, z), \pi(e', z)) < \delta$.

Suppose v, v' are words of length i . Write $v \approx_i v'$ if there are right-infinite words u, u' with prefixes v, v' such that $d(\pi(u, z), \pi(u', z)) < \delta$. Then define \sim_i to be the equivalence relation generated by \approx_i .

We claim that for all i the equivalence relation \sim_i has a single equivalence class; i.e. that *any* two words of length i can be joined by a sequence of words of length i related by \approx_i . Evidently $f \approx_1 g$ since we can choose right-infinite words fu and gu' such that $\pi(fu, z)$ and $\pi(gu', z)$ are points in $f\Lambda_z$ and $g\Lambda_z$ respectively realizing $d(\pi(fu, z), \pi(gu', z)) \leq \delta$.

If $v \sim_i v'$ for all words v, v' of length i , then $fv \sim_{i+1} fv'$ and $gv \sim_{i+1} gv'$ for all words v, v' of length i , since $v \approx_i v'$ implies $fv \approx_{i+1} fv'$ and $gv \approx_{i+1} gv'$. But if fu and gu' are as above, and w, w' are the initial words of fu and gu' of length $i+1$ then $w \approx_{i+1} w'$. So the claim is proved for all i , by induction.

Taking $i = n$ and using the defining property of n as above proves the lemma. \square

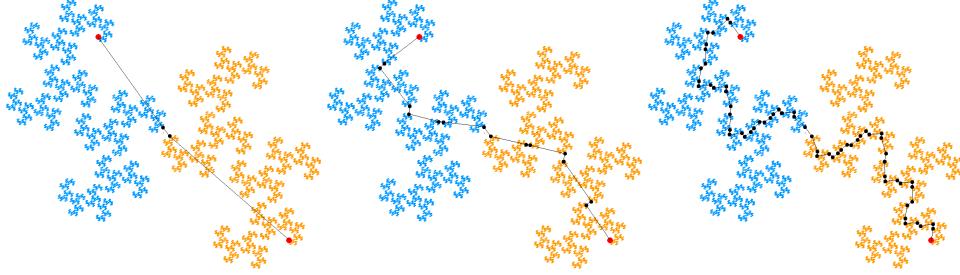


FIGURE 7. Creating a path between the red points which lies entirely within the $\delta/2$ neighborhood of Λ_z by recursively “jumping” across the pair of points in $f\Lambda_z$ and $g\Lambda_z$ which are closest, as explained in the Short Hop Lemma (Lemma 5.2.2).

Theorem 5.2.3 (Disconnected is Schottky). *The semigroup G_z has disconnected Λ_z if and only if G_z is Schottky.*

Proof. It remains to prove that if Λ_z is disconnected, then G_z admits a good disk. Since Λ_z is disconnected, by Lemma 5.2.1 the distance from $f\Lambda_z$ to $g\Lambda_z$ is some positive number δ . By Lemma 5.2.2 it follows that the closed $\delta/2$ neighborhood $\overline{N}_{\delta/2}(\Lambda_z)$ of Λ_z is path connected. So $\overline{N}_{|z|\delta/2}(f\Lambda_z)$ and $\overline{N}_{|z|\delta/2}(g\Lambda_z)$ are path connected.

Choose some ϵ with $|z|\epsilon < \delta/2 < \epsilon$. Let $L_n = \Sigma_n(z, 0)$. By Lemma 5.1.1, there is an n such that $\overline{N}_{\delta/2}(\Lambda_z) \subseteq \overline{N}_\epsilon(L_n)$. Let $E = \overline{N}_\epsilon(L_n)$. By definition, each connected component of E contains a point in Λ_z , but $\Lambda_z \subseteq \overline{N}_{\delta/2}(\Lambda_z)$, which is path connected, so there can only be a single connected component of E containing Λ_z , so E is path connected. Since E is a finite union of closed disks which is path connected, it is homeomorphic to a disk with finitely many subdisks removed.

Furthermore, fE and gE are unions of round disks of radius $|z|\epsilon$ around points of $f\Lambda_z$ and $g\Lambda_z$, and therefore are contained in $\overline{N}_{\delta/2}(\Lambda_z) \subseteq E$. Because $|z|\epsilon < \delta/2$, fE and gE are disjoint.

We now show that we can fill in the “holes” in E (if any) and obtain a good disk. By construction E has finitely many holes, so there is some hole of least diameter with boundary component γ . But then $f\gamma$ and $g\gamma$ have diameter strictly less than γ , and are contained in the interior of E , so that they must bound subdisks of E . So it follows that we can add to E the subdisk bounded by γ to obtain a new closed set E' with fE' and gE' disjoint and contained in the interior of E . Add the bounded complementary components in this way one by one until we obtain a closed topological disk D with fD and gD disjoint and contained in the interior of D . In other words, D is good for G_z , so that G_z is Schottky. \square

5.3. An algorithm to certify that Λ_z is disconnected. In this section, we describe a fast and practical algorithm to certify that the limit set Λ_z is disconnected for a given parameter z (equivalently, to certify that G_z is Schottky). Since this condition is open in z , a careful analysis of this algorithm certifies that Λ_z is disconnected on a definite open subset of parameter space. Giving a rigorous numerical certificate that Λ_z is *connected*, especially one valid in a definite open subset of parameter space, is more difficult, and is addressed in Section 7. However practically speaking, the algorithm described in this section can be used to draw fast and accurate pictures of \mathcal{M} . The algorithm we describe differs only in inessential ways from that first discussed by Bandt [1].

We give some notation. Let D_z be a round disk centered at $1/2$ with the property that fD_z and gD_z are both contained in D_z ; for example, we could take D_z to be a disk of radius $|z - 1|/2(1 - |z|)$. Let $D_n =: \Sigma_n(z, D_z)$; i.e. D_n is the union of the images of D_z under the set of words in Σ of length n . Inductively, $D_n = fD_{n-1} \cup gD_{n-1}$. By Lemma 5.1.2, $\Lambda_z = \cap_n D_n$, so Λ_z is disconnected if and only if D_n is disconnected for some n .

Lemma 5.3.1. D_n is disconnected if and only if $fD_{n-1} \cap gD_{n-1} = \emptyset$.

Proof. Obviously if $fD_{n-1} \cap gD_{n-1} = \emptyset$, then D_n is disconnected, so we must only show the converse. Suppose D_n is disconnected and $fD_{n-1} \cap gD_{n-1} \neq \emptyset$. We can take n to be minimal such that D_n is disconnected and $fD_{n-1} \cap gD_{n-1} \neq \emptyset$. Since n is minimal, for $n - 1$ we have either D_{n-1} is connected or $fD_{n-2} \cap gD_{n-2} = \emptyset$. The latter is impossible, though, because $D_{n-1} \subseteq D_{n-2}$, so $fD_{n-1} \cap gD_{n-1} \subseteq fD_{n-2} \cap gD_{n-2}$. We conclude that D_{n-1} is connected. But then fD_{n-1} and gD_{n-1} are connected, and $fD_{n-1} \cap gD_{n-1} \neq \emptyset$, so D_n is connected, a contradiction. \square

Naively, to check that fD_{n-1} is disjoint from gD_{n-1} would take exponential time, since we need to check the pairwise distances between elements of two sets, each with 2^{n-1} points. However, there is a great deal of redundancy: if u and v are words of length n starting with f and g respectively, then if $u(z, D_z)$ is disjoint from $v(z, D_z)$, then $ux(z, D_z)$ is disjoint from $vy(z, D_z)$ for all words x, y . In fact, for any fixed u and v words of length n , the images $u(z, D_n)$ and $v(z, D_n)$ are copies of D_n , scaled by z^n and translated relative to each other by $u(z, 1/2) - v(z, 1/2)$. Thus the relevant data to keep track of is the set of numbers $z^{-n}(u(z, 1/2) - v(z, 1/2))$ ranging over u, v of length n where u starts with f and v starts with g , for which $u(z, D_z)$ and $v(z, D_z)$ intersect — equivalently, for which there is an inequality $|z^{-n}(u(z, 1/2) - v(z, 1/2))| < R := 2 \text{ radius}(D_z)$.

This discussion justifies Algorithm 1 to test for connectedness of Λ_z . We briefly explain the recursion in the context of the above observations. First, the algorithm initializes the set V to contain the single number

$$z^{-1}(f(z, 1/2) - g(z, 1/2)) = z^{-1}(z/2 - (z(1/2 - 1) + 1)) = 1 - z^{-1}.$$

Next, recall that for any word u of length n , we can write $u(z, x) = z^n x + p_u(z)$, where $p_u(z)$ is a polynomial in z . Therefore, $z^{-n}(u(z, 1/2) - v(z, 1/2)) = z^{-n}(p_u(z) - p_v(z))$. So if we are given $\alpha = z^{-n}(u(z, 1/2) - v(z, 1/2)) = z^{-n}(p_u(z) - p_v(z))$, we can compute (for clarity, we write $u1/2$ in place of $u(z, 1/2)$):

$$\begin{aligned} z^{-(n+1)}(uf1/2 - vf1/2) &= z^{-1}z^{-n}(z^n z/2 + p_u(z) - z^n z/2 - p_v(z)) = z^{-1}\alpha \\ z^{-(n+1)}(ug1/2 - vg1/2) &= z^{-1}z^{-n}(z^n(1 - z/2) + p_u(z) - z^n(1 - z/2) - p_v(z)) = z^{-1}\alpha \\ z^{-(n+1)}(uf1/2 - vg1/2) &= z^{-1}z^{-n}(z^n z/2 + p_u(z) - z^n(1 - z/2) - p_v(z)) = z^{-1}(\alpha + z - 1) \\ z^{-(n+1)}(ug1/2 - vf1/2) &= z^{-1}z^{-n}(z^n(1 - z/2) + p_u(z) - z^n z/2 - p_v(z)) = z^{-1}(\alpha - z + 1) \end{aligned}$$

So given the set of differences of the form $z^{-n}(u(z, 1/2) - v(z, 1/2))$ which are less than R , where u and v may range over all words of length n , we can compute the set of differences of words of length $n + 1$, discarding those which are larger than R .

Algorithm 1 Disconnected(z , depth)

```

 $V \leftarrow \{1 - z^{-1}\}$ 
 $d \leftarrow 0$ 
while  $V \neq \emptyset$  or  $d < \text{depth}$  do
     $W \leftarrow \emptyset$ 
    for all  $\alpha \in V$  do
        if  $|z^{-1}\alpha| < R$  then  $W \leftarrow W \cup z^{-1}\alpha$ 
        if  $|z^{-1}(\alpha + z - 1)| < R$  then  $W \leftarrow W \cup z^{-1}(\alpha + z - 1)$ 
        if  $|z^{-1}(\alpha - z + 1)| < R$  then  $W \leftarrow W \cup z^{-1}(\alpha - z + 1)$ 
     $V \leftarrow W$ 
     $d \leftarrow d + 1$ 
if  $V = \emptyset$  then
    return true
else
    return false

```

If this algorithm returns true, then Λ_z is disconnected. If it returns false, then Λ_z might still be disconnected, but this would not be discovered without increasing the “depth” parameter.

Algorithm 1 is very fast, and has been implemented in our program `schottky`, available from [5]. In practice, we can check connectedness to depths exceeding 60. The algorithm is faster in certain regions than others; in particular, it is quite slow near the real axis. We follow this point up in Section 10.

5.4. Paths in Λ_z . In this section, we show how to construct paths inside the limit set Λ_z and show that it is connected if and only if it is path connected. We will not explicitly need the results in this section; however, it serves to further introduce the structure of Λ_z , and we will use very similar ideas in Section 11. These are not new results and can be derived from the general theory of IFSs, but this direct

approach is illuminating. Note that this is essentially a continuous version of the Short Hop Lemma 5.2.2.

The following construction essentially appears in [1]. Suppose that $f\Lambda_z \cap g\Lambda_z \neq \emptyset$, so there are $u, v \in \partial\Sigma$ with $u_1 = f$ and $v_1 = g$ and $\pi(u, z) = \pi(v, z)$. Then for any $a, b \in \partial\Sigma$, we will construct a continuous path within Λ_z between $\pi(a, z)$ and $\pi(b, z)$.

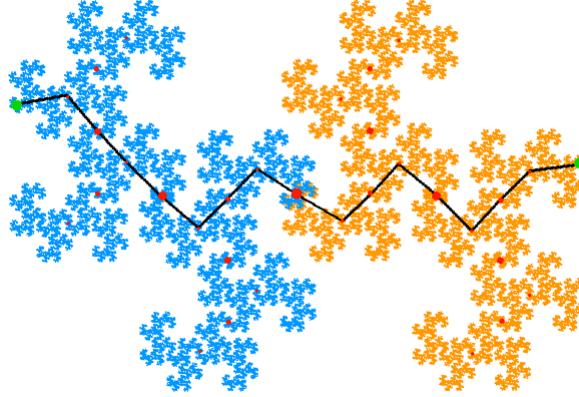


FIGURE 8. Given the words $x, y \in \partial G$ such that $\pi(z)(x) = \pi(z)(y)$, we show an approximation of a path in Λ_z between $\pi(z)(a)$ and $\pi(z)(b)$ for a given $a, b \in \partial G$. The large red disk in the middle indicates a point in the intersection $f\Lambda_z \cap g\Lambda_z$, and the other red disks show the image of this point under words of length less than or equal to 4. This method is completely analogous to Figure 7.

First, let us narrate Figure 8 to explain the construction graphically. Suppose that the IFS takes the disk of radius R inside itself. We will use this fact to bound distances between points. Suppose that $a_1 = f$ and $b_1 = g$. Note $\pi(a, z)$ and $\pi(b, z)$ are distance at most $2R$. Consider the pair of words (u, v) . By assumption Since $x_1 = f = a_1$ and $y_1 = g = b_1$, note that both $\pi(a, z)$ and $\pi(u, z) = \pi(v, z)$ lie in $f\Lambda_z$, so they are distance at most $2|z|R$ apart. Similarly, $\pi(b, z)$ and $\pi(u, z) = \pi(v, z)$ lie in $g\Lambda_z$, so they are also distance at most $2|z|R$ apart. That is, the point $\pi(u, z) = \pi(v, z)$ coarsely interpolates between $\pi(a, z)$ and $\pi(b, z)$. Next consider the words a and u . For illustrative purposes, suppose $a_2 = f$ and $u_2 = g$. Then the pair (fu, fv) interpolates between a and u : because fu agrees with a to depth 2, and fv agrees with v to depth 2, we have $\pi(fu, z) = \pi(fv, z)$, and

$$|\pi(a, z) - \pi(fu, z)| < 2|z|^2R \quad \text{and} \quad |\pi(fv, z) - \pi(u, z)| < 2|z|^2R$$

We can continue inductively producing points in Λ_z coarsely between points in our path. Figure 8 shows the coarse path which is the result of stopping the construction after 4 steps. In the limit, we will have a continuous path from the dyadic rationals into Λ_z , and because Λ_z is closed, this path extends continuously to $[0, 1]$.

We now describe the construction precisely. Given u, v as above, we define an interpolation function on infinite words $\Phi_{(u,v)} : \partial\Sigma \times \partial\Sigma \rightarrow \partial\Sigma \times \partial\Sigma$, as follows. Given two right-infinite words s, t , we may rewrite them as $s = ws'$ and $t = wt'$,

where w is the maximal common prefix of s and t . Then

$$\Phi_{(u,v)}(s,t) = \begin{cases} (wu,wv) & \text{if } s'_1 = f \text{ (and thus } t'_1 = g) \\ (wv,wu) & \text{if } s'_1 = g \text{ (and thus } t'_1 = f) \end{cases}$$

Note that $\pi(wu,z) = \pi(wv,z)$ by Lemma 3.1.8. Furthermore if the maximal common prefix of s,t has length n , i.e. $|w| = n$, then the maximal common prefix of s and $(\Phi_{(u,v)}(s,t))_1$ (denoting the first coordinate) is $n+1$; similarly, the maximal common prefix of t and $(\Phi_{(u,v)}(s,t))_2$ is also $n+1$.

Now we define a set $W \subseteq \partial\Sigma \times \partial\Sigma$. The set W will be indexed by the dyadic rational numbers, with W_r denoting the element at r ; that is $W_r = (W_{r,1}, W_{r,2})$. First set $W_0 = (a,a)$ and $W_1 = (b,b)$. Then recursively define

$$W_{k2^{-i-1} + (k+1)2^{-i-1}} = \Phi_{(u,v)}(W_{k2^{-i},2}, W_{(k+1)2^{-i},1})$$

In other words, to get the pair between $k2^{-i}$ and $(k+1)2^{-i}$, apply the interpolation function $\Phi_{(u,v)}$ to the second word at $k2^{-i}$ and the first word at $(k+1)2^{-i}$. Observe that $\pi(W_r, z)$ is well-defined because $\pi(\Phi_{(u,v)}(\cdot)_1, z) = \pi(\Phi_{(u,v)}(\cdot)_2, z)$.

Lemma 5.4.1. *Suppose that*

$$k2^{-i} \leq r_1, r_2 \leq (k+1)2^{-i}$$

Then for the construction above, we have the estimate

$$|\pi(W_{r_1}, z) - \pi(W_{r_2}, z)| < 2 \frac{|z|^{i+1}}{|1-z|}$$

Proof. By construction, $W_{k2^{-i},2}$ and $W_{(k+1)2^{-i},1}$ have a common prefix w of length i , and thus for all r with $k2^{-i} \leq r \leq (k+1)2^{-i}$, the pair of words comprising W_r also has the prefix w . Therefore, the difference $\pi(W_{r_1}, z) - \pi(W_{r_2}, z)$ is given by a power series in z with coefficients in $\{-2, 0, 2\}$ whose first nonzero coefficient has degree at least $i+1$. The estimate follows. \square

Proposition 5.4.2. *Suppose $u, v \in \partial\Sigma$ with $u_1 = f$ and $v_1 = g$ and z is such that $\pi(u, z) = \pi(v, z)$. Let $a, b \in \partial\Sigma$, and let Dy be the set of dyadic rational numbers. Using u, v, a, b as input, construct W as above, and define the map $w : \text{Dy} \rightarrow \mathbb{C}$ given by $w : r \mapsto \pi(W_r, z)$. Then w extends continuously to $[0, 1]$, and satisfies*

$$w(0) = a, \quad w(1) = b, \quad \text{and} \quad w([0, 1]) \subseteq \Lambda_z.$$

Hence, Λ_z is path connected iff it is connected iff $f\Lambda_z \cap g\Lambda_z \neq \emptyset$.

Proof. It follows from Lemma 5.4.1 that $w : \text{Dy} \rightarrow \mathbb{C}$ is continuous, and its image is contained in Λ_z , so w extends continuously as a map $[0, 1] \rightarrow \Lambda_z$.

To get the last assertion, observe that we have shown $(3) \Rightarrow (1)$; the other two implications $(1) \Rightarrow (2)$ and $(2) \Rightarrow (3)$ are obvious. \square

6. LIMIT SETS OF DIFFERENCES

6.1. Differences in Λ_z . As we saw in Section 5, the topology and geometry of Λ_z is controlled by the intersection $f\Lambda_z \cap g\Lambda_z$. In particular, $f\Lambda_z \cap g\Lambda_z \neq \emptyset$ if and only if Λ_z is connected. Said another way, the limit set is path connected if and only if 0 lies in the set of differences between points in Λ_z . It turns out that the set of differences between points in Λ_z is itself the limit set of an IFS. This limit set features prominently in [12] and in Section 9.

Define Γ_z to be the limit set of the IFS generated by the three functions

$$x \mapsto z(x+1) - 1 \quad x \mapsto zx \quad x \mapsto z(x-1) + 1$$

That is, the IFS generated by dilations by z centered at the points $-1, 0, 1$.

Lemma 6.1.1. *We have $\Gamma_z = \{a - b \mid a, b \in \Lambda_z\}$.*

Proof. Each point in Λ_z is given by a power series in z and associated with an infinite word in $\partial\Sigma$. Thus, the set of differences of points in Λ_z is associated with pairs of infinite words. Given two words $x, y \in \partial\Sigma$, it is straightforward to compute the power series giving the value $\pi(x, z) - \pi(y, z)$ recursively, as follows. Suppose that x begins with f , so $x = fx'$ and y begins with g , so $y = gy'$. Then

$$\pi(x, z) - \pi(y, z) = f(\pi(x', z)) - g(\pi(y', z)) = z(\pi(x', z) - \pi(y', z) + 1) - 1$$

In other words, the difference associated to the pair of words fx' and gy' is obtained from the difference associated to the pair of words x' and y' by the transformation $d \mapsto z(d+1) - 1$. Similarly, prefixing the pair of words (x', y') with the pair of letters $(f, f), (g, g), (g, f)$ transforms the differences by

$$d \mapsto zd, \quad d \mapsto zd, \quad z \mapsto z(d-1) + 1,$$

respectively. Note that two of these transformations are the same. Therefore, the limit set of the semigroup generated by these three transformations is precisely the set of differences in Λ_z . But that limit set is Γ_z . \square

Notice that the proof of Lemma 6.1.1 effectively shows that the set of differences of any IFS generated from a regular language as in Section 4 is itself an IFS generated from a regular language.

6.1.1. Iterated Mandelbrot sets. The set of differences between points in Λ_z is Γ_z . We can iterate this procedure by taking the set of differences in Γ_z , and so on. Let Γ_z^k be the limit set of the IFS generated by $\{f_{-k}, \dots, f_k\}$, where f_i is a dilation centered at i . The set of differences of points in Λ_z is Γ_z^1 .

Lemma 6.1.2. *The set of differences between points in $\Gamma_z^{2^k}$ is $\Gamma_z^{2^{k+1}}$.*

Proof. This is just a computation in the generators analogous to the proof of Lemma 6.1.1. We note $(z(x-m)+m)-(z(y-n)+n) = z((x-y)-(m-n))+(m-n)$, so acting by the dilations at m and n on two points acts on their difference by the dilation at $m-n$. \square

If we define

$$\mathcal{M}^k := \{z \in \mathbb{D}^* \mid \Gamma_z^k \text{ is connected}\}$$

and

$$\mathcal{M}_0^k := \{z \in \mathbb{D}^* \mid 0 \in f_1 \Gamma_z^k\},$$

then

Proposition 6.1.3. $\mathcal{M}^{2^k} = \mathcal{M}_0^{2^{k+1}}$.

Proof. To see that $\mathcal{M}_0^{2^{k+1}} \subseteq \mathcal{M}^{2^k}$, suppose that $0 \in f_1 \Gamma_z^{2^{k+1}}$, so there is a pair of generators f_n, f_{n+1} of $\Gamma_z^{2^k}$ such that $f_n \Gamma_z^{2^k} \cap f_{n+1} \Gamma_z^{2^k} \neq \emptyset$, and thus this holds for all n , so the limit set $\Gamma_z^{2^k}$ is connected. Conversely, if $\Gamma_z^{2^k}$ is connected, then $f_{2^k} \Gamma_z^{2^k}$ intersects $f_j \Gamma_z^{2^k}$ for some j . But the images $f_j \Gamma_z^{2^k}$ are translates of the same path

connected set by multiples of the same vector, so it must be that $f_{2^k}\Gamma_z^{2^k}$ intersects the translate which is closest, i.e. $f_{2^k}\Gamma_z^{2^k} \cap f_{2^k-1}\Gamma_z^{2^k} \neq \emptyset$. Hence $0 \in f_1\Gamma_z^{2^{k+1}}$. \square

Remark 6.1.4. Note that the last step in the proof of Proposition 6.1.3 is essentially the same as Bousch's proof of Proposition 4.1.3.

In general, if our IFS is a set of dilations by z at points $\{c_1, \dots, c_k\}$, then the IFS which generates the differences in our IFS is the set of dilations by z with centers at all differences of the c_i . The fact that f_1 appears in the definition of \mathcal{M}_0^k (as opposed to f_i for another i) is natural because the number 1 is always the generator of the lattice of centers.

Question 6.1.5. *What sequences of sets arise as iterated differences? What properties do these iterated IFS have?*

7. INTERIOR POINTS IN \mathcal{M}

We have already seen that \mathcal{M} contains many interior points; in fact, the entire annulus $1/\sqrt{2} \leq |z| \leq 1$ is in \mathcal{M} . In this section we develop the method of *traps* to certify the existence of many interior points in \mathcal{M} , and examine the closure of the set of interior points. The result is quite surprising: the closure of the interior is all of \mathcal{M} ... except for some subset of the two real whiskers!

This assertion is Theorem 7.2.7 below, which is the affirmation of Bandt's Conjecture (i.e. Conjecture 2.6.3). In Section 8 these techniques are used to certify the existence of (infinitely many) small holes in \mathcal{M} — i.e. exotic components of Schottky space.

7.1. Short hop paths and Traps. In this subsection we give a method to certify the existence of open subsets of \mathcal{M} . Abstractly, to certify that z is an interior point of \mathcal{M} is to give a proof that $z \in \mathcal{M}$ that depends on properties of z which are stable under perturbation. Showing that $z \in \mathcal{M}$ is equivalent to showing that $f\Lambda_z$ intersects $g\Lambda_z$, so our strategy is to show that this intersection is inevitable for some *topological reason* (depending on z). Proving that sets intersect in topology is accomplished by homology (or, more crudely, separation or linking properties). But the homological properties of Λ_z depend on its connectivity, which is what we are trying to establish! So our method is first to consider precisely chosen *neighborhoods* of Λ_z (which may be presumed to be connected for some open set of z), and then to consider homological properties of the configuration of the images of these neighborhoods under f and g which force an intersection.

The key to our method is the existence of *short hop paths* and *traps*.

Definition 7.1.1 (Short hop path). Let $p, q \in \Lambda_z$, let $\epsilon > 0$ and let D be a disk containing p and q . An (ϵ, D) -*short hop path* from p to q is a sequence e_0, e_1, \dots, e_m in $\partial\Sigma$ with $\pi(e_0, z) = p$ and $\pi(e_m, z) = q$ so that $d(\pi(e_i, z), \pi(e_{i+1}, z)) < \epsilon$ and $\pi(e_i, z) \in D$ for all i .

The existence of Short Hop Paths is guaranteed by the Short Hop Lemma; in particular, we have:

Proposition 7.1.2 (Short hop paths exist). *Let u and v be right-infinite words with a common prefix w of length n , and suppose that there are points in $f\Lambda_z$ and*

$g\Lambda_z$ which are distance at most δ apart. Let D be a disk containing the $|z|^n \delta / 2$ -neighborhood of $w\Lambda_z$. Then there is a $(|z|^n \delta, D)$ -short hop path from $\pi(u, z)$ to $\pi(v, z)$.

Proof. Let $u = wu'$ and $v = wv'$, and let D' be any disk containing the $\delta/2$ -neighborhood of Λ_z . By Lemma 5.2.2 there is a (δ, D') -short hop path from $\pi(u', z)$ to $\pi(v', z)$. Now apply w to this short hop path. \square

We now give the definition of a trap:

Definition 7.1.3 (Trap). Let D be a closed topological disk containing Λ_z in its interior. We say that a pair of words $u, v \in \Sigma$ are a *trap* for (z, D) if the following are true:

- (1) u starts with f and v starts with g ;
- (2) there are points p^\pm in $u\Lambda_z - vD$ and q^\pm in $v\Lambda_z - uD$ such that for some paths $\alpha \subseteq uD$ with endpoints p^\pm and $\beta \subseteq vD$ with endpoints q^\pm the algebraic intersection number of α and β is nonzero; and
- (3) there are points in $f\Lambda_z$ and $g\Lambda_z$ within distance ϵ of each other, where the $\epsilon/2$ neighborhood of Λ_z is contained in D .

The definition of a trap depends on a choice of paths α and β which intersect; but a homological argument shows that the property does not depend on the choice:

Lemma 7.1.4 (Any paths suffice). Suppose u, v are a trap for (z, D) , and let $p^\pm \in u\Lambda_z - vD$ and $q^\pm \in v\Lambda_z - uD$ be as in Definition 7.1.3. Then any paths $\alpha \subseteq uD$ with endpoints p^\pm and $\beta \subseteq vD$ with endpoints q^\pm must intersect.

Proof. Any two paths α, α' joining p^\pm and contained in uD are freely homotopic relative to endpoints in the complement of q^\pm , and similarly for any two β, β' joining q^\pm and contained in vD . Thus the classes $[\alpha] \in H_1(\mathbb{C} - q^\pm, p^\pm)$ and $[\beta] \in H_1(\mathbb{C} - p^\pm, q^\pm)$ are well-defined, and therefore so is their intersection product. \square

Example 7.1.5. Figure 9 shows a trap in Λ_z with $z = -0.43 + 0.54i$ which is visible to the naked eye. We have drawn D_{12} for a disk D with $fD, gD \subseteq D$, so it is guaranteed that (1) $f\Lambda_z$ and $g\Lambda_z$ are contained inside the blue and orange sets, respectively and (2) there are points in Λ_z inside every disk drawn. The computer also runs Algorithm 1 to verify that D_{12} is connected. These facts, and the (visually evident) fact that the highlighted disks satisfy the linking condition, proves the existence of points p^\pm and q^\pm inside the disks which give a trap.

The next Proposition shows that the existence of a trap for z shows that z is in the interior of \mathcal{M} .

Proposition 7.1.6 (Traps in \mathcal{M}). Let u, v be a trap for (z, D) for some disk D . Then z is in the interior of \mathcal{M} .

Proof. Suppose that δ is the distance from $f\Lambda_z$ to $g\Lambda_z$. Then any two points in Λ_z can be joined by a (δ, D) -short hop path. This is a sequence of points with gaps of size at most δ ; such a sequence is necessarily contained in the $\delta/2$ -neighborhood of Λ_z . It follows that p^+ can be joined to p^- by a path α in uD , every point on which is within distance $\delta|z|^n/2$ of some point in $u\Lambda_z$, where u has length n . Similarly, q^+ can be joined to q^- by a path β in vD , every point on which is within distance $\delta|z|^m/2$ of some point in $v\Lambda_z$, where v has length m . But α and β must intersect, by the defining property of a trap, and Lemma 7.1.4. Thus the distance from $u\Lambda_z$

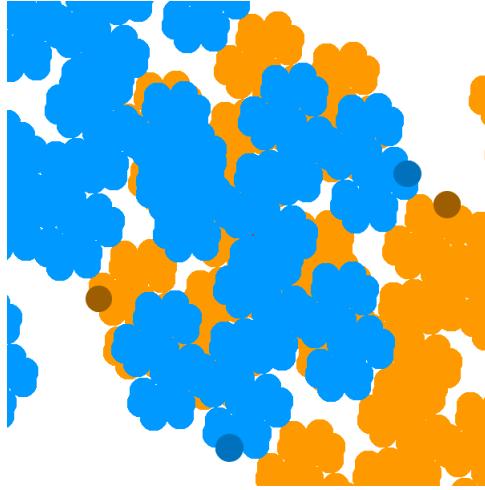


FIGURE 9. An excerpt from Λ with $z = -0.43 + 0.54i$. This picture proves the existence of a trap for this parameter, as explained in Example 7.1.5.

to $v\Lambda_z$ is at most $\delta|z|^{\min(n,m)}$. But $\delta|z|^{\min(n,m)} < \delta$ because $n, m \geq 1$. This is contrary to the definition of δ unless $\delta = 0$. \square

It is interesting to note that while the existence of a trap for z certifies that z is in the interior of \mathcal{M} and thus that $f\Lambda_z \cap g\Lambda_z \neq \emptyset$, it is difficult to use it to algorithmically produce a point of intersection: as we decrease δ , the intersecting δ -short hop paths need not converge or intersect “nicely”.

7.2. Traps are (almost!) dense. In this subsection we demonstrate the theoretical utility of traps by proving that traps are dense in \mathcal{M} away from the real axis. Since traps have nonempty interior, it follows that the interior of \mathcal{M} is dense in \mathcal{M} , again away from the real axis. This was conjectured by Bandt in [1], p. 7 and some partial results were obtained by Solomyak-Xu [13], who proved the conjecture for points in a neighborhood of the imaginary axis.

It is interesting that the proof depends on a complete analysis of the set of z for which the limit set Λ_z is convex (Lemma 7.2.3). It turns out that the z with this property are exactly the union of dyadic “spikes” — points of the form $re^{\pi ip/q}$ for coprime integers p, q and r real with $r \geq 2^{-1/q}$. For $q > 1$ these spikes are already in the interior of the solid annulus $r \geq 2^{-1/2}$ which is entirely contained in \mathcal{M} ; only the real “whiskers” protrude from this annulus, and this is why these are the only points in \mathcal{M} which are not in the closure of the interior.

Definition 7.2.1 (Cell-like, trap-like). A compact connected subset $X \subseteq \mathbb{C}$ is *cell-like* if its complement is connected. Let X be cell-like. A complex number w is *trap-like* for X if the following hold:

- (1) the union $X \cup (X + w)$ is connected (equivalently, X intersects $X + w$); and
- (2) there are 4 points in the outermost boundary of $X \cup (X + w)$ that alternate between points in $X - (X + w)$ and points in $(X + w) - X$.

Lemma 7.2.2 (Nonconvex cell has trap). *Let X be cell-like. There there is some trap-like vector w for X if and only if X is nonconvex.*

Proof. If X is convex, then the set of points in the boundary of $X \cup X + w$ in $X - (X + w)$ is connected, and similarly for those points in $(X + w) - X$, so no w is trap-like.

Conversely, suppose X is nonconvex, and let ℓ be a supporting line such that $\ell \cap X = P \cup Q$ both nonempty (not necessarily connected), and separated by an open interval I . The existence of such P and Q is guaranteed precisely by the hypothesis that X is not convex. After composing with an isometry of the plane, we can assume that ℓ is the horizontal axis, oriented positively, so that X is on the side of ℓ with negative y coordinates.

Let V be a small open disk in $\mathbb{C} - X$ containing the midpoint of I , and choose $v \in V - \ell$ on the side of ℓ with negative coordinates (i.e. the side containing X). Let p denote the point of P with biggest x coordinate. Then $w = v - p$ is trap-like for X . See Figure 10.

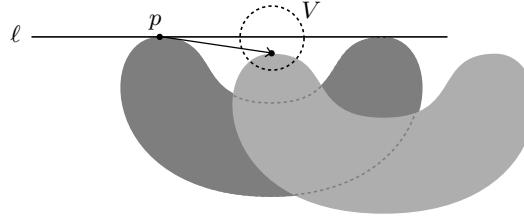


FIGURE 10. A nonconvex compact full set has a trap-like vector.
See the proof of Lemma 7.2.2.

This can be seen just by looking at the foliation of \mathbb{C} by vertical lines $\mu(t)$ with x -coordinate t , and for each line finding the point of $X \cup (X + w)$ with largest y -coordinate (where this is nonempty). Let $q \in Q$ be arbitrary, let t be the maximum number such that $\mu(t) \cap X$ is nonempty, and let r be the point on $\mu(t) \cap X$ with largest y -coordinate. Then the four points $p, v, q, r + w$ are the highest points of $X \cup (X + w)$ on their respective vertical lines $\mu(t_1), \mu(t_2), \mu(t_3), \mu(t_4)$ for $t_1 < t_2 < t_3 < t_4$, and alternate between the sets X and $X + w$. \square

We would like to apply Lemma 7.2.2 to the cell-like set X_z one obtains from a limit set Λ_z . Thus, it is important to characterize z for which the cell-like set X_z obtained from Λ_z is convex.

Lemma 7.2.3 (Convex polygon). *Let z be in \mathcal{M} , and let X_z be obtained from Λ_z by filling in bounded complementary components, so that X_z is the smallest cell-like set containing Λ_z . Then X_z is convex if and only if $z = re^{\pi ip/q}$ for coprime integers p, q and r real with $r \geq 2^{-1/q}$, in which case $X_z = \Lambda_z$ is a convex polygon.*

Proof. We make use of the following two facts: first, that X_z has rotational symmetry of order 2 about the point $1/2$; and second, that Λ_z is the union of $f\Lambda_z$ and $g\Lambda_z$, obtained from Λ_z by scaling by z and translated relative to each other by $1 - z$. Suppose X_z is convex, and consider the collection of straight segments in the boundary of X_z . This collection is nonempty; for, if p is an extremal point

for X_z tangent to the supporting line in the direction $(1 - z)/z$ then fp and gp are extremal points for X_z tangent to the same supporting line in the direction $(1 - z)$, and then the entire segment between these points is in the line. Now, if σ is a straight segment in the boundary in the direction w , then if $w \neq 1 - z$, there is a straight segment in the boundary of the form $f^{-1}\sigma$ or $g^{-1}\sigma$ of length $|z|^{-1}\sigma$. It follows that there is a chain of straight segments

$$\sigma_0, \sigma_1, \sigma_2, \dots, \sigma_{q-1}$$

where each σ_j is in a direction z^j relative to σ_0 , and has length $|z|^j|\sigma_0|$. But then $f(\sigma_{q-1})$ and $g(\sigma_{q-1})$ must be in the $1 - z$ or $z - 1$ direction, so that their union is either equal to σ_0 or the image of σ_0 under the symmetry of order 2. It follows that the argument of z is of the form $\pi p/q$ for some integers p/q , and furthermore that $|z|^q \geq 2$. This proves one direction of the claim.

The converse direction — that limit sets Λ_z for z of this kind really are convex — can be seen directly. In fact, these limit sets are zonohedra, the shadows of a linear semigroup acting in high dimensional space. Let R_q be the parallelepiped in \mathbb{R}^q consisting of vectors $v := (v_0, \dots, v_{q-1})$ whose coordinates satisfy $0 \leq v_{pj} \leq r^j$ with indices taken mod q . Note this is simply a rectangular box inside the positive orthant with one corner at the origin and edges along the coordinate axes. Let $f : \mathbb{R}^q \rightarrow \mathbb{R}^q$ be the composition $f : v \mapsto \sigma^p(rv)$ where rv means multiply the coordinates of v by r , and σ is the finite order rotation $\sigma : v \mapsto (v_{q-1}, v_0, \dots, v_{q-2})$. So f rotates and scales the box R_q to another box along the coordinate axes. Similarly, let $g : v \mapsto \sigma^p(rv) + t$ where t is the vector $(t_0, 0, 0, \dots, 0)$ for which $r^q + t_0 = 1$. The map g acts in the same way as f , except it translates the box up along the first coordinate by t_0 . The height of the box in the first coordinate is 1, and the heights of the acted-upon boxes fR_q and gR_q in the first coordinate are both r^q . Hence, providing $r^q \geq 1/2$, $fR_q \cup gR_q = R_q$, so the parallelepiped R_q is the limit set of the contracting semigroup $\langle f, g \rangle$. See Figure 11.

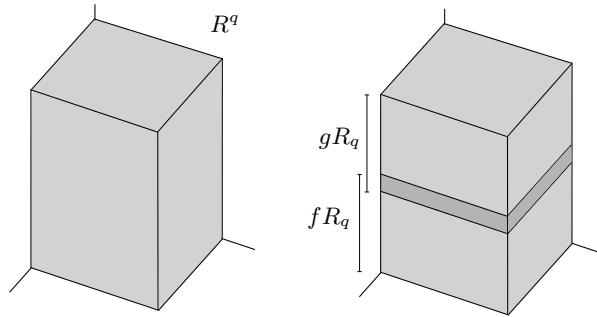


FIGURE 11. The box R_q in the proof of Lemma 7.2.3. The projection of R_q to the plane is the limit set Λ_z , so Λ_z is convex and, in particular, a zonohedron.

Projecting R_q to the plane so that the vectors $(0, 0, \dots, 1, \dots, 0)$ are projected to the $2q$ th roots of unity defines a semiconjugacy from this semigroup to G_z where $z = re^{\pi ip/q}$, taking R_q to Λ_z . \square

Example 7.2.4 (Hexagonal limit set). Take $z = 2^{-1/3}e^{2\pi i/3} \approx 0.396157 + 0.687364i$ then Λ is a hexagon with angles 120° , and side lengths in the ratio $1 : 2^{1/3} : 2^{2/3}$. See Figure 12.

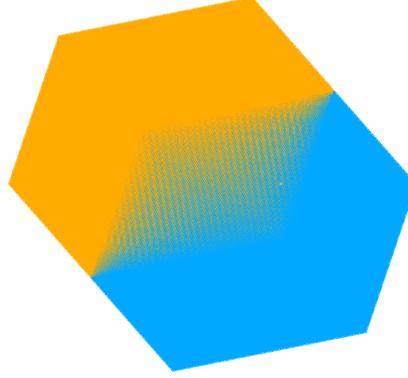


FIGURE 12. A hexagonal limit set for $z = 2^{-1/3}e^{2\pi i/3}$.

Lemma 7.2.5 (Surjective perturbation). *Let z_0 be in \mathcal{M} , and let $u, v \in \partial\Sigma$ be such that $\pi(u, z_0) = \pi(v, z_0)$. Then for any $\epsilon > 0$ there is $\delta > 0$ and integer M so that if $u_m, v_m \in \Sigma$ denote the prefixes of length m for any $m \geq M$, and T_m denotes the map*

$$T_m : z \mapsto u_m(z, 1/2) - v_m(z, 1/2)$$

then for any complex w with $|w| < \delta$ there is z_1 with $|z_1 - z_0| < \epsilon$, and $T_m(z_1) = w$.

Proof. The functions T_m converge uniformly to the limit $T_\infty : z \mapsto \pi(u, z) - \pi(v, z)$, which is holomorphic in z . Moreover, this limit could be constant only if $u = f^\infty$ and $v = g^\infty$, in which case $z \mapsto \pi(f^\infty, z)$ is identically 0 and $z \mapsto \pi(g^\infty, z)$ is identically 1, so $T_\infty(z) \equiv -1$; however, $T_\infty(z_0) = 0$. The limit is therefore nonconstant. Thus T_∞ takes the ball of radius ϵ about z_0 to a set containing the ball of radius 2δ about 0 for some positive δ , and the conclusion of the lemma is satisfied for sufficiently big m . \square

Corollary 7.2.6. *Suppose $u, v \in \partial\Sigma$ and $z_0 \in \mathcal{M}$ such that $\pi(u, z_0) = \pi(v, z_0)$. Then for any complex number w , and any positive ϵ , we can find an m and a z_1 with $|z_1 - z_0| < \epsilon$ so that*

$$z_1^{-m} (u_m(z_1, 1/2) - v_m(z_1, 1/2)) = w.$$

Proof. For sufficiently large m , the map $z \mapsto u_m(z, 1/2) - v_m(z, 1/2)$ is surjective onto a neighborhood of 0, and the claim follows. \square

We now complete the proof of Bandt's conjecture:

Theorem 7.2.7 (Interior is almost dense). *The set of interior points is dense in \mathcal{M} away from the real axis; that is*

$$\mathcal{M} = \overline{\text{int}(\mathcal{M})} \cup (\mathcal{M} \cap \mathbb{R}).$$

Proof. Let z_0 be in \mathcal{M} , and suppose the limit set Λ_{z_0} is not convex. Let X_{z_0} be the region bounded by Λ_{z_0} , so that X_{z_0} is cell-like. Since Λ_{z_0} is not convex, neither is X_{z_0} , and by Lemma 7.2.2 there is some w which is trap-like for X_{z_0} . Let $p_1, p_2 \in X_{z_0}$, and let $q_1, q_2 \in X_{z_0} + w$ be the four points from part (2) in Definition 7.2.1. Since $\partial X_{z_0} \subseteq \Lambda_{z_0}$, the points p_i, q_i lie in Λ_{z_0} . There is an ϵ so that the closed ϵ -neighborhood of X_{z_0} is connected, $p_1, p_2 \in \Lambda_{z_0} - \overline{N}_\epsilon(X_{z_0} + w)$, and $q_1, q_2 \in (\Lambda_{z_0} + w) - \overline{N}_\epsilon(X_{z_0})$. Furthermore, these conditions are open, so there is a $\delta > 0$ such that they hold for X_z for all z with $|z - z_0| < \delta$.

Now, since z_0 is in \mathcal{M} , there are $u, v \in \partial\Sigma$ starting with f and g respectively with $\pi(u, z_0) = \pi(v, z_0)$. By Corollary 7.2.6, we can find some u_m, v_m prefixes of u and v of length m , and z_1 with $|z_1 - z_0| < \delta$ so that $z_1^{-m}(u_m(z_1, 1/2) - v_m(z_1, 1/2)) = w$. We obtain a trap for (z_1, D) where $D = u_m(z_1)^{-1}(\overline{N}_\epsilon(X_{z_1}))$. This follows from the three conditions above.

We therefore find interior points of \mathcal{M} within distance δ of z_0 . Since z_0 was arbitrary, we are done in the case that Λ_{z_0} is not convex.

If $\Lambda = \Lambda_z$ is convex and $|z| < 2^{-1/2}$ then z is totally real, by Lemma 7.2.3. If $|z| > 2^{-1/2}$ then we are already in the interior, by Lemma 3.2.1. This completes the proof. \square

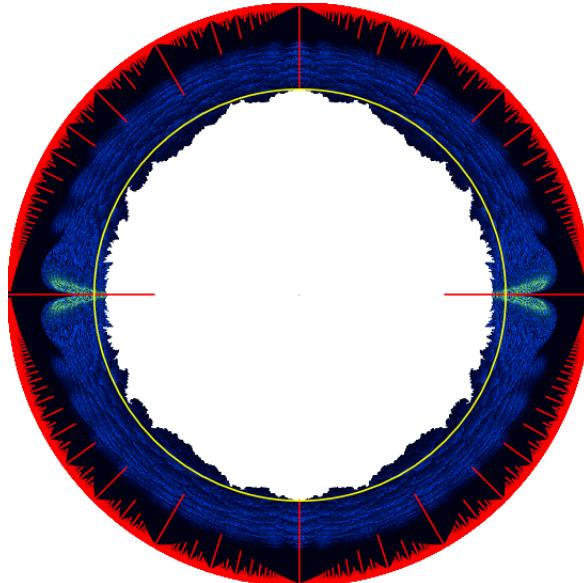


FIGURE 13. The set of z with Λ_z convex (in red) overlaid on \mathcal{M} . The yellow circle indicates $|z| = 2^{-1/2}$.

Figure 13 shows the set of z with convex Λ_z overlaid on \mathcal{M} . The picture of \mathcal{M} in a neighborhood of the real axis is surprisingly complicated; partial progress in understanding it was made by Shmerkin-Solomyak [10]; we describe some of their results in Section 10.1, and explain how the method of traps can be modified to certify the existence of interior points in $\overline{\mathcal{M}} - \overline{\mathbb{R}} \cap \mathbb{R}$.

8. HOLES IN \mathcal{M}

In this section we rigorously certify the existence of holes in \mathcal{M} (i.e. exotic components of Schottky space). Holes in \mathcal{M} were first observed experimentally by Barnsley and Harrington [2], and the existence of one hole was rigorously proved by Bandt [1]. However, our technique is quite different from Bandt's and our proof of the existence of holes is new. Furthermore, we shall show in Section 9 that our techniques generalize to prove the existence of *infinitely many* holes in \mathcal{M} .

8.1. An example. In this section, we give an example of an apparent hole in \mathcal{M} , an intuitive explanation of why the hole is truly an exotic component of Schottky space, and the output of our program rigorously certifying the hole. In the next section, we give a careful justification of the algorithm.

Figure 14 depicts an apparent collection of holes in \mathcal{M} centered at $0.459650 + 0.459654i$. The diameter of the large hole is approximately 0.000002.

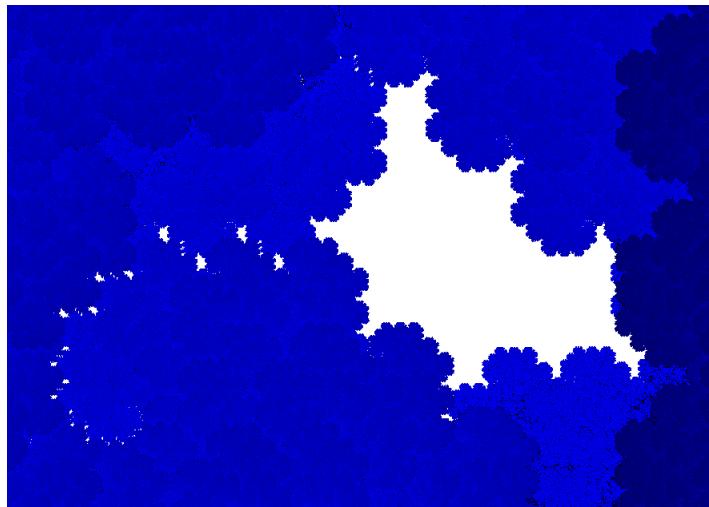


FIGURE 14. Apparent holes in \mathcal{M} centered at $0.459650 + 0.459654i$.

The limit set corresponding to a parameter inside the large hole is shown in Figure 15. The sets $f\Lambda_z$ and $g\Lambda_z$ are indeed disjoint, but they come very close. If one imagines that $f\Lambda_z$ and $g\Lambda_z$ are rigid, connected objects, then it is clear that one cannot unlink them by a rigid motion without the two sets intersecting at some intermediate step. However, movement in parameter space does not produce exactly rigid motion of the limit set, so in order to prove that this “hole” in \mathcal{M} is not, in fact, part of the large component of Schottky space, we need a more careful analysis.

Recall that the existence of a trap for parameter z is an open condition — there is some $\delta > 0$ so that a trap for z persists in $B_\delta(z)$. We call this a *ball of traps*. Our program certifies a putative hole in \mathcal{M} by producing overlapping balls of traps along a closed path encircling the hole, as shown in Figure 16. This proves that the closed path lies completely inside the interior of \mathcal{M} . Some technical remarks are in order. First, to complete the proof of the existence of a hole, we must certify some parameter z on the inside of the loop as Schottky. But since the connectedness

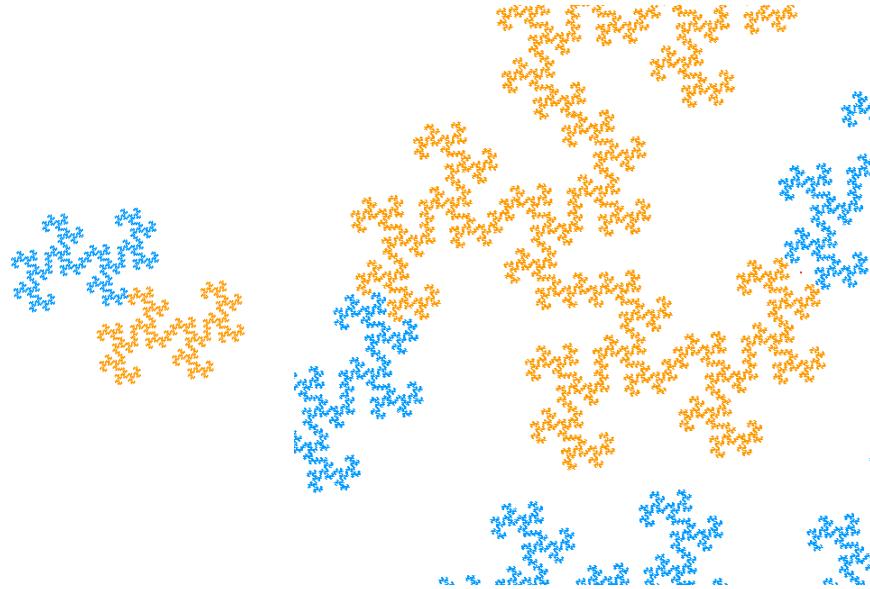


FIGURE 15. The limit set for a parameter inside the large hole shown in Figure 14, left, and a zoomed view, right. The two components $f\Lambda_z$ and $g\Lambda_z$ (blue and orange, respectively) cannot be unlinked with a rigid motion without intersecting.

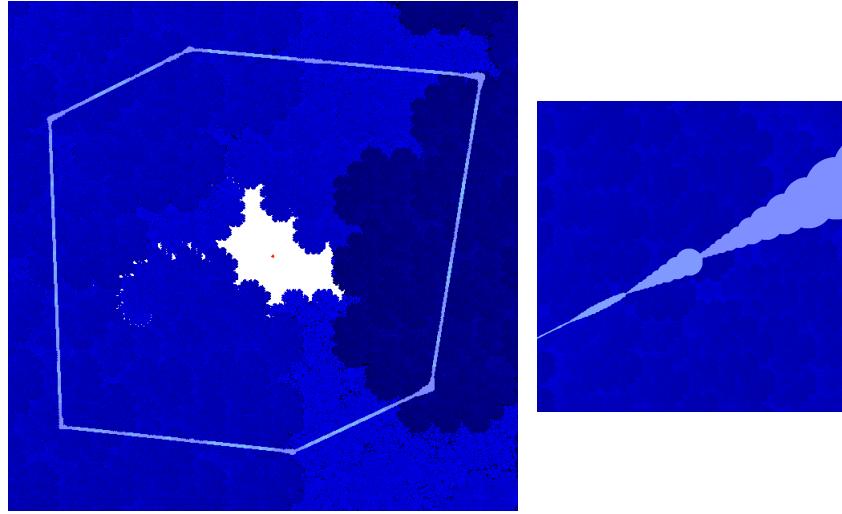


FIGURE 16. A loop of traps encircling the holes from Figure 14. A zoomed view of part of the loop shows how the program overlaps rigorous trap balls to produce a path inside the interior of \mathcal{M} .

of a pixel in Figure 16 is decided using Algorithm 1 applied to some parameter inside that pixel, a white pixel is guaranteed to contain some parameter which is Schottky, so this step is complete. Also, we note that the loop of trap balls in

Figure 16 appears to encircle many separate holes, but the output of this particular run of the program says nothing about whether these holes are actually distinct. We would need to run the program separately on loops encircling each of the holes we wished to rigorously separate. In Section 9, we extend our algorithm to prove the existence of infinitely many holes.

8.2. Numerical trap finding and loop certification. In this section, we describe our trap-finding algorithm in detail, including various numerical estimates. The details are important, because the output of a particular run of this algorithm serves as a rigorous certificate that there are multiple connected components of Schottky space, or equivalently, holes in \mathcal{M} .

The program will typically be required to produce a sequence of trap balls along a loop. Thus, we will be interested in finding a large number of traps in a given small region of parameter space. The algorithm takes advantage of this by separating the work into two pieces: a more computationally intensive piece of one-time work to find *trap-like balls*, and a fast check to produce a single ball of traps. Note that a trap-like ball (of vectors) and a ball of traps are not the same thing.

8.2.1. Finding trap-like balls.

Remark 8.2.1. This section is full of messy definitions and computations. These are necessary because we are in search of trap-like vectors similar to those found in Definition 7.2.1 but which work for *all* z in a given region, so we need to carefully estimate how the limit set changes as we change z . As a reward for this tedium, we get to compute these trap-like vectors (which is hard) only once, but we get to use them over an entire region.

Suppose that we will be searching for traps in a square region $B \subseteq \mathbb{D}^*$ of parameter space centered at z_0 and with side length $2d$. Let n , the *hull depth*, be given. Let $r_z = |z - 1|/2(1 - |z|)$; this is the minimal radius such that a disk of radius r_z centered at $1/2$ is mapped inside itself under both f and g . Let $D(p)$ denote the disk of radius p centered at $1/2$. Typically, we compute $\Sigma_n(z, D(r_z))$, the union of images of $D(r_z)$ under all words of length n , to study Λ_z . However, we need to control $\Sigma_n(z, D(r_z))$ over all $z \in B$, so we need to understand how it changes as we vary z . For this, we need some constants. The reader might consult Lemma 8.2.2 and Figure 17 for motivation before working through the technical details.

- (1) Let K be an upper bound for r_z in B . We can assume $|z| \leq 1/\sqrt{2}$, so the value $K = 2.92 > \sup_{|z|=1/\sqrt{2}} r_z$ will always work.
- (2) Let C be such that for any word $u \in \partial\Sigma$ and $z \in B$, we have $|\pi(u, z_0) - \pi(u, z)| < C|z_0 - z|$. Since u can be expressed as a power series in z with coefficients in $\{0, \pm 1\}$, an upper bound for the derivative in terms of z is given by $\sum_{i=1}^{\infty} i|z|^{i-1} = 1/(|z| - 1)^2$, so a valid value of C is given by $\sup_{z \in B} 1/(|z| - 1)^2$. As previously mentioned, we can assume that B lies within the disk of radius $1/\sqrt{2}$ by Lemma 3.2.1, so the uniform value of $C = 11.67$ will always work.
- (3) Let A be an upper bound for $|z|/|z_0|$ over B . Because $1/2 < |z|, |z_0| < 1/\sqrt{2}$, we have $|z|/|z_0| < \sqrt{2}$. For the previous two constants, a uniform upper bound like this is acceptable. In this case, though, we will be raising A to a large power, so it is critical to make A as close to 1 as possible.

Next set:

$$R_{z_0} = A^n K + 4K + 3|z_0|^{-n} C \sqrt{2} d$$

and for $z \in B$, define

$$R_z = \frac{|z_0|^n}{|z|^n} R_{z_0} - |z|^{-n} C |z - z_0|.$$

Lemma 8.2.2. *Suppose that $\Sigma_n(z_0, D(R_{z_0}))$ is connected. Then for any $z \in B$, we have*

- (1) $\Sigma_n(z, D(R_z)) \subseteq \Sigma_n(z_0, D(R_{z_0}))$.
- (2) $\Sigma_n(z, D(R_z))$ contains an ϵ -neighborhood of Λ_z for some ϵ such that there are two points $p_1 \in f\Lambda_z$, $p_2 \in g\Lambda_z$ such that $|p_1 - p_2| < \epsilon$.

Note that Algorithm 1 shows that (2) implies $\Sigma_n(z, D(R_z))$ is connected.

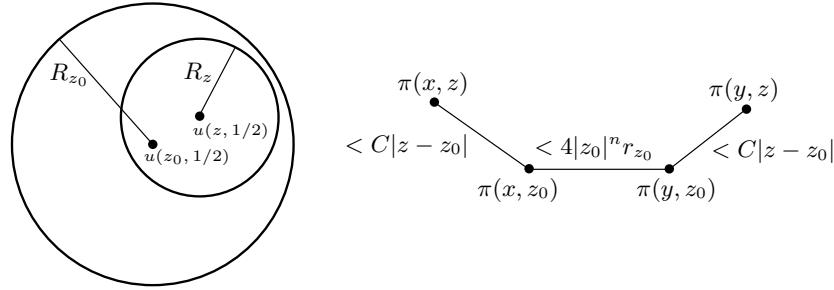


FIGURE 17. The proof of Lemma 8.2.2 just verifies that when we change the parameter z_0 to z , each disk $u(z, D(R_z))$ lies inside $u(z_0, D(R_{z_0}))$ and Λ_z still contains points that are close together. The figure on the right shows that we can prove (2) by proving that for each word u , $u(z, D(R_z))$ contains a $(4|z_0|^n r_{z_0} + 2C|z - z_0|)$ -neighborhood of $u(z, D(r_z))$. In the figure, $u \in \Sigma_n$ and $x, y \in \partial\Sigma$.

Proof. The set $\Sigma_n(z, D(R_z))$ is the union of disks of radius $|z|^n R_z$ centered at the images $u(z, 1/2)$ over all words u of length n , and the set $\Sigma_n(z_0, D(R_{z_0}))$ is a similar union of disks of radius $|z_0|^n R_{z_0}$ centered at the images $u(z_0, 1/2)$. We prove (1) by showing that for each $u \in \Sigma_n$, the disk of radius $|z|^n R_z$ at $u(z, 1/2)$ lies inside the disk of radius $|z_0|^n R_{z_0}$ at $u(z_0, 1/2)$; we just compute from the definition of R_z :

$$|z|^n R_z + C|z - z_0| = |z_0|^n R_{z_0}$$

and by the definition of C , we have $|u(z, 1/2) - u(z_0, 1/2)| < C|z - z_0|$, so (1) follows.

To prove (2), first note that since $\Sigma_n(z_0, D(R_{z_0}))$ is connected, by Algorithm 1, there are words u', v' starting with f, g , respectively, such that the disks of radius $|z_0|^n R_{z_0}$ centered at $u'(z_0, 1/2)$ and $v'(z_0, 1/2)$ intersect. Therefore, since these disks contain points in Λ_{z_0} , there are right-infinite words u, v starting with f, g , respectively, such that $|\pi(u, z_0) - \pi(v, z_0)| < 4|z_0|^n r_{z_0}$. Therefore,

$$|\pi(u, z) - \pi(v, z)| < 4|z_0|^n r_{z_0} + 2C|z - z_0|,$$

since $\pi(u, z)$ and $\pi(v, z)$ can each move by at most $C|z - z_0|$. So there are two points in Λ_z which are closer than ϵ , where ϵ is the right hand side of the inequality.

Now we must show that $\Sigma_n(z, D(R_z))$ contains an ϵ -neighborhood of Λ_z . We know that $\Sigma_n(z, D(r_z))$ contains Λ_z , so it suffices to show that the difference between the radii of the disks in $\Sigma_n(z, D(r_z))$ and the disks in $\Sigma_n(z, D(R_z))$ is at least ϵ . We compute

$$\begin{aligned} |z|^n R_z - |z|^n r_z &= |z_0|^n R_{z_0} - C|z - z_0| - |z|^n r_z \\ &= |z_0|^n \left(A^n K + 4K + 3|z_0|^{-n} C\sqrt{2}d \right) - C|z - z_0| - |z|^n r_z \\ &\geq ((|z_0|A)^n - |z|^n) r_z + 4|z_0|^n r_{z_0} + 2C|z - z_0| \\ &\geq 4|z_0|^n r_{z_0} + 2C|z - z_0| \\ &= \epsilon \end{aligned}$$

Where we have used $r_z, r_{z_0} < K$ and $|z - z_0| < \sqrt{2}d$. Also, because $|z_0|A > |z|$, we have $(|z_0|A)^n - |z|^n > 0$. \square

Let T be a component of the complement of $\Sigma_n(z_0, D(R_{z_0}))$ inside the convex hull of $\Sigma_n(z_0, D(R_{z_0}))$. Note that the boundary of T contains a line segment along a supporting hyperplane for the convex hull (the “outside” boundary of T). There are two distinguished disks in $\Sigma_n(z_0, D(R_{z_0}))$ which lie on either end of this line segment and are centered at images of $1/2$ under two words in Σ_n . Let these disks be centered at $p_1 = u_1(z_0, 1/2)$ and $p_2 = u_2(z_0, 1/2)$. Next, let q be a point in T which is distance α' from $\Sigma_n(z_0, D(R_{z_0}))$, and suppose that $\alpha' > 0$, where

$$\begin{aligned} \alpha &= \alpha' - (|z_0|A)^n K - 4|z_0|^n K - 5C\sqrt{2}d \\ &= \alpha' - |z_0|^n R_{z_0} - 2C\sqrt{2}d \end{aligned}$$

Definition 8.2.3. In the above notation, the balls $B_\alpha(p_1 - q)$ and $B_\alpha(p_2 - q)$ are *trap-like balls* for the region B .

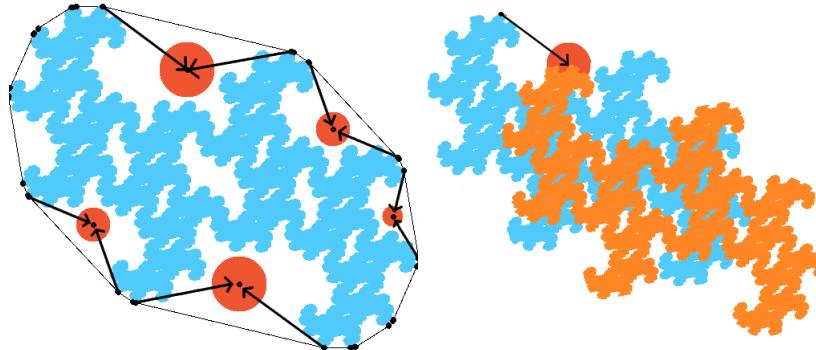


FIGURE 18. A supporting hyperplane of the convex hull of $\Sigma_n(z_0, D(R_{z_0}))$ intersects $\Sigma_n(z_0, D(R_{z_0}))$ in two balls. Vectors which translate these balls *inside* the convex hull but *outside* $\Sigma_n(z_0, D(R_{z_0}))$ are trap-like (left). Translating by a trap-like vector moves $\Sigma_n(z_0, D(R_{z_0}))$ transverse to itself and produces a trap (right).

That is, a trap-like ball for B is a ball of vectors which translate a disk at a vertex of the convex hull of $\Sigma_n(z_0, D(R_{z_0}))$ an appreciable amount into the region inside the convex hull but outside the set. See Figure 18.

Remark 8.2.4. One might wonder whether we should expect any trap-like balls to exist at all, since it's not immediately clear why α should be positive. Recall that $|z_0| < 1/\sqrt{2}$, so for n large enough, $\alpha \approx \alpha' - 5C\sqrt{2}d$, and d is probably tiny compared to the scale of $\Sigma_n(z_0, D(R_{z_0}))$ (which is approximately the limit set Λ_{z_0}).

Lemma 8.2.5. *If $B_\alpha(v)$ is a trap-like ball for B , then $B_\alpha(-v)$ is also a trap-like ball.*

Proof. The set $\Sigma_n(z_0, D(R_{z_0}))$ is rotationally symmetric under a rotation of order 2 about the point $1/2$. A trap-like ball is taken to a trap-like ball under this rotation, and in the above notation, it will negate the vectors $p_1 - q$ and $p_2 - q$. \square

8.2.2. Finding a ball of traps centered at a parameter z . In this section, we show how to use the trap-like balls produced in the previous section to verify the existence of a ball of traps at z . We fix notation as in the previous section, so we have a square region B in parameter space with side length $2d$ and centered at z_0 . We let K be an upper bound for r_z over B and C be such that $|u(z, 1/2) - u(z_0, 1/2)| < C|z - z_0|$ for all $u \in \Sigma_n$ and $z \in B$.

Lemma 8.2.6. *Let $u, v \in \Sigma_m$ be such that u starts with f and v starts with g and $z^{-m}(u(z, 1/2) - v(z, 1/2)) \in B_\alpha(p)$, where $B_\alpha(p)$ is a trap-like ball for B . Let Z be a lower bound for $|z|$ over B . Then there exists a trap for every $z' \in B_\epsilon(z) \cap B$, where*

$$\epsilon = \frac{Z^m}{2C}(\alpha - |z^{-m}(u(z, 1/2) - v(z, 1/2)) - p|)$$

Proof. We will check the 3 hypotheses of Definition 7.1.3 on the words u and v with the topological disk $\Sigma_n(z, D(R_z))$. First, u, v start with f, g by construction, so the first condition is verified. The third condition, that there are points in $f\Lambda_z$ and $g\Lambda_z$ within distance ϵ , where the $\epsilon/2$ neighborhood of Λ_z is contained in $\Sigma_n(z, D(R_z))$, is conclusion (2) of Lemma 8.2.2.

Now we need to verify condition (2) in the definition of a trap. This is the more difficult verification. After a suitable rescaling, the problem becomes more tractable. Consider the unions $z^{-m}(u\Sigma_n(z, D(R_z)))$ and $z^{-m}(v\Sigma_n(z, D(R_z)))$. These sets have a pair of intersecting paths as in condition (2) if and only if the original sets $u\Sigma_n(z, D(R_z))$ and $v\Sigma_n(z, D(R_z))$ do. Furthermore, the sets $z^{-m}(u\Sigma_n(z, D(R_z)))$ and $z^{-m}(v\Sigma_n(z, D(R_z)))$ are exactly the same, up to translation, as $\Sigma_n(z, D(R_z))$ and the translated set

$$\Sigma_n(z, D(R_z)) + z^{-m}(u(z, 1/2) - v(z, 1/2)).$$

In other words, we translate the set $\Sigma_n(z, D(R_z))$ off of itself by the vector $w = z^{-m}(u(z, 1/2) - v(z, 1/2))$. If we can find the interlocking paths of condition (2), we are done.

We start by considering $\Sigma_n(z_0, D(R_{z_0}))$ and then thinking about what can happen as we change z_0 to z . By hypothesis, w lies in a trap-like ball for B . These are four distinguished disks associated with w , as follows. The vector w is associated with a component of the complement of $\Sigma_n(z_0, D(R_{z_0}))$ inside the convex hull of it. This component has one side which lies along a supporting hyperplane H of

the convex hull, and H intersects two disks P_1, P_2 which sit on either side of the component. By the definition of the trap-like balls, the disk P_1 , which has radius $|z_0|^n R_{z_0}$, is translated by w to a disk Q_1 , which is distance at least $2C\sqrt{2d} + 2C\epsilon$ away from both H and $\Sigma_n(z_0, D(R_{z_0}))$. Also note that P_1, P_2 are this same distance away from the translated set. Let H' be a hyperplane perpendicular to w , and translate it so it supports $\Sigma_n(z_0, D(R_{z_0}))$. It lies tangent to some disk P_3 . Now P_3 is translated by w to a disk Q_2 which is distance at least $2C\sqrt{2d} + 2C\epsilon$ away from the slid H' , and thus that distance away from $\Sigma_n(z_0, D(R_{z_0}))$. The pairs of disks (P_1, P_2) and (Q_1, Q_2) are linked. See Figure 19.

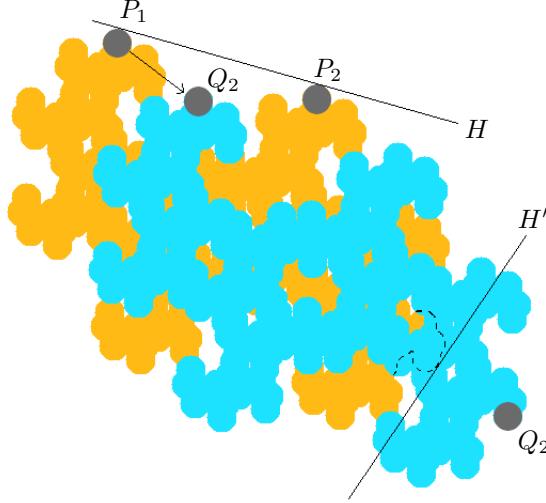


FIGURE 19. The picture of the proof of Lemma 8.2.6. If we change the parameter slightly, the marked balls P_1, P_2, Q_1, Q_2 cannot move much and thus still give a trap.

Now change the parameter from z_0 to z , and consider $\Sigma_n(z, D(R_z))$. Because $|z|^n R_z < |z_0|^n R_{z_0}$, the disks P_1, P_2, Q_1, Q_2 can only shrink. And every disk, and the supporting hyperplanes H, H' , can move at most distance $C|z - z_0| < C\sqrt{2d}$. Therefore, these four disks are still disjoint from the opposing copy of $\Sigma_n(z, D(R_z))$, and each contains points in Λ_z , and by Lemma 8.2.2, $\Sigma_n(z, D(R_z))$ remains connected, so these points can be connected by paths with algebraic intersection number 1. This verifies condition (2) of the trap definition.

This shows that there exists a trap for parameter z , but recall that we desire a trap for every z in $B_\epsilon(z)$. To see that this is true, observe that at the point z , the disks P_1, P_2, Q_1, Q_2 are still distance at least $2C\epsilon|z|^{-m}$ away from the opposing copy of $\Sigma_n(z, D(R_z))$. So we can change z again by at most ϵ while retaining this trap. All of this is contingent on the parameter z remaining in B , so the final ball of traps produced is $B_\epsilon(z) \cap B$. \square

8.3. Certifying holes. We now summarize this section. To certify a hole in \mathcal{M} which lies completely within some square region B , we compute R_{z_0} and the set $\Sigma_n(z_0, D(R_{z_0}))$ for a reasonable-sized n (say, 15). Then we compute the convex hull and some trap-like balls. Then, at the initial point of a path γ encircling the

hole, we apply Lemma 8.2.6 to produce a ball B_1 of traps. Then we go along γ to the edge of B_1 , and find another ball of traps B_2 , and so on. The balls overlap, so together they produce an open set inside set A containing γ .

We do all computations to double precision, which has a precision of at least 15 decimal digits. Therefore, as long as no number in the computation ever requires more than, say, 10 digits of precision, this is rigorous. In practice, this is never an issue.

Question 8.3.1. *Is there a combinatorial way to distinguish holes in \mathcal{M} ?*

9. INFINITELY MANY HOLES IN \mathcal{M} AND RENORMALIZATION

In this section we describe a certain family of natural operators on the parameter plane which account for much of the observed self-similarity in the structure of \mathcal{M} and \mathcal{M}_0 . Similar ideas and some similar results already appear in the work of Solomyak [12], although our approach is sufficiently different (and enough of the results we obtain are new) that it is worth including here.

The first main result we obtain is the existence of *infinitely many holes* in \mathcal{M} , arranged in certain spirals. The proof of this fact does not technically need the theoretical apparatus of renormalization; but the phenomenon is not properly explained without it. We defer the explanation until after a description and rigorous proof of the phenomenon, so that the techniques and definitions we then introduce are sufficiently motivated.

9.1. Infinitely many holes. Numerical exploration of \mathcal{M} quickly reveals many interesting phenomena, of which one of the most interesting is the appearance of apparent “spirals” of holes. One of the most prominent is centered at the point $\omega \sim 0.371859 + 0.519411i$. See Figure 20. The figure also illustrates part of \mathcal{M}_0 (in purple), and exhibits the limit as the “tip” of a spiral of \mathcal{M}_0 . Techniques of Solomyak [12] certify that \mathcal{M}_0 is self-similar at the limit, and is asymptotically similar to the limit set Λ_ω .

The main theorem we prove in this section is the following:

Theorem 9.1.1 (Limit of holes). *Let $\omega \sim 0.371859 + 0.519411i$ be the root of the polynomial $1 - 2z + 2z^2 - 2z^5 + 2z^8$ with the given approximate value. Then*

- (1) ω is in \mathcal{M} , \mathcal{M}_0 and \mathcal{M}_1 ; in fact, the intersection of $f\Lambda_\omega$ and $g\Lambda_\omega$ is exactly the point $1/2$;
- (2) there are points in the complement of \mathcal{M} arbitrarily close to ω ; and
- (3) there are infinitely many rings of concentric loops in the interior of \mathcal{M} which nest down to the point ω .

Thus, \mathcal{M} contains infinitely many holes which accumulate at the point ω .

We refer informally to the holes accumulating on ω as *hexaholes* (because of their approximate shape), and to ω itself as the *hexahole limit*.

The first step in the proof is to give a more combinatorial description of the hexahole limit ω , and prove the first bullet of Theorem 9.1.1.

Lemma 9.1.2. *The set of z for which $\pi(fgfffgggf^\infty, z) = \pi(gfgggffffg^\infty, z)$ are exactly the roots of $1 - 2z + 2z^2 - 2z^5 + 2z^8$.*

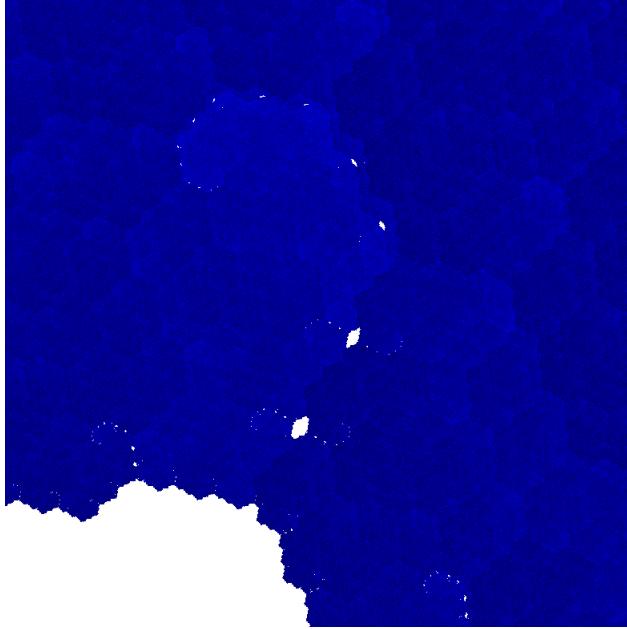


FIGURE 20. Spiral of holes converging to $\omega \sim 0.371859 + 0.519411i$.

Proof. By Proposition 4.2.1, the power series associated to these two infinite words are actually finite polynomials; equating them gives the identity

$$z - z^2 + z^5 - z^8 = 1 - z + z^2 - z^5 + z^8$$

so that $1 - 2z + 2z^2 - 2z^5 + 2z^8 = 0$. \square

The next step in the proof requires us to certify the existence of points in the complement of \mathcal{M} , arbitrarily close to ω . Any given value of z can be numerically certified as being in the complement of \mathcal{M} by Algorithm 1, but we would like to apply this algorithm uniformly to an infinite collection of z of a particular form.

First recall the form of the algorithm: given z as input, and a cutoff depth, we first load the number $1 - z^{-1}$ into a “stack” V , and then recursively replace the content of the stack at each stage with a set of *viable children*. More precisely, for each $\alpha \in V$, there are three children $z^{-1}\alpha$, $z^{-1}(\alpha + z - 1)$, and $z^{-1}(\alpha - z + 1)$. A child is *viable* if its absolute value is less than a constant R depending only on the initial value z (in fact, we can take R to be fixed throughout a neighborhood of a given z), and at each stage of the algorithm we replace each number in V with the set of its viable children. The algorithm halts whenever the stack V is empty (in which case we certify that z is in the complement of \mathcal{M}) or if we exceed the “run time” (i.e. the cutoff depth) allocated in advance.

Let’s imagine running our algorithm on an ideal machine without imposing a cutoff depth, so that the algorithm halts if and only if Λ_z is disconnected. At each successive time step d , the stack V consists of a finite list of numbers. If it happens that the content of V is eventually *periodic* (and nonempty) as a function of d , then of course the algorithm never halts — certifying that z is in fact in \mathcal{M} . Now,

the numbers in V at each finite stage are all Laurent polynomials in z of degree bounded by d , so if V is eventually periodic, then z must be algebraic.

If we apply the algorithm to the number ω defined above, then we indeed can certify that V is eventually periodic, and in fact becomes *constant* after $d = 12$.

Example 9.1.3 (Stack contents for ω). For $\omega \sim 0.371859 + 0.519411i$ the root of $1 - 2z + 2z^2 - 2z^5 + 2z^8 = 0$ we can take $R = 2.257$. Unfortunately, the full stack over all 12 steps is somewhat unwieldy, so we do not list it here. However, on step 9, the stack contains the number 1, and this is the only stack entry with viable descendants. The tree diagram of the algorithm at this point becomes periodic; we show it in Figure 21.

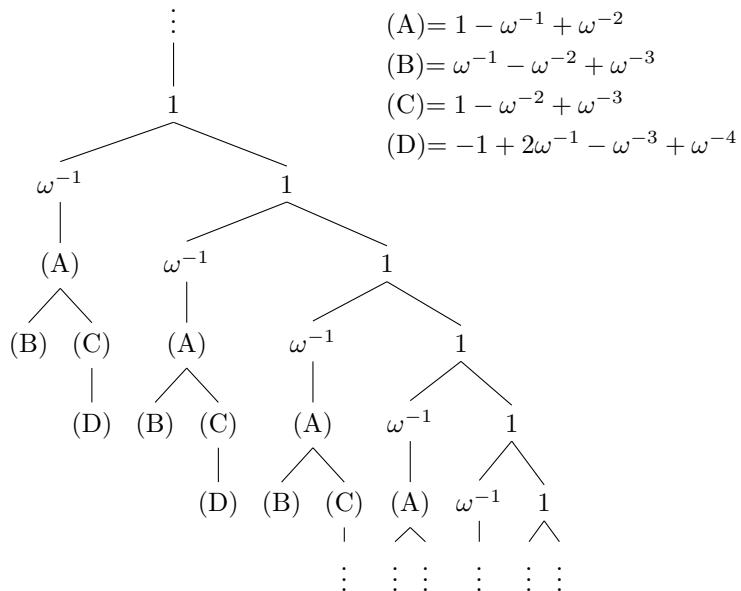


FIGURE 21. The periodic stack of the disconnectedness algorithm on the input ω .

It is clear from the tree diagram that the stack becomes constant at step 12.

By analyzing precisely which children have indefinitely viable descendants, we get a precise description of the intersection $f\Lambda_\omega \cap g\Lambda_\omega$. In this case, we can readily observe that there is a unique pair of infinite words u, v where u starts with f and v with g so that $\pi(u, \omega) = \pi(v, \omega)$; these words are in fact related under the canonical involution, so that the intersection consists exactly of the point $1/2$. This proves the first bullet in Theorem 9.1.1.

The second step in the proof of Theorem 9.1.1 is to certify the existence of points in the complement of \mathcal{M} arbitrarily close to ω . These points will all be of the form $\omega + C\omega^\ell$ for sufficiently big ℓ , and for a fixed constant $C = 0.29946137 - 0.48972405i$.

Proposition 9.1.4 (ω limit of Schottky). For $C = 0.29946137 - 0.48972405i$, there is ℓ so that point $z = \omega + C\omega^\ell$ is Schottky for sufficiently large ℓ .

In order to prove Proposition 9.1.4, we are going to formally run the disconnectedness algorithm on z and show that we can understand the contents of the stack as long as ℓ is large enough. The stack will essentially be the same as the stack for ω for a long time, followed by a uniformly bounded (in ℓ) number of steps which prove disconnectedness. This discussion is elementary, but it requires taking things to infinity in a careful order.

We first prove a general lemma which provides the stack contents; proving the proposition then reduces to doing a numerical computation for the given C value. To set up the lemma, we need to do a computation. Recall that when running the disconnectedness algorithm on ω , at every step there is a single stack entry (i.e. “1”) which has infinitely many descendants. The true, unsimplified version of this entry in the stack at step n is

$$p_{1,n}(z) = 1 + z^{-n}(-1 + 2z - 2z^2 + 2z^5 - 2z^8)$$

At every time step, there are 5 other polynomials on the stack, which are the finitely many children of $p_{1,n-1}(z)$, $p_{1,n-2}(z)$, and $p_{1,n-3}(z)$ which have not yet died. These polynomials are:

$$\begin{aligned} p_{2,n}(z) &= z^{-n}(-1 + 2z - 2z^2 + 2z^5 - 2z^8 + z^{n-4} - z^{n-3} + 2z^{n-1} - z^n) \\ p_{3,n}(z) &= z^{-n}(-1 + 2z - 2z^2 + 2z^5 - 2z^8 + z^{n-3} - z^{n-2} + z^{n-1}) \\ p_{4,n}(z) &= z^{-n}(-1 + 2z - 2z^2 + 2z^5 - 2z^8 + z^{n-3} - z^{n-2} + z^n) \\ p_{5,n}(z) &= z^{-n}(-1 + 2z - 2z^2 + 2z^5 - 2z^8 + z^{n-2} - z^{n-1} + z^n) \\ p_{6,n}(z) &= z^{-n}(-1 + 2z - 2z^2 + 2z^5 - 2z^8 + z^{n-1}) \end{aligned}$$

It is important to note that obtaining these polynomials involves *no* simplification at any stage. Thus, heuristically, for fixed n and z values close to ω , these polynomials should give the stack contents. This is the idea Lemma 9.1.5 explores in detail.

We compute the values of the $p_{i,n}(z)$ polynomials at the point $z = \omega + C\omega^{m+k}$ (it is pedagogically helpful to split ℓ in Proposition 9.1.4 into the two variables $\ell = m+k$). This computation is just an expansion, simplified using the polynomial of which ω is a root. For compactness, we denote $p_{i,n}(\omega + C\omega^{m+k})$ by $p_{i,n}^{m+k}$. All of these polynomials have a large “remainder” term, which we will denote by

$$R_{n,m,k} = \frac{\omega^{m+k}}{(\omega + C\omega^{m+k})^n} (2C - 4C\omega + 10C\omega^4 - 16C\omega^7 + O(\omega^m))$$

Now we list the polynomials:

$$\begin{aligned} p_{1,n}^{m+k} &= 1 + R_{n,m,k} \\ p_{2,n}^{m+k} &= -1 + \frac{1}{(\omega + C\omega^{m+k})^2} - \frac{1}{(\omega + C\omega^{m+k})^3} + \frac{1}{(\omega + C\omega^{m+k})^4} + R_{n,m,k} \\ p_{3,n}^{m+k} &= \frac{1}{\omega + C\omega^{m+k}} - \frac{1}{(\omega + C\omega^{m+k})^2} + \frac{1}{(\omega + C\omega^{m+k})^3} + R_{n,m,k} \end{aligned}$$

$$\begin{aligned} p_{4,n}^{m+k} &= 1 - \frac{1}{(\omega + C\omega^{m+k})^2} + \frac{1}{(\omega + C\omega^{m+k})^3} + R_{n,m,k} \\ p_{5,n}^{m+k} &= 1 - \frac{1}{\omega + C\omega^{m+k}} + \frac{1}{(\omega + C\omega^{m+k})^2} + R_{n,m,k} \\ p_{6,n}^{m+k} &= \frac{1}{\omega + C\omega^{m+k}} + R_{n,m,k} \end{aligned}$$

Lemma 9.1.5. *For any C , there are constants k and M such that for all $m > M$ and $12 < n \leq m$, the contents of the stack of the disconnectedness algorithm at step n when run on $\omega + C\omega^{m+k}$ is exactly the set of $p_{i,n}^{m+k}$ for $1 \leq i \leq 6$.*

Proof. In order to prove this lemma, first consider running the algorithm on ω : the stack beyond step 12 is constant at

$$\left\{1, -1 + \frac{1}{\omega^2} - \frac{1}{\omega^3} + \frac{1}{\omega^4}, \frac{1}{\omega} - \frac{1}{\omega^2} + \frac{1}{\omega^3}, 1 - \frac{1}{\omega^2} + \frac{1}{\omega^3}, 1 - \frac{1}{\omega} + \frac{1}{\omega^2}, \frac{1}{\omega}\right\} = \{p_{i,n}(\omega)\}_{i=1}^6$$

Now think of varying the input from ω to $\omega + C\omega^{m+k}$. In order to prove that the contents of the stack are as claimed, we need to show two things (1) the polynomials $p_{i,n}^{m+k}$ stay on the stack for all $n \leq m$ and (2) every child which was discarded for ω through step n still gets discarded.

First note that for $n \leq m$,

$$\left| \frac{\omega^{m+k}}{(\omega + C\omega^{m+k})^n} \right| \leq \left| \frac{\omega^{m+k}}{(\omega + C\omega^{m+k})^m} \right|,$$

and furthermore, $\omega^{m+k}/(\omega + C\omega^{m+k})^m$ converges to ω^k from below as $m \rightarrow \infty$. Therefore, the absolute value $|R_{n,m,k}|$ is uniformly (in n) bounded above by the “worst case” $|R_{m,m,k}|$ where $n = m$:

$$|R_{n,m,k}| \leq |R_{m,m,k}| \leq |\omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7 + O(\omega^m))|$$

So for example, the variation $|p_{6,n}(\omega) - p_{6,n}^{m+k}|$ is uniformly (in n) bounded by the expression:

$$\left| \frac{1}{\omega} - \frac{1}{\omega + C\omega^{m+k}} \right| + |\omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7 + O(\omega^m))|.$$

There are similar expressions for $|p_{i,n}(\omega) - p_{i,n}^{m+k}|$ for each i . Now if we make k large and bound m from below, we can make all these expressions as small as we like, and hence small enough so that the $p_{i,n}^{m+k}$ remain on the stack for all $n \leq m$.

To prove that these are the *only* things on the stack, we compute expressions for the children of $p_{i,n}^{m+k}$ and do exactly the same thing to prove that a large enough k and m make the worst-case deviation from the children of $p_{i,n}(\omega)$ small for $n \leq m$ and hence these children will leave the stack exactly as the children of $p_{i,n}(\omega)$ do. This computation is the same, so we omit it. \square

We note that for a specific value of C (such as the one given in the proposition), it is possible to actually numerically compute k . As an example, we show how to compute a value of k which ensures $p_{1,n}^{m+k}$ remains on the stack for sufficiently large m . As $m \rightarrow \infty$, the difference $|p_{1,n}(\omega) - p_{1,n}^{m+k}|$ is bounded above by the limit

$$|\omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7)|$$

Hence if we compute $p_{1,n}(\omega)$ and observe how far away it is from getting cut off the stack (remember things get removed if their absolute value is too large), we can

choose k so that the expression above is small enough that $p_{1,n}^{m+k}$ remains on the stack for $m > M$ and $n \leq m$, (where M can depend on k). In order to compute a value of k which actually works for Lemma 9.1.5, it is necessary to consider $p_{i,n-1}(\omega)$ over all i and all their children and make sure that k is large enough to accept or reject them appropriately.

Proof of Proposition 9.1.4. Doing the computation above for the specified value $C = 0.29946137 - 0.48972405i$ (this is an exact value) shows that Lemma 9.1.5 holds with $k = 12$. Therefore, for all m sufficiently large, the contents of the stack at time $n \leq m$ will be as claimed in the lemma. By taking m large, we can get the stack contents at step $n = m$ as close as we like to the limits, which we denote by $p_{i,\infty}^k$.

$$\begin{aligned} p_{1,\infty}^k &= 1 + \omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7) \\ p_{2,\infty}^k &= -1 + \frac{1}{\omega^2} - \frac{1}{\omega^3} + \frac{1}{\omega^4} + \omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7) \\ p_{3,\infty}^k &= \frac{1}{\omega} - \frac{1}{\omega^2} + \frac{1}{\omega^3} + \omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7) \\ p_{4,\infty}^k &= 1 - \frac{1}{\omega^2} + \frac{1}{\omega^3} + \omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7) \\ p_{5,\infty}^k &= 1 - \frac{1}{\omega} + \frac{1}{\omega^2} + \omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7) \\ p_{6,\infty}^k &= \frac{1}{\omega} + \omega^k (2C - 4C\omega + 10C\omega^4 - 16C\omega^7) \end{aligned}$$

We want to continue running the disconnectedness algorithm at this point. Recall we require a radius outside which we discard children. By taking m large, we may assume this radius is the one for ω , i.e. $2|\omega - 1|/2(1 - |\omega|) < 2.26$ and that the algorithm replaces α with the three children

$$\omega^{-1}\alpha, \quad \omega^{-1}(\alpha + \omega - 1), \quad \omega^{-1}(\alpha - \omega + 1).$$

Now start this algorithm with the given (numerical, with $k = 12$) stack contents $\{p_{i,\infty}^k\}_{i=1}^6$; it terminates (with an empty stack) in 20 steps. Therefore, for $k = 12$, there is some M such that for all $m > M$, the disconnectedness algorithm run on the input $\omega + C\omega^{m+k}$ certifies disconnectedness at step $m + 20$. This completes the proof. \square

This proves the second bullet in Theorem 9.1.1. Note that by means of this method we can numerically certify any $C \in \mathbb{C}$ for which the points $\omega + C\omega^n$ are Schottky for all sufficiently large n . However, when this method of certification fails, we cannot conclude that the corresponding points are all (eventually) in \mathcal{M} ; a different method is necessary for that.

The last step in the proof of Theorem 9.1.1 is to certify the existence of infinitely many rings of concentric loops in the interior of \mathcal{M} which nest down to the point ω . This depends on an analysis of how trap vectors transform under certain combinatorial and numerical operations. We discuss this in the remainder of the section.

Let $R \subseteq \mathbb{C}$ be a small region containing ω . Recall from Section 8 that we can produce a collection of trap-like balls for the region R such that if $z \in R$ and there exist $u, v \in \Sigma_m$ starting with f, g , respectively, such that $z^{-m}(u(z, 1/2) - v(z, 1/2))$ lies in a trap-like ball, then there exists a trap at z , and z lies in the interior of

\mathcal{M} . We will use this to show that for z of the form $\omega + C\omega^n$, we can certify the existence of a trap for z for *all* sufficiently large n .

Given two words $u, v \in \Sigma_m$, not necessarily starting with f, g , recall that we can write

$$u(z, x) = xz^m + p_u(z) \quad \text{and} \quad v(z, x) = xz^m + p_v(z)$$

for some polynomials $p_u(z)$ and $p_v(z)$ in z . For example, $p_g(z) = -z + 1$ because $g(z, x) = z(x - 1) + 1$. Define words

$$U_n = fgfffgggf^n u \quad \text{and} \quad V_n = gfgggffffg^n v$$

Lemma 9.1.6. *In the above notation, if the vector*

$$t = 2C\omega^{-m-8}(1 - 2\omega + 5\omega^4 - 8\omega^7) + \omega^{-m}(p_u(\omega) - p_v(\omega) + 1)$$

lies in a trap-like ball $B_\alpha(p)$ for the region R , then for sufficiently large n , the words U_n, V_n give a trap for $\omega + C\omega^n$. Furthermore, if we let

$$\epsilon = \frac{|\alpha - (p - t)|}{|2\omega^{-m-8}(1 - 2\omega + 5\omega^4 - 8\omega^7)|},$$

then for any compact subset S of $B_\epsilon(C)$, there is an N such that for any $C' \in S$ and $n > N$, the words U_n, V_n give a trap for $\omega + C'\omega^n$.

Proof. The proof is primarily a computation. By applying the definitions of f and g , we compute:

$$U_n(z, 1/2) = z - z^2 + z^5 - z^8 + \frac{1}{2}z^{m+n+8} + z^{n+8}p_u(z)$$

$$V_n(z, 1/2) = 1 - z + z^2 - z^5 + z^8 - z^{n+8} + \frac{1}{2}z^{m+n+8} + z^{n+8}p_v(z),$$

so

$$U_n(z, 1/2) - V_n(z, 1/2) = p_\omega(z) + z^{n+8}(p_u(z) - p_v(z) + 1),$$

where $p_\omega(z) = -1 + 2z - 2z^2 + 2z^5 - 2z^8$. Recall from the definition of ω that $p_\omega(\omega) = 0$. Since U_n and V_n have length $m + n + 8$, to show that this pair gives a trap-like vector for some z , we'll be considering the expression

$$z^{-m-n-8}(U_n(z, 1/2) - V_n(z, 1/2)) = \frac{p_\omega(z)}{z^{m+n+8}} + z^{-m}(p_u(z) - p_v(z) + 1).$$

We now show how to certify that this vector is trap-like for z of the form $\omega + C\omega^n$, for sufficiently large n . We therefore consider

$$\frac{p_\omega(\omega + C\omega^n)}{(\omega + C\omega^n)^{m+n+8}} + (\omega + C\omega^n)^{-m}(p_u(\omega + C\omega^n) - p_v(\omega + C\omega^n) + 1).$$

Note that the right summand converges to $\omega^{-m}(p_u(\omega) - p_v(\omega) + 1)$ as $n \rightarrow \infty$. We claim the left summand converges as well. To see this, we expand it out using the

definition of p_ω :

$$\begin{aligned} \frac{p_\omega(\omega + C\omega^n)}{(\omega + C\omega^n)^{m+n+8}} &= (-1 + 2\omega - 2\omega^2 + 2\omega^5 - 2\omega^8) \frac{1}{(\omega + C\omega^n)^{m+n+8}} \\ &\quad + 2C\omega^{-m-8}(1 - 2\omega + 5\omega^4 - 8\omega^7) \frac{\omega^n}{(\omega + C\omega^n)^n} \\ &\quad + \frac{-2C^2\omega^{2n} + 2C^5\omega^{5n} - 2C^8\omega^{8n} + 20C^2\omega^{3+2n} - 56C^2\omega^{6+2n}}{(\omega + C\omega^n)^{m+n+8}} \\ &\quad + \frac{20C^3\omega^{2+3n} - 112C^3\omega^{5+3n} + 10C^4\omega^{1+4n} - 140C^4\omega^{4+4n}}{(\omega + C\omega^n)^{m+n+8}} \\ &\quad + \frac{-112C^5\omega^{3+5n} - 56C^6\omega^{2+6n} - 16C^7\omega^{1+7n}}{(\omega + C\omega^n)^{m+n+8}} \end{aligned}$$

The first line is 0 because, recall, $p_\omega(\omega) = 0$, and it's straightforward to see that $\lim_{n \rightarrow \infty} \omega^n / (\omega + C\omega^n)^n = 1$, so the last three lines converge to 0 and the second line converges to $2C\omega^{-m-8}(1 - 2\omega + 5\omega^4 - 8\omega^7)$.

Therefore, if the hypothesis of the lemma holds, then for sufficiently large n , the words U_n and V_n are trap like for $\omega + C\omega^n$, as claimed.

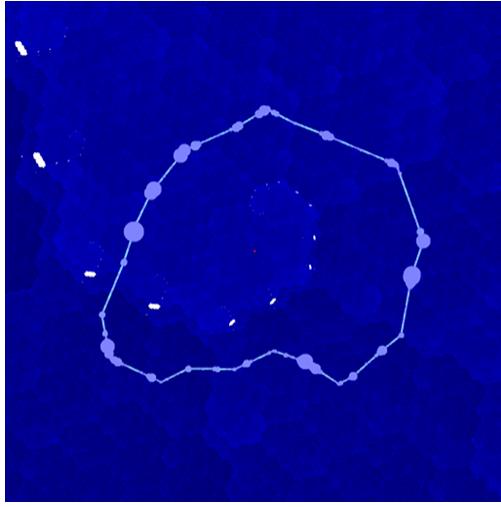
To get the last statement of the lemma, observe that the vector t varies linearly with C , so certainly for any $C' \in B_\epsilon(C)$, the hypotheses of the lemma are satisfied. But note that all the expressions above are uniformly continuous in C on compact subsets, so given any compact subset, there is a uniform bound on the value of n required. \square

To complete the proof of Theorem 9.1.1, then, it suffices to exhibit a loop of overlapping balls output by Lemma 9.1.6 encircling ω . Because there are finitely many balls, there is a uniform N such that for $n > N$, there exists a trap for $\omega + C\omega^n$ for every C in every ball in this loop. In other words, the image of this loop under the map $x \mapsto \omega(x - \omega) + \omega$ lies in the interior of \mathcal{M} for all sufficiently large iterates. Figure 22 shows the loop of trap balls which we computed.

Remark 9.1.7. Lemma 9.1.6 only states that this loop is *eventually* in the interior of \mathcal{M} (under a large enough iterate of the map $x \mapsto \omega(x - \omega) + \omega$). However, experimentally, this loop lies in the interior for all iterates. The primary evidence for this is that a picture of limit traps near ω looks the same as a picture of regular traps.

9.2. Renormalization.

9.2.1. Introduction. In this section, we place the above example in a more formal context and explain the relationship with the work of Solomyak in [12]. We first give a heuristic explanation of some of our definitions. We would *like* to define a renormalization operator $R : \Sigma \times \Sigma \times \mathbb{D}^* \rightarrow \mathbb{D}^*$ such that $R(u, v, z)$ is the parameter w such that the limit set for w is the same, in some sense, as the union $u(z, \Lambda_z) \cup v(z, \Lambda_z)$. The right definition for this operator is elusive. However, we show below that we can understand what the fixed points of renormalization should be, and at these fixed points, there is a sensible definition of a limiting trap. For certain renormalization points, we give a new interpretation of a result of Solomyak [12].

FIGURE 22. A loop of limit trap balls encircling ω .

Let u and v be given words of the same length. These will be our prefixes. Let s and t be two other words of the same length. We are interested in the appearance of the set $us^n(z, \Lambda_z) \cup vt^n(z, \Lambda_z)$ and renormalization with respect to the words us^n and vt^n as $n \rightarrow \infty$. As n gets large, renormalization at us^n , vt^n should converge to a locally-defined holomorphic function which, abusing notation, we'll call renormalization at (the now infinite words) us^∞ , vt^∞ . Parameters z for which $\pi(us^\infty, z) = \pi(vt^\infty, z)$ should be the fixed points of this renormalization.

Therefore, we say that a parameter z is a *renormalization point* if there are words u, v, s, t as above such that $\pi(us^\infty, z) = \pi(vt^\infty, z)$. We will show that there is a notion of a limit trap at a renormalization point and that this can sometimes give an asymptotic self-similarity.

9.2.2. A computation. This section is essentially concerned with the behavior of the limit set Λ_z at infinitesimal scales for renormalization points. That ω is a renormalization point means that $f\Lambda_\omega \cap g\Lambda_\omega \neq \emptyset$ and in fact there are two eventually periodic words u, v so $\pi(u, \omega) = \pi(v, \omega)$. We want to zoom in on this point of intersection. Recall that for a finite (or infinite) word $u \in \Sigma_n$, we can write $u(z, x) = xz^n + p_u(z)$, where p_u is a polynomial of degree n (if u is infinite, $p_u(z)$ is the power series $\pi(u, z)$). We take the convention that if u has length 0, then $p_u(z) \equiv 0$. If $u, v \in \Sigma_n$, then $u(z, \Lambda_z)$ and $v(z, \Lambda_z)$ are translates of each other, and the displacement vector is $p_u(z) - p_v(z)$. A more useful quantity turns out to be the displacement relative to the sizes of the sets $u(z, \Lambda_z)$ and $v(z, \Lambda_z)$, that is

$$z^{-n}(p_u(z) - p_v(z)).$$

We have already encountered this expression several times. As in the proof of Theorem 9.1.1, we will need to compute its value for parameters of the form $\omega + C\omega^n$ for long words. This section contains a rather tedious computation which will be necessary for its generalization.

Lemma 9.2.1. *Let u, v have length a ; let s, t have length b ; and let x, y have length c . Let ω be a renormalization point for u, v, s, t . Write $P(z) = p_{us^\infty}(z) - p_{vt^\infty}(z)$,*

so $P(\omega) = 0$. Then as $n \rightarrow \infty$, the quantity

$$(\omega + C\omega^{nb})^{-(a+bn+c)} (p_{us^n x}(\omega + C\omega^{nb}) - p_{vt^n y}(\omega + C\omega^{nb}))$$

converges to

$$\omega^{-a-c}(p_u(\omega) - p_v(\omega)) + \omega^{-c}(p_x(\omega) - p_y(\omega)) + \omega^{-a-c}CP'(\omega)$$

Proof of Lemma 9.2.1. First, some notation. Write d_i for the coefficients of the power series $P(z)$, so $P(z) = \sum_{i=1}^{\infty} d_i z^i$. Note that d_i is periodic with period b for large enough i ; write $P_a(z)$ to mean the eventually periodic part of $P(z)$, shifted by a , so

$$P_a(z) = \sum_{i=0}^{\infty} d_{i+a+bn} z^i.$$

Where n is taken large enough that the coefficients are constant in n . If we take a finite power for s and t , the resulting polynomial (which has degree $a+bn$) will agree with $P(z)$ to the term with degree $a+bn-1$, so define $r \in \{\pm 2, \pm 1, 0\}$ so that

$$p_{us^n}(z) - p_{vt^n}(z) = \sum_{i=0}^{a+bn-1} d_i z^i + r z^{a+bn}.$$

Observation of the power series P shows the facts (the third following from the first two):

$$\begin{aligned} P_a(z) &= p_{s^\infty}(z) - p_{t^\infty}(z) + r \\ P(\omega) &= 0 = p_u(\omega) - p_v(\omega) + \omega^a(p_{s^\infty}(\omega) - p_{t^\infty}(\omega)) \\ r - P_a(\omega) &= \omega^{-a}(p_u(\omega) - p_v(\omega)) \end{aligned}$$

We will soon encounter some rather large expressions, and it will be helpful to use some small notation. We denote the expression in the lemma by E_n , so

$$E_n = (\omega + C\omega^{nb})^{-(a+bn+c)} (p_{us^n x}(\omega + C\omega^{nb}) - p_{vt^n y}(\omega + C\omega^{nb})),$$

and we denote $\omega + C\omega^{bn}$ by Ω_n . Recall that $\lim_{n \rightarrow \infty} \omega^{bn}/\Omega_n^{bn} = 1$. We expand using the fact that $p_{us^n x}(z) = z^{a+bn} p_x(z) + p_{us^n}(z)$:

$$\begin{aligned} E_n &= \Omega_n^{-(a+bn+c)} (\Omega_n^{a+bn} (p_x(\Omega_n) - p_y(\Omega_n)) + p_{us^n}(\Omega_n) - p_{vt^n}(\Omega_n)) \\ &= \Omega_n^{-c} (p_x(\Omega_n) - p_y(\Omega_n)) + r \Omega_n^{-c} + \Omega_n^{-(a+bn+c)} \sum_{i=0}^{a+bn-1} d_i \Omega_n^i \end{aligned}$$

The first part trivially converges to $\omega^{-c}(p_x(\omega) - p_y(\omega)) + r\omega^{-c}$ as $n \rightarrow \infty$. We will show that

$$\Omega_n^{-(a+bn+c)} \sum_{i=0}^{a+bn-1} d_i \Omega_n^i \quad \longrightarrow \quad -\omega^{-c} P_a(\omega) + C\omega^{-(a+c)} P'(\omega).$$

To do this, we expand the term $\Omega_n^i = (\omega + C\omega^{bn})^i$ using the binomial theorem:

$$\begin{aligned}
 (1) \quad & \Omega_n^{-(a+bn+c)} \sum_{i=0}^{a+bn-1} d_i \Omega_n^i = \\
 (2) \quad & \Omega_n^{-(a+bn+c)} \sum_{i=0}^{a+bn-1} d_i \omega^i \\
 & + \Omega_n^{-(a+bn+c)} \sum_{i=1}^{a+bn-1} d_i i C \omega^{bn+i-1} \\
 (3) \quad & + \Omega_n^{-(a+bn+c)} \sum_{i=2}^{a+bn-1} d_i \sum_{j=0}^{i-2} \binom{i}{j} C^{i-j} \omega^{bn(i-j)+j}
 \end{aligned}$$

We handle these summand-by-summand. First, we rewrite (1) using the fact that $P(\omega) = 0$ so $\sum_{i=0}^{a+bn-1} d_i \omega^i = -\sum_{i=a+bn}^{\infty} d_i \omega^i$, so

$$\begin{aligned}
 \Omega_n^{-(a+bn+c)} \sum_{i=0}^{a+bn-1} d_i \omega^i &= -\Omega_n^{-(a+bn+c)} \sum_{i=a+bn}^{\infty} d_i \omega^i \\
 &= -\frac{\omega^a}{\Omega_n^{a+c}} \frac{\omega^{bn}}{\Omega_n^{bn}} \sum_{i=0}^{\infty} d_{i+a+bn} \omega^i \\
 &\rightarrow -\omega^{-c} P_a(\omega)
 \end{aligned}$$

Next, summand (2):

$$\begin{aligned}
 \Omega_n^{-(a+bn+c)} \sum_{i=1}^{a+bn-1} d_i i C \omega^{bn+i-1} &= \Omega_n^{-(a+c)} \frac{\omega^{bn}}{\Omega_n^{bn}} C \sum_{i=1}^{a+bn-1} d_i i \omega^{i-1} \\
 &\rightarrow \omega^{-(a+c)} C P'(\omega)
 \end{aligned}$$

Finally, summand (3). We will prove that it converges to 0. First, we bound the absolute value of the innermost sum. To do this, we pull out terms from the binomial coefficient to re-express it as a different binomial coefficient, so we can collapse the sum into a power. In the first line, we use the fact that $\binom{i}{j} = i(i-1)(i-j)(i-j-1)\binom{i-2}{j}$, and $i-1, i-j, i-j-1 \leq i$:

$$\begin{aligned}
 \left| \sum_{j=0}^{i-2} \binom{i}{j} C^{i-j} \omega^{bn(i-j)+j} \right| &\leq i^4 |C|^2 |\omega|^{2bn} \sum_{j=0}^{i-2} \binom{i-2}{j} |C^{i-2-j} \omega^{bn(i-2-j)+j}| \\
 &= i^4 |C|^2 |\omega|^{2bn} (|\omega| + |C\omega^{bn}|)^{i-2}
 \end{aligned}$$

So the entire summand (3) is bounded in absolute value by

$$\begin{aligned} & |\Omega_n|^{-(a+bn+c)} \sum_{i=2}^{a+bn-1} |d_i| i^4 |C|^2 |\omega|^{2bn} (|\omega| + |C\omega|^{bn})^{i-2} \\ &= |\Omega_n|^{-(a+c)} \frac{|\omega|^{bn}}{|\Omega_n|^{bn}} |\omega|^{bn} |C|^2 \sum_{i=2}^{a+bn-1} |d_i| i^4 (|\omega| + |C\omega|^{bn})^{i-2} \end{aligned}$$

Let $H(z) = \sum_{i=2}^{\infty} |d_i| i^4 z^{i-2}$. Using the root test, it is easy to see that $H(z)$ is uniformly convergent for $|z| < 1$, so H is uniformly convergent in a neighborhood of $|\omega|$. Therefore, as $n \rightarrow \infty$, the above expression converges to

$$\begin{aligned} & \rightarrow |\omega|^{-(a+c)} \left(\lim_{n \rightarrow \infty} \frac{|\omega|^{bn}}{|\Omega_n|^{bn}} \right) \left(\lim_{n \rightarrow \infty} |\omega|^{bn} \right) |C|^2 H(|\omega|) \\ &= |\omega|^{-(a+c)} (1)(0) |C|^2 H(|\omega|) \\ &= 0 \end{aligned}$$

We have now shown that as $n \rightarrow \infty$

$$E_n \rightarrow \omega^{-c} (p_x(\omega) - p_y(\omega)) + r\omega^{-c} - \omega^{-c} P_a(\omega) + C\omega^{-(a+c)} P'(\omega).$$

Using the observations about P at the beginning of the proof, this expression rearranges into the statement of the lemma. \square

9.2.3. Similarity. Recall from Section 6 that the set of differences between points in Λ_z is Γ_z , the limit set generated by the three contractions

$$x \mapsto z(x+1) - 1 \quad x \mapsto zx \quad x \mapsto z(x-1) + 1$$

Theorem 9.2.2 (Renormalizable traps). *Suppose that ω is a renormalizable point for u, v, s, t , where s, t have length b . Let $P(z) = p_{us^\infty}(z) - p_{vt^\infty}(z)$. Let T_ω denote $\frac{p_u(\omega) - p_v(\omega)}{P'(\omega)} - \frac{\omega^a}{P'(\omega)} \Gamma_\omega$, the translated, scaled copy of Γ_ω*

- (1) *If $C \in T_\omega$, then for all $\epsilon > 0$, there is a C' such that $|C - C'| < \epsilon$ and for all sufficiently large n , there is a trap for $\omega + C'\omega^{bn}$.*
- (2) *If there is a unique pair of infinite words $U, V \in \partial\Sigma$ such that $p_U(\omega) = p_V(\omega)$ (i.e. $U = us^\infty, V = vt^\infty$), then there is $\delta > 0$ such that for all $C \notin T_\omega$ with $|C| < \delta$, the limit set for the parameter $\omega + C\omega^{bn}$ is disconnected for all sufficiently large n .*

Remark 9.2.3. A version of part (2) of Theorem 9.2.2 still holds if there are finitely many such infinite U, V , as long as they are eventually periodic. In this case, we need to replace T_ω with a union of multiple scaled, translated copies of Γ_ω .

Remark 9.2.4. We can think of Theorem 9.2.2 as the verification of a kind of “Renormalized Bandt’s Conjecture”. It says that at a renormalizable point ω , there is an increasing union of open subsets of renormalizable traps, limiting to the asymptotically scaled copy of \mathcal{M} centered at ω . It implies (but is stronger than) one of the main consequences of Theorem 2.3 from Solomyak [12], that suitable neighborhoods of zero in T_ω converge in the sense of Hausdorff distance to suitably scaled neighborhoods of ω in \mathcal{M} .

In contrast to Solomyak, our argument is more closely expressed in the language of algorithms, since one of our aims was always to use this theorem to obtain

numerical certificates of the existence of hole spirals. This is stated carefully in Lemma 9.2.5.

Lemma 9.2.5. *Let u, v have length a ; let s, t have length b , and let x, y have length c . Let ω be a renormalization point for u, v, s, t . Write $P(z) = p_{us^\infty}(z) - p_{vt^\infty}(z)$. Suppose that C is such that the vector*

$$\omega^{-a-c}(p_u(\omega) - p_v(\omega)) + \omega^{-c}(p_x(\omega) - p_y(\omega)) + \omega^{-a-c}CP'(\omega)$$

is trap-like for ω . Then the words $us^n x, vt^n y$ give a trap for $\omega + C\omega^{bn}$ for all sufficiently large n .

Proof. This is essentially immediate from Lemma 9.2.1, which says that the vector which determines whether $us^n x, vt^n y$ give a trap converges to the above expression as n gets large. Hence, if the above is trap like, we get a trap for $\omega + C\omega^{bn}$ for all n large enough. \square

If Lemma 9.2.5 holds for some point ω and C , we say that C admits a *limit trap* for ω .

Proof of Theorem 9.2.2. We first prove part (1). Let us be given $C \in T_\omega$. By Lemma 9.2.5, if the vector:

$$\omega^{-a-c}(p_u(\omega) - p_v(\omega)) + \omega^{-c}(p_x(\omega) - p_y(\omega)) + \omega^{-a-c}KP'(\omega)$$

is trap-like for ω , then K admits a limit trap. Let T be a trap-like vector. Then we can solve for the associated value C' which admits a limit trap:

$$C' = \omega^{a+c} \frac{T}{P'(\omega)} - \frac{p_u(\omega) - p_v(\omega)}{P'(\omega)} - \frac{\omega^a}{P'(\omega)}(p_x(\omega) - p_y(\omega))$$

As c grows and x and y vary over all words of length c , the first summand goes to zero, and the second two together converge (in the Hausdorff topology, say, but quite regularly) to T_ω . Hence if $C \in T_\omega$, then for any $\epsilon > 0$, there are words $x, y \in \Sigma_c$ so that C' admitting a limit trap as above has $|C - C'| < \epsilon$. This completes the proof of part (1).

Now we prove part (2). When we run Algorithm 1 on ω , the stack entries at stage $a + bn$ are exactly the scaled differences $\omega^{-a-bn}(p_x(\omega) - p_y(\omega))$ between centers of words x, y of length $a + bn$ (when these differences are small enough to remain on the stack). If there is a unique pair of words U, V such that $p_U(\omega) = p_V(\omega)$, then there is a single stack entry with infinitely viable children, and it is $\omega^{-a-bn}(p_{us^n}(\omega) - p_{vt^n}(\omega))$. Rewriting this as in the proof of Lemma 9.2.1, we see that by making n large, this expression is as close as we'd like to $\omega^{-a}(p_u(\omega) - p_v(\omega))$.

When we vary ω to $\omega + C\omega^{a+bn}$, and make n large, then by Lemma 9.2.1, we can make this stack entry as close as we like to

$$\omega^{-a}(p_u(\omega) - p_v(\omega)) + \omega^{-a}CP'(\omega)$$

Therefore, there is a $\delta > 0$ as in the statement of the theorem such that if $|C| < \delta$, then when we run the disconnectedness algorithm on the input $\omega + C\omega^{a+bn}$, the stack at step $a + bn$ has the entry (as close as we want to) $\omega^{-a}(p_u(\omega) - p_v(\omega)) + \omega^{-a}CP'(\omega)$, and every other stack entry has children which are eliminated in finite time. The value for δ can be found by checking how far the limiting entry $\omega^{-a}(p_u(\omega) - p_v(\omega))$ is from the cutoff; then make δ small enough so that adding the term $\omega^{-a}CP'(\omega)$ does not push anything off or onto the stack.

Now, compute all possible children after c more steps; by Lemma 9.2.1, we get

$$X_{x,y} = \omega^{-a-c}(p_u(\omega) - p_v(\omega)) + \omega^{-c}(p_x(\omega) - p_y(\omega)) + \omega^{-a-c}CP'(\omega),$$

where x, y vary over all words of length c . We rearrange:

$$\omega^a \frac{X_{x,y}}{P'(\omega)} = \omega^{-c} \left(C - \left(-\frac{p_u(\omega) - p_v(\omega)}{P'(\omega)} - \frac{\omega^a}{P'(\omega)}(p_x(\omega) - p_y(\omega)) \right) \right)$$

However, the fact that C is not in T_ω means that as we increase c , the minimum value of quantity on the right above goes to infinity. Thus, $\min_{x,y} X_{x,y} \rightarrow \infty$. Hence, at some finite c , every one of these children has left the stack.

Recall the stack entries above are limits of the real stack entries we see for step $a + bn + c$, but by choosing n large enough, we can make the computation valid (because c is some finite number, so there are finitely many quantities to bring close to their limits). Hence for n large enough, the disconnectedness algorithm certifies that the limit set for $\omega + C\omega^{bn}$ is disconnected. \square

Figure 23 shows an example of T_ω near 0 for the renormalization point in Theorem 9.1.1. See also the pictures in [12].

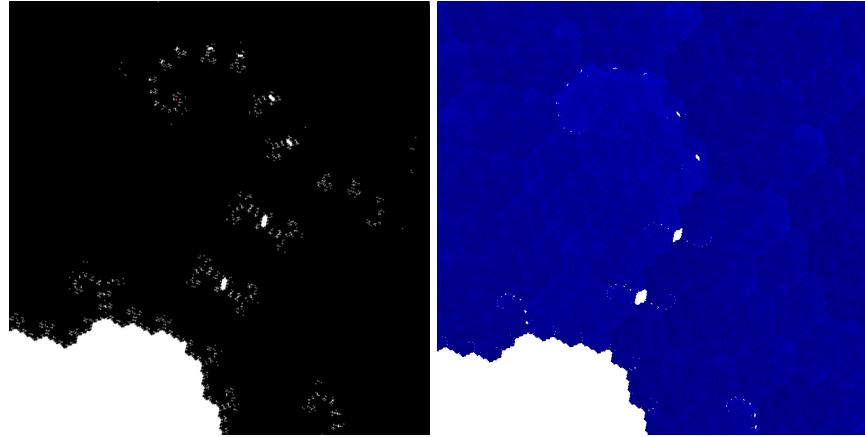


FIGURE 23. A portion of the limit set T_ω near 0 (left) for $\omega \approx 0.371859 + 0.519411i$ and set \mathcal{M} near ω (on right).

We end this section by proposing two (related) conjectures:

Conjecture 9.2.6. *The algebraic points in $\partial\mathcal{M}$ are dense in $\partial\mathcal{M}$.*

Conjecture 9.2.7. *Every point in $\partial\mathcal{M}$ not on the real axis is a limit of a sequence of holes with diameters going to zero.*

We believe that fixed points of renormalization are the key to both conjectures; such fixed points are on the one hand algebraic, and on the other hand points where \mathcal{M} is asymptotically self-similar, and asymptotically similar to the limit set of a 3-generator IFS. It is very easy for a connected limit set of a 3-generator IFS to fail to be simply-connected: irregularities in the frontiers of the translates overlap each other in complicated ways, cutting off tiny holes. Once there is one tiny hole, there will be infinitely many, accumulating densely in the boundary of the limit set; thus one expects the corresponding point in \mathcal{M} to be a limit of tiny holes.

The experimental evidence for Conjecture 9.2.7 is ambiguous. On the one hand, a computer-aided search (using `schottky`) will only reveal the holes at any scale that are big enough to see, so one must develop heuristics to identify promising regions for exploration. On the other hand, failure to find holes near some given frontier point does not rule out the possibility that they might exist, but be very elusive.

In a private communication, Boris Solomyak suggested that there might be no tiny holes accumulating at the point $i/\sqrt{2}$ in $\partial\mathcal{M}$; this is an especially good candidate counterexample to Conjecture 9.2.7, since although it is algebraic — and in fact a fixed point of renormalization — the limit set of the corresponding 3-generator IFS is full, and in fact *convex*. Thus one could not hope to prove the existence of a renormalization sequence of holes, certified by loops of limit traps, limiting to $i/\sqrt{2}$. On the other hand, very small holes *can* be found by hand, as close to $i/\sqrt{2}$ as the resolution allows — the (numerically certified) hole at $0.02269108 + 0.70320806i$ is a good example.

10. WHISKERS

In this section we discuss the subtle problem of the structure of \mathcal{M} and \mathcal{M}_0 near the real axis.

10.1. Whiskers are isolated. In light of Theorem 7.2.7 it might be surprising that the structure of \mathcal{M} and \mathcal{M}_0 near the real axis can be very complicated. In fact, as was already observed by Barnsley-Harrington [2], there is an open neighborhood of the points $\pm 1/2$ in $\pm[1/2, 1/\sqrt{2}]$ in which \mathcal{M} is totally real. We give an elementary proof of this fact, using the description of the limit set Λ_z as the values of certain power series in z , as described in Section 4.2. Getting a better estimate depends on analyzing a real 2-dimension IFS introduced by Shmerkin-Solomyak [10] which we discuss and study in Section 10.2.

Lemma 10.1.1 (Whiskers isolated). *There is some $\alpha > 1/2$ so that the intersection of \mathcal{M} with some open subset of \mathbb{C} is equal to the interval $[1/2, \alpha]$.*

Proof. Recall that for $e \in \partial\Sigma$ the image $\pi(e, z) \in \Lambda_z$ is the value of the power series $\pi(e, z) := a_0 + a_1 z + a_2 z^2 + \dots$ where the coefficients a_i are determined recursively from the infinite word e by the method in Proposition 4.2.1. The key point is that the nonzero coefficients alternate between 1 and -1 , starting with 1.

Let $z = 1/2 + \epsilon$ be real, for some small positive ϵ . The limit set Λ_z is exactly equal to the unit interval, and $f\Lambda_z = [0, z]$, $g\Lambda_z = [1 - z, 1]$ so that the intersection is exactly the interval $[1/2 - \epsilon, 1/2 + \epsilon]$. The words e with $\pi(e, z)$ in the overlap all start with fg^n or gf^n for some big n (depending only on n) so that the power series are of the form $z - z^{n+1} + \dots$ or $1 - z^{n+1} + z^m - \dots$ depending whether e starts with f or g , and in the second case $m > n+1$ (we include the possibility that $m = \infty$). In the first case, $d\pi(e, z)/dz = 1 - (n+1)z^{n+1} + \dots > 0.1$, while in the second case $d\pi(e, z)/dz = -(n+1)z^n + mz^{m-1} - \dots < 0$ for big n and any fixed $z < 1$. Since the derivative is holomorphic in z , this means that if we perturb z to $z + i\delta$ for some small positive δ , the imaginary part of $\pi(e, z)$ becomes *positive* for e beginning with f , and *negative* for e beginning with g (at least for $\pi(e, z)$ close to the interval $[1/2 - \epsilon, 1/2 + \epsilon]$), so that the two sets $f\Lambda_z$ and $g\Lambda_z$ are disjoint, and we are in the complement of \mathcal{M} . \square

10.2. A 2-dimensional IFS. We will push this argument further by analyzing the pairs $(\pi(u, z), d\pi(u, z)/dz)$ and $(\pi(v, z), d\pi(v, z)/dz)$ for left-infinite words $u, v \in \partial\Sigma$ starting with f and g respectively, and showing that for all real z in the interval $[0.5, 0.6684755]$ the pairs are disjoint.

Shmerkin-Solomyak [10] introduce a 2-dimensional *real* IFS acting on \mathbb{R}^2 whose limit set is precisely the pairs $(\pi(u, z), d\pi(u, z)/dz)$ for $u \in \partial\Sigma$. Explicitly, for real $z \in (-1, 1)$, define

$$f^{(1)} : (x, y) \rightarrow (zx, x + zy), \quad g^{(1)} : (x, y) \rightarrow (z(x - 1) + 1, x - 1 + zy)$$

and let L_z denote the limit set of the IFS generated by $f^{(1)}$ and $g^{(1)}$ (the notation is supposed to suggest the action of our familiar f and g on 1-jets). Analogous to our standard notation, we will write $u(z, x)$ for the action of the word $u \in \Sigma$ on $x \in \mathbb{R}^2$ for a parameter $z \in \mathbb{R}$. Also, we write $\pi(u, z) = \lim_{n \rightarrow \infty} u_n(z, x)$, where the limit does not depend on x .

Lemma 10.2.1. *Let $z \in \mathbb{R}$ and suppose $f^{(1)}(z, L_z)$ and $g^{(1)}(z, L_z)$ are disjoint. Then \mathcal{M} is totally real in an open neighborhood of z .*

Proof. This is the same argument as that in used in the proof of Lemma 10.1.1. \square

Since this condition is open, it can be certified numerically. Thus, if we define Ω_2 to be the subset of $z \in (-1, 1)$ for which L_z is connected, then $\overline{\mathcal{M}} - \mathbb{R} \cap \mathbb{R} \subseteq \Omega_2$. One can characterize Ω_2 as the set of real numbers z of absolute value at most 1 for which there is some power series $\zeta(z) := 1 + \sum_{n=1}^{\infty} a_n z^n$ where each $a_n \in \{-1, 0, 1\}$ for which $\zeta(z) = \zeta'(z) = 0$. We discuss later the question of whether there are points in Ω_2 which do not lie in the closure of the interior of \mathcal{M} .

Analogous to Ω_2 , one can study the subset $\Xi_2 \subseteq (-1, 1)$ consisting of z for which L_z contains the point $(1/2, 0)$, and then $\overline{\mathcal{M}_0} - \mathbb{R} \cap \mathbb{R} \subseteq \Xi_2$.

Shmerkin-Solomyak [10] define α to be the smallest positive real number in Ω_2 , and $\tilde{\alpha}$ to be the smallest real number such that $[\tilde{\alpha}, 1) \subseteq \Omega_2$. Experimentally they obtained estimates

$$\alpha \sim 0.6684755, \quad \tilde{\alpha} \sim 0.67$$

We improved the estimate of α to

$$\alpha \sim 0.6684755322100605954110550451436814$$

Getting a rigorous estimate of $\tilde{\alpha}$ is much harder, but experimentally we obtain $\tilde{\alpha} \sim 0.6693556$.

To obtain these estimates, we used an algorithm which is perfectly analogous to Algorithm 1, and is proved in essentially the same way. To describe this algorithm, we use the following shorthand:

$$A := \begin{pmatrix} z^{-1} & 0 \\ -z^{-2} & z^{-1} \end{pmatrix}, \quad Z := \begin{pmatrix} 1-z \\ -1 \end{pmatrix}$$

Furthermore, for a 2×1 column vector X we say X is *small* if $|X_1| < 1$ and $|X_2| < \sup_{k \geq 1} 2k|z|^{k-1}$, where k is an integer. Note that for z real with $|z| < 1$, this latter inequality reduces to the analysis of a small fixed number of cases for k . In the regime in which we are interested, z will be quite close to 0.66, so the relevant cases are $k = 2$ and $k = 3$, and in practice the inequality reduces to $|X_2| < 2.681165$.

The justification for this algorithm is essentially the same as that of Algorithm 1. To ask whether L_z is connected is equivalent to asking whether $f^{(1)}(z, L_z) \cap$

Algorithm 2 No Multiple Roots(z, depth)

```

 $V \leftarrow \{AZ\}$ 
 $d \leftarrow 0$ 
while  $V \neq \emptyset$  or  $d < \text{depth}$  do
     $W \leftarrow \emptyset$ 
    for all  $X \in V$  do
        if  $A(X - Z)$  is small then  $W \leftarrow W \cup A(X - Z)$ 
        if  $AX$  is small then  $W \leftarrow W \cup AX$ 
        if  $A(X + Z)$  is small then  $W \leftarrow W \cup A(X + Z)$ 
     $V \leftarrow W$ 
     $d \leftarrow d + 1$ 
    if  $V = \emptyset$  then
        return true
    else
        return false

```

$g^{(1)}(z, L_z) = \emptyset$, which is equivalent to asking whether the set of differences contains 0. Just as in Section 6, the set of differences between points in L_z is a limit set itself. We denote the set of differences by L'_z , and we note it is the limit set of the IFS generated by the three maps

$$F_{-1} : X \mapsto BX - Z, \quad F_0 : X \mapsto BX, \quad F_1 : X \mapsto BX + Z,$$

where

$$B := \begin{pmatrix} z & 0 \\ 1 & z \end{pmatrix}.$$

Note $B = A^{-1}$. We obtain these maps by looking at how pairs of maps $(f, f), (f, g), (g, f)$, and (g, g) act on differences of points; there are only three distinct maps. Since $F_1 L'_z$ consists of differences between points in L_z whose corresponding infinite words begin with $g^{(1)}$ and $f^{(1)}$, respectively, to check whether L_z is connected it suffices to check whether $0 \in F_1 L'_z$.

To determine whether $0 \in F_1 L'_z$, we start with a box R centered at $(0, 0)$, which is sent inside itself under the three maps F_{-1}, F_0, F_1 . We want to consider $F_1 L'_z$, so first we apply F_1 . Next, we subdivide $F_1 R$ into its three subboxes, which are $F_1 F_{-1} R, F_1 F_0 R$, and $F_1 F_{-1} R$, and discard those which cannot contain 0. We then subdivide again, and so on. Suppose that X is the center (image of $(0, 0)$) of an image of R under a word of length n . Since the centers of $F_{-1,0,1} L'_z$ are at $-Z, 0, Z$, respectively, the centers of the children of X will be at the points $X - B^n Z, X, X + B^n Z$. For simplicity, it makes sense to rescale the problem at every step by $A = B^{-1}$. Hence, we initialize the algorithm with the rescaled Z , i.e. AZ . Then we add $-Z, 0, +Z$, and rescale by A again, and so on. Any child which lies too far from the origin can be discarded, which is exactly what the smallness condition guarantees. The precise constants in the smallness condition follow from an analysis of how the rectangle $|X_1| \leq a, |X_2| \leq b$ behaves under the maps $F_{-1,0,1}$. It is easy to see that the infinite strip $|X_1| \leq 1$ is sent inside itself, so L'_z lies inside this strip. Then if we consider the images of the four points $(-1, -b), (1, -b), (1, b), (-1, b)$, we find that the image with the largest second coordinate is $(1, b)$ under the word $F_{-1}^k F_1^\infty$, and this image has second coordinate $2k|z|^{k-1}$. Therefore, if we find the

k maximizing that expression and set $b = 2k|z|^{k-1}$, then the limit set must lie in the rectangle $[-1, 1] \times [-b, b]$.

If we run Algorithm 2 on our numerical value for α , the output is quite interesting. For the correct theoretical value of α , the set V of children viable to each depth will never be empty, and the same must be true for our numerical approximation (of course, this is how we find the approximation in the first place). But what is not obvious from the definition (although it is intuitively plausible) is the experimental fact that the size of $|V|$ is uniformly bounded independently of the depth d , and there is apparently a *unique* lineage viable to infinite depth. If we denote the children $A(X - Z)$, AX , $A(X + Z)$ of the vector X by L, M, R respectively, then the (numerically) unique viable descendent of the initial vector AZ to 194 generations is of the form

$$L^3 \prod_i (R^i M) \text{ for } i = 1223327566251816335432839221548243362315$$

i.e. the first few terms are $LLLRMRRMRRMRRRM\cdots$. One can think of the values of i as analogs of the terms in the continued fraction expansion of a number. In fact, the analogy is quite good: if any viable sequence for an initial vector $AZ := AZ(t)$ is eventually periodic, we obtain an identity of the form $p_1(A)Z = p_2(A)Z$ for distinct polynomials p_1, p_2 with coefficients in $\{-1, 0, 1\}$, and therefore deduce that t^{-1} is a root of $p_1 - p_2$ and is therefore algebraic. The branching of the algorithm is shown in Figure 24.

In view of our experimental evidence, it seems reasonable to make the following conjecture:

Conjecture 10.2.2 (Unique lineage). *For α as above, there is a unique child at every stage with viable descendants to all future depths. Furthermore, this viable lineage consists of the initial segments in the sequence $L^3 \prod_i (R^i M)$ for some sequence $i = 122 \cdots$ as above, where the terms are uniformly bounded.*

In a similar vein, we define β to be the smallest positive real number in Ξ_2 , and $\tilde{\beta}$ to be the smallest real number such that $[\tilde{\beta}, 1) \subseteq \Xi_2$. Using similar methods we obtain the following estimates

$$\beta \sim 0.67133041244176126776, \quad \tilde{\beta} \sim 0.728781$$

(the same caveat about $\tilde{\beta}$ applies). It is easy to modify Algorithm 2 to determine, for a given real z , when there are infinite words u, v so that $\pi(u, z) = \pi(v, z) = 1/2$ and $d\pi(u, z)/dz = d\pi(v, z)/dz = 0$; we need only consider children $AX - Z$ and $AX + Z$ for each X in the stack V , and otherwise the algorithm runs in exactly the same way.

Figure 25 gives numerical plots of the subset of the intervals $[\beta, \tilde{\beta}] \cap \Xi_2$ and $[\alpha, \tilde{\alpha}] \cap \Omega_2$.

This figure strongly suggests that $\Xi_2 \cap [\beta, \tilde{\beta}]$ is totally disconnected, while $\Omega_2 \cap [\alpha, \tilde{\alpha}]$ appears to contain many solid intervals. In fact, our method of traps can be easily adapted to this more complicated IFS, and in Section 10.3 we give a method to certify interior points in Ω_2 .

10.3. Intervals in Ω_2 . Recall from Section 10.2 that Ω_2 is the set of positive real numbers $z < 1$ for which the IFS $L_z \subseteq \mathbb{R}^2$ generated by the affine linear maps

$$f^{(1)} : (x, y) \mapsto (zx, x + zy), \quad g^{(1)} : (x, y) \mapsto (z(x - 1) + 1, x - 1 + zy)$$

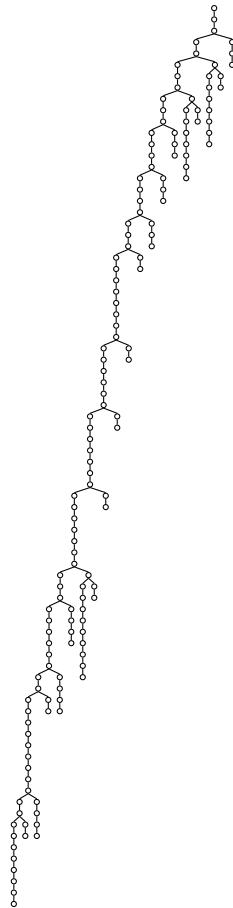


FIGURE 24. The branching of Algorithm 2 on the (numerical) input α . The long vertical chains are all R , so reading down the left edge produces strings of R 's of lengths 1,2,2,3,3,2,7, etc, agreeing with the “continued fraction” expansion of α .

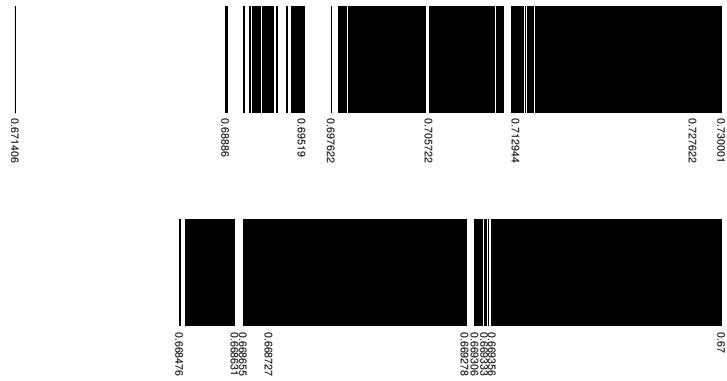


FIGURE 25. Numerical plots of Ξ_2 (top) and Ω_2 (bottom).

is connected. By abuse of notation, we denote these generators by f and g throughout this section. Note that both generators have constant Jacobian

$$B(z) := \begin{pmatrix} z & 0 \\ 1 & z \end{pmatrix}$$

Throughout this section we restrict attention to real z in the interval $[0.668, 0.67]$. The analog of Lemma 5.2.2 is the following:

Lemma 10.3.1 (Affine Short Hop Lemma). *With $z \in [0.668, 0.67]$, suppose that fL_z and gL_z contain points at distance δ apart in the L^1 metric on \mathbb{R}^2 . Then for any word u of length at least 6, the $0.9006 \cdot \delta/2$ neighborhood of $u(z, L_z)$ in the L^1 metric is path connected.*

Proof. The proof is identical to that of Lemma 5.2.2, except that one must take into account the fact that $B(z)$ does not uniformly contract the L^1 metric. However, for z in the interval in question, $B(z)^n$ multiplies the L^1 metric by at most $0.67^n + n \times 0.67^{n-1}$ which is < 0.9006 for $n \geq 6$. \square

The analog of Proposition 7.1.6 is the following:

Proposition 10.3.2 (Affine traps). *Suppose for some $z \in \Omega_2$ that there are words u, v beginning with f and g of length at least 6 so that $u(z, L_z)$ and $v(z, L_z)$ cross transversely. Then z is an interior point in Ω_2 .*

Proof. Since $u(z, L_z)$ and $v(z, L_z)$ cross transversely, the same is true for their ϵ -neighborhoods, for some sufficiently small fixed ϵ . Thus the same is true for the $\epsilon/2$ -neighborhoods of $u(z', L_{z'})$ and $v(z', L_{z'})$ whenever $|z - z'|$ is small enough, depending on z and ϵ . Thus, we choose such a z' , and suppose δ is the L^1 distance from $f(z', L_{z'})$ to $g(z', L_{z'})$, where $\delta \ll \epsilon$. Then the $0.9\delta/2$ neighborhoods of $u(z', L_{z'})$ and $v(z', L_{z'})$ are path connected, so by transversality, there is some point within L^1 distance $0.9\delta/2$ from both $u(z', L_{z'})$ and $v(z', L_{z'})$, and consequently the L^1 distance from $u(z', L_{z'})$ to $v(z', L_{z'})$ is at most 0.9δ . But then $\delta \leq 0.9\delta$ so that $\delta = 0$ and $z' \in \Omega_2$, as claimed. \square

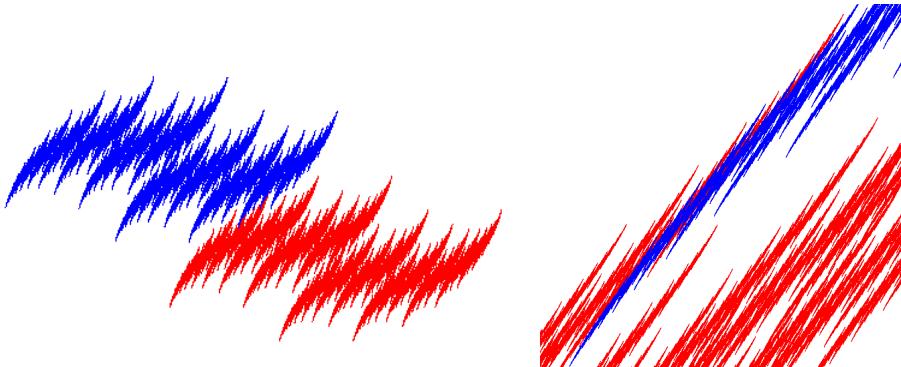


FIGURE 26. The limit set for $z = 0.669027$. The visually evident trap on the right certifies that this point lies in the interior of Ω_2 .

Such affine traps may be found and certified numerically; for example $z = 0.669027$ satisfies the proposition for the words

$$fgggf g f f f f f g f f f f f f \quad \text{and} \quad g f f f g g g g f g g g g g g f g g,$$

and we deduce that 0.669027 is an interior point in Ω_2 ; see Figure 26. One might hope to prove an analog of Bandt's Conjecture (i.e. Theorem 7.2.7) for the set Ω_2 ; that is, that the interior is dense in Ω_2 . However our proof of Theorem 7.2.7 uses in several ways the fact that points in the limit set are holomorphic functions of the parameter, which of course can no longer be true for the real parameter z . Nevertheless, such a proof does not seem beyond reach, and we comfortably conjecture:

Conjecture 10.3.3. *Affine traps are dense in Ω_2 , and hence the interior of Ω_2 is dense in Ω_2 .*

Recall that $\overline{\mathcal{M}} - \mathbb{R} \cap \mathbb{R} \subseteq \Omega_2$. It is not known whether there are any points in Ω_2 which do not lie in the closure of the interior of \mathcal{M} . However, the following lemma relates this to Conjecture 10.3.3. For clarity, we write $u(z) = \pi(u, z)$.

Lemma 10.3.4. *For every $u \in \partial\Sigma$ and $b \in \mathbb{D}^* \cap \mathbb{R}$, we have that*

$$\lim_{x+iy \rightarrow b} \frac{1}{y} \operatorname{Im}(u(x+iy)) \rightarrow u'(b).$$

The rate of this convergence does not depend on u . Consequently, the limit set Λ_{x+iy} scaled vertically by $1/y$ converges in the Hausdorff topology to L_b .

Proof. It is an easy calculus exercise to show the lemma if $x+iy$ approaches a vertically, i.e. if x is fixed at b . However, we desire convergence in general, so we will need to look at power series. Write $u(z)$ as the power series $u(z) = \sum_{k=0}^{\infty} a_k z^k$. Then

$$u(x+iy) = \sum_{k=0}^{\infty} a_k (x+iy)^k = \sum_{k=0}^{\infty} a_k \sum_{j=0}^k \binom{k}{j} (iy)^j x^{k-j}$$

The terms which contribute to the imaginary part of this sum are exactly those for which j is odd. Hence

$$\operatorname{Im}(u(x+iy)) = \sum_{k=0}^{\infty} a_k \sum_{\ell=1}^{2\ell+1=k} \binom{k}{2\ell+1} (-1)^{\ell-1} y^{2\ell+1} x^{k-(2\ell+1)}$$

The top limit on the inner sum indicates that we should run the inner sum until $2\ell+1$ is larger than k . Also note we are recording the imaginary part of $u(x+iy)$, so the $\pm i$ terms have disappeared. Therefore,

$$\frac{1}{y} \operatorname{Im}(u(x+iy)) = \sum_{k=0}^{\infty} a_k \left(kx^{k-1} + \sum_{\ell=2}^{2\ell+1=k} \binom{k}{2\ell-1} (-1)^{\ell-1} y^{2\ell} x^{k-(2\ell+1)} \right)$$

The entire sum is controlled in absolute value by $\sum_{k=0}^{\infty} |a_k| |x+iy|^k \leq \sum_{k=0}^{\infty} |x+iy|^k$, which is uniformly convergent for $x+iy \in \mathbb{D}^*$. Therefore, as $x+iy \rightarrow b$, the entire sum converges, at a rate controlled independently of u , to $\sum_{k=0}^{\infty} a_k kb^{k-1} = u'(b)$. The point in L_b associated with u has coordinates $(u(b), u'(b))$ in \mathbb{R}^2 , and the point in the vertically scaled copy of Λ_{x+iy} has coordinates $(\operatorname{Re}(u(x+iy)), \frac{1}{y} \operatorname{Im}(u(x+iy)))$, and the lemma follows. \square

If affine traps are dense in Ω_2 , then near any point in Ω_2 , by Lemma 10.3.4 there are nonreal parameters which have a trap and therefore lie in the interior of \mathcal{M} . So every point in Ω_2 would be in the closure of the interior of \mathcal{M} ; i.e. Conjecture 10.3.3 implies $\overline{\mathcal{M}} - \mathbb{R} \cap \mathbb{R} = \Omega_2$.

11. HOLES IN \mathcal{M}_0

\mathcal{M}_0 is path connected [3], but Bousch's proof is somewhat indirect. His strategy is to show that every point can be joined by a path to some parameter with absolute value close to 1. Since \mathcal{M}_0 contains an annulus around the unit circle, this gives path connectedness. He does not directly address what the paths in \mathcal{M}_0 actually look like, or when a (polygonal) path near \mathcal{M}_0 can be approximated by a path contained in \mathcal{M}_0 .

In this section, we show how to certify the existence of a point in \mathcal{M}_0 in a neighborhood of a given point and how to certify a path in \mathcal{M}_0 in a neighborhood of a given polygonal path. If we can certify paths, we can certify loops, and thus exotic holes in \mathcal{M}_0 . As with \mathcal{M} , by a *hole* in \mathcal{M}_0 , we mean a connected component of the complement which is distinct from the connected component of the complement which contains 0.

Just as \mathcal{M} closely resembles the limit set Γ_z at many points, \mathcal{M}_0 closely resembles Λ_z . Thus the methods in this section are closely related to the methods we developed in Section 5.4 to construct paths in Λ_z (e.g. Proposition 5.4.2).

11.1. Complex analysis. In this section, we prove a lemma in complex analysis, but we first motivate it. Suppose we have a holomorphic function $h(z)$, and we find that $h(z_0)$ is quite close to a value we desire c . We would like to conclude that there is a z_1 near z_0 so that $h(z_1) = c$. If the derivative of h is bounded away from 0, and does not vary much near z_0 , then h can be well approximated by a linear function, and z_1 can be found.

Thus to certify the existence of such a z_1 , and to prove the validity of the certificate, is not technically difficult. However, Lemma 11.1.1 is organized carefully to be of use to us later, and it can be confusing to read. One should understand the lemma as saying "if there are four constants r, C, C', δ which satisfy the hypotheses, then the conclusion holds". Do not worry about where the constants come from at this stage. This lemma is very similar to Lemma 3.1 in [12].

Lemma 11.1.1. *Let h be a holomorphic function and $z_0, c \in \mathbb{C}$ with $|h(z_0) - c| < \epsilon$. Suppose there are $r, C, C' > 0$ and $0 < \delta < 1$ such that $C' \leq |h'(z)| \leq C$ for all z with $|z - z_0| < r$, and*

$$r \geq \frac{\epsilon}{\delta} \frac{1 + \frac{\delta^2}{1-\delta}}{C' - C \frac{\delta}{1-\delta}} = \frac{\epsilon(1 - \delta + \delta^2)}{\delta((1 - \delta)C' - \delta C)}$$

Then there exists a unique $z_1 \in \mathbb{C}$ with $|z_0 - z_1| \leq \epsilon \frac{(1-\delta+\delta^2)}{(1-\delta)C'-\delta C}$ such that $h(z_1) = c$.

Proof. First, it suffices to prove the theorem with $c = 0$ by translation, so we will make that assumption.

Write $h_a(z)$ for the affine part of the power series for h , centered at z_0 , i.e. $h_a(z) = h(z_0) + h'(z_0)(z - z_0)$. Under h_a , the circle $z_0 + de^{i\theta}$ of radius d is mapped to the circle $h(z_0) + |h'(z_0)|de^{i\theta}$. Therefore, if $d|h'(z_0)| > \epsilon$, the image circle will

enclose 0, and hence $0 \in h_a(B_d(z_0))$, or equivalently h_a will have a zero within $B_d(z_0)$, the ball of radius d centered at z_0 .

Now consider h ; it might not be affine, and we record the remainder term as R_1 :

$$h(z) = h_a(z) + R_1(z) = h(z_0) + h'(z_0)(z - z_0) + R_1(z).$$

Suppose that there were a radius d such that for all $0 \leq \theta \leq 2\pi$, we had $|h'(z_0)|d - |R_1(z_0 + de^{i\theta})| \geq \epsilon$. In other words, the error in the affine approximation is smaller than the radius of the affine image circle minus ϵ . Then the image of the circle $z_0 + de^{i\theta}$ under h would have to contain $B_\epsilon(h(z_0))$, and hence h would have a zero in $B_d(z_0)$. Additionally, this follows immediately from Rouche's theorem, which also gives the claimed uniqueness.

To prove the lemma, then, it suffices to find a d such that $|h'(z_0)|d - \epsilon \geq |R_1(z_0 + de^{i\theta})|$ for all $0 \leq \theta \leq 2\pi$. From Taylor's theorem and Cauchy's derivative estimates, there is an inequality

$$|R_1(z_0 + de^{i\theta})| \leq \frac{M_r d^2}{r^2 - rd} \leq \frac{(\epsilon + Cr)d^2}{r^2 - rd},$$

where $M_r = \max_\theta |h(z_0 + re^{i\theta})|$, and the estimate is valid whenever $d < r$, and we can also estimate $M_r \leq \epsilon + Cr$.

Set $d = \delta r$. Rearranging the inequality in the hypothesis of the lemma, we have

$$C' \delta r - \epsilon \geq \frac{(\epsilon + Cr)\delta^2 r^2}{r^2 - \delta r^2}.$$

Since $|h'(z_0)| \geq C'$, and plugging in $d = \delta r$, we have

$$|h'(z_0)|d - \epsilon \geq \frac{(\epsilon + Cr)d^2}{r^2 - rd} \geq \frac{M_r d^2}{r^2 - rd}$$

Therefore, $d = \delta r$ satisfies the necessary inequality, so there is $z_1 \in B_{\delta r}(z_0)$ with $h(z_1) = c$. Since making r smaller maintains the validity of the bounds C, C' for $|h'(z)|$, we may shrink r until the inequality in the lemma is an equality, so the claimed bound on $|z_0 - z_1|$ holds. \square

Remark 11.1.2. The hypotheses of Lemma 11.1.1 may seem somewhat technical, but in fact they are not difficult to check in practice. We set r to be quite small but still large compared to ϵ , and we get bounds on the derivative. Then δ can be found by trial and error or any minimum-finding algorithm. In fact, Mathematica produces an explicit formula for the δ which minimizes the expression on the right of the inequality for r ; this formula is rather large and unedifying, so we omit it.

One feature we will make use of is that Lemma 11.1.1 can be checked for large collections of elements in $\partial\Sigma$ at the same time, since two words with a large common prefix will satisfy the same C, C' bounds with similar values of ϵ .

Remark 11.1.3 (Derivative bounds). Lemma 11.1.1 requires good derivative bounds on $h'(z)$ a given ball $B_{z_0}(r)$. A naive way to approach this is to get a universal upper bound K on the second derivative and then state that $|h'(z)| < |h'(z_0)| + Kr$ on $B_{z_0}(r)$. This is typically a bad estimate because r can be large compared to the potential change in $h'(z)$. Here is a better way. Since $|h'(z)|$ is holomorphic, its maximum will lie on the boundary of $B_{z_0}(r)$. Cover the boundary circle of $B_{z_0}(r)$ with many (say, 100) small balls, use the naive approach on these small balls, and take the maximum. Because the radius on which we apply the naive approach is now quite small, our error will be much less.

11.2. Paths in \mathcal{M}_0 . In this section, we explain how to find paths in \mathcal{M}_0 . These paths will be rather short, but by piecing them together, we can produce loops and thus certify holes in \mathcal{M}_0 .

We now give some initial observations about paths in \mathcal{M}_0 to clarify the construction to follow. To each point z in \mathcal{M}_0 , there is a set of distinguished words in $\partial\Sigma$; namely, the words x such that $\pi(x, z) = 1/2$. Therefore, if we have a path $\gamma : [0, 1] \rightarrow \mathbb{D}^*$ such that the image of γ lies in \mathcal{M}_0 , there is a combinatorial map $\lambda : [0, 1] \rightarrow \partial\Sigma$ such that $\pi(\lambda(t), \gamma(t)) = 1/2$. Of course, λ is not uniquely defined, as there may be more than one word mapping to $1/2$ for a given parameter.

In order to build paths in \mathcal{M}_0 , we essentially go in the other direction. Given two words $a, b \in \partial\Sigma$, we first build a nice combinatorial path interpolating between a and b . Then, provided we are close enough to \mathcal{M}_0 , we show how apply Lemma 11.1.1 to produce a path of parameters which drags this combinatorial path along $1/2$.

In this lemma, we recall the notation $p_w(z) = \pi(w, z)$, the power series associated with $w \in \partial\Sigma$.

Lemma 11.2.1. *Suppose there are $\epsilon, r, C, C' > 0$, $0 < \delta < 1$, and $z_0 \in \mathbb{C}$ such that*

$$(1) |z_0| + r < 1$$

$$(2) r \geq \frac{\epsilon(1 - \delta + \delta^2)}{\delta((1 - \delta)C' - \delta C)}.$$

$$(3) \text{ For all } v \in u\partial\Sigma \text{ we have } |p_v(z_0) - 1/2| < \epsilon.$$

$$(4) \text{ For all } v \in u\partial\Sigma \text{ and } z \in B_v(z_0) \text{ we have } C' < |p'_v(z)| < C.$$

Then for all $v \in u\partial\Sigma$, there is a unique $Z(v) \in B_{\delta r}(z_0)$ such that $p_v(Z(v)) = 1/2$. Consequently, there is a map $Z : u\partial\Sigma \rightarrow \mathcal{M}_0 \cap B_{\delta r}(z_0)$ such that $p_v(Z(v)) = 1/2$. Furthermore, Z is uniformly continuous and the image $Z(u\partial\Sigma)$ is path connected.

Proof. That the map Z exists and is well-defined (single-valued) follows immediately from Lemma 11.1.1, so the content of the lemma is the uniform continuity and path connectedness. We first address the former. This is with respect to the Cantor metric, so it suffices to show that if two words $w_1, w_2 \in u\partial\Sigma$ have a sufficiently long common prefix, then their images under Z are close (independent of what the prefix is).

Let K be equal to $|z_0| + r$. We claim that there exists a constant I such that if $w_1, w_2 \in u\partial\Sigma$ have a common prefix w of length at least I , then

$$|Z(w_1) - Z(w_2)| < \frac{2K^{|w|}}{|1 - K|} \frac{(1 - \delta + \delta^2)}{((1 - \delta)C' - \delta C)}.$$

We now prove this claim. We remark that u is a prefix of w , since w_1, w_2 already have the common prefix u . By Lemma 3.1.1, for a given z , the limit set Λ_z is contained in a ball of radius $|1 - z|/2(1 - |z|) < 1/(1 - |z|)$ centered at $1/2$, so if u is a word of length n , then $u(z, \Lambda_z)$ is contained in a ball of size $|z|^n/(1 - |z|)$ centered at $u(z, 1/2)$. In our situation, then, the limit set $w(Z(w^\infty), \Lambda_{Z(w^\infty)})$ is contained inside a ball of radius $\frac{K^{|w|}}{|1 - K|}$. Therefore, we have

$$|p_{w_1}(Z(w^\infty)) - 1/2|, |p_{w_2}(Z(w^\infty)) - 1/2| < \frac{K^{|w|}}{|1 - K|}.$$

We are going to apply Lemma 11.1.1 to w_1 and w_2 to get nearby roots, but there is a slight subtlety. We have derivative bounds on all words in $u\partial\Sigma$ and $z \in B_r(z_0)$,

but to apply Lemma 11.1.1, we need derivative bounds in a ball centered at $Z(w^\infty)$. We can achieve these bounds in the following way. Since $Z(w^\infty) \in B_{\delta r}(z_0)$, the derivative bounds C' and C must be valid over $B_{(1-\delta)r}(Z(w^\infty))$. So if $|w| > I$ for I sufficiently long enough, then

$$(1-\delta)r \geq \frac{K^{|w|}}{|1-K|} \frac{(1-\delta+\delta^2)}{\delta((1-\delta)C'-\delta C)},$$

so we can apply Lemma 11.1.1 to the words w_1, w_2 at the point $Z(w^\infty)$ with radius $(1-\delta)r$ and $\epsilon = \frac{K^{|w|}}{|1-K|}$; this gives nearby z_1, z_2 so $\pi(w_1, z_1) = 1/2$ and $\pi(w_2, z_2) = 1/2$. But Z is uniquely defined, so $Z(w_1) = z_1$ and $Z(w_2) = z_2$, and hence

$$|Z(w_1) - Z(w^\infty)|, |Z(w_2) - Z(w^\infty)| < \frac{K^{|w|}}{|1-K|} \frac{(1-\delta+\delta^2)}{((1-\delta)C'-\delta C)}.$$

The claim that Z is uniformly continuous follows from the triangle inequality, and therefore the image of Z is compact. It remains to show that the image $Z(u\partial\Sigma)$ is path connected.

Analogous to the set W we constructed to build paths through Λ_z in Section 5, given any two words $a, b \in u\partial\Sigma$, we will construct a combinatorial path through $u\partial\Sigma$ interpolating between them, and then show that applying Z to this path gives a continuous path in \mathcal{M}_0 . Given a finite word w , denote by \bar{w} the word obtained from w by swapping f and g . Note that if w is finite and there is a parameter z such that $w(z, 1/2) = 1/2$, then $\bar{w}(z, 1/2) = 1/2$, so $p_{w^\infty}(z) = 1/2$ and $p_{\bar{w}^\infty}(z) = 1/2$. Additionally, for any infinite word w_*^∞ obtained by taking an infinite power of w and swapping arbitrary copies of w for \bar{w} , we have $p_{w_*^\infty}(z) = 1/2$. Therefore, $Z(w^\infty) = Z(w_*^\infty)$.

Now let H be a set of pairs of elements of $u\partial\Sigma$ indexed by the dyadic rationals and constructed inductively as follows. First set $H_0 = (a, a)$ and $H_1 = (b, b)$. Next, given $H_{k2^{-i}}$ and $H_{(k+1)2^{-i}}$, let v be the maximal common prefix of $H_{k2^{-i}, 2}$ and $H_{(k+1)2^{-i}, 1}$, and let

$$H_{k2^{-i}+2^{-(i+1)}} = \Phi_{(v^\infty, \bar{v}^\infty)}(H_{k2^{-i}, 2}, H_{(k+1)2^{-i}, 1})$$

That is, $H_{k2^{-i}+2^{-(i+1)}}$ is either $(vv^\infty, v\bar{v}^\infty)$ or $(v\bar{v}^\infty, vv^\infty)$ depending on the first letters of $H_{k2^{-i}, 2}$ and $H_{(k+1)2^{-i}, 1}$ after the initial prefix. By the observation above, the map Z is well-defined on the pairs in H because each pair consists of two words of the form w_*^∞ for the same w .

By induction, if $k2^{-i} \leq r_1 \leq r_2 \leq (k+1)2^{-i}$, then H_{r_1} and H_{r_2} have a common prefix of length at least $|u| + i$. Here we say H_{r_1} and H_{r_2} have a common prefix of length n if at least one of the four possible pairings of a word in H_{r_1} and H_{r_2} has a common prefix of length n . Since Z is uniformly continuous, this means that $Z(H_r)$ is continuous as a function of the dyadic rational r , so $Z(H_r)$ extends continuously to $r \in [0, 1]$, and $Z(u\partial\Sigma)$ is compact, so the image $Z(H_r)$ is contained in $Z(u\partial\Sigma)$ and is a path beginning at $Z(a)$ and ending at $Z(b)$, and the lemma is proved. \square

11.3. Holes in \mathcal{M}_0 . By a *hole* in \mathcal{M}_0 , we mean a connected component of the complement which is distinct from the “obvious” large connected component containing the point 0. Lemma 11.2.1 shows how to find a map Z which takes a set of words $u\partial\Sigma$ into \mathcal{M}_0 in a nice way. In order to find a hole in \mathcal{M}_0 , we will find words $u_0, \dots, u_{n-1} \in \partial\Sigma$ satisfying Lemma 11.2.1, thus giving maps $Z_i : u_i\partial\Sigma \rightarrow \mathcal{M}_0$. The images $Z(u_i\partial\Sigma)$ are path connected, and we will show, for all i with i taken

modulo n , that we have $Z(u_i\partial\Sigma) \cap Z(u_{i+1}\partial\Sigma) \neq \emptyset$. Thus, there is a path passing through each image in turn. Furthermore, we'll show that the images encircle a point which is not in \mathcal{M}_0 . This will complete the proof of the existence of a hole in \mathcal{M}_0 .

Lemma 11.2.1 does not say what the images $Z(u_i\partial\Sigma)$ will look like; it only gives balls which are guaranteed to contain them. To get a more precise picture, we do the following: enumerate all words Σ_m of some length m , and apply Lemma 11.2.1 to $Z(u_i x \partial\Sigma)$ for every $x \in \Sigma_m$. If all these computations succeed, we obtain 2^m balls, and we know that (1) there is a point in $Z(u_i\partial\Sigma) \subseteq \mathcal{M}_0$ inside each ball and (2) these points are connected by paths inside $Z(u_i\partial\Sigma)$.

Therefore, if we can use this technique to exhibit, for each i , that the sets $Z(u_i\partial\Sigma)$ and $Z(u_{i+1}\partial\Sigma)$ lie transverse to each other, in the sense of traps, then they intersect.

Theorem 11.3.1 (Holes in \mathcal{M}_0). *There is a hole in \mathcal{M}_0 .*

Proof. After the discussion above, this proof reduces to showing the pictures shown in Figure 27 and asserting that they were produced using the method above. Note that this produces a *loop* in \mathcal{M}_0 , and checking if a parameter is *not* in \mathcal{M}_0 is rigorous, so it suffices to exhibit a single pixel in the middle of the putative hole which is not in \mathcal{M}_0 . Many such pixels are easily visible. \square

REFERENCES

- [1] C. Bandt, *On the Mandelbrot set for pairs of linear maps*, Nonlinearity **15** (2002), no. 4, 1127–1147
- [2] M. Barnsley and A. Harrington, *A Mandelbrot set for pairs of linear maps*, Phys. D. **15** (1985), no. 3, 421–432
- [3] T. Bousch, *Paires de similitudes*, preprint, 1988; available from the author's webpage
- [4] T. Bousch, *Connexité locale et par chemins höldériens pour les systèmes itérés de fonctions*, preprint, 1992; available from the author's webpage
- [5] D. Calegari and A. Walker, *schottky*, software available from
<https://github.com/dannycalegari/schottky>
- [6] V. Climenhaga and Y. Pesin, *Lectures on fractal geometry and dynamical systems*, Student Mathematical Library, **52** AMS, Providence, RI, 2009.
- [7] D. Epstein, J. Cannon, D. Holt, S. Levy, M. Paterson and W. Thurston, *Word processing in groups*, Jones and Bartlett, Boston, 1992
- [8] P. Mercat, *Semi-groupes fortement automatiques*, Bull. Soc. Math. France **141** (2013), no. 3, 423–479
- [9] A. Odlyzko and B. Poonen, *Zeros of polynomials with 0, 1 coefficients*, L'Enseignement Math. **39** (1993), 317–348
- [10] P. Shmerkin and B. Solomyak, *Zeros of $\{-1, 0, 1\}$ Power series and connectedness loci for self-affine sets*, Experimental Math. **15** (2006), no. 4, 499–511
- [11] B. Solomyak, *Mandelbrot set for a pair of line maps: the local geometry*, Analysis Thy. Appl. **20** (2004), 149–157
- [12] B. Solomyak, *On the ‘Mandelbrot set’ for pairs of linear maps: asymptotic self-similarity*, Nonlinearity **18** (2005), no. 5, 1927–1943
- [13] B. Solomyak and H. Xu, *On the ‘Mandelbrot set’ for a pair of linear maps and complex Bernoulli convolutions*, Nonlinearity **16** (2003), no. 5, 1733–1749
- [14] W. Thurston, *Entropy in dimension one*, boundarys in Complex Dynamics: In Celebration of John Milnor's 80th Birthday, pages 339–384. Princeton University Press, 2014
- [15] G. Tiozzo, *Galois conjugates of entropies of real unimodal maps*, preprint, arXiv:1310.7647

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CHICAGO, CHICAGO, IL, 60637
E-mail address: `dannyc@math.uchicago.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MI, 48109
E-mail address: `kochsc@umich.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CHICAGO, CHICAGO, IL, 60637
E-mail address: `akwalker@math.uchicago.edu`

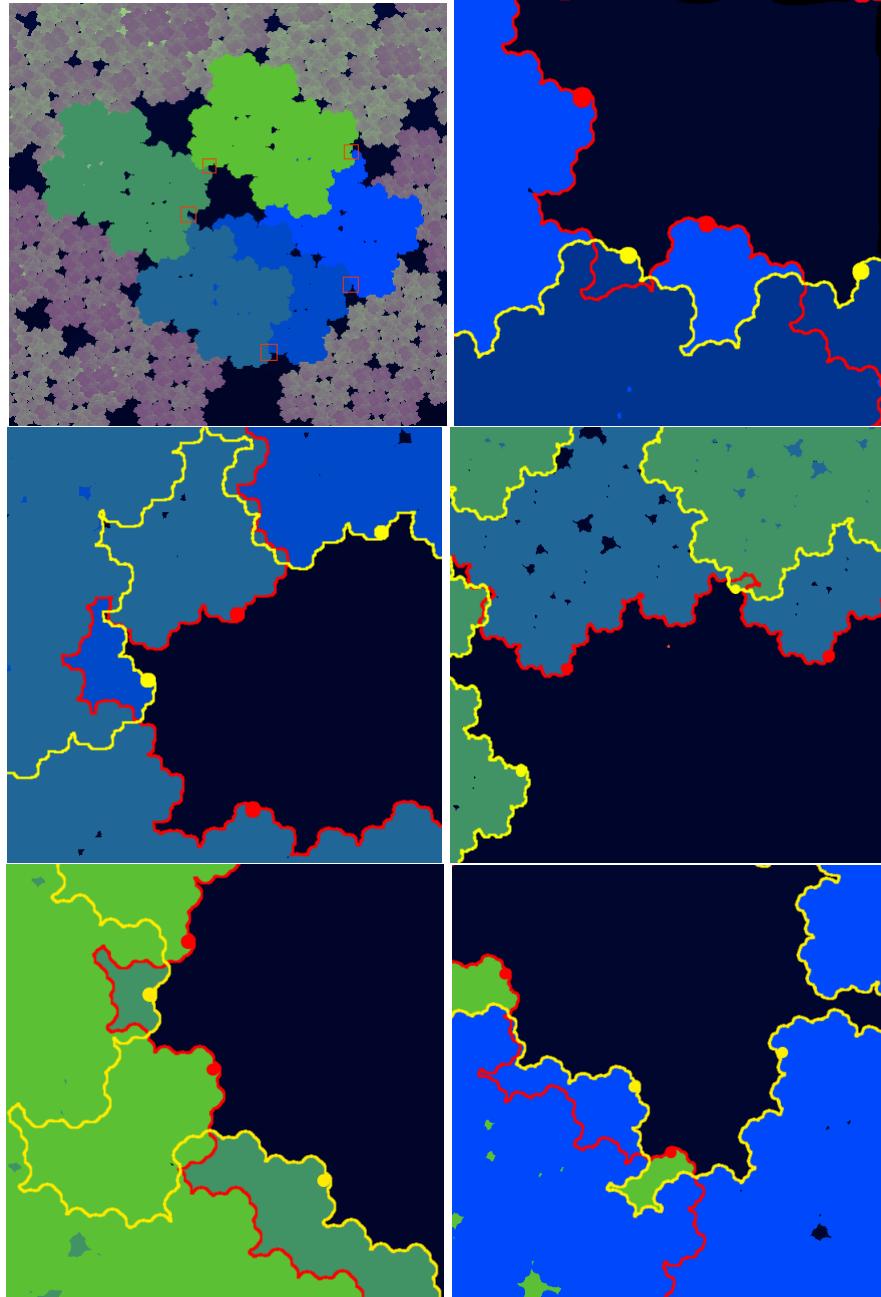


FIGURE 27. The upper left picture shows the images of $Z(u_i\partial\Sigma)$ for $0 \leq i \leq 4$, and the red boxes indicate the zoomed regions shown in the following pictures. Each picture is made up of many small disks guaranteed to contain points in \mathcal{M}_0 . Four linked disks are highlighted in each picture to show that the various images of Z must intersect, and each image is path connected, so there is a loop in \mathcal{M}_0 .