ROTATIONAL SUBSETS OF THE CIRCLE UNDER z^d

ALEXANDER BLOKH, JAMES M. MALAUGH, JOHN C. MAYER, LEX G. OVERSTEEGEN, AND DANIEL K. PARRIS

Abstract. This paper is a study of invariant sets that have "geometric" rotation numbers, which we call *rotational* sets, for the angle-tripling map $\sigma_3 : \mathbb{T} \to \mathbb{T}$, and more generally, the angle-d-tupling map $\sigma_d : \mathbb{T} \to \mathbb{T}$ for $d \geq 2$. The precise number and location of rotational sets for σ_d is determined by $d-1$, $\frac{1}{d}$-length open intervals, called *holes*, that govern, with some specifiable flexibility, the number and location of *root gaps* (complementary intervals of the rotational set of length $\geq \frac{1}{d}$). In contrast to σ_2, the proliferation of rotational sets with the same rotation number for σ_d, $d > 2$, is elucidated by the existence of canonical operations allowing one to *reduce* σ_d to σ_{d-1} and *construct* σ_{d+1} from σ_d by, respectively, removing or inserting "wraps" of the covering map that, respectively, destroy or create/enlarge root gaps.

1. Introduction

Why would one study the properties of certain invariant sets, that we call *rotational sets*, of the complex analytic map z^d restricted to the unit circle? One reason involves the connection between the map z^d and a complex polynomial map f with connected Julia set J. Let $\mathbb{C}$ denote the complex plane and $\mathbb{C}_\infty = \mathbb{C} \cup \{\infty\}$ the Riemann sphere. Suppose f is a complex polynomial map on $\mathbb{C}$. Let K denote its filled Julia set (J union its bounded complementary domains) and $\mathbb{U}_\infty = \mathbb{C}_\infty \setminus K$. Denote the complement of the closed unit disk $\overline{\mathbb{D}}$ in $\mathbb{C}_\infty$ as $\mathbb{D}_\infty$. Then there exists a conformal isomorphism, $\phi : \mathbb{D}_\infty \to \mathbb{U}_\infty$, the Böttcher uniformization, that conjugates f on $\mathbb{U}_\infty$ to z^d on $\mathbb{D}_\infty$. In this setting, points on the unit circle correspond uniquely, via ϕ, to *prime ends* of $\mathbb{U}_\infty$.

In complex dynamics prime ends are associated with *external rays*, ϕ-images of radial rays of $\mathbb{D}_\infty$, and the external rays are used extensively to study polynomial Julia sets. For example, that periodic rays "land" on repelling or rationally indifferent periodic points of the Julia set, and that every such point is a landing point of a periodic external ray, has been a very productive tool in understanding the dynamics of polynomials on Julia sets [3, 11, 5].

Depending on the complex polynomial map f, the Julia set may be simple or quite complicated, locally connected or non-locally connected, with many cyclically permuted and pre-periodic Fatou components in its complement. At one extreme, when the Julia set is locally connected the external rays all "land" at points in the Julia set, and the Julia set is a topological (and dynamical) quotient space of the circle. When the Julia set is not locally connected the *impressions* (a kind of "super-closure") of some of these external rays are non-degenerate and may be used to understand the topological and dynamical structure

Date: September 2, 2003.

1991 *Mathematics Subject Classification.* Primary: 54F20; Secondary: 30C35.

Key words and phrases. complex dynamics, circle dynamics, angle tripling, rotation number, covering map.

Corresponding author is John C. Mayer. The first author was supported by NSF research grant number DMS-0140349. The second and fifth authors were supported in part by the UAB Mathematics Fast-Track Program, supported by NSF infrastructure grant number 9901207. The fourth author was supported by NSF research grant number DMS-0072626.

of the Julia set [7, 8]. At the other extreme (though as yet unrealized), it may be that the Julia set is an indecomposable continuum, and every external ray has as its impression the entire Julia set [10].

Of particular interest in the non-locally connected, but not indecomposable, case are the external rays that correspond to irrational rotational sets. This paper may be seen as the first step in extending the results of Kiwi in [8], where irrationally indifferent points in the dynamics are explicitly excluded in order to prove the theorems, in the direction of [7]. In the latter, one seeks to understand the correspondence between certain irrational rotational sets in the circle at infinity and the boundaries of Siegel disks and Cremer points in the Julia set.

Other reasons one may want to study the properties of rotational sets are from purely geometric and combinatorial standpoints. What do they look like? How are they situated on the circle? What properties must they have? Can they be categorized? Parameterized? If so, how many are in each category? How do they correspond to parameters?

This paper will be separated into three main parts. First there is an introduction to the maps and basic definitions that are used throughout. Next, the main results, which we divide into *structural theorems* and *counting theorems*, are presented in Section 2, inspired by the paper by Bullet and Sentenac [2] about z^2 on the unit circle. Finally, proofs of the results are given, first in Section 3 of the structural theorems and then of applications in Sections 4 and 5. Parameterization of rotational sets will be covered in a subsequent paper by the second author [9].

Let us note here that for the sake of brevity some known or easy to prove statements are not supplied with proofs.

1.1. **Basic Definitions.** Let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ be the unit circle coordinatized by $[0, 1)$. We orient the circle in the positive direction: counterclockwise. For $A \subset \mathbb{T}$ let $\widehat{A}$ denote the maximal subset of $\mathbb{R}$ such that $A = \widehat{A}/\mathbb{Z}$. Let $\sigma_d : \mathbb{T} \to \mathbb{T}$ be the map given by $\sigma_d(t) = dt \pmod 1$, induced by the complex-valued map $z \to z^d$ on the unit circle in the complex plane. Often σ_2 is called *angle-doubling*, σ_3 *angle-tripling*, and so on. We will assume henceforth that $d \geq 2$. All *maps* are continuous functions.

Definition 1.1. A map $f : \mathbb{T} \to \mathbb{T}$ is *topologically exact* iff given any interval $I \subset \mathbb{T}$, there is a positive integer n such that $f^n(I) = \mathbb{T}$.

The following theorem is well-known.

Theorem 1.2. *A topologically exact covering map from the circle to itself is conjugate to σ_d for some $d \geq 2$.*

Definition 1.3. A map $f : X \to Y$ is called *monotone* if $f^{-1}(y)$ is connected for all $y \in f(X)$.

Definition 1.4 (Lift)**.** Let $e : \mathbb{R} \to \mathbb{T}$ be the natural projection map defined by $e(x) = e^{2\pi i x}$. A *lift* of a circle function $f : \mathbb{T} \to \mathbb{T}$ is a function $\widehat{f} : \mathbb{R} \to \mathbb{R}$ such that $e \circ \widehat{f}(x+m) = f(x) \circ e$ for all $x \in [0, 1]$ and all $m \in \mathbb{Z}$.

Definition 1.5 (Degree 1 and Order-Preserving)**.** We say a map $f : \mathbb{R} \to \mathbb{R}$ is *degree 1* iff $f(x + 1) = f(x) + 1$ for all $x \in \mathbb{R}$. A map $f : \mathbb{T} \to \mathbb{T}$ is *degree 1* iff f has a degree 1 lift $\widehat{f} : \mathbb{R} \to \mathbb{R}$. We say a map $f : \mathbb{T} \to \mathbb{T}$ is *order-preserving* iff f is monotone and degree 1. Let A be a closed subset of $\mathbb{T}$. We say $f : \mathbb{T} \to \mathbb{T}$ is *order-preserving on A* iff $f|_A$ can be extended to an order-preserving map $F : \mathbb{T} \to \mathbb{T}$ where $F|_A = f|_A$.

Figure 1 illustrates an order-preserving function and some of its degree 1 monotone lifts.

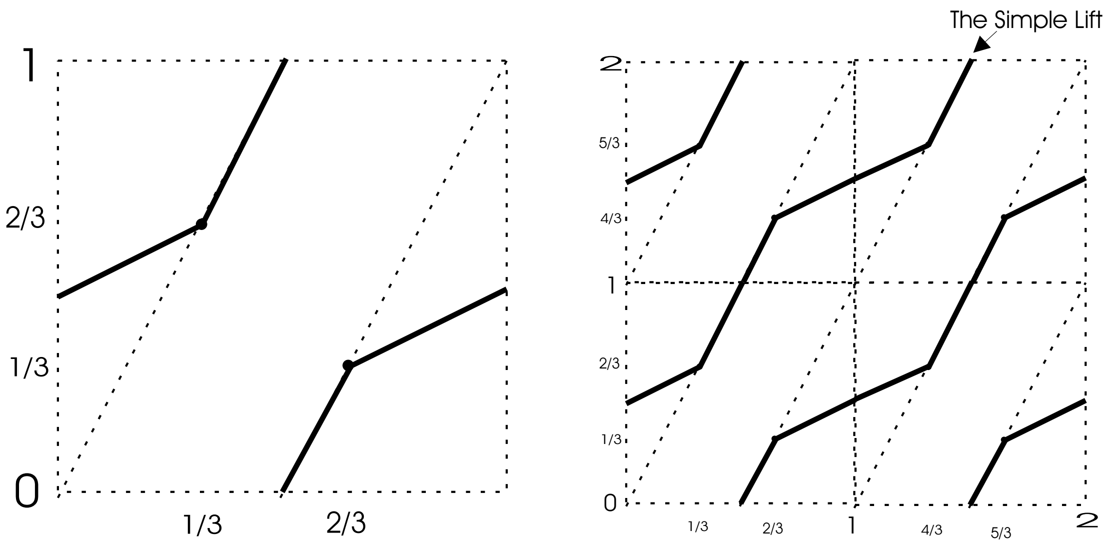

FIGURE 1. Left: "connect-the-dot" extension of the periodic orbit $\{\frac{1}{3}, \frac{2}{3}\}$. Right: illustration of several lifts of that extension, including the simple lift.

Definition 1.6 (Rotational). A subset $A \subset \mathbb{T}$ is *invariant* under $f : \mathbb{T} \to \mathbb{T}$ iff $f(A) = A$. A subset $A \subset \mathbb{T}$ is *subinvariant* under $f : \mathbb{T} \to \mathbb{T}$ iff $f(A) \subset A$. A closed invariant set A is *minimal* iff no proper closed subset of A is invariant. An invariant set A is *rotational* iff A is closed and f is order-preserving on A.

A degree 1 monotone map $f : \mathbb{T} \to \mathbb{T}$ is order-preserving. Also, f is an order-preserving map iff $\mathbb{T}$ is rotational with respect to f. If A is a proper subset of $\mathbb{T}$, invariant under σ_d (like a rotational or minimal invariant set), then A is a totally disconnected subset of $\mathbb{T}$. A rotational set may or may not be minimal. A minimal invariant set may or may not be rotational.

Definition 1.7. A degree 1 monotone circle map f has a *rotation number* $\rho(f)$ defined by

$$\rho(f) = \lim_{n\to\infty} \frac{(\widehat{f})^n(\widehat{x})}{n} \pmod 1,$$

where $x \in \mathbb{T}$, $\widehat{x}$ is any pre-image of x under the natural projection map e, and $\widehat{f}$ is any lift of f.

Proposition 1.8. *For any order-preserving map f we have that $\rho(f)$ is well-defined.*

Proof. Let $\widehat{f}$ be any lift of f. It can be shown, since $\widehat{f}$ is monotone, that $\lim_{n\to\infty} \frac{(\widehat{f})^n(\widehat{x})}{n}$ (mod 1), is independent of the point $\widehat{x}$. This value is called the lift rotation number of $\widehat{f}$ and is denoted $\rho(\widehat{f})$. Next, it can be shown that this lift rotation number, $\rho(\widehat{f})$, is the same for any lift $\widehat{f}$ of f. Hence, f has a well-defined rotation number $\rho(f)$. (See [1].) $\square$

The proof of the above proposition tells us that if we have an order-preserving map f of the circle, then we can use any lift $\widehat{f}$, and any point $\widehat{x} \in \mathbb{R}$, to compute $\rho(f)$. However, it will be often useful to consider a certain lift, called the *simple* lift, defined below (see Figure 1).

Definition 1.9. Given a function $f : \mathbb{T} \to \mathbb{T}$, then its *simple lift* will be the unique lift $\widehat{f} : \mathbb{R} \to \mathbb{R}$ of f such that $\widehat{f}(0) \in [0, 1)$.

Corollary 1.10. *If $A \subset \mathbb{T}$ is rotational under $f : \mathbb{T} \to \mathbb{T}$, then A has a well-defined rotation number $\rho(A)$ given by*

$$\rho(A) = \rho(\widehat{F}) = \lim_{n\to\infty} \frac{(\widehat{F})^n(\widehat{a})}{n} \pmod 1 \qquad \widehat{a} \in \widehat{A},$$

where $\widehat{F}$ is any lift of an order-preserving extension F of $f|_A$.

Proof. For any order-preserving extension of $f|_A$, we know that its rotation number exists and is given by the rotation number of any of its lifts by Proposition 1.8. All we need show then is that the rotation number is the same for any order-preserving extension of $f|_A$. This can be easily seen by noting that, no matter which lifts of the order-preserving extensions we choose, we can pick $\widehat{a} \in \widehat{A}$, which, with the invariance of $\widehat{A}$, forces the lifts, and thus the extensions, to have the same rotation number by Definition 1.6. □

With the above definition we can endow any rotational subset of the circle with its well-defined rotation number. It may seem that this applies to all periodic orbits and thus one should be able to compute the rotation number for any periodic orbit. This is not true, however; the periodic orbit must satisfy Definition 1.6, i.e., be rotational, in order to have a geometric rotation number, and it is easy to give examples of periodic orbits of, say, σ_2 which are not rotational. For example, the σ_2-periodic orbit $\{\frac{1}{5}, \frac{2}{5}, \frac{4}{5}, \frac{3}{5}\}$ is not rotational; one cannot find an order-preserving extension of σ_2 restricted to this set.

Theorem 1.11 (Lemma 18.8 [11])**.** *If A is a compact set invariant under σ_d such that σ_d is one-to-one on A, then A is finite.*

A consequence of Theorem 1.11 is that any invariant set on which σ_d is one-to-one must be a finite union of periodic orbits (Compare with Propositions 1.12 and 2.2).

Proposition 1.12. *If A is a minimal invariant set under σ_d, then A is either a periodic orbit (in which case, if A is rotational, a well-defined rational rotation number can be associated with it) or A is a Cantor set (in which case, if A is rotational, a well-defined irrational rotation number can be associated with it).*

The notion of a gap in a set is fairly intuitive. It is just an interval in the complement of a set. The complement of a periodic orbit is a finite union of gaps, which are just the intervals between any two spatially adjacent points. The complement of a Cantor set is a countably infinite union of disjoint open intervals. *Gaps* and *root gaps* are formally defined below in Definition 1.13. This definition and the following proposition will be used later in Sections 4 and 5.

Definition 1.13 (Gaps)**.** Let A be a closed, invariant set under σ_d. The components of $\mathbb{T} \setminus A$ are called *gaps*. A gap G of A is a *root* gap iff the *length* of G satisfies $l(G) \geq \frac{1}{d}$. A root gap G is *loose* iff $\frac{n}{d} < l(G) < \frac{n+1}{d}$ for some $n \in \mathbb{N}$, whereas, G is *taut* iff $l(G) = \frac{n}{d}$ for some $n \in \mathbb{N}$. The number n is called the *root number* of G.

Proposition 1.14. *All root gaps of a finite rotational set are loose.*

1.2. **Rotational Sets for σ_2.** This paper is inspired by what Bullett and Sentenac [2] proved about σ_2-rotational sets.

In their paper, Bullett and Sentenac completely characterize minimal rotational sets under σ_2. Not only does their work shed some new light on the topic of rotational sets under σ_2, it serves as a source in which many previously known results on the topic have been pulled together and discussed in a cohesive manner. These results include the fact that all σ_2-rotational sets must be contained in a semicircle, and that no semicircle contains more than one minimal rotational set. Consequently, they parameterize the minimal rotational sets by the semicircles that contain them. Thus, a minimal rotational set must contain an open semicircle in its complement. This concept of an interval in the complement of a rotational set is fundamental to the characterization of rotational sets under σ_d for any $d \geq 2$. Therefore, the following definition is now introduced.

Definition 1.15. Given arbitrary $d \geq 2$, and a σ_d-rotational set A, define a *hole* of A as an open interval of length $\frac{1}{d}$ in $\mathbb{T} \setminus A$. A co-existing *set of holes* will be any set of intervals of length $\frac{1}{d}$ in $\mathbb{T} \setminus A$ that are disjoint.

It will be seen later in Theorem 2.1 why the disjoint requirement is included in the definition of a set of holes. From now on, any reference to a set of holes will mean a set of holes as defined above.

Below is a list of some of the results found in Bullett and Sentenac [2]; observe that remarks made in parentheses are made exclusively in the context of the map σ_2.

(1) Each σ_2-rotational set is contained in some closed semicircle. (Note that this is equivalent to the rotational set having a hole, whether it be unique or not.)
(2) Each closed semicircle contains exactly one minimal rotational set.
(3) Given the starting point μ of any closed semicircle, there is a combinatorial algorithm allowing one to determine the minimal rotational set in that semicircle.
(4) For each rotation number in $[0, 1)$ there exists exactly one minimal rotational set with that rotation number.
(5) For any rotational Cantor set one, and only one, semicircle contains it (see Proposition 1.11). (Note this is equivalent to containing one unique hole.)
(6) For any rotational periodic orbit there is an interval of semicircles containing it (see Proposition 1.14). (Note that this is equivalent to having more than one hole.)
(7) Rotational sets can be parameterized by μ, where μ is the starting point of any semicircle in the circle.
(8) The parameter μ leads to the definition of a rotation function $\rho_2 : [0, \frac{1}{2}] \to [0, 1]$ (we consider only the semicircles whose starting point belongs to $\Delta_2 = [0, \frac{1}{2}]$. The graph of this function is a topological and measure theoretic Devil's staircase (see Definition 1.16, Definition 1.17, and Figure 2).

The following examples are given to connect some of the above results to the previous definitions. As mentioned earlier, the periodic orbit $\{\frac{1}{5}, \frac{2}{5}, \frac{4}{5}, \frac{3}{5}\}$ is minimal invariant, but not rotational. Note that it is not contained in a semicircle and fails to satisfy the last part of Definition 1.6. The set $\{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, ..., \frac{1}{2^n}, ..., 0\}$ is rotational, because it is contained in a semicircle, but not minimal.

Definition 1.16. A map $f : X \to [0, 1]$ is said to be a *Devil's staircase in topology* if f is onto, monotone, and there exists a dense open set $U \subset X$ such that f is locally constant on U.

Definition 1.17. A map $f : X \subset \mathbb{R}^n \to [0, 1]$ is said to be a *Devil's staircase in measure* if it satisfies Definition 1.16, there exists a set $A \subset [0, 1]$ such that $\lambda(A) = 1$, $f^{-1}(A) \subset X \setminus U$, and $\lambda^*(f^{-1}(A)) = 0$, where λ is one-dimensional Lebesgue measure and λ^* is the appropriate Lebesgue measure in $\mathbb{R}^n$.

1.3. **Motivating Questions.** The purpose of this paper is to determine the aspects of Bullett and Sentenac's work that can be generalized to σ_d for $d > 2$. Some specific questions that arise include:

- How are the rotational sets situated on the circle?
- How do minimal rotational sets correspond to rotation number?
- How many minimal rotational sets are there per rotation number?

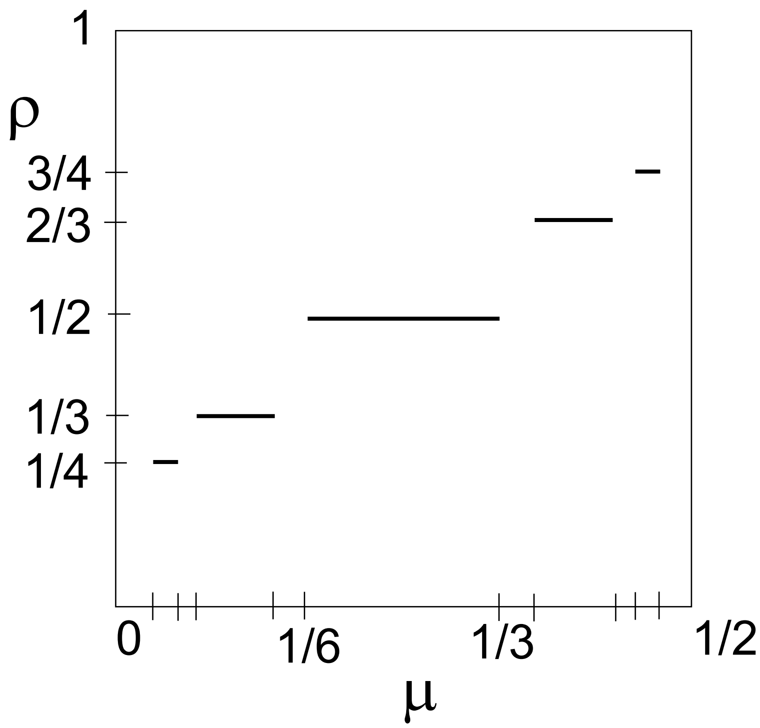

FIGURE 2. Partial construction of the Devil's staircase $\rho_2 : \Delta_2 \to [0, 1]$.

- Is there a suitable parameter space for the minimal rotational sets under σ_d?

In this paper we focus on the dynamics of rotational sets. In a subsequent paper, the second author will focus on the parameter space for σ_d, $d > 2$, and measure-theoretic aspects of rotational sets.

When considering the above questions it becomes clear that the picture for the rotational sets under σ_d for $d > 2$ is markedly different from that for $d = 2$. For example, under simple inspection of σ_3 one can find three periodic orbits with rotation number $\frac{1}{2}$, whereas under σ_2 every rotation number had only one corresponding rotational set (see Theorems 2.8 and 2.11). In addition, one of these periodic orbits, $\{\frac{1}{4}, \frac{3}{4}\}$, quickly dismisses the notion that all σ_d-rotational sets are contained in an arc of the circle of length $\frac{1}{d}$, hence our focus on the holes of the complement of the rotational sets, as mentioned above (see Theorem 2.1). Another interesting feature of σ_d when $d > 2$ is that one can prove the existence of *loose* Cantor sets (see Definition 2.10) using Theorem 2.6. This result is stated in Theorem 5.7.

2. Main Theorems

In this section, we state, explain, and exemplify the main theorems of this paper without proof. The proofs of the main theorems are contained in Sections 3, 4, and 5.

2.1. **Structural Theorems.** With the notion that rotational sets under σ_d must be contained in a $\frac{1}{d}$-arc dismissed (recall that we assume that $d > 2$), another approach must be taken to determine how they are located on the circle. It was noted earlier that the existence of holes in the set would be an important approach to this question. The precise number and location of minimal rotational sets for σ_d are determined by the $d-1$, $\frac{1}{d}$-length *holes* that govern, with some specifiable flexibility, the number and location of root gaps.

The proliferation of rotational sets for σ_d, $d > 2$ is elucidated by the existence of canonical operations allowing one to *reduce* σ_d to σ_{d-1} and *construct* σ_{d+1} from σ_d by, respectively, removing or inserting "wraps" of the covering map that, respectively, destroy or create/enlarge root gaps. A σ_d-rotational set reduces to a σ_{d-1}-rotational set with the same rotation number. Similarly, from a σ_d-rotational set we can construct a σ_{d+1}-rotational set with the same rotation number. Usually these operations do not produce unique rotational sets, and *where* the operations of reduction and construction are carried out within a rotational set determines *what* the resulting rotational set is. We address reduction and construction in Theorems 2.3 and 2.6.

Theorem 2.1 (Hole Theorem)**.** *Let A be a closed subset of $\mathbb{T}$. σ_d is order-preserving on A iff $\mathbb{T} \setminus A$ contains $d-1$ pairwise disjoint open intervals, each of length $\frac{1}{d}$, called* holes *(of A).*

Theorem 2.1 for the specific case $d = 3$: σ_3 is order-preserving on $A \subset \mathbb{T}$ iff $\mathbb{T} \setminus A$ contains two disjoint open intervals (holes) each of length $\frac{1}{3}$.

If an invariant set A satisfies the hypotheses of Theorem 2.1 then it must be rotational.

Proposition 2.2. *At least one root gap of any rotational Cantor set must be taut.*

Theorem 2.3 (Reduction Theorem)**.** *Let $J = (a, a + \frac{1}{d} \pmod 1)$ be any open interval of length $\frac{1}{d}$ on the circle (i.e., a hole). Then $X = \mathbb{T} \setminus \bigcup_{i=0}^{\infty} \sigma_d^{-i}(J)$ contains a maximal invariant Cantor set C. Moreover, there exists a monotone map $m : \mathbb{T} \to \mathbb{T}$ that at most two-to-one semi-conjugates $\sigma_d|_C$ to σ_{d-1}.*

Corollary 2.4. *Let A be a σ_d-rotational set. Then for any $J \subset \mathbb{T} \setminus A$, the corresponding set $m(A)$ is a σ_{d-1}-rotational set and $\rho(m(A)) = \rho(A)$.*

Theorem 2.5. *Any σ_d-rotational Cantor set is minimal.*

Theorem 2.6 (Construction Theorem)**.** *Let A be a rotational set for σ_d. Let x_1^0 be a point of $\mathbb{T}$. Then there is a rotational set $\widetilde{A}$ for σ_{d+1}, determined by x_1^0, such that $\sigma_{d+1}|_{\widetilde{A}}$ is at most two-to-one semi-conjugate to $\sigma_d|_A$. The semi-conjugacy is actually a conjugacy if and only if $x_1^0 \in \mathbb{T} \setminus A$.*

Corollary 2.7. *With the hypotheses of Theorem 2.6, $\rho(\widetilde{A}) = \rho(A)$.*

2.2. **Counting Theorems.** For σ_2, there is exactly one minimal rotational set for each rotation number. In contrast, for σ_d, $d > 2$, the situation is remarkably otherwise. For example, there are $d-1$ fixed points for σ_d. For σ_3 there are three period 2 rotational orbits and eight period 3 rotational orbits (of those eight orbits of period 3 there are four with rotation number $\frac{1}{3}$, and four with rotation number $\frac{2}{3}$).

Theorem 2.8 (Counting Rotational Periodic Orbits, Goldberg [4])**.** *Under the map σ_d, the number of periodic orbits with rotation number $\frac{p}{q}$ (in lowest terms) is C_q^{d-2+q}.*

Theorem 2.9 (Goldberg [4])**.** *A rotational set for σ_d with a given rational rotation number contains at most $d-1$ periodic orbits.*

Theorem 2.8 for σ_3: the number of periodic orbits with rotation number $\frac{p}{q}$ is $q + 1$. Theorems 2.8 and 2.9 were proved by Goldberg [4]. We include them here for completeness. Our proofs are different and more direct. Moreover, our rational rotational sets do not have to be finite (that is, a finite union of periodic rotational orbits).

Definition 2.10 (Taut and Loose Cantor sets)**.** A rotational Cantor set is *taut* iff all of its root gaps are taut. Otherwise, it is *loose*.

A σ_d-rotational Cantor set is *taut* if there exists exactly one choice for the $d-1$ holes it must have by Theorem 2.1.

By c we denote the cardinality of the continuum, that is, the cardinality of the reals $\mathbb{R}$.

Theorem 2.11 (Counting Rotational Cantor sets)**.** *For each σ_d-rotational Cantor set C_α with rotation number α, there are c-many rotational Cantor sets under σ_{d+1} which correspond to C_α in the sense of Construction Theorem 2.6. Moreover, if C_α had n taut root gaps, then countably many of these corresponding σ_{d+1} rotational Cantor sets have n taut root gaps, while uncountably many of them have $n+1$ taut root gaps.*

An interesting consequence of the above theorem, and of previously known results, is that for any fixed irrational rotation number α, there are no loose Cantor sets under σ_2, countably many under σ_3, and uncountably many under σ_d for $d > 3$ with rotation number α. On the other hand, for any irrational rotation number α, there is one taut Cantor set under σ_2, and uncountably many under σ_d for $d > 2$ with rotation number α.

Corollary 2.12. *For each irrational rotation number α, there are c-many rotational Cantor sets under σ_d, $d > 2$, with rotation number α.*

In fact, we prove a bit more in Section 5.4: *For each σ_d-rotational Cantor set C_α with rotation number α, there is a circularly ordered collection $\mathcal{S}_\alpha$ of σ_{d+1}-rotational Cantor sets with rotation number α, arising from inserting a wrap at each point or gap of C_α.* The induced order on the set $\mathcal{S}_\alpha$ corresponds to the order of the points (and gaps) of C_α on the circle $\mathbb{T}$.

Our approach is different than that of Goldberg and Tresser in [6]. They show that rotational Cantor sets of σ_d with a given (irrational) rotation number are in one-to-one correspondence with a $(d-2)$-dimensional simplex; indeed, rotational sets of a given rotation number are characterized by the proportion of their points between any two successive fixed points of σ_d.

3. Proof of Structural Theorems

Here we prove the three main structural theorems of Section 2.1 that are the foundation for the proofs of more specific results counting and categorizing rotational sets in Sections 4 and 5.

3.1. **Proof of Hole Theorem 2.1.**

Definition 3.1 (Algorithm for Detecting that σ_d is Order-Preserving on A)**.** We first note that σ_d is not order-preserving on $\mathbb{T}$ and, as a consequence, fails this algorithm. So let A be a proper closed subset of $\mathbb{T}$. Start with a point x_0 such that x_0 is the starting endpoint of any gap of $\mathbb{T} \setminus A$ (with the usual counterclockwise positive direction on the circle). We will start to construct our extension, F of $\sigma_d|_A$ from A to the entire circle. Our initial point on the graph of F is obviously $(x_0, \sigma_d(x_0))$.

¿From x_0 look for points of A in the interval $I = (x_0, x_0 + \frac{1}{d})$.

(1) If there is no point of A in I then construct the extension by crossing horizontally to the point $(x_0 + \frac{1}{d}, \sigma_d(x_0))$. Thus we are defining F to be constant on I. Our new point for the next iteration of the construction is $x_1 = x_0 + \frac{1}{d}$.

(2) If there is a point of A in $(x_0, x_0 + \frac{1}{d})$ then we climb the σ_d graph until we reach the right-most point of A in $(x_0, x_0 + \frac{1}{d}]$ (i.e., on the interval from x_0 to the rightmost point of A in $(x_0, x_0 + \frac{1}{d}]$ we are defining our extension as $F = \sigma_d$). Let x_1 be this rightmost point, and define $F|_{[x_0,x_1]} = \sigma_d$. Then x_1 becomes the starting point for the next iteration of the process.

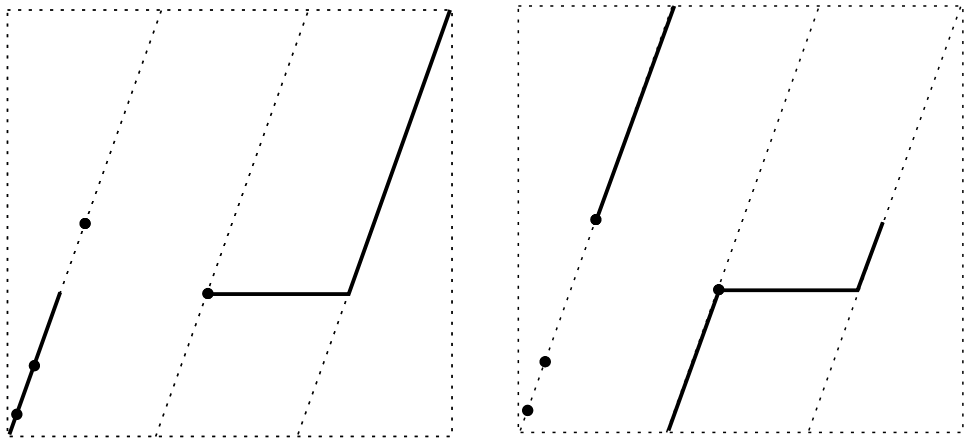

FIGURE 3. Two attempts at applying order-preserving algorithm to the periodic orbit $\{\frac{1}{20}, \frac{3}{20}, \frac{9}{20}, \frac{7}{20}\}$ of σ_3.

Note we are guaranteed on the second iteration to get past the point $x_0 + \frac{1}{d}$ in the domain of our extension. We keep repeating the process, and the above remark tells us we can keep the process finite. The resulting map F coincides with σ_d in A. In fact, it differs from σ_d exactly on gaps in A of length greater than or equal to $\frac{1}{d}$.

Clearly the process described above realizes the following more explicit algorithm:

(1) fix a root gap $I = (a, b)$;
(2) consider the root number $i = i(I)$ of I and define $F|_I$ as a constant $\sigma_d(a)$ on the arc $(a, a + \frac{i}{d})$ and as σ_d on the rest of I;
(3) repeat this construction for all root gaps of f.

We will be using one approach or the other depending on the situation. In any case, it is clear that the process leads to a new map F in which all wraps around the circle and realized on root gaps of A are "taken out". Of course this can be done for any set A, so the construction always goes through, but it does not always leads to an orientation preserving map F of the circle. We say that the process *succeeds* if at some point we will traverse the entire circle in the domain and get back to our starting point x_0, having successfully constructed an *order-preserving* extension of $\sigma_d|_A$. Otherwise the process *fails*. The way our algorithm works (and fails) is illustrated in Figure 3.

Let us now prove Hole Theorem 2.1.

Proof. First, it will be shown that if σ_d is order-preserving on A, then $\mathbb{T} \setminus A$ must contain $d-1$ holes. Indeed, suppose that there is an order-preserving extension, G, of $\sigma_d|_A$. Consider the lengths of images of arcs in more detail now (they are computed as the lengths of the curves parameterized by points of arcs). If we compute the length of the image of $\mathbb{T}$ under G it is going to be equal to 1 while for the map σ_d the length of the image of $\mathbb{T}$ is d. The difference between the two equals $d-1$ and is generated by the root gaps. Indeed, if J is a gap which is not a root gap then the fact that G is monotone implies that the length of $G(J)$ and the length of $\sigma_d(J)$ are equal. Moreover, on A both G and σ_d coincide. On the other hand for each root gap I of A its G-image is a proper arc in $\mathbb{T}$. Hence the difference between the length of $\sigma_d(I)$ and the length of $G(I)$ equals exactly the root number $i(I)$. Thus, the fact that G is monotone means that the sum of numbers $i(I)$ taken over all root gaps must be $d-1$ and therefore $\mathbb{T} \setminus A$ contains $d-1$ pairwise disjoint open intervals, each of length $\frac{1}{d}$ (they are called *holes* (of A)).

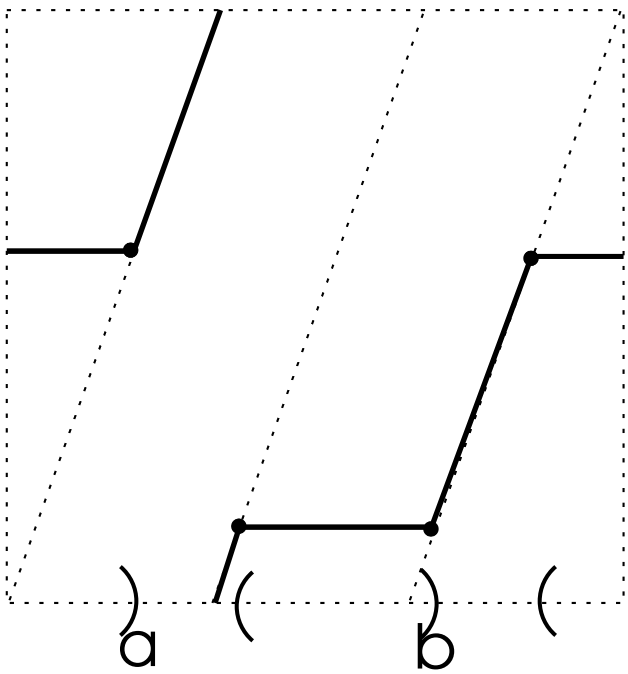

FIGURE 4. Illustration of extension of σ_d restricted on the complement to the union of holes.

On the other hand if $\mathbb{T}\setminus A$ contains $d-1$ pairwise disjoint open intervals, each of length $\frac{1}{d}$, then the process of constructing standard extension of $\sigma_d|_A$ succeeds because in this process $d-1$ "full wraps" of the circle will be taken out leaving the length of the image of $\mathbb{T}$ under F equal to 1 and thus implying that F is not just locally monotone (which it always is), but also globally monotone (as a map of $\mathbb{T}$ onto itself) and so $\sigma_d|_A$ is order preserving. (See Figure 4.) $\square$

3.2. **Proof of Proposition 2.2.** We now show that any rotational Cantor set must have a *taut* root gap.

Proof. Given a rotational Cantor set C, we know from Theorem 1.11 that there exist two points $x, y \in C$ such that $\sigma_d(x) = \sigma_d(y)$. Since C is rotational, it satisfies the Hole Theorem 2.1. From the proof of that theorem, we know that there is a monotone degree 1 map $m : \mathbb{T} \to \mathbb{T}$ that agrees with σ_d on C and is flat on the $d-1$ holes in $\mathbb{T} \setminus C$. Let I and J denote the complementary arcs to $\{x, y\}$ in $\mathbb{T}$. If there are points of C in both I and J, then m must be increasing on some points of both I and J. Hence, $\sigma_d(x) \neq \sigma_d(y)$, a contradiction. We conclude that either I or J (not both) is a taut root gap of C. $\square$

3.3. **Proof of Reduction Theorem 2.3.** We now prove that σ_d can be *reduced* to σ_{d-1} by removing a *wrap*, J, and all its pre-images.

Proof. To show that X must contain a Cantor set we show that there must exist a binary tree (in the sense of subset containment) of intervals in the construction of X. Since X must be totally disconnected because every interval is eventually onto, it then follows that this binary tree of intervals limits to a Cantor set contained in X. In the case $d \geq 4$ it is easy to show that this binary tree exists.

For all $d > 2$, $\sigma_d^{-1}(J)$ contains d evenly spaced pre-image intervals, at least $d-2$ of which must be completely contained in $\mathbb{T} \setminus J$ and therefore "break" $\mathbb{T} \setminus J$ into at least

$(d-2)+1 \geq 2$ non-degenerate components. When $d > 3$, at least two of these are *nice* components, meaning that they map homeomorphically onto $X \setminus J$. (The only components which may not are those adjacent to J.) In turn, each of theses nice pieces contains at least two more nice subpieces, components of $\mathbb{T} \setminus \sigma_d^{-2}(J)$ and $\mathbb{T} \setminus \bigcup_{i=0}^{2} \sigma_d^{-i}(J)$, that map homeomorphically onto one of the nice pieces at the previous stage. This process repeats, always with two nice subpieces in each nice piece of the previous stage, as we consider higher order pre-images of J. Therefore, we know that X must contain a Cantor set C_0.

For the case $d = 3$ the above proof holds for certain arrangements of the first pre-images of J. However, if the components of the first pre-image of J are situated on the circle such that one component is properly contained in J, and does not share an endpoint with J, the above proof does not follow through directly. For this arrangement the idea of the proof is the same. One component of $\mathbb{T} \setminus \bigcup_{i=0}^{2} \sigma_d^{-i}(J)$ is nice, but the two components of $\mathbb{T} \setminus \sigma_d^{-1}(J)$ which are adjacent to J that do not map homeomorphically onto $\mathbb{T} \setminus J$ must combine to substitute for the lacking nice piece. Our intersection in this case is not a binary tree of intervals, but a binary tree of sets where at each branch point the set at one sub-branch contains at least one interval, and the set along the other sub-branch contains at least one interval. This binary tree of sets must therefore contain a binary tree of intervals, and consequently a Cantor set.

It now follows by a standard maximum principle argument that X, which is totally disconnected and contains a Cantor set, contains a unique maximal Cantor set C.

The unique maximal Cantor set C found above is also invariant, as we now show using the fact that it is maximal. Indeed, $\sigma_d(C) \subset C$ because otherwise $C \cup \sigma_d(C)$ is a Cantor set properly containing C while, since X is subinvariant, we have that $C \cup \sigma_d(C)$ is contained in X, a contradiction of the maximality of C.

On the other hand, suppose that $\sigma_d(C)$ is a proper subset of C. Choose a small compact neighborhood V in $C \setminus (\sigma_d(C) \cup \{\sigma_d(a)\}$ (recall that $J = (a, a + \frac{1}{d})$). Then V is a Cantor set itself, and since it is disjoint from $\sigma_d(a)$ it has Cantor set pre-images disjoint from J. Each such pre-image is contained in X (indeed, by the construction they are disjoint from J and their images are contained in X). Hence, they must be contained in the unique maximal Cantor subset of X, i.e. in C, a contradiction. So, we have shown that C is invariant.

Definition 3.2 ($\overline{\sigma}_d$)**.** Given $J = (a, a + \frac{1}{d} \pmod 1)$, we define the map $\overline{\sigma}_d : \mathbb{T} \to \mathbb{T}$ by

$$\overline{\sigma}_d(x) = \begin{cases} \sigma_d(a) = \sigma_d(a + \frac{1}{d}) = b, & \text{if } x \in J; \\ \sigma_d(x), & \text{if } x \notin J. \end{cases}$$

The map $\overline{\sigma}_d$ basically collapses an interval, which would have mapped onto the circle, to a single point while allowing σ_d to act normally on the rest of the circle. Although we do not have a special notation for the map defined by any particular interval (like J in this case), it will be clear in context which interval we are referring to. Note that $J \subset \mathbb{T} \setminus C$.

Lemma 3.3. *$\overline{\sigma}_d$ maps complementary gaps of C to complementary gaps of C unless the gap is equal to J, in which case $\overline{\sigma}_d$ carries the gap to a single point in C.*

Proof. It is clear by definition of $\overline{\sigma}_d$ that if the gap is equal to J, then its image is a point. Also, since J is the gap, then its endpoints are in C. Thus, by invariance of C under $\overline{\sigma}_d$, the image of these endpoints, and thus the image of J, is a point in C.

So now let our gap G not be equal to J. By way of contradiction, suppose that $\overline{\sigma}_d(G)$ is not a gap of C. Then there exists a point of C in $\overline{\sigma}_d(G)$. In fact, there also exists a small clopen neighborhood $U \subset C$ in $\overline{\sigma}_d(G)$, and we may assume that U does not contain $\overline{\sigma}_d(a)$. Then we can take the $\overline{\sigma}_d$-pre-image V of U in G (observe that since U does not contain a then σ_d and $\overline{\sigma}_d$ coincide on the entire pre-image of U, in particular on V). Then V has to be a part of C: a) it is disjoint from J (because $a \notin U$) and all its images are disjoint from J (because $U \subset C$). However this contradicts the assumption that $V \subset G$ because G is a gap in C. □

Definition 3.4. We define $m' : \mathbb{T} \to \mathbb{T}'$ as the monotone map which shrinks endpoints of the complementary gaps of the Cantor set C found above to points.

The quotient space $\mathbb{T}'$ that remains, once the complementary gaps of C are shrunk to points, is a circle with a natural local order inherited from $\mathbb{T}$.

Definition 3.5. Define the map $\sigma'_{d-1} : \mathbb{T}' \to \mathbb{T}'$ by $\sigma'_{d-1} = m' \circ \overline{\sigma}_d \circ (m')^{-1}$

Note that σ'_{d-1} is well-defined by Lemma 3.3.

Lemma 3.6. *σ'_{d-1} is conjugate to σ_{d-1}*

Proof. By Theorem 1.2, we need to show is that σ'_{d-1} is a $(d-1)$-to-one covering map of the circle that is topologically exact. First we will show that σ'_{d-1} is an $(d-1)$-to-one covering map of the circle. Observe that if $x \in \mathbb{T}'$ then its m'-pre-image is either a point of C or the closure of a gap in C. Then by Lemma 3.3 there are $d-1$ components of $\overline{\sigma}_d^{-1}((m')^{-1}(x))$ each of which is the closure of a gap in C or a point of C. Finally, all these components project by m' onto $d-1$ points which are σ'_{d-1}-pre-images of x. On the other hand, by the construction it follows that σ'_{d-1} is a local homeomorphism. Therefore σ'_{d-1} is a $(d-1)$-to-one covering map.

To prove that σ'_{d-1} is topologically exact we first note the fact that $(\sigma'_{d-1})^k = m' \circ \overline{\sigma}_n^k \circ (m')^{-1}$ by the semi-conjugacy. Note also that $\overline{\sigma}_d|_C = \sigma_d|_C$ since C meets J at most at the endpoints of J. Let U be any open interval of $\mathbb{T}'$ intersected with C. Then $(m')^{-1}(U)$ must contain an open subset, V, of the Cantor set C. Without loss of generality, we may assume U does not meet J. Consider the "sisters" of J, that is, the intervals of the form $\overline{J} + \frac{k}{d}$, each of which maps onto $\mathbb{T}$. Evidently, pre-images of C under successive iterates of σ_d^{-1} become as small as we want them and are evenly spaced on the circle. Hence, we can see that there must exist some k such that $\sigma_d^k(V) = C$. We also have that $\overline{\sigma}_d^k((m')^{-1}(U)) \supset \overline{\sigma}_d^k(V) = C$, and therefore that $m' \circ \overline{\sigma}_d^k \circ (m')^{-1}(U) = \mathbb{T}'$. Thus we have found k such that $(\sigma'_{d-1})^k(U) = \mathbb{T}'$. □

Proposition 3.7. *There are exactly $d-1$ different homeomorphisms that conjugate σ'_{d-1} on $\mathbb{T}'$ to σ_{d-1} on $\mathbb{T}$.*

Proof. Since σ'_d is a topologically exact $(d-1)$-to-one covering map, σ'_d is conjugate to σ_d. If h and h' are two homeomorphisms which conjugate σ'_{d-1} on $\mathbb{T}'$ to σ_{d-1} on $\mathbb{T}$, then they can differ only by multiplication by a $d-1$ root of unity. In other words, they can differ only in the way in which they label the $d-1$ fixed points. Therefore, since these fixed points have a circular order to them, there are $d-1$ different homeomorphisms which conjugate σ'_{d-1} on $\mathbb{T}'$ to σ_{d-1} on $\mathbb{T}$. □

Let $h : \mathbb{T}' \to \mathbb{T}$ be a homeomorphism which conjugates σ'_{d-1} on $\mathbb{T}'$ to σ_{d-1} on $\mathbb{T}$.

Definition 3.8. Define the map $m : \mathbb{T} \to \mathbb{T}$ by $m = h \circ m'$.

Proposition 3.9. *The map m semi-conjugates $\overline{\sigma}_d$ and σ_{d-1}. Moreover, m at most two-to-one semi-conjugates $\sigma_d|_C$ and σ_{d-1}.*

This concludes the proof of Reduction Theorem 2.3. □

We will need the second of the following two lemmas in order to prove Corollary 2.4. The first lemma will be used in the proof of the second.

Lemma 3.10. *If $a, b : \mathbb{T} \to \mathbb{T}$ are degree 1 circle maps then for any lift $\widehat{a \circ b}$ of the map $a \circ b : \mathbb{T} \to \mathbb{T}$ there exists lifts $\widehat{a}$ and $\widehat{b}$ of a and b such that $\widehat{a \circ b} = \widehat{a} \circ \widehat{b}$.*

Proof. We take the rather odd approach, on the face of it, of first showing that for any lifts $\widehat{a}, \widehat{b}$ of a, and b, respectively, there exists a lift $\widehat{a \circ b}$ of $a \circ b$ such that $\widehat{a} \circ \widehat{b} = \widehat{a \circ b}$. To this end note that it is easy to prove that for any degree 1 lifts $\widehat{a}$ and $\widehat{b}$ that $\widehat{a} \circ \widehat{b}$ is a degree 1 map and that $\widehat{a} \circ \widehat{b} = a \circ b \pmod 1$ for all $x \in [0, 1)$. These two facts show that $\widehat{a} \circ \widehat{b}$ is some lift of $a \circ b$; i.e., there exists a lift $\widehat{a \circ b}$ such that $\widehat{a \circ b} = \widehat{a} \circ \widehat{b}$.

To prove the lemma let $\widehat{a \circ b}$ be any lift of $a \circ b$. Pick any lifts $\widehat{a}'$ and $\widehat{b}'$ of a and b. The previous claim then shows that there exists some lift of $a \circ b$, call it $\widehat{a \circ b}'$, such that $\widehat{a \circ b}' = \widehat{a}' \circ \widehat{b}'$. There then exists an integer m such that $\widehat{a}' \circ \widehat{b}' = \widehat{a \circ b}' = \widehat{a \circ b} + m$. Therefore, among many options, we can define lifts $\widehat{a}$ and $\widehat{b}$ by $\widehat{a} = \widehat{a}'$ and $\widehat{b} = \widehat{b}' - m$. Then we have that

$$\widehat{a} \circ \widehat{b} = \widehat{a}' \circ (\widehat{b}' - m) = \widehat{a}' \circ \widehat{b}' - m = \widehat{a \circ b},$$

and our lemma is proven. □

Lemma 3.11. *Let f and g be monotone degree 1 circle maps. Suppose there exists a degree 1 circle map ϕ and a set $A \subset \mathbb{T}$ such that $\phi \circ f|_A = g \circ \phi|_A$ and $f(A) \subset A$. Then $\rho(g) = \rho(f)$.*

Proof. First, we note that $\rho(f)$ and $\rho(g)$ exist by Proposition 1.8. Since $\phi \circ f|_A = g \circ \phi|_A$ we can fix lifts such that $\widehat{\phi \circ f}|_{\widehat{A}} = \widehat{g \circ \phi}|_{\widehat{A}}$. Applying Lemma 3.10 to these lifts, we then know there exists lifts $\widehat{f}, \widehat{g}, \widehat{\phi}_1$, and $\widehat{\phi}_2$ such that $\widehat{\phi}_1 \circ \widehat{f}|_{\widehat{A}} = \widehat{g} \circ \widehat{\phi}_2|_{\widehat{A}}$. Since $\widehat{\phi}_2 = \widehat{\phi}_1 + m$, for some integer m, we have that $\widehat{\phi}_1 \circ \widehat{f}|_{\widehat{A}} - m = \widehat{g} \circ \widehat{\phi}_1|_{\widehat{A}}$. Now consider $\widehat{g}^n \circ \widehat{\phi}_1$ for arbitrary n. We have that

$$\begin{aligned}
\widehat{g}^n \circ \widehat{\phi}_1|_{\widehat{A}} &= \widehat{g}^{n-1} \circ \widehat{g} \circ \widehat{\phi}_1|_{\widehat{A}} \\
&= \widehat{g}^{n-1}(\widehat{\phi}_1 \circ \widehat{f}|_{\widehat{A}} - m) \\
&= \widehat{g}^{n-1} \circ \widehat{\phi}_1 \circ \widehat{f}|_{\widehat{A}} - m \\
&= \widehat{g}^{n-2} \circ \widehat{g} \circ \widehat{\phi}_1 \circ \widehat{f}|_{\widehat{A}} - m \\
&= \widehat{g}^{n-2} \circ (\widehat{\phi}_1 \circ \widehat{f} - m) \circ \widehat{f}|_{\widehat{A}} - m \\
&= \widehat{g}^{n-2} \circ \widehat{\phi}_1 \circ \widehat{f} \circ \widehat{f}|_{\widehat{A}} - 2m \\
&\vdots \\
&= \widehat{\phi}_1 \circ \widehat{f}^n|_{\widehat{A}} - nm.
\end{aligned}$$

Note that we required the subinvariance of A under f, and thus of $\widehat{A}$ under $\widehat{f}$, in the above step from the substitution in the fifth to sixth equality and in all such subsequent substitutions.

Before we get to the actual computing of the rotation number of g we state one more important fact that we will use in the proof. Since $\widehat{\phi}_1$ is a degree 1 map of the reals, we have that there exists an integer M such that $|\widehat{\phi}_1(x) - x| \leq M$ for all $x \in \mathbb{R}$. Moreover, since $\widehat{f}^n(x)$ is a real number for all n and all $x \in \mathbb{R}$ we can further say that $|\widehat{\phi}_1 \circ \widehat{f}^n(x) - \widehat{f}^n(x)| \leq M$ for all $x \in \mathbb{R}$ and all positive integers n.

We now would like to calculate $\rho(g)$. Let $\widehat{a} \in \widehat{A}$ and consider $\widehat{\phi}(\widehat{a})$. We have already shown that

$$\widehat{g}^n \circ \widehat{\phi}_1(\widehat{a}) = \widehat{\phi}_1 \circ \widehat{f}^n(\widehat{a}) - nm.$$

This is equivalent to

$$\widehat{g}^n \circ \widehat{\phi}_1(\widehat{a}) = \widehat{\phi}_1 \circ \widehat{f}^n(\widehat{a}) - \widehat{f}^n(\widehat{a}) + \widehat{f}^n(\widehat{a}) - nm.$$

Therefore

$$\frac{\widehat{g}^n \circ \widehat{\phi}_1(\widehat{a})}{n} = \frac{\widehat{\phi}_1 \circ \widehat{f}^n(\widehat{a}) - \widehat{f}^n(\widehat{a})}{n} + \frac{\widehat{f}^n(\widehat{a})}{n} - \frac{nm}{n}.$$

Recalling that $|\widehat{\phi}_1 \circ \widehat{f}^n(x) - \widehat{f}^n(x)| \leq M$ for all $x \in \mathbb{R}$ and all integers n, and taking the limit as $n \to \infty$, we get that

$$\lim_{n\to\infty} \frac{\widehat{g}^n \circ \widehat{\phi}_1(\widehat{a})}{n} = \lim_{n\to\infty} \frac{\widehat{f}^n(\widehat{a})}{n} - m.$$

Therefore, since we can pick any point to compute the rotation numbers of f and g,

$$\rho(g) = \lim_{n\to\infty} \frac{\widehat{g}^n \circ \widehat{\phi}_1(\widehat{a})}{n} = \lim_{n\to\infty} \frac{\widehat{f}^n(\widehat{a})}{n} = \rho(f) \pmod 1.$$

□

3.4. **Proof of Corollary 2.4.** We now show that the rotation number is preserved when reducing A to $m(A)$.

Proof. Suppose we have a σ_d-rotational set A. Let J be any hole of A and C the maximal invariant Cantor set in X as shown above. We will now show that $m(A)$ is a σ_{d-1}-rotational set and that $\rho(m(A)) = \rho(A)$. Clearly, $m(A)$ is invariant and closed.

To show that $m(A)$ is rotational, we pick a particular extension f of $\sigma_d|_A = \overline{\sigma}_d|_A$. Let $\{H_1, H_2, ..., H_{d-1}\}$ be any $d-1$ disjoint holes of A. We want to change the holes so that they project, under m, to $d-2$ holes of $m(A)$. If only one of the holes, H_i, picked above intersects J then we replace that hole with J itself and we still have $d-1$ holes of A. If two of the above holes intersect J then we wish to adjust one of them so that it abuts the other one in J, which is certainly possible if it was a hole in the first place. In the above two cases keep in mind that $A \cap J = \emptyset$. The only other case is when none of the holes intersect J, and this can only happen when A is a fixed point together with any subset of its first pre-images. Here we can easily replace one of these holes with J. Now the reader may check that under m the projection of these modified $d-1$ holes of A under σ_d is $d-2$ holes of $m(A)$ under σ_{d-1}. (The only non-trivial case is when two holes abut in J.) Thus, from Hole Theorem 2.1, σ_{d-1} is order-preserving on $m(A)$. Since we earlier showed that $m(A)$ is closed and invariant, it now satisfies all the conditions of Definition 1.6 and is rotational.

Now let f and g be any monotone extensions of $\sigma_d|_A$ and $\sigma_{d-1}|_{m(A)}$, respectively. Since A and $m(A)$ are rotational we have that $\rho(f) = \rho(A)$ and $\rho(g) = \rho(m(A))$. Recall that $m \circ \overline{\sigma}_d = \sigma_{d-1} \circ m$. Since $f|_A = \overline{\sigma}_d|_A = \sigma_d|_A$ (not needed but true) and $g|_{m(A)} = \sigma_{d-1}|_{m(A)}$, we get that $m \circ f|_A = g \circ m|_A$. We can apply Lemma 3.11 to get that $\rho(m(A)) = \rho(A)$. □

3.5. **Proof of Construction Theorem 2.6.** We show that σ_{d+1} can be constructed from σ_d by inserting a *wrap* and its appropriate pre-images.

Proof. We label our initial point as x_1^0 for reasons we will see shortly. Let $\text{orb}^-(x_1^0) = \{x \in \mathbb{T} \mid \sigma_d^k(x) = x_1^0 \text{ for some } k \geq 0\}$. Now the points in $\text{orb}^-(x_1^0)$ are labelled as follows: x_j^{-i} where i is the minimal power such that $\sigma_d^i(x_j^{-i}) = x_1^0$ and $1 \leq j \leq d^i$. For fixed i there may be exactly d^i pre-images x_j^{-i}, or fewer, depending on the initial x_1^0, and they are labelled with subscripts from 1 up to d^i in no particular order.

We now want to construct a new space from our circle $\mathbb{T}$ using the set $\text{orb}^-(x_1^0)$. Insert an interval $I_j^{-i} = [a_j^{-i}, b_j^{-i}]$ into the circle $\mathbb{T}$ at all $x_j^{-k} \in \text{orb}^-(x_1^0)$. This insertion of intervals gives rise to another circle $\mathbb{T}^*$ and a natural monotone map $M^* : \mathbb{T}^* \to \mathbb{T}$, which collapses our inserted intervals back to the points they came from. Note that $\mathbb{T}^*$ has a natural local order inherited from $\mathbb{T}$.

Now we define a metric d^* on $\mathbb{T}^*$. The reason we define a metric is to force the map we define below to be topologically exact. The metric d^* induces the order topology on $\mathbb{T}^*$. In this new metric, we consider each interval I_j^{-i} to be an isometric copy of the interval $[0, \frac{1}{(2d)^{i+1}}]$. Hence, $d^*(a_j^{-i}, b_j^{-i})$ is the *length* of the interval I_j^{-i} in $\mathbb{T}^*$, denoted $l(I_j^{-i}) = \frac{1}{(2d)^{i+1}}$. Let $l'(B)$ denote the length of an interval $B \subset \mathbb{T}$. Let $A = [a, b] \subset \mathbb{T}^*$, where $a < b$. We now define $d^*(a, b)$. We start with $a, b \in I_j^{-i}$ for the same i and j. We simply define $d^*(a, b)$ to be the inherited isometric distance between the two.

For all other cases we define the values $J(a)$ and $J(b)$. If $a \in [a_j^{-i_1}, b_j^{-i_1}] = I_j^{-i_1}$ for some i_1, then $J(a) = d^*(a, b_j^{-i})$, else $J(a) = 0$. If $b \in [a_j^{-i_2}, b_j^{-i_2}] = I_j^{-i_2}$ for some $i_2 \neq i_1$, then $J(b) = d^*(b, a_j^{-i})$, else $J(b) = 0$. Let

$$l(A) = l'(M^*(A)) + J(a) + J(b) + \sum_{y \in \mathrm{orb}^-(x_1^0) \cap M^*(A)} l((M^*)^{-1}(y)).$$

Let A_1 and A_2 be the two intervals in $\mathbb{T}^*$ with endpoints a and b and define $d^*(a, b) = \min\{l(A_1), l(A_2)\}$. Note that the summation of length in the above definition is finite. Even if we inserted the maximum number of intervals, d^i, for each i, their total length would be

$$\sum_{i=0}^{\infty} d^i \frac{1}{(2d)^{i+1}} < 1.$$

Now define a new map $\sigma_{d+1}^* : \mathbb{T}^* \to \mathbb{T}^*$ as follows: Let $x \in \mathbb{T}^*$. If $x \notin I_j^{-i}$ for any $i, j \geq 0$ then $\sigma_{d+1}^*(x) = (M^*)^{-1}(\sigma_d(M^*(x)))$.

If $x \in I_j^{-i}$ for some $i, j \geq 0$ then we define $\sigma_{d+1}^*(x)$ more carefully. If $x \in I_1^0 = (M^*)^{-1}(x_1^0)$ and $\sigma_d(x_1^0) \notin \mathrm{orb}^-(x_1^0)$, then the value $(M^*)^{-1}(\sigma_d(x_1^0))$ is well-defined, but we do not define $\sigma_{d+1}^*(x)$ to be this value because then the whole interval $I_1^0 = [a_1^0, b_1^0]$ would have this value under σ_{d+1}^* and we do not want this. We define

$$\sigma_{d+1}^*(a_1^0) = \sigma_{d+1}^*(b_1^0) = (M^*)^{-1}(\sigma_d(x_1^0))$$

and let σ_{d+1}^* map $\mathrm{Int}(I_1^0)$ linearly around the circle $\mathbb{T}^*$, in order, with respect to the metric d^* (i.e., σ_{d+1}^* is really mapping the interval I_1^0 around the circle exactly once starting and stopping at $(M^*)^{-1}(\sigma_d(x_1^0))$).

If $x \in I_1^0$ and $\sigma_d(x_1^0) = x_j^{-i} \in \mathrm{orb}^-(x_1^0)$ for some $i, j \geq 0$ then

$$(M^*)^{-1}(\sigma_d(x_1^0)) = (M^*)^{-1}(x_j^{-i}) = [a_j^{-i}, b_j^{-i}]$$

for some i and $1 \leq j \leq d^i$. In this case we define

$$\sigma_{d+1}^*(a_1^0) = a_j^{-i}, \sigma_{d+1}^*(b_1^0) = b_j^{-i}$$

and let σ_{d+1}^* map $\mathrm{Int}(I_1^0)$ linearly around the circle $\mathbb{T}^*$ in order, with respect to the metric d^*. That is, σ_{d+1}^* maps the interval I_1^0 around the circle a little more than once, starting at a_j^{-i}, passing b_j^{-i} and a_j^{-i} again, and finally stopping at b_j^{-i}. We say I_1^0 *overlaps* the interval I_j^{-i}.

In the final case that $x \in I_j^{-i}$ for $i > 0$, we have that $\sigma_d(M^*(x)) = x_{j'}^{-i+1} \in \mathrm{orb}^-(x_1^0)$ for some j'. Here we define σ_{d+1}^* to send the interval $[a_j^{-i}, b_j^{-i}]$ linearly onto $[a_{j'}^{-i+1}, b_{j'}^{-i+1}]$ in order.

Note that $M^* \circ \sigma^*_{d+1} = \sigma_d \circ M^*$. Also note that we have shown that σ^*_{d+1} takes endpoints of inserted intervals to endpoints of inserted intervals, and in fact takes all inserted intervals to inserted intervals except I^0_1, which is mapped completely around the circle while overlapping an inserted interval as defined above.

Lemma 3.12. *σ^*_{d+1} is conjugate to σ_{d+1}*

Proof. By Theorem 1.2, it suffices to show that σ^*_{d+1} is a $(d+1)$-to-one covering map of the circle that is topologically exact. The proof that σ^*_{d+1} is a covering map (locally one-to-one, locally order-preserving, and locally onto) is straightforward and left to the reader.

To show σ^*_{d+1} is a $(d+1)$-to-one map it suffices to find one point in the range with exactly $d+1$ pre-images. Pick $x \notin \mathrm{orb}^-(x^0_1) \cup \{\sigma_d(x^0_1)\}$. Then $(M^*)^{-1}(x)$ is one point, call it y, and $y \notin I^{-i}_j$ for any i and j. We will show that y has $d+1$ pre-images under σ^*_{d+1}. Let $\{x_1, \dots, x_d\}$ be the d distinct pre-images of x under σ_d. Since $x \notin \mathrm{orb}^-(x^0_1)$ then $x_1, \dots, x_d \notin \mathrm{orb}^-(x^0_1)$ as well. Thus $\{(M^*)^{-1}(x_1), \dots, (M^*)^{-1}(x_d)\}$ are d distinct points in $\mathbb{T}^*$ and, by definition, $\sigma^*_{d+1}((M^*)^{-1}(x_i)) = y$ for $i = 1, \dots, d$. Where could any other pre-images of y under σ^*_{d+1} come from? All possible candidates from outside an I^{-i}_j have been found above because by the definition of σ^*_{d+1} on $\mathbb{T}^* \setminus \bigcup_{i,j} I^{-i}_j$, the image of such candidates under M^* must be in the set $\{x_1, \dots, x_d\}$. Now consider the inserted intervals, I^{-i}_j's. We know that for $I^{-i}_j \neq I^0_1$, σ^*_{d+1} carries it to another $I^{-i'}_{j'}$. Since $y \notin I^{-i}_j$ for any i, j, that leaves I^0_1 as the only remaining source of possible pre-images of y. I^0_1 contains at least one pre-image of y because it always wraps. It contains no more than one pre-image of y due to the fact that if it does double cover anything at all it is one of the inserted intervals that y is not in. Thus we have found a point with exactly $d+1$ pre-images under σ^*_{d+1}.

Finally we show that σ^*_{d+1} is topologically exact. With our metric defined on $\mathbb{T}^*$ we have that $l(\sigma^*_{d+1}(A)) \geq d \cdot l(A)$ for any $A \subset \mathbb{T}^*$.

This can be seen by thinking of A as being a union of intervals that are either part of an I^{-i}_j or not. By the definition of σ^*_{d+1}, it multiplies the length of intervals outside any I^{-i}_j by a factor of d, intervals a part of any $I^{-i}_j \neq I^0_1$ by a factor of $\frac{1}{(2d)^{i+1}} / \frac{1}{(2d)^{i+2}} = 2d$, and intervals a part of I^0_1 by a factor of at least $2d$ (more so if it overlaps). Hence, it is topologically exact.

We have shown σ^*_{d+1} is conjugate to σ_{d+1}. □

We proceed to define the σ_{d+1}-rotational set $\widetilde{A}$ *constructed* from the σ_d-rotational set A. Let $h : \mathbb{T} \to \mathbb{T}^*$ be the homeomorphism conjugating σ^*_{d+1} and σ_{d+1}. Define a monotone map $M : \mathbb{T} \to \mathbb{T}$ by $M = M^* \circ h$. We now want to consider the set $\widetilde{A} = M^{-1}(A) \setminus \mathrm{Int}(M^{-1}(A))$ where A is our rotational set under σ_d.

First we claim that $\widetilde{A}$ is a rotational set under σ_{d+1}. We need to show first that it is invariant. It follows from the semi-conjugacy, and the fact that A is invariant under σ_d, that if $x \in \widetilde{A}$ then $\sigma_{d+1}(x) \in M^{-1}(A)$. We need to show that $\sigma_{d+1}(x) \notin \mathrm{Int}(M^{-1}(A))$. To do this we note that $M^{-1}(A) \setminus \mathrm{Int}(M^{-1}(A))$ is homeomorphic to $(M^*)^{-1}(A) \setminus \mathrm{Int}((M^*)^{-1}(A))$ and consider how σ^*_{d+1} acts on $h(x)$. As we have already shown in the proof of Theorem 2.6, σ^*_{d+1} takes endpoints of inserted intervals to endpoints of inserted intervals. Based on this, and the fact that σ^*_{d+1} maps only inserted intervals into inserted intervals, we can see that σ^*_{d+1} cannot take points of $(M^*)^{-1}(A) \setminus \mathrm{Int}((M^*)^{-1}(A))$ into $\mathrm{Int}((M^*)^{-1}(A))$. Thus, $\sigma_{d+1}(x) = h^{-1} \circ M^* \circ h(x)$ cannot be in $\mathrm{Int}(M)^{-1}(A)$, and we have shown $\sigma_{d+1}(\widetilde{A}) \subset \widetilde{A}$.

Now we need to show that $\widetilde{A} \subset \sigma_{d+1}(\widetilde{A})$. Let $x \in \widetilde{A}$. If $h(x) \notin I^{-i}_j$ for any pair i, j then it will follow from the diagram, the invariance of A, and the one-to-one properties of the maps on such x, that $x \in \sigma_{d+1}(\widetilde{A})$. On the other hand, for such x that $h(x) \in I^{-i}_j$, consider how

$h(x)$ must be an endpoint of an I_j^{-i} because $x \in \widetilde{A}$. Then recall that σ_{d+1}^* carries endpoints of all I_j^{-i+1}'s for $i \geq 0$ onto the endpoints of all I_j^{-i}'s. Therefore, since $h(x)$ is an endpoint of an I_j^{-i}, we know there exists some endpoint, call it y, of an I_j^{-i} which maps to it under σ_{d+1}^*. It then follows from the invariance of A, and the commutativity of the diagram, that in fact $y \in h(\widetilde{A})$ and, again by the diagram, that $x \in \sigma_{d+1}(\widetilde{A})$.

Now that we have shown that $\widetilde{A}$ is invariant, we need to show that it is closed and that $\sigma_{d+1}|_{\widetilde{A}}$ is order-preserving for it to be rotational. It is closed because A is closed and M is continuous. If a is a limit point of $\widetilde{A}$ one can show that $M(a) \in A$ by considering the image, under M, of a sequence of points in $\widetilde{A}$ that limits to a and by using the fact that A is closed. With a little additional argument concerning the removal of the interior of inserted intervals, one can show that a must have been in $\widetilde{A}$, showing that $\widetilde{A}$ is closed. To show that $\sigma_{d+1}|_{\widetilde{A}}$ is order-preserving we use the Hole Theorem 2.1 in conjunction with the easily proven facts that any $d-1$ holes of A lift to $d-1$ holes of $\widetilde{A}$, and that another hole is contained in the interior of I_1^0. Hence $\widetilde{A}$ is rotational.

To finish the proof of Construction Theorem 2.6 it is enough to note that if $x_1^0 \in A$ then M^* maps

$$(M^*)^{-1}(\mathrm{orb}^-(x_1^0) \cap A) \setminus \mathrm{Int}((M^*)^{-1}(\mathrm{orb}^-(x_1^0) \cap A))$$

two-to-one onto $\mathrm{orb}^-(x_1^0) \cap A$. Points in $A \cap \mathrm{orb}^-(x_1^0)$ will be mapped onto in a two-to-one manner while points of A not in $\mathrm{orb}^-(x_1^0)$ will be mapped onto in a one-to-one fashion. M behaves the same way on $\widetilde{A}$. Hence, $\sigma_{d+1}|_{\widetilde{A}}$ is at most two-to-one semi-conjugate to $\sigma_d|_A$. Now, in the other case, if $x_1^0 \notin A$, then no insertions are done at any points of A and everything is kept one-to-one, hence $\sigma_{d+1}|_{\widetilde{A}}$ is conjugate to $\sigma_d|_A$. □

We now prove Corollary 2.7.

Proof. Let f and g be any monotone extensions of $\sigma_{d+1}|_{\widetilde{A}}$ and $\sigma_d|_A$, respectively. Since $\widetilde{A}$ and A are rotational we have that $\rho(f) = \rho(\widetilde{A})$ and $\rho(g) = \rho(A)$. Recall that $M \circ \sigma_{d+1} = \sigma_d \circ M$. Since $f|_{\widetilde{A}} = \sigma_{d+1}|_{\widetilde{A}}$ and $g|_A = \sigma_d|_A$ we get that $M \circ f|_{\widetilde{A}} = g \circ M|_{\widetilde{A}}$. We can apply Lemma 3.11 to get that $\rho(\widetilde{A}) = \rho(A)$. □

4. Rotational Periodic Orbits for σ_d

How many periodic orbits are there under σ_d with rotation number $\frac{p}{q}$ (in lowest terms)? We know from Bullett and Sentenac [2] that under σ_2 there is only one periodic orbit with any given rational rotation number. Upon brief inspection of σ_3, however, one finds that there are two fixed points and three periodic orbits with rotation number $\frac{1}{2}$. After looking at a few other rotation numbers one may conjecture that, under σ_3, any rational rotation number $\frac{p}{q}$ has $q+1$ periodic orbits with that rotation number. When $d \geq 4$, however, such inspection and conjecturing is not so easy, as no pattern presents itself so quickly when looking at the number of periodic orbits with any given rational rotation number under σ_d.

The main results of this section have been obtained previously by Goldberg [4]. We provide proofs for three reasons: completeness, our approach is distinctly different, and our proofs suggest generalizations to non-rotational orbits.

4.1. **Distinguishing Rotational Periodic Orbits by the Placement of Pre-images of 0.** Before we get to the proof of Theorem 2.8 we will discuss some preliminaries that will make the proof more efficient. First, given any rational rotation number $\frac{p}{q}$ in lowest terms we know that any rotational periodic orbit of rotation number $\frac{p}{q}$ must have q points, and

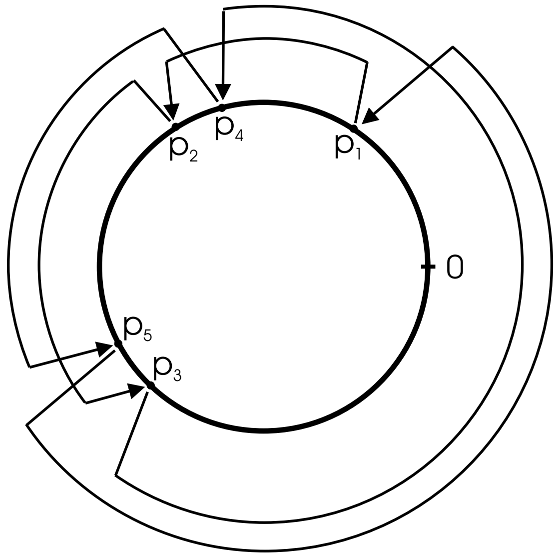

FIGURE 5. Illustration of a periodic orbit with rotation number $\frac{2}{5}$. The point p_1 maps to p_2 skipping p_4. In turn, p_2 maps to p_3 skipping p_5, and so on.

that each time the map is applied to any point in the orbit, it skips over $p-1$ points of the orbit spatially on the circle (see Figure 5).

We will also use what we know about the d-ary expansion of a periodic orbit. We can "read-off" the d-ary expansion of a periodic orbit by its placement with relation to the $d-1$ pre-images of 0. The circle is split into d sectors by 0 and its pre-images: the 0th, 1st, ..., (d-1)-th sectors. Label the trajectory of a periodic orbit as $p_1 \to p_2 \to \ldots \to p_q$. Note that these are not arranged consecutively counterclockwise on the circle unless $p = 1$. The expansion of p_1 is the repeating string of q digits where the i-th digit is whatever sector p_i is in. The expansion of p_2 will contain the same repeating pattern just shifted one place to the left, i.e., starting with the second digit of p_1's expansion. In fact, the expansion of any point in a periodic orbit is the same repeating pattern starting in different places. We therefore refer to the d-ary expansion of a periodic orbit, or the periodic expansion, as the class of expansions of p_1 through p_q (i.e., when we say two expansions are different, we mean that the repeating pattern that defines those expansions is different and one is not just the shift of the other, which would be an equivalent expansion). We also note here that the only points which can have ambiguous expansions are pre-images of 0. We do not worry about this, however, because we will only be referring to the expansions of periodic orbits, and these do not contain any pre-images of 0.

Easily established facts about the expansions of periodic orbits are that different periodic orbits have different expansions and two different periodic expansions must arise from different periodic orbits.

4.2. **Proof of Counting Theorem 2.8 for rotational periodic orbits.** The idea of our proof is that for any given rational rotation number $\frac{p}{q}$ we can place the trajectory of the periodic orbit on the circle spatially in only one way (in other words, for any two such orbits an orientation preserving homeomorphism of the circle conjugates the maps on these orbits). We then place 0 arbitrarily on the circle. It remains to determine how many different ways there are to place the pre-images of 0 on the circle with one necessary yet simple restriction that will be discussed in the proof. Each different way of placing these pre-images gives rise to a different periodic expansion, and thus a different periodic orbit with rotation number $\frac{p}{q}$.

Proof. Fix d. Let $\frac{p}{q}$ be in reduced form. Place q points on a circle to denote a periodic orbit with q points. Place 0 somewhere on the circle between two points of the periodic orbit. Label the first point counterclockwise from 0 as p_1. The rest of the points are labelled temporally, or as they are arrived at in the trajectory of p_1; i.e., $p_2 = \sigma_d(p_1), p_3 = \sigma_d(p_2)$, etc... Since we have already placed p_1 , there is only one way to do this temporal labelling. Let the first point clockwise from 0 be p_k so that 0 is in the gap between the points p_k and p_1 of the periodic orbit. We will be placing the pre-images of 0 between points of the periodic orbit, between 0 and p_1 and between p_k and 0. Since there are q points in the orbit, the periodic orbit and 0 give us $q+1$ complementary intervals, called gaps from here on, in which to place the pre-images of 0.

We are now ready to begin the process of placing the pre-images of 0 on the circle in as many ways as possible. We will unite the pre-images of 0 into groups; a group consists of all pre-images of 0 belonging to the same gap. Therefore groups are ordered on the circle and *unlinked* in the sense that their convex hulls are disjoint. We can say that groups are *finite intervals* of points. Some groups may be empty (this corresponds to the fact that there are no pre-images of 0 in a certain gap). In this way we partition all $d-1$ pre-images of 0 into $q+1$ disjoint and pairwise unlinked finite intervals of points which themselves are ordered on the circle.

Since the order of gaps is defined, it may seem that the number of ways one can divide $d-1$ pre-images of 0 into $q+1$ ordered among themselves finite intervals of points would answer the question. However there is one more necessary condition which must be satisfied here. Namely, since 0 is between p_k and p_1 then there must be at least one pre-image of 0 between p_{k-1} and p_q. This is because the gap (p_{k-1}, p_q) maps over the gap (p_k, p_1) (not necessarily one-to-one). Now, given p/q and the choice of p_1, the gap (p_{k-1}, p_q) is well-defined. So what we need to count is the number of partitions of $d-1$ pre-images of 0 into $q+1$ ordered among themselves finite intervals of points (groups) such that a particular group is non-empty.

It is easy to see that this is equivalent to dividing $d-2$ ordered points into $q+1$ groups because then we can always add one more element to exactly the group which must be non-empty to make sure that it is non-empty. Now, to divide $d-2$ ordered points into $q+1$ groups we first add q fictitious points ("dividers") to the set of $d-2$ points and then choose q points out of the just created collection of $d-2+q$ points. This yields the number C_q^{d-2+q}. We need to show that this is the number of rotational periodic orbits of z^d of rotation number p/q.

In fact one way it has already been shown: every rotational orbit of rotation number p/q of z^d must give rise to one of the C_q^{d-2+q} itineraries listed in the arguments above. Moreover, by the construction distinct itineraries correspond to distinct periodic orbits. What remains to show is that each such itinerary gives rise to a (rotational) orbit. To observe this let us consider a given "abstract" rotational periodic orbit P of rotation number p/q and follow the construction from above inserting points which will play the roles of 0 and its pre-images. This provides us with an abstract itinerary which is rotational, but which we must show can be realized by σ_d.

Having done this, construct a map of the circle into itself as follows. On the rotational periodic orbit it acts exactly as the corresponding rotation. On the arcs connecting points of the periodic orbit it acts as a version of "connect-the-dots" map except that the dots are connected on the circle. In other words, let a, b belong to the periodic orbit P and a', b' be their images. Moreover, let (a, b) be the arc in S^1 containing no points of P, and the direction from a to b is counterclockwise. Then we define our map so that it maps (a, b) onto a counterclockwise arc connecting a' and b'. Moreover, if in the above construction we assume that there are several pre-images of 0 in (a, b) then the map we construct will have to wrap around the circle exactly this number of times before b gets mapped into b'. Clearly, by the construction the map which we get will be of degree d, and the initially chosen orbit P will be its periodic orbit.

By arguments standard in one-dimensional theory one can show now that this map can actually be monotonically semi-conjugated to z^d so that the orbit in question will map onto a rotational orbit of z^d of the same rotation number, and moreover, the same itinerary as the one corresponding to the construction. Since the correspondence between periodic orbits and itineraries is one-to-one, there exists a unique rotational periodic orbit corresponding to each of C_q^{d-2+q} itineraries constructed above. This completes the proof. □

We now prove Theorem 2.9, that A rotational set for σ_d with a given rational rotation number contains at most $d-1$ periodic orbits. The following lemma is the heart of the proof of Theorem 2.9. The theorem's proof immediately follows the proof of the lemma.

Lemma 4.1. *For any σ_d-rotational set A ($d > 2$), let J be any hole in the complement of A and the map m subsequently defined as in the proof of Reduction Theorem 2.3. Then there cannot exist 3 (or more) periodic orbits of A taken to one by the map m. In fact, given any n σ_d-periodic orbits in A, they must map to at least $n-1$ σ_{d-1}-periodic orbits under m.*

At first, this may seem to be a vacuous statement. Since the rotational set must be in the complement of J and all its pre-images, and since the map m is at most two-to-one on the maximal Cantor set C (recall proof of Reduction Theorem 2.3), the lemma may appear to easily follow from Reduction Theorem 2.3. However, all we know is that $A \subset X$, and not necessarily that $A \subset C$ (here X is the set of all points avoiding J, and so $C \subset X$), hence certain arguments are necessary to cover the case when periodic orbits from A are not contained in C.

Proof. Assume, by way of contradiction, that there exists a rotational set A, and that some $J \subset \mathbb{T} \setminus A$, such that A contains three periodic orbits that are mapped to one under the map m defined by J as in Reduction Theorem 2.3. Hence, the m-pre-image of each point of this periodic orbit $P = \{p, \sigma_{d-1}(p), \dots, \sigma_{d-1}^{q-1} p\}$ is non-degenerate. So $m^{-1}(p) = I$ is a non-degenerate arc. Moreover, we can always choose p so that I is the smallest possible.

All arcs complementary to C contain arcs which are pre-images of J. Moreover, among arcs complementary to C, there is only one greater than or equal in length to $1/d$, namely the arc S, complementary to C, containing J (otherwise m cannot semi-conjugate $\overline{\sigma}_d$ and σ_{d-1}), and this complementary arc S cannot cover its image more than 2-to-1 for the same reason. Let $S = (u, v)$ and $J = (a, a + 1/d)$ with both arcs oriented counterclockwise. Then I maps onto S by some power of σ_d, and without loss of generality, we may assume that two points x, y from the periodic orbits we study belong to $[u, a]$ (recall that our periodic orbits avoid J). The arc complementary to C containing $\sigma_d(a)$ is $(\sigma_d(u), \sigma_d(v))$. It is less than $1/d$ in length, and will then be mapped onto its images until it gets mapped back onto S because p is periodic. Moreover, since along the way its length grows, we conclude that actually $(u, v) = I$. In any case, in the described situation we get $\sigma_d^q(x) = x, \sigma_d^q(y) = y$ (because x, y are periodic) which contradicts the expanding properties of σ_d.

The final part of the lemma follows by noting that even if we do have two σ_d-periodic orbits that are mapped to one σ_{d-1}-periodic orbit, we cannot have another pair mapped to one because they would also have to go through the big gap containing J, meaning we would have all four orbits being taken to one under the map m, and the above argument rules this out. □

We now complete the proof of Theorem 2.9.

Proof. Suppose, by way of contradiction, there existed d or more periodic orbits in a σ_d-rotational set. By applying Lemma 4.1, along with Reduction Theorem 2.3, $d-2$ times we would see that this σ_d-rotational set would have to project down to a σ_2-rotational set containing at least two σ_2-periodic orbits, a contradiction. □

5. Rotational Cantor Sets for σ_d

5.1. **Proof of Theorem 2.5.** We now prove that σ_d-rotational Cantor sets are minimal.

Proof. We first show that any σ_3-rotational Cantor set C is minimal. From the proof of Reduction Theorem 2.3 we have that $m(C)$ is a σ_2-rotational Cantor set, where m is the map defined as in the proof of the Reduction Theorem 2.3. Moreover, we know from the work of Bullett and Sentenac that all σ_2-rotational Cantor sets are minimal. Now take any point $c \in C$ and any point $c_1 \in C$. We will show that C is minimal by showing that the limit set of the orbit of c contains arbitrary $c_1 \in C$; i.e., the orbit of c is dense. Consider $m(c)$ in our σ_2-Cantor set. By minimality of this Cantor set, the orbit of $m(c)$ under σ_2 is dense. Hence, there exists a subsequence of the orbit of $m(c)$ which approaches $m(c_1)$. If c_1 is the only point of C which maps to $m(c_1)$ under m, then clearly this subsequence lifts to a subsequence of the orbit of c which converges to c_1. If m carries another point of C, say c_2, to $m(c_1)$, then c_1 and c_2 must be endpoints of some gap of C that is shrunk to a point by m. In this case though, $m(c_1) = m(c_2)$ cannot be an endpoint of our σ_2-Cantor set. One can then show by minimality that the orbit of $m(c)$ must approach $m(c_1)$ from both sides. When this is lifted up to σ_3, this sequence is split over the gap that was shrunk to a point, one side converging to one endpoint c_1, the other converging to c_2. Thus, we have shown that the σ_3 orbit of c is dense, so C is minimal.

The general case now follows from an induction on d in which the induction step uses the same argument as above. □

5.2. **Root Gaps of Rotational Cantor Sets.** For this section, assume C is a σ_d-rotational Cantor set contained in $\mathbb{T}$. Note that C is minimal by Theorem 2.5. A few lemmas proven below are of a technical nature. A major observation which makes their proofs more straightforward is that Corollary 2.4 can be applied not just once but a few times. This obviously results in the following conclusion: *If C is a rotational Cantor set then there exists a map m_C which semi-conjugates $\sigma_d|C$ with an irrational rotation* τ of $\mathbb{T}$. In fact, m_C simply collapses all gaps of C. However, note that a taut root gap G of C goes to a point p_G such that $m_C^{-1}(\tau(p_G))$ is a point, while a loose gap H of C goes to a point p_H such that $m_C^{-1}(\tau(p_H))$ is a gap of C.

For the sake of completeness let us suggest a sketch of an alternative proof of the existence of m_C which does not rely upon the developed techniques (well, almost). Even though this is only a sketch, it presents a different way of arguing and hence may be of interest. Given a σ_d-invariant rotational Cantor set C, we associate with it a well-defined rotation number ρ. If ρ is rational then the points of C are either periodic or preperiodic, so there are no more than countably many of them, a contradiction. Thus ρ is irrational. Take a minimal subset C' of C (recall that we do not rely upon the above developed tools!). Then C' is a Cantor set itself and $\sigma_d|C'$ is onto, and a map ϕ which collapses all gaps of C', is well-defined. Clearly, $\phi(\mathbb{T})$ is a circle. Since $\sigma_d|_C$ preserves cyclic order then σ_d maps the endpoints of a gap in C' either onto the endpoints of a gap in C', or onto the same point. Thus, ϕ semi-conjugates $\sigma_d|_{C'}$ with a map τ of the circle which does not change the cyclic order. Moreover, τ is minimal because so is $\sigma_d|_{C'}$. This easily implies that τ is an irrational rotation. Since the original set C is a Cantor set, then in fact $C = C'$, and the claim about the existence of a semi-conjugacy m_C between $\sigma_d|_C$ and an irrational rotation of the circle is proven.

Lemma 5.1. *Let M be a monotone extension of $\sigma_d|_C$. Then the following holds:*

(1) *M is one-to-one on C, except possibly at the endpoints of root gaps.*
(2) *M is two-to-one on endpoints of taut root gaps.*
(3) *M is one-to-one on endpoints of loose root gaps.*

Proof. Immediately follows from the existence of m_C and its action on points corresponding to taut and loose root gaps. □

Lemma 5.2. *If I is a complementary gap of C which is not a root gap, then $\sigma_d(I)$ is a complementary gap of C. Moreover, there exists a unique k such that $\sigma_d^k(I)$ is a root gap.*

Proof. The first part of the lemma follows from the existence of m_C. To prove the second part, observe that as we apply σ_d to non-root gaps their length grows d-fold, hence at some moment a non-root gap will be mapped onto a root gap for the first time. Since the next time its image is the entire circle, this is in fact the only time when the root gap is mapped onto a root gap as desired. □

Lemma 5.3. *Let $I' = (a', b')$ be a loose root gap of C. Let $x \in I'$ be maximal such that $\sigma_d(x) = \sigma_d(a')$. Then $\sigma_d((x, b'))$ is a complementary gap of C.*

Proof. Follows from the existence of m_C and its action on points corresponding to loose root gaps. □

Definition 5.4. Let $I' = (a', b')$ be a loose root gap. Let x be the maximal point in I' such that $\sigma_d(a') = \sigma_d(x)$. Then the gap $(\sigma_d(x), \sigma_d(b')) = (\sigma_d(a'), \sigma_d(b'))$ is called the *overshoot interval.*

Lemma 5.5. *Let $I' = (a', b')$ be a loose root gap of C. Let $K = (\sigma_d(x), \sigma_d(b'))$ be the corresponding overshoot interval. Given the unique $k \geq 0$ from Lemma 5.2 such that $\sigma_d^k(K)$ is a root gap, we have $\sigma_d^k(K) \cap I' = \emptyset$.*

Proof. First note that k could be equal to 0 because K may itself be a root gap. Since $\sigma_d^k(K)$ and I' are both root gaps, if $\sigma_d^k(K) \cap I' \neq \emptyset$, then $\sigma_d^k(K) = I'$. Moreover, $\sigma_d^k|_K$ has not reversed the order of the points of x and b'. Therefore, $\sigma^{k+1}(b') = b'$, a contradiction with minimality of C. We conclude that $\sigma_d^k(K) \cap I' = \emptyset$. □

Recall that if a root gap is such that $\frac{n}{d} \leq l(G) < \frac{n+1}{d}$ for some $n \in \mathbb{N}$ then the number n is called the *root number* of G. Recall also that we have assumed $d > 2$.

Theorem 5.6 (Root Gap Length)**.** *Each root gap of C with root number k_0 has length*

$$\frac{k_0}{d} + \sum_{i=1}^{d-2} \frac{k_i}{d^{l_i}}, \text{ where } l_i > i, l_{i+1} > l_i \text{ and } \sum_{i=0}^{d-2} k_i \leq d-1.$$

It is true that for any possible length given by the formula above there exists a Cantor set with a gap of that length. We provide an explicit proof in the case $d = 3$ in Theorem 5.7. The proof below can be utilized as a recursive recipe for σ_d-rotational Cantor sets, $d > 3$.

Proof. Note that one root gap of C must be taut by Proposition 2.2. Take a root gap $R = (a, b)$ of C with root number k_0. If R is taut then we are done as it has length $\frac{k_0}{d}$ for some $k_0 \leq d-1$. So assume R is loose. Let's track R's overshoot interval $\widetilde{R}$.

As we know $\widetilde{R}$, which is a complementary gap (Lemma 5.3), maps to a root gap R_1 (Lemma 5.2) of the Cantor set C after $m_1 \geq 0$ steps. If R_1 is taut then $l(R_1) = \frac{k_1}{d}$ where k_1 is the root number for R_1. It follows that

$$l(\widetilde{R}) = \frac{l(R_1)}{d^{m_1}} = \frac{k_1}{d^{m_1+1}}.$$

By knowing the length of $\widetilde{R}$, we then know that

$$l((x, b')) = \frac{l(\widetilde{R})}{d} = \frac{k_1}{d^{m_1+2}},$$

and therefore that

$$l(R) = l((a', x)) + l((x, b')) = \frac{k_0}{d} + \frac{k_1}{d^{m_1+2}}.$$

If R_1 is loose then we can inductively take the argument one more step further. Suppose then that $R_1's$ overshoot interval, $\widetilde{R}_1$, maps to a root gap R_2 after $m_2 \geq 0$ steps which happens to be taut with root number k_2. If we then apply the above case to R_1 (R_1 takes the role of R above and R_2 the role of R_1), we see that

$$l(R_1) = \frac{k_1}{d} + \frac{k_2}{d^{m_2+2}}.$$

This implies that

$$l(\widetilde{R}) = \frac{l(R_1)}{d^{m_1}} = \frac{k_1}{d^{m_1+1}} + \frac{k_2}{d^{m_1+m_2+2}}.$$

And getting back to R we get

$$l(R) = \frac{k_0}{d} + \frac{l(\widetilde{R})}{d} = \frac{k_0}{d} + \frac{k_1}{d^{m_1+2}} + \frac{k_2}{d^{m_1+m_2+3}}.$$

If R_2 happens to be loose also then we need to map its overshoot interval, $\widetilde{R}_2$, to another root gap R_3, determine if it is taut or not, and apply the same inductive process as mentioned above. We can keep repeating this process until R_{n-1}'s overshoot interval maps to a taut root gap R_n in $m_n \geq 0$ steps. This must occur for every root gap R of C as consequence of Lemma 5.2.

Moreover, it must occur in at most $n \leq d-2$ iterations of the above process, with the first iteration being finding R_1 (i.e. $n \leq d-2$). When the above process terminates we would have root gaps R and $\{R_1, .., R_n\}$. These root gaps must be pairwise disjoint, for otherwise two of them coincide. That would mean that a root gap mapped back to itself, making its endpoints, which are in the Cantor set, periodic, a contradiction. Therefore, we have $n+1$ disjoint root gaps $\{R, R_1, \ldots, R_n\}$ in the above construction. On the other hand, by Theorem 2.1 there are no more than $d-1$ root gaps. Hence $n \leq d-2$, and the number of terms in the root gap length summation can not exceed $d-1$.

Let us consider several cases. If we happen to start with a taut root gap, then $n = 0$. If we do not start with a taut root gap but the first overshoot interval maps to a taut root gap, then $n = 1$. Now, with $\{m_1, .., m_n\}$ and $\{k_1, .., k_n\}$ defined as above, we simply let

$$m_{n+1} = m_{n+2} = ... = m_{d-1} = 0 = k_{n+1} = k_{n+2} = .. = k_{d-2}.$$

Defining

$$l_i = \sum_{j=1}^{i} m_i + (i+1) \text{ for } i \geq 1$$

we see that we have found l_i such that $l_i > i$ and $l_{i+1} > l_i$ are satisfied, and that

$$l(R) = \frac{k_0}{d} + \sum_{i=1}^{d-2} \frac{k_i}{d^{l_i}}.$$

What is left to show is $\sum_{i=0}^{n} k_i \leq d-1$ for each root gap. This follows easily from the fact that the root gaps in the construction were pairwise disjoint and that each one is at least as big as its corresponding root number, k_i, divided by d. Therefore, if $\sum_{i=0}^{n} k_i > d-1$ then our root gaps would cover the entire circle except possibly for a finite set of endpoints, a contradiction. Hence $\sum_{i=0}^{n} k_i \leq d-1$, and since $\sum_{i=0}^{d-2} k_i = \sum_{i=0}^{n} k_i$, we have that $\sum_{i=0}^{d-2} k_i \leq d-1$. □

In the theorems below we refer to the rotational Cantor sets that arise from "lifting" a Cantor set in the Construction Theorem 2.6. We can do this because although the theorem only guarantees us that a Cantor set lifts to some rotational set with the same irrational rotation number, we know it must contain a Cantor set because the map M is one-to-one except for countably many points. Moreover, it is only strictly bigger than a Cantor set if, in the Construction Theorem 2.6, we insert intervals at endpoints of the Cantor set. If this happens then the endpoint lifts to two points of our rotational set. Of such pairs of points, one would be in our lifted Cantor set, while the other would be off on its own in a gap. We can then throw away all such isolated points without affecting the fact that our rotational set is closed and invariant. The invariance is not affected because the isolated points of our lifted set map to each other, except at the last step, the inserted I_1^0 interval. At this stage, the isolated point and its Cantor set counterpart are mapped to the same point in the Cantor set. Thus, after throwing away this countable set of isolated points, invariance of the lifted set is preserved. By throwing away these points we are left with a rotational Cantor set, which is what we will mean when we refer to the Cantor set obtained from another Cantor set via the Construction Theorem 2.6.

Incidentally, if we do not throw away the isolated point of the lifted set, but rather keep all points that are iterated pre-images of it and contained in the complement of the holes, then we produce an irrational rotational set which is not minimal. No such irrational rotational sets exist for σ_2.

5.3. **Rotational Cantor Set Construction.** In this section we explicitly classify and construct all rotational Cantor sets for σ_3. The extension to σ_d, $d > 3$, is left to the reader.

Theorem 5.7 (Cantor Set Construction)**.** *Let α be any irrational rotation number. Then, for each σ_3-rotational Cantor set $\widetilde{C}_\alpha$ with rotation number α, exactly one of the following conditions holds:*

(1) *There exists a gap of length $\frac{2}{3}$.*
(2) *There exist two disjoint gaps of length $\frac{1}{3}$.*
(3) *There exist two disjoint gaps, one of length $\frac{1}{3}$, the other of length $\frac{1}{3} + \frac{1}{3^k}$, for some $k \in \{2, 3, ...\}$.*

Moreover, for each irrational α, each of these cases is realized by some σ_3-rotational Cantor set $\widetilde{C}_\alpha$ with rotation number α.

Proof. Given a σ_3-rotational Cantor set $\widetilde{C}_\alpha$ with irrational rotation number α, we can just apply Hole Theorem 2.1 and Theorem 5.6 to prove that exactly one of the above conditions holds.

The remaining claims of the theorem follow from careful consideration of the map σ_3^* and insertion point x_1^0 defined in the proof of Construction Theorem 2.6. Consider first the unique σ_2-rotational Cantor set, C_α, with irrational rotation number α. Recall that

C_α has a taut hole of length $\frac{1}{2}$. If x_1^0 is inserted in this hole, or at its endpoints, then the corresponding gap of $\widetilde{C}_\alpha$, which we obtain by Construction Theorem 2.6, contains I_1^0 and therefore must be larger than $\frac{1}{3}$. (An isolated point is removed from the gap in the case that an insertion was done at an endpoint. See above remark.) Since σ_3^* maps this gap's endpoints together by conjugacy, it must be of length $\frac{2}{3}$ and we are in case 1.

If x_1^0 is inserted in any other gap or endpoint of a gap, then the corresponding gap of $\widetilde{C}_\alpha$ must be a root gap. The map σ_3^* does not map the endpoints of this gap together, so it is a gap of length bigger than $\frac{1}{3}$. Moreover, the gap of C_α which was exactly $\frac{1}{2}$ in length, must form the corresponding gap of $\widetilde{C}_\alpha$ that is of length $\frac{1}{3}$, because its endpoints must map together by the conjugacy. Hence we are in case 3. (Note that the gap in which x_1^0 is inserted - or rather the number of steps which is necessary to map this gap onto the taut root gap of length $\frac{1}{2}$ in C_α - determines the exact k used in the length above and hence every k can be realized. See remark following Theorem 5.6.)

Note that in the above two cases, if we insert at an endpoint, then only one of the points to which it pulls back under the map M will actually be in $\widetilde{C}_\alpha$. The other one is isolated in a gap of $\widetilde{C}_\alpha$.

Finally, if x_1^0 is inserted at a non-endpoint of C_α, then both points to which it pulls back will be in $\widetilde{C}_\alpha$ (compare previous paragraph). This will be one gap of $\widetilde{C}_\alpha$ which is of length $\frac{1}{3}$. The other will be the gap corresponding to the $\frac{1}{2}$-gap of C_α, as above. Hence, we are in case 2. □

Every σ_d-rotational set can be obtained by construction from a σ_{d-1}-rotational set. In particular, every σ_d-rotational set arises from $d-2$ consecutive appropriate applications of Theorem 2.6 to the σ_2-rotational set with the same rotation number.

5.4. **Proof of Counting Theorem 2.11 for rotational Cantor sets.**

Proof. Let C be any σ_d-rotational Cantor set for $d \geq 2$. Pick any two points x and y in C, with $x < y$, as long as they are not the endpoints of the same complementary gap of C. Consider the rotational Cantor sets C_x and C_y, obtained from the Construction Theorem 2.6 when beginning the insertion at x and y respectively. To simplify the proof, we may assume that the homeomorphism which labels the fixed points of σ_{d+1}, defined during construction (see Proposition 3.7), is in fact the one which labels the pre-image of the fixed point 0 of σ_d as the 0 under σ_{d+1}. This allows us to avoid the case where construction done at $d-1$ different points of C may give rise to the same σ_{d+1}-Cantor set if the homeomorphism permutes the fixed points in the right way. Both C_x and C_y are minimal by Theorem 2.5. We will show that C_x and C_y are different Cantor sets.

Let x_l and y_l be the least points of C_x and C_y, respectively, counterclockwise from 0. Then their orbits are dense, by minimality, and follow the pattern of the orbit of the least point of C, call it z, to which they project under the map M of the Construction Theorem 2.6 proof (note that z is the least point of C because of the particular homeomorphism chosen). Consider the first time that z maps between x and y. Then, at the same time 'upstairs' we have that x_l must have jumped over the wrap inserted at x while y_l has not yet jumped over the wrap inserted at y. Since these wraps must contain a pre-image of 0 under the map σ_{d+1}, we see that the itineraries of x_l and y_l differ at this time step. Thus, x_l and y_l are different, and we have shown that C_x and C_y are different. The fact that they have the same rotation number as C comes from Corollary 2.7.

Since we can do the insertion at c-many points of C which are not endpoints of gaps to construct different Cantor sets under σ_{d+1}, we know immediately that there are c-many σ_{d+1}-rotational Cantor sets with the same irrational rotation number as C. Moreover, we can imitate the proof of Theorem 5.7 to show that an insertion done at either end or inside a complementary gap of C leads to a corresponding loose root gap in the constructed Cantor set, and insertion at a non-endpoint leads to a taut root gap in the constructed Cantor set. These Cantor sets formed by different insertion points are indeed different, as shown above, and so the theorem is proven. □

In fact, we have proven more. Let x stand for a non-endpoint or closure of a gap of the σ_d-rotational Cantor set C_α with rotation number α. We have shown that there exists a one-to-one order-preserving correspondence between $\mathcal{S}_\alpha = m_{C_\alpha}(C_\alpha)$ and the σ_{d+1}-rotational Cantor sets that arise from lifting C_α by inserting in $\mathbb{T}$ a wrap (via the Construction Theorem 2.6) at x.

References

[1] L. Alsed, J. Llibre, and M. Misiurewicz, *Combinatorial Dynamics and Entropy in Dimension One*, Second edition, Advanced Series in Nonlinear Dynamics 5, World Scientific Publishing Co., Inc., River Edge, NJ, 2000.

[2] S. Bullett and P. Sentenac, Ordered orbits of the shift, square roots, and the devil's staircase, *Math. Proc. Camb. Phil. Soc.* 115: 451–481, 1994.

[3] A. Douady and J.H. Hubbard, Etude dynamique des polynomes complexes, *Publications Mathematiques d'Orsay* Parties I et II, 1984–85.

[4] L.R. Goldberg, Fixed points of polynomial maps I. rotation subsets of the circles, *Annales Scientifiques de l'Ecole Normale Superieure, 4e serie* 25: 679–685, 1992.

[5] L.R. Goldberg and J. Milnor, Fixed points of polynomial maps II. fixed point portraits, *Annales Scientifiques de l'Ecole Normale Superieure, 4e serie* 26: 51–98, 1993.

[6] L.R. Goldberg and C. Tresser, Rotation orbits and the Farey tree, *Ergodic Theory and Dynamical Systems* 16: 1011–1029, 1996.

[7] J. Grispolakis, J.C. Mayer, and L.G. Oversteegen, Building blocks for quadratic Julia sets, *Transactions of the AMS* 351(3): 1171–1201, 1999.

[8] J. Kiwi, Real laminations and the topological dynamics of complex polynomials, preprint, 2000.

[9] J.M. Malaugh, Parameterizing rotational subsets of the circle under z^d, preprint 2003.

[10] J.C. Mayer and J.T. Rogers, Jr., Indecomposable Julia sets of polynomials, *Proceedings of the AMS* 117(3): 795–802, 1993.

[11] J. Milnor, *Dynamics in One Complex Variable*, Vieweg, Braunschweig, Germany, 2000.

E-mail address, Alexander Blokh: `ablokh@math.uab.edu`

E-mail address, James M. Malaugh: `malaugh@math.uab.edu`

E-mail address, John C. Mayer: `mayer@math.uab.edu`

E-mail address, Lex G. Oversteegen: `overstee@math.uab.edu`

(Alexander Blokh, James M. Malaugh, John C. Mayer, Lex G. Oversteegen, and Daniel K. Parris) Department of Mathematics, University of Alabama at Birmingham, Birmingham, AL 35294-1170