Chapter 2
Rotation Sets

Throughout this chapter d will be a fixed integer ≥ 2. We study certain invariant sets for the multiplication by d map $m_d : \mathbb{T} \to \mathbb{T}$ defined by

$$m_d(t) = dt \qquad (\text{mod } \mathbb{Z}).$$

The low-degree cases m_2 and m_3 are often referred to as the ***doubling*** and ***tripling*** maps.

Definition 2.1 A non-empty compact set $X \subset \mathbb{T}$ is called a ***rotation set*** for m_d if

- X is m_d-invariant in the sense that $m_d(X) = X$,[1] and
- the restriction $m_d|_X$ can be extended to a degree 1 monotone map of the circle.

Roughly speaking, the latter condition means that m_d preserves the cyclic order of all triples in X, except that it may identify some pairs.

If X is a rotation set for m_d and g, h are degree 1 monotone extensions of $m_d|_X$, then $g = h$ on every orbit in X, so $\rho(g) = \rho(h)$ by Theorem 1.8. This quantity, which therefore depends on X only, is called the ***rotation number*** of X and is denoted by $\rho(X)$. We refer to X as a rational or irrational rotation set according as $\rho(X)$ is rational or irrational.

2.1 Basic Properties

Since multiplication by d commutes with the rigid rotation $r : t \mapsto t + 1/(d-1)$ (mod $\mathbb{Z}$), if X is a rotation set for m_d, so are its $d-2$ rotated copies

$$X + \frac{1}{d-1}, \; X + \frac{2}{d-1}, \; \ldots, \; X + \frac{d-2}{d-1} \quad (\text{mod } \mathbb{Z}).$$

[1]Thus, our notion of invariance is stronger than *forward invariance* $m_d(X) \subset X$ and weaker than *full invariance* $m_d^{-1}(X) = X$.

S. Zakeri, *Rotation Sets and Complex Dynamics*, Lecture Notes in Mathematics 2214, https://doi.org/10.1007/978-3-319-78810-4_2

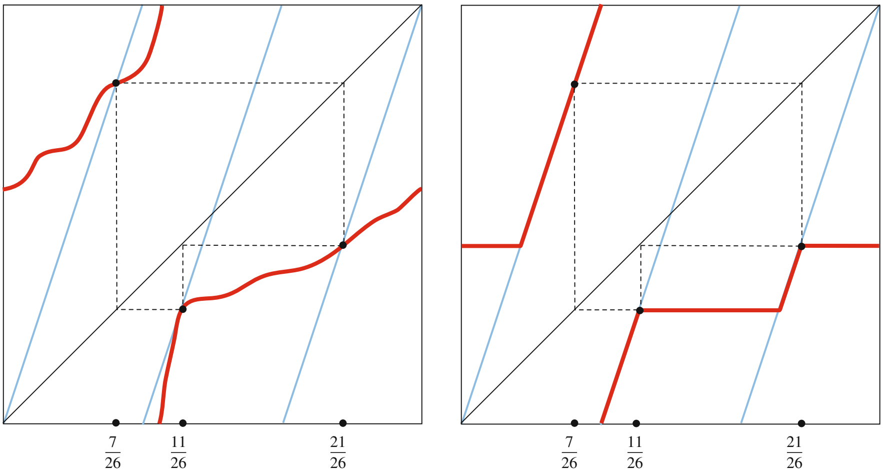

Fig. 2.1 The cycle $X : \frac{7}{26} \mapsto \frac{21}{26} \mapsto \frac{11}{26}$ under tripling is a rotation set with $\rho(X) = \frac{2}{3}$. Left: A generic monotone extension of $m_3|_X$. Right: The "standard" monotone extension of $m_3|_X$ (see the discussion leading to Eq. (2.1))

Moreover, all these sets have rotation number $\rho(X)$ since if g is a monotone extension of $m_d|_X$, then the conjugate map $r^{\circ i} \circ g \circ r^{-i}$ will be a monotone extension of the restriction of m_d to $X + i/(d-1)$.

Example 2.2 The 3-cycle $X : \frac{7}{26} \mapsto \frac{21}{26} \mapsto \frac{11}{26}$ under tripling is a rotation set with rotation number $\frac{2}{3}$. Two possible monotone extensions of m_3 restricted to this cycle are shown in Fig. 2.1. The 180°-rotation of X produces the new rotation set $X + \frac{1}{2} : \frac{8}{26} \mapsto \frac{24}{26} \mapsto \frac{20}{26}$ with the same rotation number. On the other hand, the 4-cycle $\frac{1}{5} \mapsto \frac{3}{5} \mapsto \frac{4}{5} \mapsto \frac{2}{5}$ under tripling is *not* a rotation set since it fails to have a combinatorial rotation number (compare Corollary 1.16).

A rotation set containing periodic orbits is clearly rational. Conversely, every orbit in a rational rotation set is eventually periodic. Here is a more precise statement:

Theorem 2.3 *Suppose X is a rational rotation set for m_d, with $\rho(X) = p/q$ in lowest terms. Then, every forward orbit in X under m_d is finite. More precisely, for every $t \in X$ there is an integer $i \geq 0$ such that $m_d^{\circ i}(t)$ is periodic of period q. In particular, X is at most countable.*

Proof Take any $t \in X$ and any degree 1 monotone extension g of $m_d|_X$. We know from Theorem 1.14 that the sequence $\{g^{\circ nq}(t) = m_{d^q}^{\circ n}(t)\}$ tends to a periodic point $t' \in X$ of period q as $n \to \infty$. Since the map m_{d^q} is uniformly expanding on the circle, its fixed point t' is repelling. Hence $m_{d^q}^{\circ n}(t)$ cannot converge to t' unless $m_{d^q}^{\circ n}(t) = t'$ for some n. □

Remark 2.4 Most periodic orbits of m_d do not define rotation sets. For each prime number q the equation $m^{\circ q}(t) = t$ has $d^q - 1$ solutions $t = i/(d^q - 1) \pmod{\mathbb{Z}}$. Discarding the $d - 1$ fixed points of m_d, it follows that the number $(d^q - d)/q$ of distinct q-cycles of m_d grows exponentially fast as $q \to \infty$. On the other hand, the number of q-cycles of m_d that form a rotation set is precisely $(q-1)\binom{q+d-2}{q}$, which grows like the power q^{d-1} as $q \to \infty$ (see Corollary 3.11).

Every rotation set is nowhere dense since any open interval on the circle eventually maps to the whole circle under the iterations of m_d. By contrast, Lebesgue measure on the circle is ergodic for m_d,[2] so a randomly chosen point on $\mathbb{T}$ has a dense orbit almost surely. This proves the following

Theorem 2.5 *The union $\mathcal{R}_d$ of all rotation sets for m_d has Lebesgue measure zero.*

McMullen [19] has proved the sharper statement that the Hausdorff dimension of $\mathcal{R}_d$ is zero.[3] For more on the set $\mathcal{R}_d$, see Sect. 4.3.

To study of the structure of a rotation set, we first look at its complement.

Definition 2.6 Let X be a rotation set for m_d. A connected component of the complement $\mathbb{T} \smallsetminus X$ is called a ***gap*** of X. A gap of length ℓ is ***minor*** if $\ell < 1/d$ and ***major*** otherwise. The ***multiplicity*** of a major gap is the integer part of $d\ell \geq 1$. A major gap is ***taut*** or ***loose*** according as $d\ell$ is or is not an integer.

Intuitively, a minor gap is short enough so it maps homeomorphically onto its image by m_d. On the other hand, a major gap is too long and wraps around the circle by m_d as many times as its multiplicity (see Lemma 2.8 below).

It will be convenient to work with a specific degree 1 monotone extension of $m_d|_X$ which can be defined whenever X has more than one point. This map, which we call the ***standard monotone map*** of X and denote by g_X, is defined as follows: On every minor gap, set $g_X = m_d$. On every major gap $(a, a+\ell)$ of length $0 < \ell < 1$ and multiplicity n, define

$$g_X(t) = \begin{cases} m_d(a) & t \in \left(a, a + \frac{n}{d}\right] \\ m_d(t) & t \in \left(a + \frac{n}{d}, a + \ell\right) \end{cases} \tag{2.1}$$

(see Figs. 2.1 and 2.2). The map g_X is piecewise affine with derivatives 0 and d, so the total length of the plateaus of g_X is $1 - 1/d = (d-1)/d$. Since by the

[2] Assuming $m_d^{-1}(E) = E$ for some measurable set E, the characteristic function χ_E satisfies $\chi_E \circ m_d = \chi_E$. Expanding χ_E into the Fourier series $\sum c_n e^{2\pi i n t}$, it follows that $\sum c_n e^{2\pi i d n t} = \sum c_n e^{2\pi i n t}$ which implies $c_n = c_{dn}$ for all n. Since $c_n \to 0$, this can hold only if $c_n = 0$ for all $n \neq 0$.

[3] He proves the statement for the closure of the union of all *finite* rotation sets for m_d, but an inspection of his proof shows that it also works for the a priori larger set $\mathcal{R}_d$. The zero dimension statement for individual rotation sets was known much earlier [29].

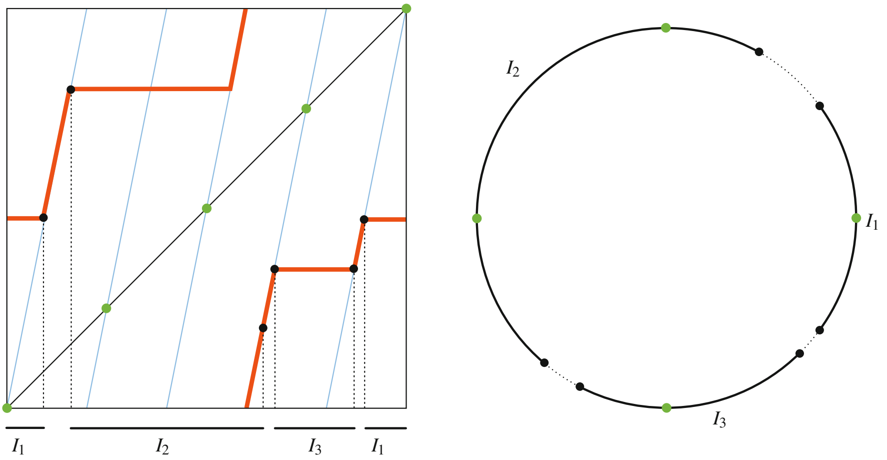

Fig. 2.2 Left: The standard monotone map g_X of some rotation set X for m_5. Counting multiplicities, X has four major gaps, two taut gaps I_1, I_3 of multiplicity 1 and a loose gap I_2 of multiplicity 2. Right: The position of major gaps around the circle. Notice that each major gap contains as many fixed points of m_d as its multiplicity, as asserted in Lemma 2.13

construction each major gap of multiplicity n contributes a plateau of length n/d, we arrive at the following fundamental fact (compare [2] and [4]):

Theorem 2.7 *Every rotation set for m_d containing more than one point has $d-1$ major gaps counting multiplicities.*

The following lemma summarizes the mapping properties of gaps:

Lemma 2.8 *Let X be a rotation set for m_d containing more than one point and $I = (a, a+\ell)$ be a gap of X. Take any degree 1 monotone extension g of $m_d|_X$.*

(i) If I is a minor gap, the interior J of $g(I)$ is a gap of length $d\ell$. Moreover, $m_d : I \to J$ is a homeomorphism.
(ii) If I is a taut gap of multiplicity n, the image $g(I)$ is the single point $m_d(a) \in X$. Under m_d, the point $m_d(a)$ has $n-1$ preimages in I, whereas every point in $\mathbb{T} \smallsetminus \{m_d(a)\}$ has n preimages in I.
(iii) If I is a loose gap of multiplicity n, the interior J of $g(I)$ is a gap of length $d\ell - n$. Under m_d, every point in J has $n+1$ preimages in I, whereas every point outside J has n preimages in I.

Proof For the standard monotone map g_X the statements follow immediately from the definition. For an arbitrary extension g, we can use the fact that g is monotone and takes the same values as g_X on the boundary of gaps to arrive at the same conclusions. The details are straightforward and will be left to the reader. □

The preceding lemma shows that the pattern of how gaps map around is independent of the choice of the monotone extension g. For any gap I, the image

$g(I)$ is either a point or a gap J modulo its endpoints. In practice, it is convenient to ignore the issue of endpoints and simply declare that I *maps to* J. With this convention in mind, we see from the above lemma that every minor gap eventually maps to a major gap I. If I is taut, it maps to a point and the process stops. If I is loose, it maps to a new gap and the process continues.

Let us collect some corollaries of this basic observation.

Theorem 2.9 *A rotation set is uniquely determined by its major gaps.*

Proof Let X, Y be rotation sets with the same collection of major gaps. We may assume neither of X, Y is a single point. Suppose there is some $t \in Y \smallsetminus X$. Then t must belong to a minor gap I of X. Take the smallest integer $i > 0$ such that $J = m_d^{\circ i}(I)$ is a major gap of X. Then $m_d^{\circ i} : I \to J$ is a homeomorphism, so $m_d^{\circ i}(t) \in J \cap Y$, which is impossible since J is a major gap of Y as well. This proves $Y \subset X$. Similarly, $X \subset Y$. □

Theorem 2.10 *Suppose X is a rotation set containing more than one point and I is a gap of X. Then either I is periodic or it eventually maps to a taut gap.*

Proof Let I_i denote the interior of $g_X^{\circ i}(I)$ and assume that I_i is not taut for any i. By Lemma 2.8 there is a sequence $i_1 < i_2 < i_3 < \cdots$ of positive integers for which I_{i_k} is loose. Since there are finitely many loose gaps, we must have $I_{i_j} = I_{i_k}$ for some $j < k$. This proves that I eventually maps to a periodic gap. Since by monotonicity of g_X every gap is the image of precisely one gap, it follows that I itself must be periodic. □

Corollary 2.11 *Every infinite rotation set has at least one taut gap.*

Conversely, all major gaps of a finite rotation set are loose since in this case m_d, being surjective, must also be injective on the rotation set.

Proof Otherwise every gap would be periodic by the previous theorem, so its endpoints would be periodic points in the rotation set. By Theorem 1.14 these infinitely many endpoints would have the same period $q > 0$ under m_d. This is impossible since m_d has only finitely many q-cycles. □

Remark 2.12 Here is an alternative approach to the above corollary (compare [2]): Lemma 2.8 applied to g_X shows that $m_d(t) = m_d(t')$ for a distinct pair $t, t' \in X$ precisely when t, t' form the endpoints of a taut gap or more generally when there is a chain $t = t_1, t_2, \ldots, t_k = t' \in X$ such that each pair t_i, t_{i+1} forms the endpoints of a taut gap. Thus, if X had no taut gap, the map $m_d : X \to X$ would be a homeomorphism. Since m_d is expanding, this would imply that X is finite [21, Lemma 18.8].

The next result establishes a connection between the major gaps of a rotation set and the $d-1$ fixed points

$$u_i = \frac{i}{d-1} \pmod{\mathbb{Z}}$$

of the map m_d. This connection will play an important role in Sects. 3.2 and 3.3.

Lemma 2.13 *Suppose X is a rotation set for m_d with $\rho(X) \neq 0$. Then each major gap of X of multiplicity n contains exactly n fixed points of m_d.*

Compare Fig. 2.2.

Proof The assumption $\rho(X) \neq 0$ tells us that each fixed point of m_d belongs to a gap, which is necessarily major since a minor gap is disjoint from its image under m_d. Let I be a major gap of multiplicity n and assume that it contains $n+1$ adjacent fixed points $u_i, \ldots, u_{i+n}$. Since each open interval (u_j, u_{j+1}) contains precisely one preimage of every fixed point under m_d, it follows that u_i has at least $n+1$ preimages in I. By Lemma 2.8, I is loose and u_i belongs to the interior of $g_X(I)$. This implies that the closure of I maps onto itself by g_X, so the endpoints of I must be fixed by m_d, which contradicts the assumption $\rho(X) \neq 0$. Thus, I contains at most n fixed points of m_d.

Now let $\{I_i\}$ be the finite collection of major gaps of X of multiplicities $\{n_i\}$. We have shown that the number k_i of fixed points in I_i satisfies $0 \leq k_i \leq n_i$. Since $\sum k_i = \sum n_i = d - 1$, we must have $k_i = n_i$ for all i. □

To each rotation set X for m_d we can assign a ***gap graph*** Γ_X which is a finite directed (not necessarily connected) graph having one vertex for each major gap of X, with an edge going from vertex I to vertex J whenever J is the first major gap in the forward orbit of I. We also assign to each vertex I a weight $w(I) \geq 1$ equal to its multiplicity. Thus, Γ_X has the following properties:

(i) $\displaystyle\sum_{\text{vertices } I} w(I) = d - 1.$

(ii) The degree of every vertex is either 0 (no edge going out or coming in), or 1 (only one edge going out or coming in), or 2 (one edge going out and one coming in, possibly a loop).

If X has no loose gaps, Γ_X is a trivial graph consisting of at most $d-1$ vertices and no edges. If X is an irrational rotation set, Theorem 2.10 tells us that every directed path in Γ_X terminates at a taut vertex and in particular there are no closed paths (see Fig. 2.3).

Let us call a finite directed graph ***admissible*** of degree d if it satisfies the conditions (i) and (ii) above. It is natural to ask the following

Question 2.14 Given an admissible graph Γ of degree d, does there exist a rotation set X for m_d whose gap graph Γ_X is isomorphic to Γ?

In Sect. 4.2 we will provide the answer to this question in the case Γ has no closed paths (see Theorem 4.6).

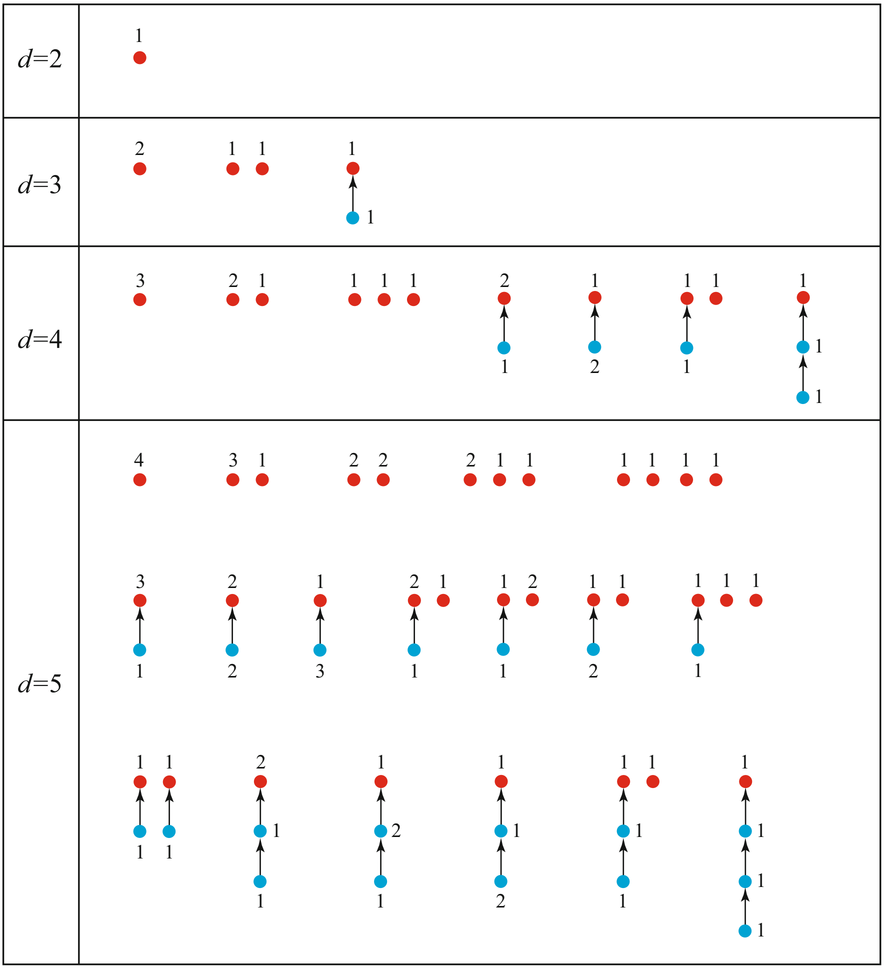

Fig. 2.3 Possible gap graphs for irrational rotation sets under m_d for $2 \leq d \leq 5$. The red and blue vertices correspond to taut and loose gaps respectively, and the weights denote multiplicities

2.2 Maximal Rotation Sets

Take any collection

$$\mathscr{I} = \{I_1, \dots, I_{d-1}\}$$

of *disjoint* open intervals on the circle, each of length $1/d$. Consider the set

$$X_{\mathscr{I}} = \{t \in \mathbb{T} : m_d^{\circ n}(t) \notin I_1 \cup \cdots \cup I_{d-1} \text{ for all } n \geq 0\}.$$

Theorem 2.15 ([4]) *$X_{\mathscr{I}}$ is a rotation set for m_d.*

Proof First we check that $X_{\mathscr{I}} \neq \emptyset$. Denote by U the open set $I_1 \cup \cdots \cup I_{d-1}$. Under m_d, every $t_0 \in \mathbb{T}$ has d preimages which are a distance $1/d$ apart, hence at least one of these preimages, say t_1, must be outside U. It follows inductively that there is a backward orbit $\cdots \mapsto t_2 \mapsto t_1 \mapsto t_0$ such that $t_n \notin U$ for every $n \geq 1$. Evidently, any accumulation point of the sequence $\{t_n\}$ belongs to $X_{\mathscr{I}}$.

It is immediate from the definition that $X_{\mathscr{I}}$ is compact and maps into itself by m_d. Of the d preimages of any point in $X_{\mathscr{I}}$, at least one lies outside U and therefore belongs to $X_{\mathscr{I}}$. This proves $m_d(X_{\mathscr{I}}) = X_{\mathscr{I}}$. Finally, m_d restricted to $X_{\mathscr{I}}$ can be extended to a degree 1 monotone map $g : \mathbb{T} \to \mathbb{T}$ by setting $g = m_d$ outside U and mapping each interval I_i to a point. □

Corollary 2.16 *A non-empty compact m_d-invariant set X is a rotation set if and only if $\mathbb{T} \smallsetminus X$ contains $d-1$ disjoint open intervals, each of length $1/d$.*

Proof Necessity follows from Theorem 2.7. For sufficiency, let $\mathscr{I}$ be the collection of the $d-1$ disjoint intervals of length $1/d$ in $\mathbb{T} \smallsetminus X$. By the above theorem $X_{\mathscr{I}}$ is a rotation set that contains X. Hence X itself is a rotation set. □

If Y is a rotation set for m_d and if $X \subset Y$ is compact and m_d-invariant, then clearly X is also a rotation set for m_d, with $\rho(X) = \rho(Y)$. We record the following simple lemma for future reference:

Lemma 2.17 *Suppose X, Y are rotation sets for m_d containing more than one point, and assume $X \subset Y$. Then each major gap of X of multiplicity n contains n major gaps of Y counting multiplicities.*

Proof Evidently each major gap of Y is contained in a major gap of X. Let $\{I_i\}$ be the collection of major gaps of X of multiplicities $\{n_i\}$. The number k_i of major gaps of Y contained in I_i satisfies $0 \leq k_i \leq n_i$. Since $\sum k_i = \sum n_i = d-1$ by Theorem 2.7, we must have $k_i = n_i$ for all i. □

Let us call a rotation set ***maximal*** if it is not properly contained in another rotation set. Theorem 2.15 provides a convenient recipe for enlarging every rotation set to a maximal one.

Lemma 2.18 *Every rotation set is contained in a maximal rotation set.*

Proof Suppose X is a rotation set for m_d. For each major gap $(a, a+\ell)$ of X of multiplicity n, consider the n disjoint subintervals $(a + (j-1)/d, a + j/d)$ for $1 \leq j \leq n$. Let $\mathscr{I}$ denote the collection of the $d-1$ disjoint open intervals of length $1/d$ thus obtained. The rotation set $X_{\mathscr{I}}$ of Theorem 2.15 clearly contains X. Moreover, the endpoints of the intervals in $\mathscr{I}$ map to X under m_d, which shows they all belong to $X_{\mathscr{I}}$. Thus, $X_{\mathscr{I}}$ has $d-1$ taut gaps of multiplicity 1. By Theorem 2.9 and Lemma 2.17, $X_{\mathscr{I}}$ is maximal. □

Corollary 2.19 *A rotation set X for m_d is maximal if and only if it has $d-1$ distinct gaps of length $1/d$. In this case $X = X_{\mathscr{I}}$, where $\mathscr{I}$ is the collection of the major gaps of X.*

The proof of Lemma 2.18 in fact gives the following improved lower bound for the number $N_{\max}(X)$ of the maximal rotation sets containing X:

Corollary 2.20 *Suppose X is a rotation set for m_d with loose gaps $I_1, \dots, I_k$ of multiplicities $n_1, \dots, n_k$. Then*

$$N_{\max}(X) \geq \prod_{j=1}^{k} (n_j + 1).$$

In particular, X is contained in at least 2^k maximal rotation sets.

Proof For each loose gap $I = (a, a + \ell)$ of X with multiplicity n, there are $n + 1$ different ways of choosing n disjoint subintervals of length $1/d$ whose endpoints map to $m_d(a)$ or $m_d(a + \ell)$ (the one in the proof of Lemma 2.18 was one of these choices). This leads to $\prod_{j=1}^{k}(n_j + 1)$ different choices for the collection $\mathscr{I}$. □

Example 2.21 The 2-cycle $X = \{\frac{1}{3}, \frac{2}{3}\}$ under doubling is contained in precisely two maximal rotation sets

$$X_{\mathscr{I}_1} = \left\{\frac{1}{3}, \frac{2}{3}\right\} \cup \left\{\frac{1}{3} - \frac{1}{3 \cdot 2^{2n-1}}\right\}_{n \geq 1} \cup \left\{\frac{2}{3} - \frac{1}{3 \cdot 2^{2n}}\right\}_{n \geq 1}$$

and

$$X_{\mathscr{I}_2} = \left\{\frac{1}{3}, \frac{2}{3}\right\} \cup \left\{\frac{1}{3} + \frac{1}{3 \cdot 2^{2n}}\right\}_{n \geq 1} \cup \left\{\frac{2}{3} + \frac{1}{3 \cdot 2^{2n-1}}\right\}_{n \geq 1}$$

corresponding to the collections $\mathscr{I}_1 = \{(\frac{2}{3}, \frac{1}{6})\}$ and $\mathscr{I}_1 = \{(\frac{5}{6}, \frac{1}{3})\}$. Note that each orbit in $X_{\mathscr{I}_i}$ eventually hits the 2-cycle X, and the intersection of $X_{\mathscr{I}_i}$ with the major gap of X is countably infinite.

The above example is a special case of a count for $N_{\max}(X)$ that we will establish in the next section for certain rational rotation sets (see Theorem 2.30). These rotation sets, however, are not typical. In fact, *when $d > 2$ there are rational rotation sets for m_d that are contained in infinitely many maximal rotation sets*. Here is an example:

Example 2.22 Consider the 2-cycle $X = \{\frac{1}{4}, \frac{3}{4}\}$ under tripling. Define the sequences

$$t_n = \sum_{j=0}^{n} \frac{1}{3^{2j+1}} + \frac{1}{3^{2n+1} \cdot 12}$$

$$s_n = m_3(t_n) = \sum_{j=0}^{n} \frac{1}{3^{2j}} + \frac{1}{3^{2n} \cdot 12}$$

for $n \geq 0$. Then $\frac{1}{3} < t_0 < t_1 < t_2 < \cdots$ with $t_n \to \frac{3}{8}$ and $\frac{1}{12} = s_0 < s_1 < s_2 < \cdots$ with $s_n \to \frac{1}{8}$. For each $n \geq 0$ the collection

$$\mathscr{I}_n = \left\{ \left(t_n, t_n + \frac{1}{3}\right), \left(\frac{3}{4}, \frac{1}{12}\right) \right\}$$

produces a rotation set $X_{\mathscr{I}_n}$ which evidently contains the 2-cycle X. The endpoints $\frac{3}{4}, \frac{1}{12}$ map to $\frac{1}{4}$ under m_3, so they both belong to $X_{\mathscr{I}_n}$. The other endpoints $t_n, t_n + \frac{1}{3}$ have the m_3-orbit

$$t_n,\ t_n + \frac{1}{3} \mapsto s_n \mapsto t_{n-1} \mapsto s_{n-1} \mapsto \cdots \mapsto t_0 \mapsto s_0 = \frac{1}{12} \mapsto \frac{1}{4}$$

which, by monotonicity of $\{t_j\}$ and $\{s_j\}$, never meets the pair of intervals in $\mathscr{I}_n$. This shows that both $t_n, t_n + \frac{1}{3}$ belong to $X_{\mathscr{I}_n}$. Thus $X_{\mathscr{I}_n}$ has a pair of major gaps of length $\frac{1}{3}$ and therefore is maximal by Corollary 2.19.

The situation in the irrational case is different and in fact simpler:

Theorem 2.23 *Every irrational rotation set X for m_d is contained in finitely many maximal rotation sets. For any maximal rotation set $Y \supset X$ and any gap I of X, the intersection $Y \cap I$ is finite (possibly empty) and eventually maps into X under the iterations of m_d.*

Proof Take any maximal rotation set $Y \supset X$. First suppose I is a major gap of X of multiplicity n. By Lemma 2.17 and Corollary 2.19, Y has exactly n taut gaps of multiplicity 1 contained in I. We distinguish two cases:

- *Case 1:* I is taut. Then I has the form $(a, a + n/d)$ and

$$Y \cap \bar{I} = \left\{ a, a + \frac{1}{d}, \ldots, a + \frac{n}{d} \right\}.$$

  This condition uniquely determines the major gaps of Y that are contained in I. Notice that the inclusion $m_d(Y \cap \bar{I}) \subset X$ holds.
- *Case 2:* I is loose. Consider the standard monotone map g_Y which is also an extension of $m_d|_X$. By Theorem 2.10, there is an $i > 0$ such that the interior J of $g_Y^{\circ i}(I)$ is a taut gap of X (there can be no periodic loose gap of X since $\rho(X)$ is irrational). Note that $m_d^{\circ i}(Y \cap I) = g_Y^{\circ i}(Y \cap I)$ is contained in $Y \cap \bar{J}$ which is uniquely determined by *Case 1*. Hence the elements of $Y \cap I$ are among the finitely many $m_d^{\circ i}$-preimages of $Y \cap \bar{J}$. This gives finitely many choices for the major gaps of Y in I.

The two cases above show that there are only finitely many choices for the major gaps of Y, hence for Y itself by Theorem 2.9.

We have shown that for any major gap I of X, the intersection $Y \cap I$ is finite and eventually maps into X. Since every minor gap of X maps homeomorphically onto a major gap under some iterate of m_d, the result must also hold when I is minor. □

The number $N_{\max}(X)$ of maximal rotation sets Y containing an irrational rotation set X depends on the structure of the gap graph Γ_X defined in the previous section. Suppose there is a maximal path in Γ_X of the form

$$I_k \to I_{k-1} \to \cdots \to I_1, \quad \text{with} \quad w(I_i) = n_i. \tag{2.2}$$

Since I_1 is taut, the major gaps of Y in I_1 are already determined. However, there are $\binom{n_1+n_2}{n_2}$ choices for the major gaps of Y in I_2. For each of these choices, there are $\binom{n_1+n_2+n_3}{n_3}$ choices for the major gaps of Y in I_3 and so on. This gives the count

$$N_{\max}(X) = \prod \binom{n_1+n_2}{n_2}\binom{n_1+n_2+n_3}{n_3}\cdots\binom{n_1+\cdots+n_k}{n_k} = \prod \frac{(n_1+\cdots+n_k)!}{n_1!\,\cdots\,n_k!}, \tag{2.3}$$

where the product is taken over all maximal paths in Γ_X of the form (2.2) (if there is no path in Γ_X, the product is taken over the empty set and is understood to be 1).

A quick inspection of Fig. 2.3 reveals that $N_{\max}(X) = 1$ for $d = 2$, $N_{\max}(X) \leq 2$ for $d = 3$, and $N_{\max}(X) \leq 6$ for $d = 4$, and $N_{\max}(X) \leq 24$ for $d = 5$. More generally, we have the following

Theorem 2.24 $N_{\max}(X) \leq (d-1)!$ *whenever* X *is an irrational rotation set for* m_d.

Proof If the gap graph Γ_X has no path, then $N_{\max}(X) = 1$ and there is nothing to prove. Otherwise, let Γ_X have $p \geq 1$ distinct maximal paths of the form (2.2), where the weights of the vertices in the i-th path add up to N_i, so $N_1 + \cdots + N_p \leq d-1$. Then, by (2.3),

$$N_{\max}(X) \leq \prod_{i=1}^{p} N_i! \leq \Big(\sum_{i=1}^{p} N_i\Big)! \leq (d-1)!$$

as required. □

2.3 Minimal Rotation Sets

A rotation set is called ***minimal*** if it does not properly contain another rotation set. This section will study the question of existence and uniqueness of minimal rotation sets that are contained in a given rotation set, in both rational and irrational cases.

Before we begin, a quick comment on topological dynamics is in order. The simple proof that minimality is equivalent to having all orbits dense requires a slight modification here, as the closure of an orbit in a rotation set is only *forward* invariant and may not be a rotation set.[4] Similarly, the standard application of Zorn's lemma

[4]In fact, it will follow from the results of this section that for rotation sets minimality is equivalent to having a single dense orbit, a property that is often called ***point transitivity***.

to show that every rotation set contains a minimal rotation set needs some care because the intersection of a linearly ordered family of rotation sets is a priori *forward* invariant. This minor problem is addressed by observing that under m_d, every compact forward invariant set contains a compact invariant set. In fact, if Z is compact and satisfies $m_d(Z) \subset Z$, the nested intersection $K = \bigcap_{n \geq 0} m_d^{\circ n}(Z)$ is easily seen to satisfy $m_d(K) = K$.

Let us first consider the rational case, where minimal rotation sets are cycles. Let $C = \{t_1, \ldots, t_q\}$ be a cycle of rotation number p/q under m_d, where the t_j are in positive cyclic order and their subscripts are taken modulo q (see Sect. 1.3). By Theorem 1.14, $dt_j = t_{j+p} \pmod{\mathbb{Z}}$ for every j. The q gaps $I_j = (t_j, t_{j+1})$ are permuted under any monotone extension g of $m_d|_C$, so $g(\bar{I}_j) = \bar{I}_{j+p}$. Recall that these gaps are either minor or loose: there can be no taut gap.

It follows from Theorem 2.3 that every rotation set X for m_d with $\rho(X) = p/q$ in lowest terms contains at least one q-cycle. But there could be several such minimal sets in X. For instance, under the tripling map m_3, the union $X = C_1 \cup C_2$ of the 3-cycles

$$C_1 : \frac{4}{26} \mapsto \frac{12}{26} \mapsto \frac{10}{26} \quad \text{and} \quad C_2 : \frac{7}{26} \mapsto \frac{21}{26} \mapsto \frac{11}{26}$$

is a rotation set with $\rho(X) = \frac{2}{3}$. This can be seen, for example, from Corollary 2.16 since $\mathbb{T} \smallsetminus X$ contains the intervals $(\frac{12}{26}, \frac{12}{26} + \frac{1}{3})$ and $(\frac{21}{26}, \frac{21}{26} + \frac{1}{3})$ on the circle. The general situation can be understood as follows.

We call a collection $C_1, \ldots, C_N$ of distinct q-cycles under m_d with the same rotation number ***compatible*** if their union $C_1 \cup \cdots \cup C_N$ is a rotation set. We say that $C_1, \ldots, C_N$ are ***superlinked*** if for every pair $i \neq j$, each gap of C_i meets C_j. Geometrically, this means that the points of C_i and C_j alternate as we go around the circle.

Lemma 2.25 *$C_1, \ldots, C_N$ are compatible if and only if they are superlinked.*

In follows in particular that a collection of cycles are compatible if and only if they are pairwise compatible.

Proof First suppose $X = C_1 \cup \cdots \cup C_N$ is a rotation set. Consider the standard monotone map $g = g_X$, which is also a monotone extension of $m_d|_{C_i}$ for each i. Pick any pair C_i, C_j. Since these cycles are distinct, there is a gap I of C_i that meets C_j at some point t. Then for every $k \geq 0$, the interior J_k of $g^{\circ k}(I)$ meets C_j at $g^{\circ k}(t) = m_d^{\circ k}(t)$. Since $J_0 = I, J_1, \ldots, J_{q-1}$ form all the gaps of C_i, we conclude that C_i, C_j are superlinked.

Conversely, suppose $C_1, \ldots, C_N$ are superlinked and consider the standard monotone map $g = g_{C_1}$. Take a gap I of C_1 and let J be the interior of $g(I)$. For $2 \leq i \leq N$, let $C_i \cap I = \{a_i\}$ and $C_i \cap J = \{b_i\}$. Using the fact that the C_i have the same rotation number, it is easy to see that $b_i = m_d(a_i)$. As the C_i are superlinked, the points a_i appear in the same order in I as the points b_i in J, so there is an orientation-preserving homeomorphism $h : I \to J$ such that

$h(a_i) = b_i$ for $2 \le i \le N$. Repeating this process for every gap of C_1 and gluing together the resulting homeomorphisms will then yield an orientation-preserving homeomorphism $h : \mathbb{T} \to \mathbb{T}$ which restricts to m_d on the union $C_1 \cup \cdots \cup C_N$. □

Theorem 2.26 *The number of distinct cycles in a rational rotation set is bounded above by the number of its distinct major gaps.*

In view of Theorem 2.7, we recover the following result of Goldberg as a special case (see [11] for the original combinatorial proof and [2] for an inductive argument reducing the problem down to $d = 2$):

Corollary 2.27 *A rational rotation set for m_d contains at most $d-1$ distinct cycles.*

The upper bound $d - 1$ can always be achieved; see Corollary 3.15.

Proof of Theorem 2.26 Let Y be a rational rotation set for m_d with $\rho(Y) = p/q$ in lowest terms. Suppose $C_1, \ldots, C_N$ are the distinct cycles in Y, all necessarily of length q. The union $X = C_1 \cup \cdots \cup C_N$ is an m_d-invariant subset of Y, so it is a rotation set. By Lemma 2.25, the C_i are superlinked. It follows that any gap I of C_1 contains precisely N gaps $J_1, \ldots, J_N$ of X. Each J_i is periodic of period q and its orbit contains at least one major gap of X. Moreover, the orbits of $J_1, \cdots, J_N$ are disjoint, so they cannot share any major gap of X. It follows that X, hence Y, has at least N distinct major gaps. □

Corollary 2.28 *Every rational rotation set under the doubling map contains a unique cycle.*

Example 2.29 Under the tripling map m_3 there are five 4-cycles of rotation number $\frac{1}{4}$:

$$\begin{aligned}
C_1 &: \frac{1}{80} \mapsto \frac{3}{80} \mapsto \frac{9}{80} \mapsto \frac{27}{80} \\
C_2 &: \frac{2}{80} \mapsto \frac{6}{80} \mapsto \frac{18}{80} \mapsto \frac{54}{80} \\
C_3 = C_3 + \frac{1}{2} &: \frac{5}{80} \mapsto \frac{15}{80} \mapsto \frac{45}{80} \mapsto \frac{55}{80} \\
C_4 = C_2 + \frac{1}{2} &: \frac{14}{80} \mapsto \frac{42}{80} \mapsto \frac{46}{80} \mapsto \frac{58}{80} \\
C_5 = C_1 + \frac{1}{2} &: \frac{41}{80} \mapsto \frac{43}{80} \mapsto \frac{49}{80} \mapsto \frac{67}{80}
\end{aligned}$$

By Corollary 2.27, at most two 4-cycles under tripling can be compatible. By Lemma 2.25, this happens precisely when the two 4-cycles are superlinked. Simple inspection shows that (C_1, C_2), (C_2, C_3), (C_3, C_4) and (C_4, C_5) are the only compatible pairs (compare Fig. 2.4).

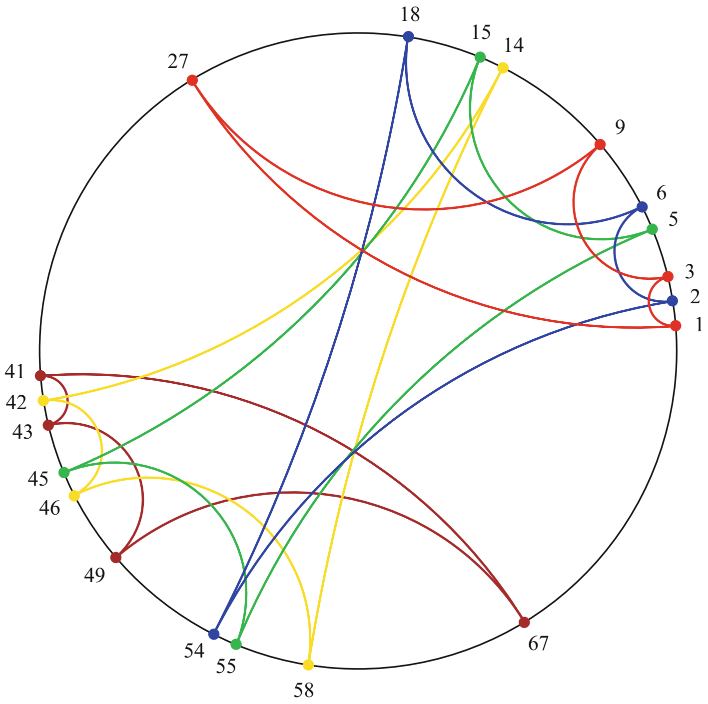

Fig. 2.4 The five 4-cycles of rotation number $\frac{1}{4}$ under tripling, shown in different colors (angles are given in multiples of $\frac{1}{80}$). Only the four superlinked pairs (red, blue), (blue, green), (green, yellow), and (yellow, brown) are compatible cycles

Before moving on to the irrational case, let us use the above ideas to show that for some rational rotation sets the lower bound of Corollary 2.20 is sharp:

Theorem 2.30 *Let X be a rational rotation set for m_d which is the union of $d-1$ distinct cycles. Then $N_{\max}(X) = 2^{d-1}$.*

Proof By Theorem 2.26 X has $d-1$ major gaps, all loose and of multiplicity 1. If $\rho(X) = p/q$, these major gaps have disjoint orbits which are periodic of period q. Let Y be any maximal rotation set containing X. Each major gap $I = (a, b)$ of X contains a single major gap J of Y of length $1/d$. We claim that $J = (a, a + 1/d)$ or $J = (b - 1/d, b)$. Otherwise $J = (t, t + 1/d)$, where $a < t < t + 1/d < b$. The standard monotone map $g = g_Y$ is also a monotone extension of $m_d|_X$, so $g^{\circ q}$ maps I onto itself fixing the endpoints a, b. Moreover, the gaps $g(I), \ldots, g^{\circ q-1}(I)$ of X are all minor, so they cannot contain major gaps of Y; as such, g acts homeomorphically on them. It follows that $g^{\circ q}$ is homeomorphic on $[a, t] \cup [t + 1/d, b]$ and collapses J to the single point $m_d^{\circ q}(t)$. This image point necessarily lies in J since $g^{\circ q} = m_d^{\circ q}$ is expanding on both $[a, t]$ and $[t + 1/d, b]$. This is a contradiction since $m_d^{\circ q}(t) \in Y$.

Thus, there are just two possibilities for each major gap of Y inside a given major gap of X, hence 2^{d-1} possibilities altogether for the major gaps of Y, and therefore for Y itself. This proves $N_{\max}(X) \leq 2^{d-1}$. The result now follows since $N_{\max}(X) \geq 2^{d-1}$ by Corollary 2.20. □

The following corollary immediately follows from the above theorem and its proof:

Corollary 2.31 *Every rotation cycle X under the doubling map is contained in exactly two maximal rotation sets. Moreover, if (a, b) is the major gap of X, then the intervals $(a, a + \frac{1}{2})$ and $(b - \frac{1}{2}, b)$ are the major gaps of these maximal rotation sets.*

Compare Example 2.21.

Example 2.32 Consider the 2-cycle $X = \{\frac{1}{4}, \frac{3}{4}\}$ under tripling. We showed in Example 2.22 that $N_{\max}(X) = \infty$. However, the enlarged rotation set $Y = \{\frac{1}{8}, \frac{1}{4}, \frac{3}{8}, \frac{3}{4}\}$, a union of two 2-cycles under tripling, has $N_{\max}(Y) = 4$ by Theorem 2.30!

We now consider minimal rotation sets in the irrational case.

Theorem 2.33 *Every irrational rotation set X for m_d contains a unique minimal rotation set K. Moreover,*

(i) K is the Cantor attractor of any monotone extension of $m_d|_X$.
(ii) Each gap of K contains at most finitely many points of X, all of which eventually map to K under the iterations of m_d.

Proof Take a monotone extension g of $m_d|_X$ and let K be the Cantor attractor of g, as in Theorem 1.20. Let Z be any non-empty compact m_d-invariant subset of X. By Theorem 1.20, $K = \omega_g(t) \subset Z$ for every $t \in Z$. It follows that K is the unique minimal rotation set contained in X.

To verify the second statement, let Y be any maximal rotation set containing X (whose existence is guaranteed by Lemma 2.18). Since Y contains K, Theorem 2.23 shows that for each gap I of K, the intersection $Y \cap I$ is at most finite and maps into K under the iterations of m_d. Hence the same must be true of $X \cap I$. □

By (the proof of) Theorem 1.20, the gaps of the Cantor attractor of g are the plateaus of the Poincaré semiconjugacy φ between g and r_θ. Thus, we have the following

Corollary 2.34 *Suppose X is a minimal rotation set for m_d with $\rho(X) = \theta$ irrational. Then there exists a degree 1 monotone map $\varphi : \mathbb{T} \to \mathbb{T}$, whose plateaus are precisely the gaps of X, which satisfies $\varphi \circ m_d = r_\theta \circ \varphi$ on X.*

Here is the converse statement. Recall that for each point $s \in \mathbb{T}$, I_s denotes the interior of the fiber $E_s = \varphi^{-1}(s)$.

Theorem 2.35 *Let θ be irrational and $\varphi : \mathbb{T} \to \mathbb{T}$ be a degree 1 monotone map with the property that $I_s \neq \emptyset$ implies $I_{s-\theta} \neq \emptyset$. Denote by X the complement of the* *union of all plateaus of φ. If*

$$\varphi \circ m_d = r_\theta \circ \varphi \quad on \quad X, \tag{2.4}$$

then X is a minimal rotation set for m_d with $\rho(X) = \theta$.

Proof The assumptions imply that φ has plateaus; otherwise $X = \mathbb{T}$ and (2.4) would exhibit a global conjugacy between the degree $d \geq 2$ map m_d and the rotation r_θ, which is impossible.

We invoke Theorem 1.22 to find a degree 1 monotone map $g : \mathbb{T} \to \mathbb{T}$ such that $\varphi \circ g = r_\theta \circ \varphi$ on $\mathbb{T}$. Then X is the Cantor attractor of g. If $t \in X$ is not an endpoint of a plateau and $s = \varphi(t)$, then $E_s = \{t\}$, so by the assumption $E_{s+\theta}$ is a singleton $\{t'\}$. The semiconjugacy relation (2.4) for m_d and the one for g then show that $m_d(t) = t' = g(t)$. Since the set of such t is dense in X, we conclude that $g = m_d$ on X. As the Cantor attractor of g, X is minimal for g and hence for m_d, and $m_d(X) = g(X) = X$. This completes the proof that X is a minimal rotation set. □

We conclude this section with characterizations of minimal rotation sets, as well as those that are both minimal and maximal.

Theorem 2.36 *A rotation set for m_d is a Cantor set if and only if it is minimal and has irrational rotation number.*

Proof The "if" part follows from Theorem 2.33. For the "only if" part, suppose X is a Cantor set. Then $\rho(X)$ is irrational since a rational rotation set is at most countable (Theorem 2.3). Let K be the unique minimal rotation set contained in X. If $K \neq X$, some gap I of K would have to meet X. But then by Theorem 2.33 the intersection $X \cap I$ would be finite, consisting of isolated points of X. This would contradict the assumption that X is a Cantor set. □

Let us call a rotation set ***exact*** if it is both minimal and maximal.[5] Evidently a rational rotation set can never be exact. In the irrational case, the following criterion follows immediately from Corollary 2.19 and Theorem 2.36:

Theorem 2.37 *An irrational rotation set for m_d is exact if and only if it is a Cantor set with $d - 1$ distinct gaps of length $1/d$.*

Corollary 2.38 *Every irrational rotation set under the doubling map is exact.*

Proof Let X be an irrational rotation set under doubling. Then X has a single major gap I of multiplicity 1 which is necessarily taut by Corollary 2.11. If K is the unique minimal rotation set contained in X, then K is a Cantor set with a single taut gap of multiplicity 1 which can only be I. It follows from Theorem 2.9 that $K = X$, and then from Theorem 2.37 that X is exact. □

---

[5]The terminology is meant to suggest that nothing can be added to or removed from such a set without losing the property of being a rotation set.

Remark 2.39 The above corollary is false in higher degrees. For example, there are minimal irrational rotation sets under tripling with a pair of major gaps of lengths $\frac{1}{3}$ and $\frac{4}{9}$ which therefore are not maximal (compare Theorem 4.31). However, every irrational rotation set under tripling is either minimal, or maximal, or both. In every degree > 3, there are irrational rotation sets that are neither minimal nor maximal.

For more on the role of exact rotation sets, see Sect. 4.3.