Orbits, periodic points, and cubic polynomials

Mateusz G. Olechnowicz
University of Waterloo

April, 2013

Abstract

This is a report on the author's explorations in arithmetic dynamics which were motivated by various conjectures regarding the preperiodic point sets of polynomials. One goal was to classify (by sagittal graph isomorphism) a large amount of data relating to cubics; to better understand graph isomorphism, functions on finite sets were studied and a formula to count endofunctions was found. We also present a product identity involving the map $z \mapsto z^2 + c$, an elementary proof of a crucial divisibility property, and a classification of some cubic polynomial pictures.

Contents

1 Introduction **3**

2 Dynamics **4**
2.1 Basics 4
2.2 Pictures 9
2.3 Hooks 10
2.4 Graph isomorphism 13
2.5 Functions on finite sets 14
2.6 Counting endofunctions 16

3 On to polynomials **21**
3.1 A fundamental identity 21
3.2 A curious identity 22
3.3 A note on conjugation 23
3.4 Linear classification 25
3.5 Lagrange interpolation polynomials 28
3.6 Quadratic 2-cycle 30

4 Cubic pictures: narrated equation-solving 31
4.1 Cubic with two fixed points and one preperiodic point 31
4.2 Cubic with a 2-cycle and antipodal preimages 32
4.3 Cubic with a fixed point and a Y tree above it 34
4.4 Cubic with a point of type 1_1 and another of type n_1 34
4.5 Cubics with a point of type 3_1 36
4.6 Other cubics 38

5 Further questions 39

6 Acknowledgements 39

List of Figures

1 The anatomy of a finite orbit 8
2 The sagittal graph of $z \mapsto \frac{-1}{60}(z^3 - 49z - 300)$ restricted to its preperiodic point set over $\mathbf{Q}$ 10
3 Picture for $\frac{1}{30}z^3 - \frac{79}{30}z + 1$ 12
4 Lattice of orbits of $\frac{1}{30}z^3 - \frac{79}{30}z + 1$ 12
5 There are four functions on the set $\{1, 2\}$... 13
6 ...but only three "types" of functions. 13
7 If τ exists, then f and g have the same picture. 14
8 A single fixed point 26
9 The identity map's picture 26
10 The most interesting case: $z \mapsto \beta - z$ 27
11 Can a quadratic have this as its picture? 30
12 A quadratic 2-cycle can never occur alone. 31
13 $1 \to 0 \to 0$, $\varepsilon \to \varepsilon$ 31
14 $\varepsilon \to 0 \to 1 \to 0$, $\delta \to 1$ 32
15 There are only two cubics with this picture. 34
16 A graph not among the data 36
17 Another graph not among the data 38
18 Only one of these was found among the data 38
19 $-6z^3 + \frac{9}{2}z + 1$ and $\frac{3}{2}z^3 - \frac{9}{2}z + 1$ 38
20 $-\frac{289}{16}z^3 + \frac{27}{4}z + 1$ and $-\frac{49}{250}z^3 + \frac{27}{10}z + 1$ 38
21 Is this picture possible? 39
22 What about this one? 39

1 Introduction

Polynomials under iteration exhibit a wealth of interesting and mysterious structures. The orbit of a point under a polynomial map may escape to infinity, wander aimlessly, or become trapped in a cycle. It is this latter case which is most curious, especially when all the concerned quantities are rational. There are many examples of degree d maps with $(d+1)$-cycles; for every β, the map $z \mapsto \beta - z$ has a 2-cycle, and there are infinitely many quadratic maps with a 3-cycle.

Question. For what values of n can we find a quadratic polynomial over $\mathbf{Q}$ with a rational n-cycle?

In [8], Patrick Morton answered that it was impossible for $n = 4$, and in [4], E. V. Flynn, Bjorn Poonen, and Edward Schaefer proved it was impossible for $n = 5$. Recently, Micheal Stoll showed it was impossible for $n = 6$, but the proof is conditional on the Birch and Swinnerton-Dyer conjecture.[14] Assuming that it was impossible for *all* $n > 3$, Poonen proved that quadratic maps have at most 9 rational preperiodic points, and classified all possible pictures of quadratics; there are 12 of them.[12, p. 12]

Less work has been done on cubic maps. In 2007 a team of mathematicians led by Rob Benedetto analyzed over twelve billion cubic polynomials, determined all the preperiodic points for each one, and tabulated the results in [2]. In the accompanying paper [3] they defined a normal form for cubic polynomials, and proved that every cubic can be put into this form via a sequence of conjugations. We reproduce their definition here.

Definition 1.1. Let K be a field, and let f in $K[z]$ be a cubic polynomial. We will say that f is *in normal form* if either

$$f(z) = az^3 + bz + 1$$

or

$$f(z) = az^3 + bz$$

If two cubics of the first form are conjugate to each other, then they are equal, whereas if two cubics of the second form are conjugate to each other, then the linear terms are equal and the quotient of their leading terms is a square.

Their search algorithm tested one cubic per class with coefficients less than 300. It found 781019 cubic polynomials with a nonempty set of preperiodic points (plus over two billion cubics of the second form which only fixed zero).

Each entry in the data includes the cubic $f(z)$ and its preperiodic points $\{x_i\}$. Drawing arrows from x_i to $f(x_i)$ results in a directed graph, which we call the *picture* of f. We grouped the cubics by picture and found 102 distinct structures, a great leap from Poonen's 12 for quadratics. Among them was an example of a cubic with a 5-cycle. These are not all the possible pictures, however; we prove the existence of at least 3 more in Section 4. A table of all the diagrams will soon be available on the author's webpage [11].

In general, graph isomorphism is hard, but picture isomorphism is not. In Section 2 we introduce the language of dynamics and apply it to the characterization and visualization of functions on finite sets. The results of that section were motivated by the need for an algorithm that would, given two finite sets S and T and two maps $f : S \to S$ and $g : T \to T$, determine if they have the same picture.

In Section 3 we introduce the arithmetic dynamics of polynomials, provide an elementary proof of the implication $m \mid n \Rightarrow \phi^m(z) - z \mid \phi^n(z) - z$, prove a curious identity also stumbled upon by Benedetto, show that linear maps admit exactly 4 possible pictures, and introduce Lagrange polynomials as a tool for generating examples. Section 4 we devote to various proofs involving cubics and their pictures, guided by the classified data. We end in Section 5 with some questions whose answers are unknown to us, and suggest other paths of research.

2 Dynamics

2.1 Basics

We begin with a few fundamental definitions, readily available from any standard text, like Silverman's [13]. Throughout the following, we assume S to be a nonempty set.

Definition 2.1. An *endofunction on S* is a map $f : S \to S$ (which sends S to itself). Such a self-map can be *iterated*; the n-fold composition of f with itself,

$$f^n = \underbrace{f \circ \cdots \circ f}_{n \text{ times}},$$

is called the nth iterate of f. By convention, we let f^0 be the identity map.

Example 2.1. The cubic polynomial $f(z) = z^3$ over $\mathbf{Q}$ has a nice closed-form expression for its nth iterate: $f^n(z) = z^{3^n}$. However, it is generally difficult to write down a closed-form expression for the nth iterate of an arbitrary polynomial map; even $g(z) = z^3 + 1$ becomes very complicated after a few iterations.

Definition 2.2. Let a be a point in S. The *orbit* of a under f is the set

$$\begin{aligned}\mathcal{O}_f(a) &= \{f^n(a) : n \geq 0\} \\ &= \{a, f(a), f(f(a)), f(f(f(a))), \ldots\}.\end{aligned}$$

Example 2.2. The orbit of 0 under $g(z) = z^3+1$ is $\mathcal{O}_g(0) = \{0, 1, 2, 9, 730, \ldots\}$.

The principal goal of dynamics is "to classify the points a in the set S according to the behavior of their orbits $\mathcal{O}_f(a)$"[13]. Orbits are useful because they capture everything we need to know about f on S. As a demonstration of their utility, we shall characterize all subsets of S which are not quitted by f.

Definition 2.3. A subset U of S is said to be f*-invariant* if $f(U) \subseteq U$—that is, if $f(u)$ is in U for all u in U.

Proposition 2.1. *Every f-invariant subset of S is of the form*

$$\bigcup_{a \in V} \mathcal{O}_f(a)$$

for some subset V of S.

Proof. First of all, for any $V \subseteq S$, we have

$$\begin{aligned}\bigcup_{a \in V} \mathcal{O}_f(a) &= \bigcup_{a \in V} \{a, f(a), f^2(a), \ldots\} \\ &= \{a_1, f(a_1), f^2(a_1), \ldots\} \\ &\cup \{a_2, f(a_2), f^2(a_2), \ldots\} \\ &\cup \{a_3, f(a_3), f^2(a_3), \ldots\} \\ &\vdots \\ &= V \cup f(V) \cup f^2(V) \cup \cdots\end{aligned}$$

(where the numbering of a few elements of V is not intended to assert any sort of countability of V, but merely serves to clarify the rearrangement of the union). Since function application distributes over union, we have

$$\begin{aligned} f\left(\bigcup_{a\in V}\mathcal{O}_f(a)\right) &= f\left(V\cup f(V)\cup f^2(V)\cup\cdots\right) \\ &= f(V)\cup f^2(V)\cup\cdots \\ &\subseteq V\cup f(V)\cup f^2(V)\cup\cdots \\ &= \bigcup_{a\in V}\mathcal{O}_f(a) \end{aligned}$$

so that any union of orbits of f is f-invariant.

On the other hand, for any f-invariant subset U of S we clearly have

$$U\subseteq\bigcup_{a\in U}\mathcal{O}_f(a)=U\cup f(U)\cup\cdots.$$

Notice that $f(U)\subseteq U$ implies $f^2(U)\subseteq f(U)\subseteq U$, so that, by transitivity, $f^2(U)\subseteq U$. In fact, $f^k(U)\subseteq U$ for all $k\geq 0$ and hence

$$\bigcup_{a\in U}\mathcal{O}_f(a)=U\cup f(U)\cup f^2(U)\cup\cdots\subseteq U. \qquad \blacksquare$$

Sometimes, the orbit of a point is finite. This happens when there are integers $n_1<n_2$ such that $f^{n_1}(a)=f^{n_2}(a)$. When $n_1=0$ we see that successive iterates will periodically jump back to a. We have special names for such points (Northcott called them "exceptional" in [10]).

Definition 2.4. The point a is *periodic* if $f^n(a)=a$ for some $n>0$; the least such n is called the period of a, and a is called *n-periodic* or *n-cyclic*, and f is said to have an n-cycle. If $n=1$, a is said to be fixed, and f is said to have a fixed point. The point a is *preperiodic* if some iterate $f^m(a)$ of a is periodic. We collect these exceptional points as follows: for any f-invariant subset U of S, let

$$\mathrm{Per}(f,U)=\{a\in U: f^n(a)=a \text{ for some } n>0\}$$

and

$$\mathrm{PrePer}(f,U)=\{a\in U: f^n(a)=f^m(a) \text{ for some } n>m\geq 0\}.$$

We often omit U when it equals the domain of f and simply write (Pre)Per(f) for the set of (pre)periodic points of f on S.

Example 2.3. The map $z \mapsto e^{2i\pi/n}z$ over $\mathbf{C}$ fixes 0; the orbit of any nonzero point z_0 is $z_0 \to e^{2i\pi/n}z_0 \to e^{4i\pi/n}z_0 \to \cdots \to e^{2ni\pi/n}z_0 = z_0$. Therefore, 0 is a fixed point and every other point is n-cyclic.

Example 2.4. The point -2 when iterated under the cubic map $z \mapsto -\frac{1}{2}z^3 + \frac{3}{2}z + 1$ behaves as follows: $-2 \to 2 \to 0 \to 1 \to 2 \to \cdots$. Therefore -2 is preperiodic, 0, 1, 2 are periodic, and the cubic has a 3-cycle.

For any f-invariant subset U of S, we have, by definition, that $\mathrm{Per}(f, U)$ is a subset of $\mathrm{PrePer}(f, U)$. A corollary of the following proposition is that if $\mathrm{Per}(f, U)$ is empty then so is $\mathrm{PrePer}(f, U)$.

Proposition 2.2. *Let $f : S \to S$ be a map and let a be a point in S. Then*

$$\bigcap_{n=0}^{\infty} \mathcal{O}_f(f^n(a)) \subseteq \mathrm{Per}(f)$$

and

$$\bigcap_{n=0}^{\infty} \mathcal{O}_f(f^n(a)) \neq \varnothing \iff a \in \mathrm{PrePer}(f).$$

Proof. If

$$x \in \bigcap_{n=0}^{\infty} \mathcal{O}_f(f^n(a))$$

then, in particular,

$$x \in \mathcal{O}_f(a)$$

so that we may write $x = f^m(a)$ for some nonnegative integer m. But because x is in *all* the orbits,

$$x \in \mathcal{O}_f(f^{m+1}(a)) = \{f(f^m(a)), f^2(f^m(a)), f^3(f^m(a)), \dots\}$$

so that $x = f^n(f^m(a)) = f^n(x)$ for some positive integer n. Therefore $x \in \mathrm{Per}(f)$ and $a \in \mathrm{PrePer}(f)$.

On the other hand, if $a \in \mathrm{PrePer}(f)$, then $\mathcal{O}_f(a)$ is finite, so that

$$\mathcal{O}_f(a) \supseteq \mathcal{O}_f(f(a)) \supseteq \mathcal{O}_f(f^2(a)) \supseteq \cdots$$

is a descending chain of finite, nonempty sets, so it must eventually stabilize: $\mathcal{O}_f(f^k(a)) = \mathcal{O}_f(f^{k+1}(a))$. Therefore

$$\bigcap_{n=0}^{\infty} \mathcal{O}_f(f^n(a)) = \mathcal{O}_f(f^k(a)) \neq \varnothing. \qquad \blacksquare$$

If the m and the n in the above proposition are picked to be the least integers of their kind, then we may describe the orbit of a under f with the following picture:

$$\begin{tikzcd}
 & & & & & f^{m+1}(a) \arrow[dr, bend left] & \\
a \arrow[r] & f(a) \arrow[r] & f^2(a) \arrow[r] & \cdots \arrow[r] & f^m(a) \arrow[ur, bend left] & & \cdots \arrow[dl, bend left] \\
 & & & & & f^{m+n-1}(a) \arrow[ul, bend left] &
\end{tikzcd}$$

Figure 1: The anatomy of a finite orbit

Hence we arrive at another piece of useful terminology.

Definition 2.5. Given a map f on a set S, a point a in S is said to be of *type* n_m if it enters an n-cycle after m applications of f. Periodic points are of type n_0.

It is not hard to see that the orbit of a point of type n_m has cardinality $n + m$.

Next we prove a small lemma which reveals when two points are periodic, and which will be useful when dealing with hooks.

Lemma 2.3. *If $\mathcal{O}_f(a) = \mathcal{O}_f(b)$ then either $a = b$ or a and b are both periodic.*

Proof. If $a \neq b$, then since $a \in \mathcal{O}_f(b)$, there exists some $n \geq 1$ such that $a = f^n(b)$; similarly, since $b \in \mathcal{O}_f(a)$, $b = f^m(a)$ for some $m \geq 1$. Then f^{n+m} fixes both a and b: $f^n(f^m(a)) = f^n(b) = a$ and $f^m(f^n(b)) = f^m(a) = b$. ∎

The following few facts are given as exercises in [13].

1. If S is finite, then so is every orbit. But points with finite orbits are preperiodic. Hence, $\mathrm{PrePer}(f) = S$.

2. If f is injective, then $f^n(a) = f^m(a) \Rightarrow f^{n-1}(a) = f^{m-1}(a) \Rightarrow \cdots \Rightarrow f^{n-m}(a) = f^{m-m}(a) = a$ so that $\mathrm{PrePer}(f) = \mathrm{Per}(f)$.

3. It follows from (1) and (2) that if S is finite and f is injective, surjective, or bijective, then $\mathrm{Per}(f) = \mathrm{PrePer}(f) = S$.

4. Conversely, if $\mathrm{Per}(f) = S$ then f is bijective; for every x in S there exists some positive integer n_x such that $f^{n_x}(x) = x$; a commuting inverse g for f can be constructed by letting $g(x) = f^{n_x - 1}(x)$ for each x in S.

Example 2.5. Consider the map $q(a,b) = (a+1, ab)$ on $\mathbf{Z} \times \mathbf{Z}$. Each application of q increments the abscissa, so that $\mathrm{Per}(q)$ (and hence $\mathrm{PrePer}(q)$) is empty. Iterating q reveals a pattern:

$$(a,b) \mapsto (a+1, ab) \mapsto (a+2, (a+1)ab) \mapsto (a+3, (a+2)(a+1)ab) \mapsto \cdots$$

so that $q^n(a,b) = (a+n, (a)_n b)$ (where $(x)_n = x(x+1)(x+2)\cdots(x+n-1)$ is the Pochhammer symbol).

This formula tells us that starting with $a < 0$ and any b value yields an orbit which inexorably advances toward the y-axis while the ordinate jumps above, jumps below the x-axis with each step (unless it is stuck to it). Having arrived at the y-axis—we are at $(0, b')$, say—we get sent to $(1, 0)$, and from then on we may never leave the x-axis.

Starting with $a > 0$ and $b \neq 0$ tells a different story. The ordinate, having seen its friends sucked into the x-axis, never once crosses it, and instead goes off to $\pm\infty$ faster and faster.

2.2 Pictures

Every function f on a set S has a special type of directed graph associated to it, called its *sagittal graph* [5, p. 174] or its *functional digraph*. It is defined as follows: the vertices are the points in S, and the arrows are drawn from x to $f(x)$.

Conversely, every directed graph (V, E) with the property that each vertex has out-degree 1 corresponds to a function $f : V \to V$ as follows: for (x, y) in E, define $f(x) = y$.

The sagittal graphs of functions on infinite sets can be quite complicated. For a polynomial p of degree $d > 1$ defined over $\mathbf{C}$ every point x has up to d pre-images, so that the sagittal graph of p is an uncountable directed graph where every vertex has positive in-degree at most d.

However, there is a well known result due to Northcott[10] which implies that a polynomial over $\mathbf{Q}$ of degree at least two has finitely many rational preperiodic points. Therefore, we can restrict our attention (and our polynomials) to preperiodic point sets.

Definition 2.6. Let $f : \mathbf{Q} \to \mathbf{Q}$ be a polynomial map of degree at least two. We shall call the sagittal graph of $f|_{\mathrm{PrePer}(f,\mathbf{Q})}$ the *picture* of f.

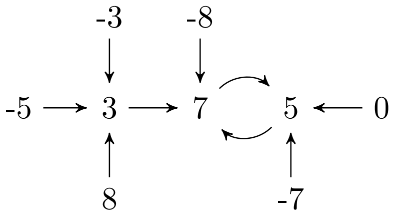

Figure 2: The sagittal graph of $z \mapsto \frac{-1}{60}(z^3 - 49z - 300)$ restricted to its preperiodic point set over $\mathbf{Q}$

We borrow a term from graph theory and say that the *sources* of f are the points of in-degree 0; they are precisely the elements of $S \setminus f(S)$. In the above figure, the sources are -8, -7, -5, -3, 0, and 8. Sources are closely related to the main topic of the next section.

2.3 Hooks

The domain of an endofunction f is always f-invariant, so by Proposition 2.1 it can be written as a union of orbits. The set V from that proposition need not be all of S, and in fact can be much smaller.

Definition 2.7. We shall call a subset H of S a *hook-set* for f if

$$\bigcup_{h \in H} \mathcal{O}_f(h) = S$$

and if for all k in H,

$$\bigcup_{h \in H \setminus \{k\}} \mathcal{O}_f(h) \neq S.$$

The elements of a hook-set are called *hooks.*

The motivation for hooks came from the realization that with a polynomial expression for f, and relatively few points preperiodic points $h_1, \ldots, h_d$, one could 'discover' the entirety of $\mathrm{PrePer}(f)$ by computing forward orbits of the h_i under f. The second requirement ensures that such a description of the full picture does not have any redundancies—every h_i contributes some new information.

We will now see that hooks map to maximal orbits in the partially ordered set $\{\mathcal{O}_f(a) : a \in S\}$.

Lemma 2.4. *Let H be a hook-set for $f : S \to S$. If $a \in H$ then $\mathcal{O}(a)$ is maximal.*

Proof. Suppose $\mathcal{O}(b) \supseteq \mathcal{O}(a)$. Then

$$b \in S = \bigcup_{h \in H} \mathcal{O}(h)$$

so $b \in \mathcal{O}(k)$ for some hook k. Now if $k \neq a$ we have $\mathcal{O}(k) \supseteq \mathcal{O}(b) \supseteq \mathcal{O}(a)$ and

$$S = \bigcup_{h \in H} \mathcal{O}(h) = \bigcup_{h \in H \setminus \{a\}} \mathcal{O}(h) \neq S,$$

a contradiction. So $k = a$ and $\mathcal{O}(b) = \mathcal{O}(a)$. ∎

For finite S, hooks always exist, but for infinite S anything can happen.

Example 2.6. Consider the map $s(n) = n + 1$ over the integers. Then $\mathcal{O}_s(n) \subsetneq \mathcal{O}_s(n-1)$, so s has no hooks. Looking at its restriction to the naturals, we see that $\mathbf{N} = \mathcal{O}_s(0)$.

If hooks map to maximal orbits, what maps to minimal orbits?

Lemma 2.5. *The point a is periodic if and only if $\mathcal{O}_f(a)$ is minimal.*

Proof. Let a be a periodic point of order n for f and suppose $\mathcal{O}_f(a) \supseteq \mathcal{O}_f(b)$ for some b in S. Then $b = f^m(a)$ for some $m \geq 0$; pick k large enough that $kn > m$. Then $f^{kn-m}(b) = f^{kn-m}(f^m(a)) = f^{kn}(a) = f^{kn}(a) = a$, so that a and hence $\mathcal{O}_f(a)$ are contained in $\mathcal{O}_f(b)$.

Conversely, if a is not periodic, then no forward iterate of a returns to a. Therefore $a \notin \mathcal{O}_f(f(a))$ and $\mathcal{O}_f(a) \supsetneq \mathcal{O}_f(f(a))$. ∎

It turns out that hook-sets are quite simple to characterize.

Theorem 2.6. *Let $f : S \to S$ be a map and let $H \subseteq S$ be a hook-set for it. Then H comprises the sources of f and one periodic point per isolated cycle.*

Proof. First we prove that $S \setminus f(S) \subseteq H$. Let v be in $S \setminus f(S)$. Then $v \neq f(u)$ for any u in S. But since the union of the orbits of the elements of H is all of S, there must be some hook k such that

$$v \in \mathcal{O}_f(k) = \{k\} \cup \{f(f^n(k)) : n \geq 0\};$$

the fact that v is a source forbids its membership in the second set. Thus $v = k$ is in H.

Now we demonstrate that H need not comprise only sources. Let h be in H and suppose $f(a) = h$ (where a need not be a hook). Then $\mathcal{O}(a) \supseteq \mathcal{O}(h)$, but since the latter is maximal, we have $\mathcal{O}(a) = \mathcal{O}(h)$. Hence either $a = h$, so that h is a fixed point, or $a, h \in \mathrm{Per}(f)$, in which case h is periodic.

This shows that if anything points to $\mathcal{O}_f(h)$ then it's actually an element of $\mathcal{O}_f(h)$. Therefore, $\mathcal{O}_f(h)$ is isolated. $\blacksquare$

Example 2.7. Consider the set of rational preperiodic points for the cubic polynomial $f(z) = \frac{1}{30}z^3 - \frac{79}{30}z + 1$. Its picture looks like this:

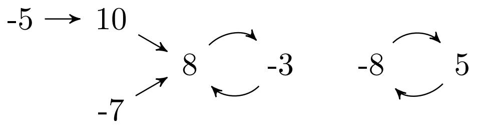

Figure 3: Picture for $\frac{1}{30}z^3 - \frac{79}{30}z + 1$

The orbits form a lattice, whose maximal elements are readily seen to be generated by -5, -7, and either 5 or -8. In this case, there are two possible hook-sets: $\{-7, -5, 5\}$ and $\{-7, -5, -8\}$.

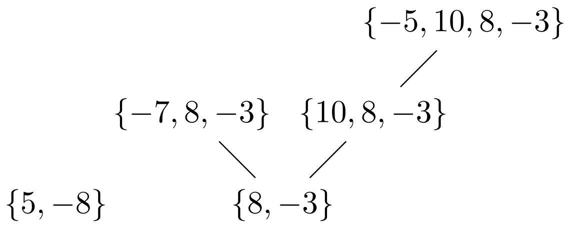

Figure 4: Lattice of orbits of $\frac{1}{30}z^3 - \frac{79}{30}z + 1$

Example 2.8. Consider the map $r(a, b) = (a - b, a + b)$ on $\mathbf{Z} \times \mathbf{Z}$. Writing it as

$$r(a,b) = \sqrt{2}\begin{pmatrix} \cos(\frac{\pi}{4}) & -\sin(\frac{\pi}{4}) \\ \sin(\frac{\pi}{4}) & \cos(\frac{\pi}{4}) \end{pmatrix}\begin{pmatrix} a \\ b \end{pmatrix}$$

we see that it is injective and fixes the origin. That it fixes nothing else is clear from the inequality $||r(a, b)||^2 = (a - b)^2 + (a + b)^2 = 2(a^2 + b^2) > ||(a, b)||^2$. Points are sent spiralling outward, performing an eighth of a rotation with each iteration. What are the hooks of r?

The previous theorem comes in handy. If some point (x, y) is mapped to by r, then

$$a - b = x$$
$$a + b = y.$$

Adding the two equations gives $x + y = 2a$; subtracting gives $x - y = -2b$. Since $x + y \equiv x - y \pmod{2}$ (i.e. they have the same parity), the point (x, y) has a preimage if and only if $2 \mid x + y$. Therefore the hook-set of r is $\{(0, 0)\} \cup \{(x, y) \in \mathbf{Z} \times \mathbf{Z} : 2 \nmid x + y\}$.

The compression afforded by storing only hooks and f is, in general, ineffective. On a set of size n, a constant map has $n - 1$ hooks, and the identity map has n. On the other hand, an n-cycle has 1 hook. In practise, because are dealing with polynomial f, the in-degree of a given vertex is bounded by the degree of f.

2.4 Graph isomorphism

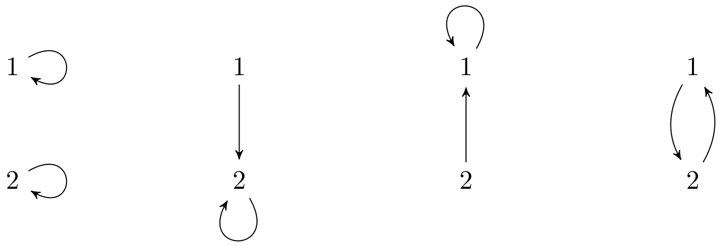

Figure 5: There are four functions on the set $\{1, 2\}$...

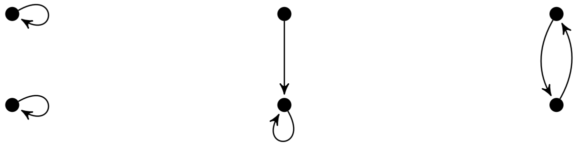

Figure 6: ...but only three "types" of functions.

Two directed graphs are said to share an isomorphism when there exists an arrow-preserving bijection between their vertex sets. Thus, two maps $f : S \to S$ and $g : T \to T$ have the same picture if and only if there exists an isomorphism $\tau : S \to T$ between their sagittal graphs.

Let's look closely at what such a map τ is really doing. In S, we have $x \to f(x)$. Therefore, in T, we must have $\tau(x) \to \tau(f(x))$. But in T, we also have $\tau(x) \to g(\tau(x))$. Since the arrow out of $\tau(x)$ only points to one thing, we see that $\tau(f(x))$ must equal $g(\tau(x))$.

This equality precisely captures what it means for two functions to have the same picture. Hence our algorithm, given $f : S \to S$ and $g : T \to T$ as input, should assert or disprove the existence of a bijection $\tau : S \to T$ which satisfies the functional equation $\tau \circ f = g \circ \tau$.

$$\begin{tikzcd}
S \arrow[r, "f"] \arrow[d, "\tau"'] & S \arrow[d, "\tau"] \\
T \arrow[r, "g"'] & T
\end{tikzcd}$$

Figure 7: If τ exists, then f and g have the same picture.

From the commutative diagram it is clear that if two functions f and g, defined on a field K, are *linearly conjugate* (that is, $g = f^{\varphi} = \varphi \circ f \circ \varphi^{-1}$ for some $\varphi = \alpha z + \beta$) then they have the same picture. However, there is no requirement that τ be a linear homeomorphism.

We remind the reader that the restriction of a cubic polynomial f to $\operatorname{PrePer}(f, \mathbf{Q})$ is quite clearly a function on a finite set. The vertex-set of the sagittal graph was defined to be the domain of $f|_{\operatorname{PrePer}(f,\mathbf{Q})}$, which is $\operatorname{PrePer}(f, \mathbf{Q})$.

2.5 Functions on finite sets

A good algorithm is designed to take advantage of all the available structure of its input, so we intend now to fully describe functions on finite sets.

Given $f : S \to S$ we may partition S into components which are connected with respect to f. Write $x \sim y$ if and only if $\mathcal{O}_f(x) \cap \mathcal{O}_f(y) \neq \varnothing$. Reflexivity and symmetry are clear; to see that this relation is transitive, suppose $x \sim y$ and $y \sim z$. Then $f^n(y) \in \mathcal{O}_f(x)$ and $f^m(y) \in \mathcal{O}_f(z)$ for some $n, m \geq 0$, so

$$\begin{aligned} f^{m+n}(y) &= f^m(f^n(y)) \\ &= f^n(f^m(y)) \in \mathcal{O}_f(x) \cap \mathcal{O}_f(z) \end{aligned}$$

In some sense orbits are like rivers: once x and y meet, they do not flow apart, so if y also meets z, then x and z will eventually meet as well.

Definition 2.8. The equivalence classes of $\sim$ on S are called the *(connected) components* of f. When there is only one class—when no two orbits are disjoint—then f is said to be *connected.*

It is not hard to see that if two functions have the same picture, then they have the same number of components, and each component of one function is isomorphic to a component of the other. Conversely, if $C_1, \ldots, C_k$ are the components of f, and $D_1, \ldots, D_k$ are the components of g, and $\sigma : \{1, \ldots, k\} \to \{1, \ldots, k\}$ is a bijection such that C_i is isomorphic to $D_{\sigma(i)}$ for all $1 \leq i \leq k$, then in fact f and g are isomorphic.

What does a connected function f on a finite set S look like? All the orbits intersect nontrivially, so let $C = \bigcap_{a \in S} \mathcal{O}(a)$. Now, because S is finite, C is nonempty, and because the set $\{\mathcal{O}(a) : a \in S\}$ is a meet-semilattice under set intersection, C is equal to some orbit $\mathcal{O}(c)$. In fact, it is the unique minimal orbit (for it is contained in every other orbit), so it is equal to $\mathcal{O}(a)$ for every a in $\mathrm{Per}(f)$ (since periodic points have minimal orbits). Finally, since the set of f's periodic points is f-invariant,

$$\mathrm{Per}(f) = \bigcup_{a \in \mathrm{Per}(f)} \mathcal{O}(a) = C.$$

Definition 2.9. Let f be a connected function on a set S. We call $\mathrm{Per}(f)$ the *central cycle* of f on S, and we call its cardinality the *cycle number* of f.

The cycle number can be used to discern graphs of small size, because while there are, say, 9 connected graphs of size 4, only 3 have cycle number 2.

For a connected function $f : S \to S$ with cycle number N, every point in S is of type N_m, so we can define $m = h(x)$ to equal the *height* of x—a natural name, for $h(x)$ is distance (measured in applications of f) between x and the central cycle.

What happens around the central cycle? If we define $a \approx b$ if and only if $f^{h(a)}(a) = f^{h(b)}(b)$, then the equivalence classes of $\approx$ (which are just elements of the coimage of the map $S \ni x \overset{\Delta}{\longmapsto} f^{h(x)}(x) \in \mathrm{Per}(f)$) partition S into regions of points, all of which enter the central cycle at the same spot. If orbits are rivers, then these regions are drainage basins, and within these watersheds, the intersection $\mathcal{O}(a) \cap \mathcal{O}(b)$ equals $\mathcal{O}(c)$ for some unique c, which we affectionately call the *conflux* of a and b.

Unfortunately, this train of thought does not hold water. The correct approach is to appropriately 'invert' f and realize that the points on the central cycle are the roots of trees.

Definition 2.10. Let $f : S \to S$ be connected. Then the *arborification* of f is the map $\mathrm{Arb}(f) : S \to P(S)$ defined by $\mathrm{Arb}(f)(x) = (f^{-1}(x)) \setminus \mathrm{Per}(f)$.

What does this map do? It regards every point on S as a tree node whose child nodes are exactly its non-periodic preimages. The childless nodes (those for which $\mathrm{Arb}(f)$ returns $\varnothing$) are precisely the sources of f and the periodic points with in-degree 1.

Tree isomorphism is very simple: two childless trees are isomorphic, and two trees T_1, T_2 with subtrees $S_{1,1}, \ldots, S_{1,j}$ and $S_{2,1}, \ldots, S_{2,j}$ are isomorphic if and only if there exists a bijection $\sigma : \{1, \ldots, j\} \to \{1 \ldots, j\}$ such that $S_{1,i}$ is tree-isomorphic to $S_{2,\sigma(i)}$ for all $1 \leq i \leq j$.

We conclude that two connected functions f and g have the same picture if and only if there is a bijection $\kappa : \mathrm{Per}(f) \to \mathrm{Per}(g)$ such that $\kappa(f(c)) = g(\kappa(c))$ and the tree above c is isomorphic to the tree above $\kappa(c)$ for all c in $\mathrm{Per}(f)$.

With this complete understanding of functions on finite sets, it is not so hard to implement an algorithm for determining whether two maps f and g have the same picture.

2.6 Counting endofunctions

To learn more about functions on finite sets we seek to count how many possible pictures there are of size n. Of course, there are n^n functions on an n-element set, but, as we have already seen, many of them are indistinguishable when viewed as unlabelled directed graphs.

Let us define $\gamma : \mathbf{N} \to \mathbf{N}$ to be the number of graphs of size n. Immediately we have the crude upper bound $\gamma(n) \leq n^n$, and the following proposition gives us a lower one.

Proposition 2.7. *For all $n \in \mathbf{N}$, $\gamma(n) \geq n$.*

Proof. Fix n and for each $1 \leq m \leq n$ define

$$f_m(k) = \begin{cases} k+1 & \text{if } k < n \\ m & \text{if } k = n \end{cases}$$

Then 1 is a point of type $(n - m + 1)_{m-1}$ for f_m, and thus f_{m_1} and f_{m_2} cannot have isomorphic sagittal graphs for $m_1 \neq m_2$ (e.g. f_1 is a bijection and $f_{m>1}$ is not). Hence there are at least n pictures of size n, and a concrete example of each is given by f_m. ∎

Let us now employ the theory of group actions to find an exact formula for $\gamma(n)$. Let $X = S^S$, the set of all functions from S to itself, and let S_n (the symmetric group on n letters) act on X by

$$\sigma \cdot f = \sigma \circ f \circ \sigma^{-1}$$

so that the orbits $[f] = \{g \in X : \sigma \circ f = g \circ \sigma\}$ correspond exactly to pictures of size n. Therefore, to count to number of possible pictures of size n, we can use the Cauchy-Frobenius lemma from group theory. The number of orbits is given by

$$\gamma(n) = \frac{1}{n!} \sum_{\sigma \in S_n} |X^\sigma|$$

where $|X^\sigma| = \{f \in X : \sigma \cdot f = f\}$ is the set of points in X fixed by σ.

Now for a simple lemma which greatly reduces the number of terms in the summation.

Lemma 2.8. *If σ is conjugate to τ (in S_n) then $|X^\sigma| = |X^\tau|$.*

Proof. Suppose $\pi\sigma\pi^{-1} = \tau$ and let f be fixed by σ. Then $\pi f \pi^{-1}$ is fixed by τ, for $\tau\pi f \pi^{-1}\tau = \pi\sigma\pi^{-1}\pi f \pi^{-1}\pi\sigma\pi^{-1} = \pi\sigma f \sigma\pi^{-1} = \pi f \pi^{-1}$. Hence $|X^\sigma| \leq |X^\tau|$; by symmetry, we have equality. ∎

The conjugacy classes of S_n are well understood: there are as many of them as there are partitions of n, and, if $\alpha = (\alpha_1, ..., \alpha_n)$ is one such partition (where $\alpha_k \geq 0$ and $\sum_{k=1}^n k\,\alpha_k = n$) the number of permutations in S_n whose cycle structures are of the form

$$\underbrace{(\cdot)\ldots(\cdot)}_{\alpha_1}\underbrace{(\cdot\,\cdot)\ldots(\cdot\,\cdot)}_{\alpha_2}\cdots$$

is given by

$$\frac{n!}{\prod_{k=1}^n k^{\alpha_k}\alpha_k!}.$$

Hence, if $\alpha(\sigma)$ denotes the cycle structure of a permutation σ, the number of orbits is given by

$$\gamma(n) = \frac{1}{n!} \sum_{\text{one } \sigma \text{ per class}} \frac{n!}{\prod_{k=1}^n k^{\alpha(\sigma)_k}\alpha(\sigma)_k!}|X^\sigma|.$$

Can we find an expression for $|X^\sigma|$ in terms of the cycle structure of σ? Let's look at two examples.

Example 2.9. Suppose $n = 7$ and $\sigma = (1\,2\,3\,4)(6\,7)$. How many possible f can σ fix? The relation $f = \sigma \cdot f = \sigma \circ f \circ \sigma^{-1}$ yields a system of seven equations:

$$\begin{aligned} f(1) &= \sigma(f(4)) \\ f(2) &= \sigma(f(1)) \\ f(3) &= \sigma(f(2)) \\ f(4) &= \sigma(f(3)) \\ f(5) &= \sigma(f(5)) \\ f(6) &= \sigma(f(7)) \\ f(7) &= \sigma(f(6)) \end{aligned}$$

First of all, notice that $f(5)$ is fixed by σ, and since the only fixed point of σ is 5, we must have $f(5) = 5$. Also, $f(7)$ is determined by $f(6)$, and each of $f(2), f(3), f(4)$ is determined by $f(1)$. Where can f send 6? It can be fixed (in which case $f(7) = \sigma(6) = 7$), sent to 7 (in which case $f(7) = \sigma(7) = 6$), or sent to 5 (in which case $f(7) = \sigma(5) = 5$). Note that the equations imply $f(6) = \sigma(f(7)) = \sigma(\sigma(f(6))) = \sigma^2(f(6))$, so that $f(6)$ cannot equal any of 1, 2, 3, 4, because none of those is fixed by σ^2. Finally, we consider $f(1)$. The only restriction on where f sends 1 is that $f(1)$ be fixed by σ^4, but since $4 = o(\sigma)$, $f(1)$ can be anything it wants to be. Thus there are $3 \cdot 7 = 21$ functions fixed by $(1\,2\,3\,4)(6\,7)$ in S_7.

Example 2.10. Suppose $n = 5$ and $\sigma = (1\,2\,3)(4\,5)$. To count how many possible functions σ can fix, let's look at the five equations

$$\begin{aligned} f(1) &= \sigma(f(3)) \\ f(2) &= \sigma(f(1)) \\ f(3) &= \sigma(f(2)) \\ f(4) &= \sigma(f(5)) \\ f(5) &= \sigma(f(4)) \end{aligned}$$

$f(2)$ and $f(3)$ are determined by $f(1)$, and $f(5)$ is determined by $f(4)$. Where can f send 1? Any of 1, 2, 3 will work, but $f(1) = 4$ yields a contradiction when we realize that $\sigma^3(f(1)) = f(1)$ but $\sigma^3(4) = 5 \neq 4$. Similarly, $f(4)$ can only be 4 or 5. Thus there are $3 \cdot 2 = 6$ functions fixed by $(1\,2\,3)(4\,5)$ in S_5.

With these observations noted we are ready to prove two lemmata.

Lemma 2.9 (Division lemma)**.** *Let f be fixed by σ. If i is a period-k point for σ, then $f(i)$ is a period-d point for σ for some d dividing k.*

Proof. We can expand k-fold the equation $f = \sigma \cdot f$ to get $f = \sigma \circ f \circ \sigma^{-1} = \sigma^2 \circ f \circ \sigma^{-2} = \ldots = \sigma^k \circ f \circ \sigma^{-k}$. Using the fact that $\sigma^k(i) = i$ we can evaluate the expansion at i to get $f(i) = \sigma^k(f(\sigma^{-k}(i))) = \sigma^k(f(i))$, which implies that $f(i)$ is a period-d point for σ for some d dividing k. ∎

Lemma 2.10. *If $\alpha = (\alpha_1, \ldots, \alpha_n)$ is the cycle structure of some permutation σ, then $|X^\sigma| = \prod_{k=1}^{n} (\sum_{d|k} d\,\alpha_d)^{\alpha_k}$.*

Proof. For each $1 \leq k \leq n$, σ has α_k k-cycles. A given k-cycle can be written as $(l_1\; l_2\; l_3 \ldots l_k)$ for some letters $1 \leq l_i \leq n$. Thus $f = \sigma \cdot f$ implies

$$f(l_i) = \sigma(f(\sigma^{-1}(l_i))) = \begin{cases} \sigma(f(l_k)) & \text{if } i = 1 \\ \sigma(f(l_{i-1})) & \text{if } i > 1 \end{cases}$$

so that whither f sends l_2 through l_k is determined by whither f sends l_1.

Now, l_1 is a period-k point for σ. By the division lemma, $f(l_1)$ must be a period-d point for σ for some d dividing k. But for any d dividing k, σ has α_d d-cycles, each with d points. Hence σ has $d\,\alpha_d$ period-d points. Thus σ provides f with $\sum_{d|k} d\,\alpha_d$ possible targets for l_1.

Whither f sends one cycle of σ does not depend on whither f sends another (cycle of σ), so we can count independently:

$$\begin{aligned} |X^\sigma| &= \prod_{\text{cycles of } \sigma} \left(\sum_{d\,|\,\text{cycle length}} d\,\alpha_d \right) \\ &= \prod_{k=1}^{n} \left(\prod_{k\text{-cycles of } \sigma} \Big(\sum_{d|k} d\,\alpha_d \Big) \right) \\ &= \prod_{k=1}^{n} \Big(\sum_{d|k} d\,\alpha_d \Big)^{\alpha_k} \end{aligned}$$ ∎

We summarize this chain of deductions in the following

Theorem 2.11. *The number of endofunctions (up to sagittal graph isomorphism) on a set of size n is given by*

$$\gamma(n) = \sum_{\alpha} \prod_{k=1}^{n} \frac{(\sum_{d|k} d\,\alpha_d)^{\alpha_k}}{k^{\alpha_k} \alpha_k!}$$

where the sum is taken over all partitions α of n.

Since $\alpha_k = 0$ implies $(\sum_{d|k} d\,\alpha_d)^{\alpha_k} = 1$, and since $\sum_{k=1}^{n} k\,\alpha_k = n$ implies not all α_k can be zero, the product may be taken over all $\alpha_k \neq 0$.

Corollary. *For all $n \in \mathbf{N}$, $\gamma(n) \geq \frac{n^n}{n!}$.*

Proof. This improvement on the previous lower bound can be derived by noting that $\alpha = (n, 0, \ldots, 0)$ is a partition of n for every $n \geq 1$. Thus

$$\begin{aligned}
\gamma(n) &\geq \prod_{\alpha_k \neq 0} \frac{(\sum_{d|k} d\,\alpha_d)^{\alpha_k}}{k^{\alpha_k} \alpha_k!} \\
&= \frac{(\sum_{d|1} d\,\alpha_d)^{\alpha_1}}{1^{\alpha_1}(\alpha_1)!} \\
&= \frac{(1 \cdot n)^n}{1^n \cdot n!} \\
&= \frac{n^n}{n!}
\end{aligned}$$

∎

Example 2.11. If $n = 4$ then the partitions of n are given by

$$\begin{aligned}
4 &= 3 + 1 \\
&= 2 + 2 \\
&= 2 + 1 + 1 \\
&= 1 + 1 + 1 + 1
\end{aligned}$$

which correspond to 4-tuples (0, 0, 0, 1), (1, 0, 1, 0), (0, 2, 0, 0), (2, 1, 0, 0), and (4, 0, 0, 0). Thus

$$\begin{aligned}
\gamma(4) = &\frac{(1 \cdot 0 + 2 \cdot 0 + 4 \cdot 1)^1}{4^1 \cdot 1!} \\
&+ \frac{(1 \cdot 1)^1}{1^1 \cdot 1!} \cdot \frac{(1 \cdot 1 + 3 \cdot 1)^1}{3^1 \cdot 1!} \\
&+ \frac{(1 \cdot 0 + 2 \cdot 2)^2}{2^2 \cdot 2!} \\
&+ \frac{(1 \cdot 2)^2}{1^2 \cdot 2!} \cdot \frac{(1 \cdot 2 + 2 \cdot 1)^1}{2^1 \cdot 1!} \\
&+ \frac{(1 \cdot 4)^4}{1^4 \cdot 4!} \\
= &\, 1 + \frac{4}{3} + 2 + 4 + \frac{32}{3} \\
= &\, 19
\end{aligned}$$

The first few values of γ are $\gamma(1) = 1$, 3, 7, 19, 47, 130; this is sequence A001372 in [7].

3 On to polynomials

We turn our attention now to polynomial maps over fields.

3.1 A fundamental identity

The main goal of this section is to prove the following

Theorem 3.1. *Let ϕ be a polynomial over a commutative unital ring R. Then for all nonnegative integers m and n,*

$$m \mid n \Rightarrow (\phi^m(z) - z) \mid (\phi^n(z) - z).$$

This result is often used in the literature (and is crucial in the definition of dynatomic polynomials, which are the dynamical analogues of cyclotomic polynomials; see [13, p. 148]) where it is accompanied by a parenthetical proof sketch which advises readers to consider roots on both sides of the bar. There exists an elementary, ring-theoretic proof of this fact, and we present it here.

First, we need to establish a few useful lemmata.

Lemma 3.2. *Let R be a commutative unital ring and let $n \geq 0$ be an integer. Then $a - b \mid a^n - b^n$.*

Proof. The result follows by induction: certainly we have $a - b \mid 0$, and $a^{n+1} - b^{n+1} = a(a^n - b^n) + ab^n - b^{n+1} = a(a^n - b^n) + (a - b)b^n$. ∎

Lemma 3.3. *Let R be a commutative unital ring. For a, b, f in $R[x]$, we have that $a - b \mid f(a) - f(b)$.*

Proof. Letting $c_0, \dots, c_n$ be the coefficients of f, we observe that

$$\begin{aligned} f(a) - f(b) &= c_n a^n + c_{n-1} a^{n-1} + \cdots + c_1 a + c_0 \\ &\quad - c_n b^n - c_{n-1} b^{n-1} - \cdots - c_1 b - c_0 \\ &= c_n(a^n - b^n) + c_{n-1}(a^{n-1} - b^{n-1}) + \cdots + c_1(a - b) \end{aligned}$$

which, by the previous lemma ($R[x]$ is certainly commutative and unital if R is), is clearly divisible by $a - b$. ∎

As a corollary we have a proof of the factor theorem which does not require R to be a field nor $R[x]$ to be a Euclidean domain: if $f(\lambda) = 0$ then $x - \lambda \mid f(x) - f(\lambda) = f(x)$.

Now we are ready to prove the main theorem of this section.

Proof. Let ϕ in $R[x]$ be any polynomial. We may evaluate ϕ at ϕ and receive another polynomial in $R[x]$; repeating this process gives us the forward iterates of ϕ. Repeatedly applying the preceding lemma with $f = \phi$, we see that

$$\begin{gathered}
\phi(z) - z \mid \phi(\phi(z)) - \phi(z) = \phi^2(z) - \phi(z) \\
\phi^2(z) - \phi(z) \mid \phi(\phi^2(z)) - \phi(\phi(z)) = \phi^3(z) - \phi^2(z) \\
\phi^3(z) - \phi^2(z) \mid \phi(\phi^3(z)) - \phi(\phi^2(z)) = \phi^4(z) - \phi^3(z) \\
\vdots
\end{gathered}$$

so that by transitivity of divisibility, $\phi(z) - z$ divides $\phi^k(z) - \phi^{k-1}(z)$ for all positive integers k.

Now suppose $n = ml$. Adding $\phi^k(z) - \phi^{k-1}(z)$ from $k = 1$ to l, the sum telescopes:

$$\phi(z) - z \mid \sum_{k=1}^{l} \phi^k(z) - \phi^{k-1}(z) = \phi^l(z) - z.$$

Finally, replacing ϕ with ϕ^m and noting that the lth iterate of ϕ^m is just $\phi^{lm} = \phi^n$, we get the desired result:

$$\phi^m(z) - z \mid \phi^n(z) - z. \qquad \blacksquare$$

3.2 A curious identity

While reading Walde's and Russo's parametrization of those values of c for which $Q_c(z) = z^2 + c$ has a 3-cycle, we noticed that the three given identities [15, p. 323]

$$x_1 + x_2 = \tau, \quad x_2 + x_3 = -\frac{\tau + 1}{\tau}, \quad x_3 + x_1 = -\frac{1}{\tau + 1}$$

multiply to 1. Was this a coincidence?

Proposition 3.4. *Let $x_1, x_2, \ldots, x_n$ be an n-cycle for $Q_c(z) = z^2 + c$ (not necessarily over $\mathbf{Q}$). Then*

$$\prod_{i=1}^{n} x_i + Q_c(x_i) = 1.$$

Proof. Since $Q_c(x_i) = x_{i+1}$ (where x_{n+1} is taken to be x_1) and $x_i \neq x_{i+1}$, we may divide and multiply by $1 = \frac{x_i - x_{i+1}}{x_i - x_{i+1}}$:

$$\begin{aligned} x_i + Q_c(x_i) &= x_i + x_{i+1} \\ &= (x_i + x_{i+1}) \cdot \frac{x_i - x_{i+1}}{x_i - x_{i+1}} \\ &= \frac{x_i^2 - x_{i+1}^2}{x_i - x_{i+1}} \end{aligned}$$

and add $0 = c - c$:

$$\begin{aligned} \frac{x_i^2 - x_{i+1}^2}{x_i - x_{i+1}} &= \frac{x_i^2 + c - x_{i+1}^2 - c}{x_i - x_{i+1}} \\ &= \frac{Q_c(x_i) - Q_c(x_{i+1})}{x_i - x_{i+1}} \\ &= \frac{x_{i+1} - x_{i+2}}{x_i - x_{i+1}} \end{aligned}$$

so that the product telescopes:

$$\begin{aligned} \prod_{i=1}^{n} x_i + Q_c(x_i) &= \prod_{i=1}^{n} \frac{x_{i+1} - x_{i+2}}{x_i - x_{i+1}} \\ &= \frac{x_2 - x_3}{x_1 - x_2} \cdot \frac{x_3 - x_4}{x_2 - x_3} \cdots \frac{x_n - x_1}{x_{n-1} - x_n} \cdot \frac{x_1 - x_2}{x_n - x_1} \\ &= 1 \end{aligned}$$

∎

Benedetto observed this result as well and generalized it in [1].

3.3 A note on conjugation

As we move toward some classification theorems, we must take note of an essential device for simplifying our task. To classify the general nth-degree polynomial $f(z) = c_n z^n + \cdots + c_1 z + c_0$ is to grapple with the $n + 1$ degrees of freedom provided by its coefficients. Linear conjugacy is a special type of equivalence which preserves all the dynamical properties of a map, including its picture. It allows us to "label" two dots on the picture, so to speak, and this reduces the number of unknown coefficients of f.

To put it concretely, if $f(a) = b \neq a$ then conjugating by $\varphi(z) = \alpha z + \beta$ where

$$\alpha = \frac{c - d}{a - b}$$

and

$$\beta = \frac{ad - bc}{a - b}$$

will yield a new map $f^\varphi = \varphi \circ f \circ \varphi^{-1}$ that sends c to d. These values for α and β are obtained by solving the system of equations

$$\varphi(a) = c, \quad \varphi(b) = d.$$

If, further, f sends b back to a, then so will f^φ send d back to c:

$$\varphi(f(\varphi^{-1}(d))) = \varphi(f(b)) = \varphi(a) = c.$$

If $f(a) = a$ and then we may conjugate by $\psi(z) = \alpha(z - a) + c$ (for any α) so that f^ψ fixes c. If f has another fixed point, b, we may conjugate by the φ above so that f^φ fixes c and d. Note that this time $a \neq b$ so α is well defined. If, instead of another fixed point, we choose to relabel a preimage b of a to d then conjugating by φ above will do the trick: f^φ will map d to c to c.

When seeking to conjugate a polynomial into a form with fewer coefficients, it's useful to have an expression for the the conjugate in terms of the original. Unfortunately, the convention we adopted earlier in this paper for the order of conjugation is a little messier than conjugating "the other way"; if $\varphi(z) = \alpha z + \beta$, compare

$$f^\varphi(z) = \varphi(f(\varphi^{-1}(z))) = \alpha(f(\frac{1}{\alpha}z - \frac{\beta}{\alpha})) + \beta$$

to

$$f^{\varphi^{-1}}(z) = \varphi^{-1}(f(\varphi(z))) = \frac{1}{\alpha}(f(\alpha z + \beta) - \beta).$$

In the following we opt for the latter, more aesthetic of the two formulae.

Proposition 3.5. *If $\varphi(z) = \alpha z + \beta$ and*

$$f(z) = \sum_{k=0}^{n} c_k z^k$$

is a general nth-degree polynomial, then

$$f^{\varphi^{-1}}(z) = \sum_{m=0}^{n} \Big(\alpha^{m-1} \sum_{k=m}^{n} c_k \binom{k}{m} \beta^{k-m}\Big) z^m - \frac{\beta}{\alpha}.$$

Proof. First we evaluate f at $\varphi(z)$:

$$\begin{aligned} f(\alpha z+\beta) =& \sum_{k=0}^{n} c_k(\alpha z+\beta)^k \\ =& \sum_{k=0}^{n} c_k \sum_{j=0}^{k} \binom{k}{j} \alpha^j \beta^{k-j} z^j \\ =& c_0 \left[\binom{0}{0} \alpha^0 \beta^0 z^0 \right] \\ &+ c_1 \left[\binom{1}{0} \alpha^0 \beta^1 z^0 + \binom{1}{1} \alpha^1 \beta^0 z^1 \right] \\ &+ c_2 \left[\binom{2}{0} \alpha^0 \beta^2 z^0 + \binom{2}{1} \alpha^1 \beta^1 z^1 + \binom{2}{2} \alpha^2 \beta^0 z^2 \right] \\ &\vdots \end{aligned}$$

Suppose we want the z^m coefficient (for some m between 0 and n). Each inner sum features some j equal to m if and only if k allows it (by being at least m). Adding vertically (down those $k \geq m$ and $j = m$) we get that the z^m coefficient is

$$\sum_{k=m}^{n} c_k \binom{k}{m} \alpha^m \beta^{k-m}$$

so that

$$f(\alpha z+\beta) = \sum_{m=0}^{n} \left(\alpha^m \sum_{k=m}^{n} c_k \binom{k}{m} \beta^{k-m} \right) z^m.$$

Subtracting β and dividing by α yields the desired result. ∎

Corollary.

- *The leading term of $f^{\varphi^{-1}}$ is $\alpha^{n-1}c_n$.*
- *The next term is $\alpha^{n-2}(c_{n-1} + c_n n\beta)$, which vanishes if and only of $\beta = -\frac{c_{n-1}}{nc_n}$; this is a Tschirnhaus transformation.*
- *The linear term is unaffected by α and hence by pure scalings $z \mapsto \alpha z$.*

3.4 Linear classification

The maps $z \mapsto \alpha z + \beta$ where α and β are rational are very simple to classify because iteration does not increase the degree of the map. If $f(z) = \alpha z + \beta$

then $f^n(z) = \alpha^n z + (\alpha^{n-1} + \cdots + \alpha + 1)\beta$. The n-cyclic points are among the solutions to

$$0 = f^n(z) - z = (\alpha^n - 1)z + \frac{\alpha^n - 1}{\alpha - 1}\beta.$$

If $\alpha^n \neq 1$ then this equation reduces to

$$\begin{aligned} 0 &= z + \frac{1}{\alpha - 1}\beta \\ \frac{\beta}{1-\alpha} &= z \end{aligned}$$

but $f(\frac{\beta}{1-\alpha}) = \frac{\alpha\beta}{1-\alpha} + \beta = \frac{\alpha\beta+\beta-\alpha\beta}{1-\alpha} = \frac{\beta}{1-\alpha}$ so that f's "n-cycle" is actually fixed. Thus for α not a root of unity, $z \mapsto \alpha z + \beta$ has a single fixed point.

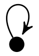

Figure 8: A single fixed point

If, on the other hand, $\alpha^n = 1$, then $\alpha = \pm 1$ (over $\mathbf{Q}$). In the first case, $f(z) = z + \beta$, which is of no dynamical interest unless $\beta = 0$, in which case f is the identity map and it fixes everything.

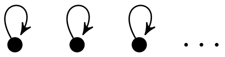

Figure 9: The identity map's picture

In the second case, $f(z) = -z + \beta$, so that the n-cyclic points are among the solutions to

$$\begin{aligned} 0 &= f^n(z) - z \\ &= ((-1)^n - 1)z + \frac{(-1)^n - 1}{-2}\beta \\ \frac{\beta}{2}((-1)^n - 1)z &= ((-1)^n - 1)z \end{aligned}$$

If n is odd, we get $\frac{\beta}{2}(-2) = (-2)z$; the solution, $z = \frac{\beta}{2}$, is actually a fixed point. If n is even, we get $0 = 0z$ so that all rational z are solutions; indeed, $f(f(z)) = f(-z+\beta) = -(-z+\beta) + \beta = z - \beta + \beta = z$. Hence for $\alpha = -1$, the picture consists of a fixed point and infinitely many 2-cycles.

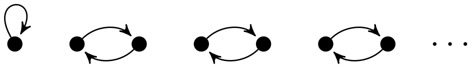

Figure 10: The most interesting case: $z \mapsto \beta - z$

It is of pedagogical interest to note that these linear maps provide a natural example of a semidirect product. Given a commutative unital ring R, the subset G of $R[z]$ consisting of invertible linear maps

$$G = \{z \mapsto \alpha z + \beta : \alpha \in R^\times, \beta \in R\}$$

forms a group under composition

$$\big(z \mapsto \alpha z + \beta\big) \circ \big(z \mapsto \gamma z + \delta\big) = z \mapsto (\alpha\gamma)z + (\alpha\delta + \beta).$$

We may identify two intuitive subgroups of G: the translations

$$T = \{z \mapsto z + \beta : \beta \in R\}$$

and the scalings

$$S = \{z \mapsto \alpha z : \alpha \in R^\times\}.$$

It is clear that both are indeed subgroups and that T is further a normal subgroup. The only translation that is also a scaling is the identity map, so S and T intersect trivially. Finally, every element $z \mapsto \alpha z + \beta$ in G actually decomposes uniquely as

$$\big(z \mapsto z + \beta\big) \circ \big(z \mapsto \alpha z\big)$$

so that G is the semidirect product of T and S.

A closer look into the structure of T reveals an obvious isomorphism with $(R, +)$:

$$\big(z \mapsto z + \beta_1\big) \circ \big(z \mapsto z + \beta_2\big) = z \mapsto z + (\beta_1 + \beta_2),$$

whereas composition in S looks like

$$\big(z \mapsto \alpha_1 z\big) \circ \big(z \mapsto \alpha_2 z\big) = z \mapsto (\alpha_1\alpha_2)z$$

so that S is isomorphic to $(R^\times, \cdot)$. Therefore $G = T \rtimes S \cong R^+ \rtimes R^\times$. The homomorphism from $R^\times$ into $\mathrm{Aut}(R^+)$ is given by $r \mapsto (z \mapsto rz)$.

3.5 Lagrange interpolation polynomials

We enrich this section with a useful method to generate examples.

Let K be a field. Given a finite set of distinct indeterminates $x_1, \ldots, x_n$ we may define

$$f_i(x) = \prod_{\substack{j=1 \\ j \neq i}}^{n} \frac{x - x_j}{x_i - x_j} \in K(x_1, \ldots, x_n)[x]$$

so that

$$f_i(x_k) = \prod_{j \neq i} \frac{x_k - x_j}{x_i - x_j} = \begin{cases} 1 & \text{if } k = i \\ 0 & \text{if } k \neq i \end{cases} = \delta_{ik}$$

(the Kronecker delta). For a permutation σ in S_n (or, generally, for any self-map on $\{1, \ldots, n\}$) we may define an interpolant polynomial through the points $(x_i, x_{\sigma(i)})$ by

$$f_\sigma(x) = \sum_{i=1}^{n} x_{\sigma(i)} f_i(x) \in K(x_1, \ldots, x_n)[x].$$

When restricted to the indeterminates, f_σ "implements" σ. Choosing actual, pairwise distinct constants c_i in the ground field K we get a concrete polynomial. For example, if σ is an n-cycle, then f_σ is a degree-$(n-1)$ polynomial with an n-cycle. Of particular interest is the leading term

$$L_\sigma(x_1, \ldots, x_n) = \sum_{i=1}^{n} x_{\sigma(i)} \prod_{j \neq i} \frac{1}{x_i - x_j}$$

which, if made to vanish by appropriate choice of (pairwise distinct) x_i, will yield a degree-$(n-2)$ polynomial which implements σ.

For $n = 1$, $L_{(1)}(x_1) = x_1$ which also equals $f_{(1)}(x)$.

For $n = 2$,

$$L_\sigma(x_1, x_2) = \frac{x_{\sigma(1)} - x_{\sigma(2)}}{x_1 - x_2}.$$

This simplifies to $L_{(1)}(x_1, x_2) = 1$ if σ is the identity (indeed, $f_{(1)}(x) = x$); if σ is a transposition, then $L_{(1\,2)}(x_1, x_2) = -1$ (indeed, $f_{(1\,2)}(x) = x_1 + x_2 - x$). Leaving S_n, we see that if σ is constant, then $L_\sigma(x_1, x_2) = 0$, and $f_\sigma(x) = x_{\sigma(1)}$ is of degree $2 - 2 = 0$.

For $n \geq 3$, the equations get pretty hairy; we can take advantage of conjugation and choose $x_1 = 0$ and $x_2 = 1$. Since we are interested in periodic points, we pick σ to be some 3-cycle, say, $\sigma = (1\,2\,3)$:

$$\begin{aligned} L_{(1\,2\,3)}(0, 1, x_3) &= \frac{1}{(0-1)(0-x_3)} + \frac{x_3}{(1-0)(1-x_3)} \\ &= \frac{1}{x_3} + \frac{x_3}{1-x_3} \\ &= \frac{x_3^2 - x_3 + 1}{x_3(1-x_3).} \end{aligned}$$

Note that the numerator $x_3^2 - x_3 + 1$ does not vanish over $\mathbf{Q}$ so that a rational 3-cycle cannot be exhibited by anything smaller than a quadratic.

Less trivial is the fact that $L_{(1\,2\,3\,4\,5)}(0, 1, x_3, x_4, x_5)$ vanishes for $x_3 = -1$, $x_4 = 2$, and $x_5 = -2$, corresponding to our only known example (up to conjugacy) of a cubic with a 5-cycle: $\frac{x^3}{3} - \frac{x^2}{2} - \frac{11x}{6} + 1$, conjugate to $\frac{x^3}{12} - \frac{25x}{12} + 1$ in normal form.

We can use Lagrange polynomials to show that the existence of a quadratic with a rational 4-cycle is equivalent to the existence of non-trivial rational points on a certain curve. Let $\sigma = (1\,2\,3\,4)$ and suppose $(x_1, x_2, x_3, x_4) = (0, 1, a, b)$, where a and b are distinct from each other and from 0 and 1. Then the interpolant is of degree at most 3, with leading term

$$\begin{aligned} L_\sigma(0, 1, a, b) &= \frac{1}{(0-1)(0-a)(0-b)} \\ &\quad + \frac{a}{(1-0)(1-a)(1-b)} \\ &\quad + \frac{b}{(a-0)(a-1)(a-b)} \\ &= \frac{a^3b - a^2b^2 - a^2b + a^2 + ab^2 - a + b^3 - 2b^2 + b}{ab(a-1)(b-1)(a-b)}, \end{aligned}$$

which vanishes if and only if the numerator does. A plot of the numerator set to zero reveals $(a, b) = (0, 1)$ to be a critical point. The pencil of lines through $(0, 1)$ is given by $b = at + 1$ as t ranges; making the substitution and dividing by $-a^2$ (remember, $a \neq 0$) yields the following quadratic in a:

$$a^2(t^2 - t) + a(-t^3 - t^2 + 3t - 1) + (-t^2 - 2t + 1).$$

Its solutions are given by

$$a = \frac{(t-1)(t^2 + 2t - 1) \pm \sqrt{(t-1)(t+1)(t^2+1)(t^2+2t-1)}}{2(t-1)t}$$

so that a rational 4-cycle exists if and only if the radicand is a square, which is to say the curve

$$s^2 = t^6 + 2t^5 - t^4 - t^2 - 2t + 1$$

has non-trivial rational points.

Of course there are some obvious solutions, namely the roots of the radicand, but $t = 1$ is forbidden lest our "quadratic in a" simplify to -2 (which is never 0); and $t = -1$ is forbidden because then $a = 1$.

If we decide to venture beyond $\mathbf{Q}$ and into $\mathbf{Q}(i)$, then $t = \pm i$ corresponds to interpolant polynomials $z \mapsto \pm iz + 1$, which aren't even quadratics (but they do get the job done: $0 \mapsto 1 \mapsto 1 \pm i \mapsto \pm i \mapsto 0$). Finally, we cannot use the roots of $t^2 + 2t - 1$ (which live in $\mathbf{Q}(\sqrt{2})$ because they will cause a itself to vanish.

A disadvantage of Lagrange polynomials is that we have no control over the other preperiodic points of the interpolant. For example, the linear map which implements a 2-cycle also happens to have infinitely many other 2-cycles and a fixed point (see figure 10 on page 27), and the quadratic which implements a 3-cycle always has purely preperiodic preimages of that cycle (this is a consequence of part (1) of Theorem 3 in [12, p. 16]).

3.6 Quadratic 2-cycle

As a warm up to the analyses which follow, let's see if a general quadratic can have only a 2-cycle as its picture.

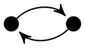

Figure 11: Can a quadratic have this as its picture?

Without loss of generality, let's assume $f(z) = az^2 + bz + c$ sends 0 to 1 and 1 to 0, so that $c = 1$ and $a+b+1 = 0$. Therefore, $f(z) = az^2 - (a+1)z + 1$. The pre-images of 0 are easy to find, because we already know one: $z = 1$. The other is the solution to

$$0 = \frac{f(z)}{z-1} = az - 1;$$

thus we have $f(1/a) = 0$. The pre-images of 1 are likewise easily found:

$$0 = \frac{f(z) - 1}{z} = az - (1 + a);$$

thus we have $f(1 + 1/a) = 1$. In order for this quadratic's picture to be exactly a 2-cycle, we need a value of a for which $1/a = 1$ and $1 + 1/a = 0$—this is impossible, because the first equation implies the second equation reduces to $2 = 0$, which is absurd.

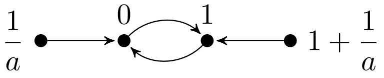

Figure 12: A quadratic 2-cycle can never occur alone.

4 Cubic pictures: narrated equation-solving

Our analysis of Benedetto's data has revealed over 102 possible pictures that cubics can yield, a number which shatters any hopes of a complete classification. Nonetheless, certain statements can be made. Having classified the cubic data, we were able to observe the frequency with which certain pictures occurred—some very often, some quite rarely. By computing all possible pictures of a certain size (using the results from the previous section) we were able to note that some pictures did not occur at all *in the data*. These observations motivated us to try to prove whether these pictures really are abundant, rare, unique, or non-existent.

4.1 Cubic with two fixed points and one preperiodic point

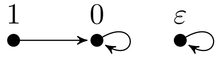

Figure 13: $1 \to 0 \to 0$, $\varepsilon \to \varepsilon$

This is the only size-3 graph which does not occur in the classified data. We endeavoured to see if it was possible for a cubic to have this as its picture.

If a general cubic $f(z) = az^3 + bz^2 + cz + d$ sends 1 to 0, fixes 0, and fixes some other point, then $f(0) = 0 \Rightarrow d = 0$ and $f(1) = 0 \Rightarrow c = -(a + b)$. (Recall that we have the freedom of relabeling two points in $\mathrm{PrePer}(f)$ because of conjugation.)

Notice that $f(z)$ has two known roots—0 and 1—so that it splits completely over $\mathbf{Q}$. This third root cannot be anything new, because we are supposing that f have *one* point of type 1_1. Hence this third root must actually be 0 or 1, and so $f(z)$ has a double root. This means that its discriminant, $(a+b)^2(2a+b)^2$, vanishes. Case 1 below deals with $b=-a$; Case 2 with $b=-2a$.

Another restriction is given by looking at $f(z)-z$. This polynomial also has two known roots—0 and some other point—so that it likewise splits completely over $\mathbf{Q}$. Now, because $f(z)-z$ has two "pictorial" fixed points but three roots counting multiplicity, the polynomial $f(z)-z$ must also have a double root. This means that its discriminant, $(a+b+1)^2(4a^2+4ab+4a+b^2)$, also vanishes.

There are very few cubic polynomials for which both $f(z)$ and $f(z)-z$ have a double root. The two simultaneously vanishing discriminants provide enough information to solve for a and b exactly.

Case 1: $b=-a$. The discriminant of $f(z)-z$ simplifies to $(a-a+1)^2(4a^2+4a(-a)+4a+(-a)^2) = 4a-a^2 = a(4-a)$ so that $a=-4$, in which case $b=4$ and $c=0$. Indeed, $f(z)=-4z^3+4z^2$ has the desired picture; the other fixed point is $\frac{1}{2}$.

Case 2: $b=-2a$. The discriminant of $f(z)-z$ simplifies to $(a-2a+1)^2(4a^2+4a(-2a)+4a+(-2a)^2) = (1-a)^2(4a) = 4a(a-1)^2$ so that $a=1$, in which case $b=-2$ and $c=1$. Indeed, $f(z)=z^3-2z^2+z$ has the desired picture; the other fixed point is 2.

These two polynomials do not appear in the data because when conjugated into normal form the denominators of their leading terms are both 729, which exceeds the bound of 300 used in the search.

4.2 Cubic with a 2-cycle and antipodal preimages

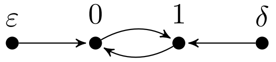

Figure 14: $\varepsilon \to 0 \to 1 \to 0$, $\delta \to 1$

This small and symmetric picture occurs only once in the classified data. We sought to determine if its existence is unique.

If a general cubic sends 0 to 1 and 1 to 0, then it must be of the form $f(z)=az^3+bz^2-(1+a+b)z+1$. If we further require that both 0 and 1 have *only one* non-periodic preimage each (say $f(\varepsilon)=0$ and $f(\delta)=1$), then, because the equations

$$\frac{f(z)-0}{(z-1)(z-\varepsilon)}, \quad \frac{f(z)-1}{(z-0)(z-\delta)}$$

are both linear, the points 0 and 1 actually have three rational preimages. But we don't want these preimages to be anything new, so $f^{-1}(0)$ must equal $\{0, \varepsilon\}$ and $f^{-1}(1)$ must equal $\{1, \delta\}$. Therefore, both $f(z)$ and $f(z)-1$ have double roots.

The discriminant of $f(z)$ is $(2a+b-1)^2(a^2+4a+2ab+b^2)$, and the discriminant of $f(z)-1$ is $(a+b+1)^2(4a^2+4a+4ab+b^2)$. These must simultaneously vanish; the case analysis follows.

Case 1: $b = 1-2a$. The discriminant of $f(z)-1$ simplifies to $(2-a)^2(1+4a)$. If $a=2$, then $f(z) = 2z^3-3z^2+1$. If $a = -\frac{1}{4}$, then $f(z) = -\frac{1}{4}z^3+\frac{3}{2}z^2-\frac{9}{4}z+1$.

Case 2: $b = -1-a$. The discriminant of $f(z)$ simplifies to $(a-2)^2(4a+1)$. If $a=2$, then b happens to equal -3 again. On the other hand, if $a=-\frac{1}{4}$, then $f(z) = -\frac{1}{4}z^3-\frac{3}{4}z^2+1$, which is conjugate to the second thing in Case 1 via $z \mapsto 1-z$.

Case 3: $a^2+4a+2ab+b^2 = 0$ and $4a^2+4a+4ab+b^2=0$. This system of equations actually has a swift solution; setting them equal immediately yields $a^2+2ab = 4a^2+4ab$. Recalling that $a \neq 0$, we get that

$$b = -\frac{3}{2}a.$$

With this restriction, both discriminants simplify to constant multiples of $a(a-2)(a+16)$. If $a=2$ then $b = -\frac{3}{2}(2) = -3$ and we have found nothing new. The other solution, $a=-16$, yields $f(z) = -16z^3+24z^2-9z+1$.

In summary, the algebra has revealed three possible candidate cubics, but only $-\frac{1}{4}z^3+\frac{3}{2}z^2-\frac{9}{4}z+1$ has precisely the desired picture. The fact that

$$\frac{f^n(z)-z}{f(z)-z}$$

has no unforeseen rational roots for $n = 2, 3, 4, 5$ (verified with a computer) makes us highly confident this picture is unique.

The other two maps each have a fixed point at $z = \frac{1}{2}$, but no other (pre)periodic points. They are not conjugate to each other, because 0's double-preimage under $2z^3-3z^2+1$ is 1 (the other period-2 point) whereas 0's double-preimage under $-16z^3+24z^2-9z+1$ is $\frac{1}{4}$ (a purely preperiodic point). Since nothing in our above analysis precluded figure 15 being the target picture, we conclude that these two cubics are the only two cubics with this picture.

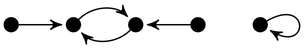

Figure 15: There are only two cubics with this picture.

4.3 Cubic with a fixed point and a Y tree above it

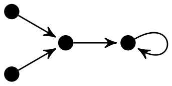

If a general cubic $f(z) = az^3 + bz^2 + cz + d$ sends 1 to 0 and fixes 0, then $f(z) = az^3 + bz^2 - (a+b)z$. If it further has the property that both $f(z)$ and $f(z) - 1$ have double roots, then both their discriminants must vanish. Initially, $\mathrm{disc}_z(f(z)-1)$ does not factor, but after setting $\mathrm{disc}_z f(z) = (a+b)^2(2a+b)^2$ to zero and getting b in terms of a, it does.

Case 1: $b = -a$. The discriminant of $f(z) - 1$ simplifies to $-a^2(4a+27)$, which implies $a = -\frac{27}{4}$ and $f(z) = -\frac{27}{4}z^3 + \frac{27}{4}z^2$. The preimages of 1 are $\frac{2}{3}$ and $-\frac{1}{3}$.

Case 2: $b = -2a$. The discriminant of $f(z) - 1$ simplifies to $a^2(4a - 27)$, which implies $a = \frac{27}{4}$ and $f(z) = \frac{27}{4}z^3 - \frac{27}{2}z^2 + \frac{27}{4}z$. The preimages of 1 are $\frac{4}{3}$ and $\frac{1}{3}$.

Neither of these cubics has any fixed points, 2-cycles, or 3-cycles, nor do any of their points have more preimages—their pictures are exactly as desired.

4.4 Cubic with a point of type 1_1 and another of type n_1

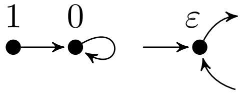

We will begin with an assumption on $f(z) = az^3 + bz^2 + cz + d$ which gives some freedom in the possible pictures for f. A cubic with a graph as above must satisfy the following two conditions:

1. $f(z)$ has a double root
2. $f(z) - \varepsilon$ has a double root

but what's happening around $\varepsilon \neq 0, 1$ is a little uncertain; we only ask that it be part of an n-cycle for some small n and that it have exactly two distinct preimages.

First, the equations $f(1) = 0$ and $f(0) = 0$ imply (as before) that $f(z) = az^3+bz^2-(a+b)z$. Condition 1 tells us that f's discriminant, $(a+b)^2(a+2b)^2$, equals zero, so that our cubic is of one of the following two forms:

$$\begin{aligned} f_1(z) &= az^3 - az^2 = az^2(z-1) \\ f_2(z) &= az^3 - 2az^2 + az = az(z-1)^2. \end{aligned}$$

Case 1: $\operatorname{disc}_z(f_1(z) - \varepsilon) = -a^2\varepsilon(4a + 27\varepsilon)$. The only interesting solution is $\varepsilon = -\frac{4a}{27}$, whence $f_1(z) - (-\frac{4a}{27}) = \frac{a}{27}(3z-2)^2(3z+1)$.

Since we are assuming that ε is part of some cycle, some forward iterate of it must map to one of its preimages: $\frac{2}{3}$ or $\frac{-1}{3}$. The first few iterates of ε are

$$\begin{aligned} f_1^0\left(-\frac{4a}{27}\right) &= -\frac{4a}{27}, \\ f_1^1\left(-\frac{4a}{27}\right) &= -\frac{16a^3(4a+27)}{27^3}, \end{aligned}$$

and

$$f_1^2\left(-\frac{4a}{27}\right) = -\frac{256a^7(4a+27)^2(64a^4+432a^3+19683)}{27^9}$$

after which the equations become unmanageable.

It turns out that the only rational solution to be found among the three equations $f_1^i(\varepsilon) = \frac{2}{3}$ for $i = 0, 1, 2$ (not simultaneously, just whatsoever) is $a = -\frac{9}{2}$, yielding the cubic $-\frac{9}{2}z^3 + \frac{9}{2}z^2$, conjugate to $-2z^3 + \frac{3}{2}z$ in normal form. This map fixes $\frac{1}{3}$.

The only rational solution among all the three equations $f_1^i(\varepsilon) = -\frac{1}{3}$ is $a = \frac{9}{4}$, which yields a cubic conjugate to $\frac{9}{16}z^3 - \frac{3}{4}z + 1$ in normal form. This map fixes $\frac{4}{3}$.

Both these cubics have the following picture:

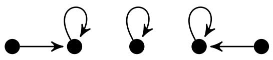

Case 2: $\operatorname{disc}_z(f_2(z)-\varepsilon) = a^2\varepsilon(4a-27\varepsilon)$; this time, $\varepsilon = \frac{4a}{27}$. The multiplicities of its preimages are readily seen from the factorization $f_2(z) - (\frac{4a}{27}) = \frac{a}{27}(3z-4)(3z-1)^2$. Again, we assume ε is trapped in some n-cycle, so

compute the forward iterates of ε:

$$f_2^0\left(\frac{4a}{27}\right) = \frac{4a}{27},$$

$$f_2^1\left(\frac{4a}{27}\right) = \frac{4(4a^2-27a)^2}{27^3},$$

and

$$f_2^2\left(\frac{4a}{27}\right) = \frac{4a^3(4a-27)^2(64a^4-864a^3+2916a^2-19683)^2}{27^9}.$$

This case is much more fruitful. $f_2^0(\varepsilon) = \frac{4}{3}$ implies $a = 9$, which yields $9z^3 - 18z^2 + 9z$, conjugate to $z^3 - 3z$. $f_2^0(\varepsilon) = \frac{1}{3}$ implies $a = \frac{9}{4}$, which yields a cubic conjugate to $\frac{9}{16}z^3 - \frac{3}{4}z + 1$, seen before.

At this point, we have exactly three cubics with two double-rooted fixed points.

The equation $f_2(\varepsilon) = \frac{4}{3}$ has a new rational solution: $a = -\frac{9}{4}$, which corresponds to a 2-cycle between $\varepsilon = -\frac{1}{3}$ and $\frac{4}{3}$. This map is conjugate to $-\frac{25}{16}z^3 + \frac{3}{4}z + 1$.

The equation $f_2(\varepsilon) = \frac{1}{3}$ has a new rational solution as well: $a = -\frac{9}{2}$, which corresponds to a 2-cycle between $\varepsilon = \frac{2}{3}$ and its double-preimage $\frac{1}{3}$. This map is conjugate to $\frac{1}{2}z^3 - \frac{3}{2}z + 1$.

Both these cubics have the following picture:

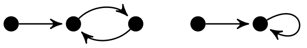

Finally, neither of the equations $f_2^2(\varepsilon) = \frac{1}{3}, \frac{4}{3}$ has novel rational solutions, which proves that no cubic with a fixed point with one preimage and a periodic point with one non-cyclic preimage can have a 3-cycle. Further, only two can have a 2-cycle, and only three can have another fixed point.

4.5 Cubics with a point of type 3_1

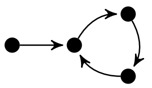

Figure 16: A graph not among the data

This picture is one of the few size-4 graphs which does not occur in the classified data, although there is no obvious reason why it shouldn't.

If a general cubic f sends 1 to 0, the linear coefficient may be eliminated so that $f(z) = az^3+bz^2-(a+b+d)z+d$. Since we want 0 to have exactly two preimages, we can use discriminants again. Setting $\operatorname{disc}_z(f(z)) = (2a+b-d)^2(a^2+2ab+4ad+b^2)$ to zero yields two cases. We only analyze the simpler one (equation-wise): $d = 2a+b$. Here, $f(z) = az^3+bz^2-(3a+2b)z+2a+b = (z-1)^2(az+2a+b)$.

The equation $f(2a+b) = 1$ can be solved by parametrizing the constant term using t (despite it already equalling d—perhaps the author prefers unvoiced variables in his notes). Trudging through much more arithmetic than in the above cases, we come to

$$f(z) = \frac{(z-1)^2((-t^3+2t^2-t+1)z+t^4-2t^3+t^2)}{t(t-1)^2}$$

where $t \in \mathbf{Q} \setminus \{0,1\}$.

Though we are unable to prove anything about this general form, a few concrete examples reveal pictures not found in the data.

Letting $t = \frac{1}{3}$ yields $\frac{23}{4}z^3 - \frac{67}{6}z^2 + \frac{61}{12}z + \frac{1}{3}$, which has $-\frac{4}{69}$ as a point of type 3_1 and (as far as we can tell) no other preperiodic points. It is conjugate to

$$\frac{12158811289}{103456682352}z^3 - \frac{5329}{2484}z + 1.$$

Letting $t = \frac{2}{3}$ yields $\frac{25}{2}z^3 - \frac{73}{3}z^2 + \frac{67}{6}z + \frac{2}{3}$, which has $-\frac{4}{75}$ as a point of type 3_1 and no other preperiodic points. It is conjugate to

$$\frac{38962417321}{16607531250}z^3 - \frac{6241}{1350}z + 1.$$

Both these cubics have figure 16 as their picture; both were missed by Benedetto's cubic search because their coefficients have denominators quite larger than 300.

The equation $f(2a+b) = -2 - \frac{b}{a}$ is a little trickier, but can be solved with the same substitution $t = 2a+b$. What results is a quadratic in a with the discriminant equal to $(t-3)(t+1)$.

With $t = 3$ we get $a = -\frac{1}{2}$ and $f(z) = -\frac{1}{2}z^3 + 4z^2 - \frac{13}{2}z + 3$, which sends $1 \to 0 \to 3 \to 6 \to 0$ and fixes 2. It is conjugate to

$$-\frac{2809}{1458}z^3 + \frac{25}{6}z + 1$$

and has the following picture

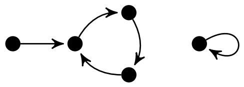

Figure 17: Another graph not among the data

With $t = 1$ we get $a = -\frac{1}{2}$ again (of course) and $f(z) = -\frac{1}{2}z^3 + \frac{3}{2}z - 1$; conjugating by $z \mapsto -z$ flips the sign on the constant term and puts it into normal form. This cubic *was* found among the data, and there was only one instance of it. The points -1 and -2 are of type 3_1; it has the following picture:

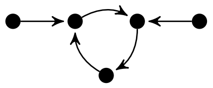

Figure 18: Only one of these was found among the data

We need not only evaluate t at roots of the discriminant; it appears under a radical in the formula for a, so we seek values of t for which it is a rational square. The equation $s^2 = (t-3)(t+1)$ determines a hyperbola in the plane, which has infinitely many rational points. A pencil of lines through either $(-1, 0)$ or $(3, 0)$ does the trick. One of the cubics churned out ($t = \frac{7}{2}$) is conjugate to

$$-\frac{1442401}{2000}z^3 + \frac{363}{20}z + 1$$

and has figure 16 as its picture (the hook is $-\frac{110}{1201}$).

4.6 Other cubics

Some simple pictures with relatively few instances eluded proof of uniqueness; some simple pictures which did not appear in the data eluded proof of existence. We leave them here for future reference.

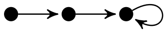

Figure 19: $-6z^3 + \frac{9}{2}z + 1$ and $\frac{3}{2}z^3 - \frac{9}{2}z + 1$

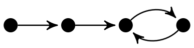

Figure 20: $-\frac{289}{16}z^3 + \frac{27}{4}z + 1$ and $-\frac{49}{250}z^3 + \frac{27}{10}z + 1$

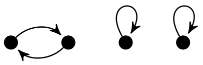

Figure 21: Is this picture possible?

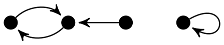

Figure 22: What about this one?

5 Further questions

A natural step is to proceed toward quartics. What is the longest cycle a quartic can have? Lagrange polynomials guarantee this number is at least 5, though we have found no examples of a quartic with a 6- or 7-cycle. A search like Benedetto's could be done for quartic polynomials; it is not hard to see that the general quartic belongs to one of the following three classes of conjugacy classes:

$$az^4 + cz^2 + dz + 1, \quad az^4 + z^2 + dz, \quad az^4 + dz.$$

Instead of increasing the degree of the polynomial, one can increase the degree of the field extension over $\mathbf{Q}$. For example, in [14] it is shown that $z^2 - \frac{71}{48}$ has a 6-cycle over $\mathbf{Q}(\sqrt{33})$; one point in the cycle is $z = -1 + \frac{1}{12}\sqrt{33}$. What is the longest cycle a polynomial over a quadratic extension can have?

Another direction one can take is to follow Narkiewicz in [9] where he investigates cycles of polynomials over certain rings, in particular over $\mathbf{Z}[\frac{1}{N}]$.

One may also study polynomial maps over finite fields, where every point is preperiodic; in [6], Gilbert et al. conjecture that there are exactly p possible pictures modulo p, except for $p = 17$.

6 Acknowledgements

The results of this paper came from a USRA in Winter 2013, held at the University of Waterloo, supported by NSERC and the University of Waterloo. The author gratefully acknowledges the advice and support of his supervisors Kevin Hare and David McKinnon, who were unfailing in their patience, wise with their explanations, and always willing to put their fingers to a keyboard or a pencil to some paper.

References

[1] Rob Benedetto. An elementary product identity in polynomial dynamics. *The American Mathematical Monthly*, 108:860–864, 2001.

[2] Rob Benedetto. Small height points for cubic polynomials. `http://www.cs.amherst.edu/~rlb/cubicdata/`, 2007.

[3] Rob Benedetto, Ben Dickman, Sasha Joseph, Ben Krause, Dan Rubin, and Xinwen Zhou. Computing points of small height for cubic polynomials. *Involve*, 2:37–64, 2009.

[4] E.V. Flynn, Bjorn Poonen, and Edward Schaefer. Cycles of quadratic polynomials and rational points on a genus 2 curve. *Duke Mathematics Journal*, 90:435–463, 1997.

[5] Harald Fripertinger and Peter Schöpf. Endofunctions of a given cycle type. *Annales des sciences mathématiques du Québec*, 23(2):173–188, 1999.

[6] Christie L. Gilbert, Joseph D. Kolesar, Clifford A. Reiter, and John D. Storey. Function digraphs of quadratic maps modulo p. *Fibonacci Quarterly*, 39:32–49, 2001.

[7] OEIS Foundation Inc. The on-line encyclopedia of integer sequences. `http://oeis.org/A123456`, 2011.

[8] Patrick Morton. Arithmetic properties of periodic points of quadratic maps, II. *Acta Arithmetica*, 87(2):89–102, 1998.

[9] Władysław Narkiewicz. Polynomial cycles in certain rings of rationals. *Journal de théorie des nombres de Bordeaux*, 14(2):529–552, 2002.

[10] D. G. Northcott. Periodic points on an algebraic variety. *Annals of Mathematics*, 51(1), January 1950.

[11] Mateusz G. Olechnowicz. Table of cubic pictures. `http://csclub.uwaterloo.ca/~mgolechn/cubics.html`, 2013.

[12] Bjorn Poonen. The complete classification of rational preperiodic points of quadratic polynomials over $\mathbf{Q}$: a refined conjecture. *Mathematische Zeitschrift*, 228:11–29, 1998.

[13] Joseph H. Silverman. *The Arithmetic of Dynamical Systems.* Springer-Verlag, 1st edition, 2007.

[14] Michael Stoll. Rational 6-cycles under iteration of quadratic polynomials. `arXiv:0803.2836v2 [math.NT]`, April 2009.

[15] Ralph Walde and Paula Russo. Rational periodic points of the quadratic function $Q_c(x) = x^2 + c$. *The American Mathematical Monthly*, 101(4):318–331, April 1994.