Arithmetic dynamics

Dragos Ghioca and Thomas Tucker

A simple question

- Let A be an invertible matrix in $\mathrm{GL}_r(\mathbb{C})$, acting in the usual way on $\mathbb{C}^r$.
- Let V be a subspace of $\mathbb{C}^r$.
- Let $\mathbf{z}$ be a point in $\mathbb{C}^r$.

Question

Is there any simple obvious pattern describing the set of n such that $A^n\mathbf{z} \in V$?

Exercise

Suppose that we are in two dimensions and that V is a line through the origin. Show that one of the two following holds:

1. There is at most one n such that $A^n\mathbf{z} \in V$; or
2. There is an entire coset $i + m\mathbb{Z}$ of $\mathbb{Z}$ such that $A^n\mathbf{z} \in V$ for all n in this coset.

This can be proved using simple linear algebra and group theory. [Hint: Show that if two iterates pass V, then $A^m V = V$ for some m.]

One more very simple example

Let's just do one very simple example.

Example

Let $A = \begin{pmatrix} -2 & 0 \\ 0 & 2 \end{pmatrix}$ and let $s = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$. Let V be the subspace consisting of all vectors of the form $\begin{pmatrix} x \\ -x \end{pmatrix}$. Then $A^n s \in V$ if and only if n is odd, i.e. if $n \in 1 + 2\mathbb{Z}$.

Note that A^2 sends V to itself, so once you get some $A^i s \in V$, you must get $A^{i+2k} s \in V$ for any k (this is why you do not get any "singleton cosets" here).

Other examples

The case of lines through the origin is particularly simple, both because the proof is obvious and because the formulation is so simple: Things are not quite so simple in general. For example, when V is a line that does not pass through the origin, it is easy to see that you can have $s \neq t$ with $A^s\mathbf{z} \in V$ and $A^t\mathbf{z} \in V$ without getting an entire coset, just by drawing a line.

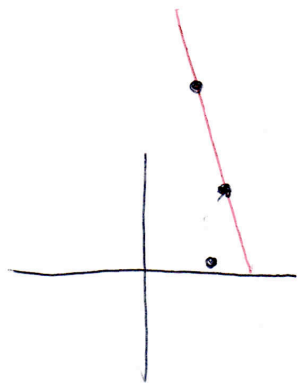

But in the case of lines not passing through the origin, once you have *a large* finite number of n such that $A^n\mathbf{z} \in V$, you must have an entire coset of such n. (In fact, there is an explicit bound on that number, 61, due to Beukers-Schlickewei, which is likely nowhere near sharp.)

Linear recurrence sequences

The first results on the problem above were proved by Skolem, in the context of closely related questions on *linear recurrence sequences*.

We call $\{a_n\}_{n\in\mathbb{N}} \subset \mathbb{C}$ a linear recurrence sequence if for some $k \in \mathbb{N}$ there exist $c_0, \ldots, c_{k-1} \in \mathbb{C}$ such that for each $n \in \mathbb{N}$ we have

$$a_{n+k} = c_{k-1}a_{n+k-1} + \cdots + c_1 a_{n+1} + c_0 a_n.$$

The associated characteristic equation for this sequence is

$$x^k - c_{k-1}x^{k-1} - \cdots - c_1 x - c_0 = 0.$$

We let $r_1, \ldots, r_m$ be the *distinct, nonzero* roots of the above characteristic equation. Then there exist polynomials $f_1, \ldots, f_m \in \mathbb{C}[z]$ ($\deg(f_i)$ is less than the order of multiplicity of the root r_i) such that for each $n \in \mathbb{N}$ we have

$$a_n = f_1(n)r_1^n + \cdots + f_m(n)r_m^n. \tag{1}$$

It's easy to see formula (1) in the case $m = k$ (i.e., all roots of the characteristic equation are distinct and nonzero). In that case the polynomials f_i are simply constants which we find by solving the k-by-k system of equations obtained from verifying the above formula for $n = 1, \ldots, k$. Then, note that since each λ_i is a root of $x^k - c_{k-1}x^{k-1} - \cdots - c_1 x - c_0 = 0$, we have

$$\lambda_i^{n+k} = c_{n+k-1}\lambda_i^{n+k-1} + \cdots + c_1\lambda_i^{n+1} + c_0\lambda^n$$

so (1) then continues to hold by induction. When λ_i is a double root, then it also satisfies the derivative of $(x^n)(x^k - c_{k-1}x^{k-1} - \cdots - c_1 x - c_0)$, so we have

$$(n+k)\lambda_i^{n+k} = (n+k-1)c_{n+k-1}\lambda_i^{n+k-1} + \cdots + (n+1)c_1\lambda_i^{n+1} + c_0\lambda^n,$$

which allows relations involving linear polynomials $f_i(n)$ to be preserved by induction. The same reasoning applies to higher order $f_i(n)$ when λ_i is a higher-order root.

Example

Let $\{a_n\}_{n\geq 1}$ be the sequence defined by

$$a_1 = -3;\ a_2 = 0;\ a_3 = -15 \text{ and}$$

$$a_{n+3} = 3a_{n+1} - 2a_n.$$

The characteristic equation is

$$x^3 - 3x + 2 = 0$$

whose roots are $r_1 = 1$ (twice) and $r_2 = -2$. Thus we search for a formula

$$a_n = (An + B)r_1^n + Cr_2^n,$$

with $A, B, C \in \mathbb{C}$. We find that $A = -3$, $B = 2$ and $C = 1$; so

$$a_n = -3n + 2 + (-2)^n.$$

Question

Given a linear recurrence sequence $\{a_n\}$, what is the set of $n \in \mathbb{N}$ such that $a_n = 0$?
Looking at the above example, this leads to solving the equation

$$-3n + 2 + (-2)^n = 0$$

and it is easy to check that the only solution is $n = 2$. However, in general the above question might be more challenging if there are more roots r_i of the characteristic equation which have the same largest absolute value.

"Analytic functions"

Assume that each r_i is a positive real number; then

$$F(x) := \sum_{i=1}^{m} f_i(x) r_i^x$$

is a real analytic function. So, the question is when $F(x) = 0$ (especially, for which values x which are positive integers). Still this doesn't solve the problem, but it provides the motivation for the Skolem's method which will solve the problem. Rather than work over $\mathbb{R}$, we will work over "p-adic fields" where:

1. The integers $\mathbb{Z}$ are actually in the unit disc (very different from $\mathbb{R}$ or $\mathbb{C}$.)
2. Non-zero convergent power series have finitely many zeros (as over $\mathbb{R}$ or $\mathbb{C}$).

Then writing $F(x)$ as a p-adic analytic power series will solve our problem completely.

p-adic valuation

Let p be a prime number. For each nonzero integer n we denote by $v_p(n)$ the exponent of p in n. We extend the function v_p to all nonzero rational numbers a/b by letting $v_p(a/b) = v_p(a) - v_p(b)$.
Example: $v_3(7/9) = -2$

p-adic norm

For a prime number p and any nonzero rational number x we let

$$|x|_p := p^{-v_p(x)}$$

be the p-adic norm of x. It's easy to show that

$$|x + y|_p \leq \max\{|x|_p, |y|_p\}$$

with equality if $|x|_p \neq |y|_p$. This allows us to define a metric on $\mathbb{Q}$ by letting the distance between two rational numbers x and y be $|x - y|_p$.

p-adic numbers

We let $\mathbb{Q}_p$ be the completion of $(\mathbb{Q}, |\cdot|_p)$. These are the p-adic numbers. The set of all $x \in \mathbb{Q}_p$ such that $|x|_p \leq 1$ is the set of p-adic integers $\mathbb{Z}_p$. Each p-adic number z can be uniquely written as $p^a x$, where $x \in \mathbb{Z}_p$ and $a := v_p(z)$. If $|x|_p = 1$, then x is called a p-adic unit. For any p-adic unit x there exists a positive integer m less than p such that $|x^m - 1|_p < 1$.
Each $x \in \mathbb{Z}_p$ is uniquely written as an infinite series in powers of p:

$$x = \sum_{i=0}^{\infty} c_i p^i,$$

where each $c_i \in \{0, 1, \ldots, p-1\}$. With the above notation, we note that m may be obtained such that $c_0^m \equiv 1 \pmod{p}$.
Using the above expansion, we obtain that $\mathbb{Q}_p$ has uncountably many numbers. Hence, $\operatorname{trdeg}_{\mathbb{Q}} \mathbb{Q}_p$ is uncountable.

p-adic analytic functions

Let $a \in \mathbb{Q}_p$ and $r > 0$ be a positive real number. We say that $f : \mathbb{Q}_p \longrightarrow \mathbb{Q}_p$ is a p-adic analytic function on the open ball $D(a; r)$ if it has a power series expansion

$$\sum_{i=0}^{\infty} c_i (x - a)^i$$

which is convergent for all $x \in D(a; r)$.

Theorem

If f is p-adic analytic on $D(a; r)$, then its zeros cannot accumulate at a point inside $D(a; r)$.

The p-adic exponential and logarithmic function

The function

$$\exp_p(x) = \sum_{n=0}^{\infty} \frac{x^n}{n!}$$

is convergent p-adically for all $x \in p^2 \cdot \mathbb{Z}_p$. Its inverse is the p-adic logarithmic function

$$\log_p(1+x) = \sum_{n=1}^{\infty} (-1)^{n-1} \frac{x^n}{n}$$

which is convergent on $p \cdot \mathbb{Z}_p$. This allows us to define the p-adic analytic function $f(x) = (1+a)^x$ for each $x \in p \cdot \mathbb{Z}_p$, where a is a given number in $p \cdot \mathbb{Z}_p$. Indeed, we let

$$(1+a)^x = \exp_p(x \cdot \log_p(1+a)).$$

What is our goal?

For a given *finite* set S of nonzero complex numbers, we want to find a prime number p and an embedding σ of the finitely generated field $\mathbb{Q}(S)$ into $\mathbb{Q}_p$ such that for each $x \in S$, we have $\sigma(x) \in \mathbb{Z}_p$.

Why?

We want this since then we can choose the set S be the finite set containing all (nonzero) coefficients of the polynomials f_i and all r_i's, and also the inverses of all of these numbers. We also choose S to contain the inverses of these numbers so that the embedding we found actually sends each number of S into a p-adic unit.

So, assume we obtained the above embedding. We can let then M be a positive integer such that for each $i = 1, \ldots, m$ we have

$$|r_i^M - 1|_p < 1.$$

Then for each $j = 0, \ldots, M-1$, the function

$$F_j(x) := \sum_{i=1}^{m} f_i(Mx + j) r_i^{Mx+j} = \sum_{i=1}^{m} f_i(Mx + j) r_i^j \cdot \left(r_i^M\right)^x$$

is a p-adic analytic function. Hence we split the sequence $\{a_n\}$ into M subsequences and for each one of them we found a parametrization by a p-adic analytic function. Using the Theorem on zeros of p-adic analytic functions we conclude that the set of integers n for which $a_n = 0$ is a finite union of arithmetic progressions, where (for us) a single number is the arithmetic progression of ratio 0.

An example of embedding

Say that we want to embed into a suitable $\mathbb{Z}_p$ the following set of numbers $S := \left\{\frac{2}{51}, \pi, e, \sqrt[3]{7}, i\right\}$. We proceed one number from S at a time.

As long as p does not divide 51, we can always embed $\frac{2}{51}$ into $\mathbb{Z}_p$; so all prime numbers p work except for 3 and 17.

For the numbers π and e, we can easily find two algebraically independent (over $\mathbb{Q}$) elements u and v of $\mathbb{Q}_p$ (for *any* prime p) since $\mathrm{trdeg}_{\mathbb{Q}}\, \mathbb{Q}_p$ is infinite. Then we can find an isomorphism $\sigma : \mathbb{Q}(\pi, e) \longrightarrow \mathbb{Q}(u, v)$ such that $\sigma(\pi) = u$ and $\sigma(e) = v$.

It's a bit harder to find a suitable prime p for the desired embedding of $\sqrt[3]{7}$ and of i.

Embedding of algebraic numbers into $\mathbb{Q}_p$

If $i \in \mathbb{Z}_p$ then it has a p-adic expansion

$$i = c_0 + c_1 \cdot p + \cdots .$$

In particular, $c_0^2 + 1 \equiv 0 \pmod{p}$. So we are searching for prime numbers p such that the congruence equation

$$x^2 \equiv -1 \pmod{p}$$

is solvable in integers. Immediately this yields that $p \equiv 1 \pmod{4}$. Conversely, if $p \equiv 1 \pmod{4}$, then we can solve for $c_0, c_1, \ldots$ so that $i \in \mathbb{Z}_p$. For proving this we use Hensel's Lemma.

Hensel's Lemma

Theorem

Let $f \in \mathbb{Z}[x]$, *let* $x_0 \in \mathbb{Z}$ *such that* $f(x_0) \equiv 0 \pmod{p}$ *and* $f'(x_0) \not\equiv 0 \pmod{p}$. *Then there exists* $z_0 \in \mathbb{Z}_p$ *such that* $z_0 - x_0 \in p \cdot \mathbb{Z}_p$ *and* $f(z_0) = 0$.

The proof constructs progressively better approximations of z_0, i.e., we construct $\{x_n\}_{n \geq 1}$ such that for each $n \geq 1$ we have

$$x_n \equiv x_{n-1} \pmod{p^n} \text{ and } f(x_n) \equiv 0 \pmod{p^{n+1}}.$$

Then z_0 is the limit (in $\mathbb{Z}_p$) of the above sequence $\{x_n\}$.

So, going back to finding a suitable prime number p so that both i and $\sqrt[3]{7}$ embed into $\mathbb{Z}_p$, we deal now with the latter.
If $\sqrt[3]{7} = c_0 + c_1 p + \cdots \in \mathbb{Z}_p$, then

$$c_0^3 \equiv 7 \pmod{p}.$$

Furthermore, as long as p is neither 3 nor 7, then Hensel's Lemma can be applied to show that $\sqrt[3]{7} \in \mathbb{Z}_p$. So all we need is to find a prime number p such that the congruence equation

$$x^3 \equiv 7 \pmod{p}$$

is solvable. If $p \equiv 2 \pmod 3$, then *any* integer is a perfect cube modulo p.

Conclusion of the example

So, as long as $p \notin \{3, 7, 17\}$ and in addition,

$$p \equiv 1 \pmod 4 \text{ and } p \equiv 2 \pmod 3$$

then each of the numbers from $S = \left\{ \frac{2}{51}, \pi, e, \sqrt[3]{7}, i \right\}$ can be embedded into $\mathbb{Z}_p$. Next we have to show that this can be done simultaneously.
Yes, we can do this since for each distinct $x, y \in S$, $\mathbb{Q}(x) \cap \mathbb{Q}(y) = \mathbb{Q}$ (the intersection being taken in $\mathbb{C}$).

Skolem-Mahler-Lech Theorem

The above always works, we can always embed any arithmetic sequence in $\mathbb{Q}_p$. So we obtain the following.

Theorem

Let $(a_n)_{n=1}^{\infty}$ be a linear recurrence sequence. Then the set of positive integers n such that $a_n = 0$ is a finite union of arithmetic sequences.

An arithmetic sequence is a set of integers of the form $\ell, \ell + M, \ell + 2M, \ldots..$ We allow $M = 0$, which gives a "singleton sequence". The M here is the power we must raise elements to so that we can take their p-adic logarithm.

Algebraic dynamics

For any algebraic variety X endowed with a self-map Φ, we denote by Φ^n the n-th iterate of Φ (for each $n \in \mathbb{N}$). For any $\alpha \in X$ we denote by $\mathrm{Orb}_\Phi(\alpha)$ the orbit of α under Φ, i.e., the set of all $\Phi^n(\alpha)$ for nonnegative integers n.

Example. $X = \mathbb{A}^1$, $\Phi(x) = x^2 + 1$ and $\alpha = 0$. Then

$$\mathrm{Orb}_\Phi(\alpha) = \{0, 1, 2, 5, 26, 676, \dots\}$$

If $\mathrm{Orb}_\Phi(\alpha)$ is finite, then we call the point α preperiodic.

One geometric reformulation of Skolem-Mahler-Lech

The following comes directly from the proof of the Skolem-Mahler-Lech theorem.

Theorem

Let $\Phi : \mathbb{C}^N \longrightarrow \mathbb{C}^N$ be a linear map, let $\mathbf{z} \in \mathbb{C}^N$, and let V be a subvariety of $\mathbb{C}^N$. Then the set of n such that $\Phi^n(\mathbf{z}) \in V$ is a finite union of arithmetic sequences.

Recall that a subvariety of $\mathbb{C}^N$ is simply the zero set of a set of n-variable polynomials with coefficients in $\mathbb{C}$. In this context we usually write $\mathbb{C}^N$ as $\mathbb{A}^N(\mathbb{C})$.

Another geometric reformulation of Skolem's result

Theorem

Let $\Phi : \mathbb{A}^N \longrightarrow \mathbb{A}^N$ *be given by*

$$\Phi(x_1, \ldots, x_N) = (f_1(x_1), \ldots, f_N(x_N))$$

where each f_i *is an affine linear map. Then for each* $\alpha \in \mathbb{A}^N(\mathbb{C})$ *and for each subvariety* V *of* $\mathbb{A}^N$*, the set of all* $n \in \mathbb{N}$ *such that* $\Phi^n(\alpha) \in V(\mathbb{C})$ *is a finite union of arithmetic progressions.*

It is easy to find a formula for the n-th iterate of a number $c \in \mathbb{C}$ under a linear map $f(x) = ax + b$.

Case 1. If $a = 1$, then $f^n(c) = c + nb$.

Case 2. If $a \neq 1$, then $f^n(c) = c \cdot a^n + b \cdot \frac{a^n - 1}{a - 1}$.

Now, a point $(\gamma_1, \ldots, \gamma_N)$ in $\mathbb{A}^N$ lies on a given affine subvariety V if for each $F(x_1, \ldots, x_N)$ in the vanishing ideal of V, we have $F(\gamma_1, \ldots, \gamma_N) = 0$.
Using the fact that we have explicit formulas for the n-th iterate of a number under a linear map, we immediately get that the equation

$$F\left(f_1^n(\alpha_1), \ldots, f_N^n(\alpha_n)\right) = 0$$

translates into an equation of the form

$$\sum_{i=1}^{m} g_i(n) r_i^n = 0$$

for some polynomials g_i and some numbers r_i. Then Skolem's Theorem finishes the problem. Note that during our proof we actually obtain a p-adic parametrization of the orbit of α under Φ.

Extension to automorphisms of $\mathbb{A}^N$

The above morphisms given by the coordinatewise action of a linear map are automorphisms of $\mathbb{A}^N$. However there are more complicated automorphisms of $\mathbb{A}^N$; for example

$$\Phi(x, y) = (y + x^2, -2x)$$

is an automorphism of $\mathbb{A}^2$ with inverse

$$\Psi(x, y) = (-y/2, x - y^2/4).$$

It's still possible to find suitable p-adic parametrizations of the orbit of any point in $\mathbb{A}^N$ under an automorphism Φ of $\mathbb{A}^N$. The key is that the determinant of the Jacobian of an automorphism of $\mathbb{A}^N$ is a nonzero **constant**.

Preliminaries

So, we have an automorphism of $\mathbb{A}^N$ given by

$$\Phi(x_1, \ldots, x_N) := (F_1, \ldots, F_N)$$

where each $F_i \in \mathbb{C}[x_1, \ldots, x_N]$. We choose a *suitable* prime number p such that

(1) each coefficient of each F_i embeds into $\mathbb{Z}_p$; and

(2) under this embedding, the Jacobian of Φ is mapped to a p-adic unit.

Strategy – informal

Suppose that there is a prime p and a modulus M such that for each congruence class i modulo M, there is a p-adic analytic map

$$\theta_i : \mathbb{Z}_p \longrightarrow \mathbb{A}(\mathbb{Q}_p) \quad \text{such that} \quad \theta_i(k) = \Phi^{i+Mk}(\mathbf{z}).$$

For each polynomial G that vanishes on V, we have that

$$G \circ \theta_i : \mathbb{Z}_p \longrightarrow \mathbb{Z}_p$$

is a one variable analytic power series. Thus $G \circ \theta_i$ either has at most finitely many zeroes or is itself *identically zero*. Since for $n = Mk + i$, we have $A^n(\mathbf{z}) \in V$ if and only if $G \circ \theta_i(k) = 0$ for all G vanishing on V, we see that either

- $\Phi^{i+Mk}(\mathbf{z}) \in V$ **at most for finitely many** k (these give "singleton" cosets); *or*
- $\Phi^{i+Mk}(\mathbf{z}) \in V$ **for all** k (these give infinite cosets).

Construction of the p-adic analytic functions

Lemma

Let $\Phi := (F_1, \ldots, F_N) : \mathbb{Z}_p^N \longrightarrow \mathbb{Z}_p^N$ be a surjective polynomial map whose Jacobian has constant determinant which is a p-adic unit. Then there exists a positive integer k such that $\Phi^k := (H_1, \ldots, H_N)$ has the following two properties:

(i) *$H_i(z_1, \ldots, z_N) \equiv z_i \pmod{p}$ for all $z_1, \ldots, z_N \in \mathbb{Z}_p$; and*

(ii) *for each $\overline{z} := (z_1, \ldots, z_N) \in \mathbb{Z}_p^N$, the Jacobian of Φ^k at $\overline{z}$ is of the form $I_N + p \cdot M_{\overline{z}}$, for some matrix $M_{\overline{z}}$ with entries in $\mathbb{Z}_p$.*

The first part is easy since Φ induces a bijective action on $\mathbb{F}_p^N$, and thus a suitable power Φ^j induces the identity map on $\mathbb{F}_p^N$. Now, we let m be the order of the reduction of $J(\Phi^j, \cdot)$ modulo p seen as an element of $\mathrm{GL}_N(\mathbb{F}_p)$. We have

$$J(\Phi^{jm}, \overline{z}) = J(\Phi^j, \overline{z}) \cdot J(\Phi^j, \Phi^j(\overline{z})) \cdots J(\Phi^j, \Phi^{j(m-1)}(\overline{z})).$$

By our choice of j, we have $J(\Phi^j, \Phi^{ji}(\overline{z})) \equiv J(\Phi^j, \overline{z}) \pmod{p}$ for all i; hence $k = j \cdot m$ works for the above Lemma.

Construction of the p-adic analytic functions (continued)

Lemma

Let $\Psi := (H_1, \ldots, H_N) : \mathbb{Z}_p^N \longrightarrow \mathbb{Z}_p^N$ *be a polynomial map satisfying:*

(i) $H_i(z_1, \ldots, z_N) \equiv z_i \pmod{p}$ *for all* $z_1, \ldots, z_N \in \mathbb{Z}_p$*; and*

(ii) *for each* $\overline{z} := (z_1, \ldots, z_N) \in \mathbb{Z}_p^N$*, the Jacobian of* Ψ *at* $\overline{z}$ *is of the form* $I_N + p \cdot M_{\overline{z}}$*, for some matrix* $M_{\overline{z}}$ *with entries in* $\mathbb{Z}_p$*.*

Then for each given point $\overline{z_0} \in \mathbb{Z}_p^N$*, there exist* $g_1, \ldots, g_N \in \mathbb{Q}_p[[x]]$ *which are analytic on* $\mathbb{Z}_p$ *such that*

(1) $(g_1(0), g_2(0) \cdots, g_N(0)) = \overline{z_0}$*; and*

(2) $g_i(z+1) = H_i(g_1(z), \cdots, g_N(z))$ *for each* $i = 1, \ldots, N$ *and for each* $z \in \mathbb{Z}_p$*.*

The above lemma shows that

$$\Psi^n(\overline{z_0}) = (g_1(n), g_2(n), \cdots g_N(n)),$$

which is the desired p-adic analytic parametrization for our orbit.

Generalized Skolem-Mahler-Lech

The above method can be generalized to any unramified self-map on any quasiprojective variety. A map is said to be unramified if the Jacobian never vanishes.

Theorem

Let X be a quasiprojective variety defined over $\mathbb{C}$, let Φ be an unramified endomorphism of X, let $\alpha \in X(\mathbb{C})$ and let V be a subvariety of X. Then the set of all $n \in \mathbb{N}$ such that $\Phi^n(\alpha) \in V(\mathbb{C})$ is a finite union of arithmetic progressions.

Examples: endomorphisms of group varieties, automorphisms of K3 surfaces.... What about when a map is not unramified?

Why do constant Jacobians help?

Here's the (admittedly oversimplified) idea...
If the Jacobian is constant, one can take its p-adic logarithm at all but finitely many p, after raising to a sufficient power. Consider

$$A = \begin{pmatrix} 4 & 0 \\ 0 & 7 \end{pmatrix}$$

Then we can take p-adic logarithms of suitable powers of A at all primes but 4 and 7. For example, at 3 we obtain

$$A^n = \begin{pmatrix} \exp_3(n \log_3 4) & 0 \\ 0 & \exp_3(n \log_3 7) \end{pmatrix}$$

Something more sophisticated will work whenever the Jacobian is non-vanishing.

p-adic dynamics in one variable

Let $f(x) \in \mathbb{Z}$. Choose p prime to the degree of f such that all the nonzero coefficients of f are p-adic units.
Then the p-adic dynamics of f are surprisingly simple.
First look at the level of "residue classes modulo p".
Within each periodic residue class, there are three possible behaviors: indifferent, attracting, and superattracting Each gives different sorts of functions. (Note: there are no "repelling periodic points" so no chaos.)

Analytic uniformization of the orbit of a single variable analytic function

Definition

Let p be a prime. Let $U = p\mathbb{Z}_p$. Let $f : U \to U$ be a function such that

$$f(t) = \sum_{i \geq 0} c_i t^i \in \mathbb{Q}_p[[t]],$$

with $|c_0|_p < 1$, $|c_1|_p = 1$, and $|c_i|_p \leq 1$ for all $i \geq 1$. Then we say U is a quasiperiodicity disk *for f.*

The conditions on f in the above Definition mean precisely that f is p-adic analytic and maps U bijectively onto U. U is called a quasiperiodicity domain of f (in the sense of Rivera-Letelier).
In particular, the preperiodic points of f in U are in fact periodic.

The indifferent case – theorem

Although quasiperiodicity domains are complicated in general, if we put in one extra assumption – that c_1 is congruent to 1 modulo p, we get the following very simple theorem.

Theorem

(Rivera-Letelier, BGT) Pick $z \in U$ as above. There exists a bijective analytic map $\sigma : p\mathbb{Z}_p \longrightarrow U$ such that

$$f^k(z) = \sigma(\sigma^{-1}(z) + k).$$

In other words, z moves around as if being translated!
Note that σ depends on z.

Example of constructing the parametrization in the indifferent case

Say that $f(z) = \lambda z$, where $|\lambda - 1|_p < 1$. We want to find a function $u(z)$ such that $f(u(z)) = u(z+1)$. We let $\mu = \log(\lambda)$ and $u(z) = \exp(\mu z)$ would work.

Example of the indifferent case

Let $f(x) = 7 + 3x + 4x^2 - 5x^4 + 2x^5$ and consider the orbit of 0 under $f(x)$.
Then for $p = 7$, $D(0;1)$ is a domain of quasiperiodicity for $f(x)$ and therefore there exists a 7-adic uniformization map $u(x)$ as above.
Furthermore, since $f^6(x)$ has the linear term $729 \equiv 1 \pmod 7$, we get that $f^6(x)$ is conjugated through a 7-adic analytic map $u : \mathbb{Z}_7 \longrightarrow \mathbb{Z}_7$ to $x \mapsto x + 1$, i.e.,

$$f(u(x)) = u(x + 1).$$

So, for each $n \in \mathbb{N}$ we have

$$f^n(0) = u(n + \alpha_0),$$

where $u(\alpha_0) = 0$.

Example of the attracting case

Let $f(x) = 3x + 4x^2 - 5x^4 + 2x^5$ and consider the orbit of 1 under $f(x)$. 7-adically we are in the previous scenario, but what about 3-adically?
One can prove the existence of a point $y \in 3 \cdot \mathbb{Z}_3$ such that $f(y) = y$; this is Hensel's Lemma again. Then we can show that $f(x)$ is conjugated through a 3-adic analytic function $u(x)$ to $x \mapsto 3x$, i.e.,

$$f(u(x)) = u(3x).$$

Hence for each $n \in \mathbb{N}$, we have

$$f^n(1) = u(3^n \cdot \alpha_0),$$

where $u(\alpha_0) = 1$.

Example of the superattracting case

Let $f(x) = 4x^2 - 5x^4 + 2x^5$ and consider the orbit of 3 under $f(x)$. It is immediate to see that the orbit converges to 0 in the 3-adic topology.
One can prove that $f(x)$ is conjugated through a 3-adic analytic function $u(x)$ to $x \mapsto x^2$, i.e.,

$$f(u(x)) = u(x^2).$$

Hence for each $n \in \mathbb{N}$, we have

$$f^n(3) = u\left(\alpha_0^{2^n}\right),$$

where $u(\alpha_0) = 3$.

Parametrization at an attracting fixed point – theorem

Let $f(z) = a_0 + a_1 z + a_2 z^2 + \cdots \in \mathbb{Z}_p[[z]]$ be a nonconstant power series with $|a_0|_p, |a_1|_p < 1$. Then there is a point $|y|_p < 1$ such that $f(y) = y$, and $\lim_{n\to\infty} f^n(z) = y$ for all $|z|_p < 1$. In other words, an attracting point. . Write $\lambda = f'(y)$; then $|\lambda|_p < 1$. We may change coordinates to make the fixed point 0. Let $U = p\mathbb{Z}_p$. Suppose that $\lambda \neq 0$. Then we have

Theorem

There is a bijective analytic map $\sigma : p\mathbb{Z}_p \longrightarrow U$ such that

$$f^k(z) = \sigma(\lambda^k \sigma^{-1}(z))$$

for all $z \in U$.

Sketch of constructing the uniformization at an attracting point

We have

$$f(z) = \lambda z \cdot (1 + g(z)), \text{ where } g(z) = O(z).$$

So, for each positive integer n we have

$$u_n(z) := \lambda^{-n} f^n(z) = z \cdot (1+g(z))(1+g(f(z))) \cdots (1+g(f^{n-1}(z))).$$

One can prove that $\{u_n\}$ converges to an analytic function u and since

$$\lambda u_{n+1} = u_n \circ f$$

we conclude that indeed f is p-adically analytically conjugate to $z \mapsto \lambda z$ through the function $u(z)$.

Parametrization at a superattracting fixed point – theorem

Let $f(z) = a_0 + a_1 z + a_2 z^2 + \cdots \in \mathbb{Z}_p[[z]]$ be a nonconstant power series with $|a_0|_p, |a_1|_p < 1$. Then there is a point $y \in p\mathbb{Z}_p$ such that $f(y) = y$, and $\lim_{n\to\infty} f^n(z) = y$ for all $z \in D(0,1)$. In other words, an attracting point. . Write $\lambda = f'(y)$; then $|\lambda|_p < 1$. We may change coordinates to make the fixed point 0. Let $U = p\mathbb{Z}_p$. Suppose now that $\lambda = 0$, and let a_m be the first nonzero coefficient of f. Then we have

Theorem

There is a bijective analytic map $\sigma : p\mathbb{Z}_p \longrightarrow U$ such that

$$f^k(z) = \sigma\left((\sigma^{-1}(z))^{m^k}\right).$$

for all $z \in U$.

Sketch of constructing the uniformization at a superattracting point

After conjugating by a suitable map $z \mapsto cz$ we may assume that

$$f(z) = z^m \cdot (1 + g(z)), \text{ where } g(z) = O(z).$$

We compute then for each positive integer n:

$$f^n(z) = z^{m^n} \cdot (1+g(z))^{m^{n-1}} \cdot (1+g(f(z)))^{m^{n-2}} \cdots (1+g(f^{n-1}(z))).$$

One can prove that the function

$$z \mapsto \sum_{n \geq 0} \frac{\log(1 + g(f^n(z)))}{m^{n+1}}$$

converges uniformly to another function $h(z)$.

Sketch of constructing the uniformization at a superattracting point (continued)

Since

$$h(f(z)) = mh(z) - \log(1 + g(z))$$

we get that

$$\log\left(\frac{f(z)}{z^m}\right) = mh(z) - h(f(z))$$

and letting $u(z) = z \cdot \exp(h(z))$ yields

$$u(f(z)) = u(z)^m.$$

Indifferent is good

In the above, the indifferent U are much more complicated, but...they do give us p-adic analytic uniformization.
Look at the underlying maps:

1. Indifferent: we have $f^k(z)$ looks like $z + k$ (obviously p-adic analytic in k).
2. Attracting: we have $f^k(z)$ looks like $\lambda^k z$, where $|\lambda|_p < 1$, so we cannot take a logarithm, so not p-adic analytic.
3. Superattracting, we have $f^k(z)$ looks like (z^{m^k}), also not analytic in k when $|z|_p < 1$.

One final note: the translation for indifferents comes from "taking a second log". You start with $\lambda^k z$, but taking $\log_p$ makes this $k \log_p \lambda + z$. Then rescale to get 1.

The Dynamical Mordell-Lang Conjecture

Conjecture

Let X be a quasiprojective variety defined over $\mathbb{C}$, let Φ be an endomorphism of X, let $\alpha \in X(\mathbb{C})$ and let V be a subvariety of X. Then the set of all $n \in \mathbb{N}$ such that $\Phi^n(\alpha) \in V(\mathbb{C})$ is a finite union of arithmetic progressions.

We've seen that the result holds for all unramified maps. The key lies in finding a suitable prime p and a p-adic parametrization of the orbit. This works when the orbit ends up in an indifferent cycle modulo p.

When we deal with an attracting cycle modulo p, the method *might* work as long as we have a *single* multiplier, i.e., the Jacobian of Φ at a point of the orbit is a multiple λ of the identity matrix, where $|\lambda|_p < 1$.

Also the superattracting case works under restricted hypothesis.

Why Mordell-Lang?

Assume $X = \mathbb{A}^N$, $\alpha = (1, \ldots, 1)$ and

$$\Phi(x_1, \ldots, x_N) = (a_1 x_1, \ldots, a_N x_N)$$

for some nonzero numbers a_i. Then the above Conjecture reduces to the problem of understanding what is the intersection between an affine subvariety V and a cyclic multiplicative group Γ spanned by $(a_1, \ldots, a_N)$. This is a special case of a Theorem of Laurent.

Laurent's Theorem

For each field K, $\mathbb{G}_m(K) = (K^*, \cdot)$ is the set of its nonzero numbers, with the usual multiplication operation.

Theorem

Let $F \in \mathbb{C}[x_1, \ldots, x_N]$ and let Γ be a finitely generated subgroup of $\mathbb{G}_m^N$. Then the set of all $\gamma \in \Gamma$ such that $F(\gamma) = 0$ is a finite union of cosets of subgroups of Γ.

One can reduce to the case F is a linear polynomial, but still the problem remains difficult. For example, Laurent's Theorem implies that there are *only finitely many* positive integers m and n such that

$$3^m - 2^n = 1.$$

This is the case since the curve $x - y = 1$ contains no positive dimensional algebraic subgroup of $\mathbb{G}_m^2$.

$3^m - 2^n = 1$

Such diophantine exponential equations are difficult since they involve more than 1 variable. There are some obvious solutions: $m = n = 1$ and $(m, n) = (2, 3)$. **Are there any other solutions? No.**

But what about the equation

$$3^m - 2^n = 115?$$

Are there any other solutions besides $m = 5$ and $n = 7$?

The beginnings

Laurent's Theorem is a special case of the classical Mordell-Lang Conjecture. On the other hand, everything started with Mordell's Conjecture 83 years ago.
Essentially, Mordell's Conjecture asks that a curve which is not sufficiently simple (i.e., its genus is larger than 1) it cannot have infinitely many rational points.
Note that curves of genus 0 are parametrizable and thus they contain infinitely many rational points, while curves of genus 1 (which are the elliptic curves) *sometimes* contain infinitely many rational points.
If $F(x, y) \in \mathbb{Q}[x, y]$ is a *generic* polynomial of degree larger than 3, then the corresponding curve has genus larger than 1. For example,

$$x^4 + y^4 = 1$$

has genus 3.

Already Mordell's Conjecture implies that an equation such as

$$2^m - 3^n = 115$$

has finitely many solutions $(m, n) \in \mathbb{N} \times \mathbb{N}$. Indeed, if there were infinitely many solutions, then there would be infinitely many solutions $(m, n) \in \mathbb{N} \times \mathbb{N}$ with both m and n in fixed residue classes modulo 4 (or even some larger moduli, if needed). This yields that for some constants c_1 and c_2, the equation

$$c_1 x^4 - c_2 y^4 = 115$$

would contain infinitely many solutions $(x, y) \in \mathbb{Q} \times \mathbb{Q}$. If the genus of the above curve is larger than 1 this contradicts Mordell's conjecture.

Mordell's Conjecture

Mordell's conjecture was proven by Faltings:

Theorem

Let K be a number field, and C be a curve of genus greater than 1 defined over K. Then $C(K)$ is a finite set.

Lang's generalization

One can interpret Mordell's Conjecture as follows: we view the curve C embedded into its Jacobian J (which is an abelian variety of dimension equal to the genus of C). Then finding $C(K)$ reduces to finding the intersection between the curve C (inside J) with the finitely generated group $J(K)$ of K-rational points of J.
Mordell's Conjecture predicts that $C(\overline{K}) \cap J(K)$ is finite. Lang generalized this problem to the case of intersecting an arbitrary subvariety X of a semiabelian variety G with a finitely generated subgroup of $G(\overline{K})$. This is called the Mordell-Lang Conjecture and it was proven by Faltings in the case G is an abelian variety, and by Vojta in the general case of a semiabelian variety.

Conjecture (Mordell-Lang)

Let G be a semiabelian variety defined over $\mathbb{C}$, let $V \subset G$ be a subvariety, and let $\Gamma \subset G(\mathbb{C})$ be a finitely generated subgroup. Then $V(\mathbb{C}) \cap \Gamma$ is a finite union of cosets of subgroups of Γ.

The above conjecture is indeed a generalization of Mordell's Conjecture. One sees this by noting that if we assume the intersection $C(\overline{K}) \cap J(K)$ (taken inside the Jacobian J of the curve C) is infinite, then Mordell-Lang's Conjecture predicts that the above intersection contains a coset of a *infinite* subgroup H of $J(K)$. Since the Zariski closure of H (in J) is a (infinite) algebraic subgroup, we conclude that C is a translate of a 1-dimensional algebraic subgroup of J. However, this forces C be isomorphic to an elliptic curve, which contradicts th assumption that its genus is larger than 1.

Special points and special subvarieties

In particular, the above analysis shows that if the intersection in the Mordell-Lang problem is infinite, then the subvariety V contains a translate of an infinite algebraic subgroup of G.
This phenomenon fits into the philosophy that only *special subvarieties* may contain a Zariski dense set of *special points*. For the Mordell-Lang Conjecture, the special points are the points of Γ, while the special subvarieties are finite unions of translates of algebraic subgroups of of G. This principle of special points and special subvarieties is fundamental in other important conjectures from arithmetic geometry, such as: Manin-Mumford, Bogomolov, André-Oort, Pink-Zilber.

Limitations to possible extensions of the Mordell-Lang problem

One cannot replace G by an arbitrary algebraic subgroup defined over $\mathbb{C}$. For example, the *Mordell-Lang principle* fails in $\mathbb{G}_a^2$ due to the existence of infinitely many solutions to the Pell's equation:

$$x^2 - 2y^2 = 1.$$

Then the plane curve C given by the above Pell's equation contains infinitely many points in common with the rank-2 subgroup $\mathbb{Z} \times \mathbb{Z}$ of $\mathbb{G}_a^2(\mathbb{C})$. Also, C is not a coset of 1-dimensional algebraic subgroup of $\mathbb{G}_a^2$, and thus the intersection is not a finite union of cosets of subgroups of $\mathbb{Z} \times \mathbb{Z}$.
Similarly, we cannot (naively) extend the Mordell-Lang principle in positive characteristic. Indeed, the curve $C \subset \mathbb{G}_m^2$ defined by the equation $x + y = 1$ over $\mathbb{F}_p(t)$ contains infinitely many points in common with the cyclic subgroup $\Gamma \subset \mathbb{G}_m^2$ generated by $(t, 1-t)$. However, the intersection is not a finite union of cosets of subgroups of Γ, and also C is not a translate of a tori.

"Special points" and "special subvarieties in the Dynamical Mordell-Lang Conjecture

Conjecture

Let X be a quasiprojective variety defined over $\mathbb{C}$, let Φ be an endomorphism of X, let $\alpha \in X(\mathbb{C})$ and let V be a subvariety of X. Then the set of all $n \in \mathbb{N}$ such that $\Phi^n(\alpha) \in V(\mathbb{C})$ is a finite union of arithmetic progressions.

The special points are the points in the orbit of α under Φ, while the special subvarieties are the irreducible, Φ-periodic subvarieties of X which intersect the given orbit.

A dynamical generalization

In the Mordell-Lang Conjecture, the subgroup Γ of the semiabelian variety G can be viewed as the orbit of the identity 0 of G under a finitely generated group of translations $\Phi_1, \ldots, \Phi_r$ acting on G.
Then the conclusion would be that for any subvariety V of G, the set of all $(n_1, \ldots, n_r) \in \mathbb{Z}^r$ such that

$$\Phi_1^{n_1} \cdots \Phi_r^{n_r}(0) \in X(\mathbb{C})$$

is a finite union of cosets of subgroups of $\mathbb{Z}^r$.
It is natural to ask what happens if one replaces:

- G with an arbitrary quasiprojective variety X defined over $\mathbb{C}$;
- $\Phi_1, \ldots, \Phi_r$ with arbitrary commuting endomorphisms of G;
- 0 by any point in $X(\mathbb{C})$.

Dynamical Mordell-Lang problem Let X be a quasiprojective variety defined over $\mathbb{C}$, let V be a subvariety of X, let $\Phi_1, \ldots, \Phi_r$ be commuting endomorphisms of X, and let $\alpha \in X(\mathbb{C})$. Is the set of all $(n_1, \ldots, n_r) \in \mathbb{N}^r$ such that $\Phi_1^{n_1} \cdots \Phi_r^{n_r}(\alpha) \in V(\mathbb{C})$ a finite union of sets of the form $\gamma + H \cap \mathbb{N}^r$, where $\gamma \in \mathbb{N}^r$ and H is a subgroup of $\mathbb{Z}^r$?
As shown by the example of the Pell's equation the answer is sometimes "NO". It is interesting to note that the answer is also "NO" for certain commuting semigroups of linear transformations acting on $\mathbb{A}^N$. But we do have a positive answer in an important case of linear transformations.

Theorem

If $M_1, \ldots, M_r$ are commuting matrices which are diagonalizable, then the Dynamical Mordell-Lang problem has a positive answer.

The case of linear transformations

For the Dynamical Mordell-Lang (DML) problem, assume $X = \mathbb{A}^N$ and that each endomorphism Φ_i is linear, given by some matrix M_i; we assume $M_i M_j = M_j M_i$ for each i and j. Then the DML problem for (linear) varieties V of $\mathbb{A}^N$ reduces to finding all tuples $(n_1, \ldots, n_r) \in \mathbb{N}^r$ such that for a given vector $\alpha \in \mathbb{A}^N$ we have

$$M_1^{n_1} \cdots M_r^{n_r}(\alpha) \in V.$$

Since the matrices M_i commute, they share the same Jordan blocks, and thus without loss of generality we may assume they are all in Jordan form (at the expense of replacing each M_i by $B^{-1} M_i B$ for some suitable invertible matrix B). Furthermore, we may assume each M_i consists of a single Jordan block, i.e. there exist strictly upper triangular matrices U_i such that

$$M_i = \lambda \cdot I_N + U_i$$

where λ is the common eigenvalue of the M_i's.

Powers of a matrix

For a matrix $M := \lambda \cdot I_N + U$ where U is nilpotent, we immediately see that there exist polynomials $f_1, \ldots, f_{N-1}$ such that for all $n \in \mathbb{N}$ we have

$$M^n = \lambda^n \cdot I_N + \sum_{i=1}^{N-1} f_i(n) \cdot U^i.$$

Hence the equation

$$M_1^{n_1} \cdots M_r^{n_r}(\alpha) \in V$$

translates into a polynomial-exponential equation of *several* variables:

$$F\left(r_1^{n_1}, \ldots, r_m^{n_m}, n_1, \ldots, n_m\right) = 0$$

where the r_i's are the eigenvalues of the matrices M_i, and F is a polynomial. We've seen that when $n_1 = \cdots = n_m$, the Skolem's method applies and yields a result compatible with the Mordell-Lang principle. However, in the general case when the n_i's are distinct, then we can get much more different answers.

An example

Let

$$M_1 = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -2 \\ 0 & 0 & 1 \end{pmatrix}$$

and

$$M_2 = \begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & 4 \\ 0 & 0 & 1 \end{pmatrix}.$$

Then

$$M_1 M_2 = M_2 M_1 = \begin{pmatrix} 1 & 1 & -4 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}.$$

We let $\alpha = (3, 2, 1)$ and let V be the line given by the equations $x = 3$ and $z = 1$. We can compute

$$M_1^m M_2^n = \begin{pmatrix} 1 & 2n - m & 4n(n-1) - 4mn + m(m-1) \\ 0 & 1 & 4n - 2m \\ 0 & 0 & 1 \end{pmatrix}.$$

So, we have to find nonnegative integers m and n such that

$$4n - 2m + 4n(n-1) - 4mn + m(m-1) = 0.$$

The solutions are

$$m = 3k^2 \text{ and } n = \frac{3(k^2 \pm k)}{2}, \text{ where } k \in \mathbb{N}.$$

This doesn't fit into the Mordell-Lang principle.

The diagonal case

However when all the matrices M_i are diagonalizable, then the above equation is only exponential; the polynomial part appears only in the presence of Jordan blocks of dimension larger than 1. So, in the diagonalizable case, the DML problem reduces to Laurent's Theorem, i.e., solving equations of the form

$$F\left(r_1^{n_1}, \ldots, r_m^{n_m}\right) = 0$$

where $F \in \mathbb{C}[x_1, \ldots, x_m]$.

The linear case

An interesting special case for the DML problem is when $X = \mathbb{A}^2$, $\alpha = (x_0, y_0)$, and V is the diagonal line given by the equation $x = y$, and moreover $\Phi_1(x, y) = (f(x), y)$ while $\Phi_2(x, y) = (x, g(y))$ for some polynomials f and g with complex coefficients. Then the Dynamical Mordell-Lang problem reduces to asking for which positive integers m and n we have

$$f^m(x_0) = g^n(y_0)?$$

If one of the polynomials is linear, then the answer will not necessarily fit the Mordell-Lang principle. Indeed, if

$$f(x) = x + 1;\ g(y) = 2y;\ (x_0, y_0) = (0, 1)$$

then $f^m(x_0) = g^n(y_0)$ if and only if $m = 2^n$. If instead $g(y) = y^2$ and $y_0 = 2$, then we obtain that

$$f^m(x_0) = g^n(y_0) \text{ if and only if } m = 2^{2^n}.$$

However, if neither f nor g is a linear polynomial we can show that the Mordell-Lang principle holds for lines in the affine space.

A simple version

Theorem

(GTZ) Let $f, g \in \mathbb{C}[x]$ have degree greater than 1. If the orbit $\mathrm{Orb}_f(x_0)$ intersects the orbit $\mathrm{Orb}_g(y_0)$ in infinitely many points then there exist positive integers m and n such that $f^m = g^n$.

The full case of lines

Theorem

If $X = \mathbb{A}^N$ and $V \subset \mathbb{A}^N$ is a line, and for each $i = 1, \ldots, r$, we have

$$\Phi_i(x_1, \ldots, x_N) = (x_1, \ldots, x_{i-1}, f_i(x_i), x_{i+1}, \ldots, x_N)$$

for some polynomials f_i of degree larger than 1, then the Dynamical Mordell-Lang problem has a positive answer.

In particular, going back to the special case: $N = 2$ and V is the diagonal line, we get that if the orbit $\mathrm{Orb}_f(x_0)$ intersects the orbit $\mathrm{Orb}_g(y_0)$ in infinitely many points then there exist positive integers m and n such that $f^m = g^n$. Using our method, together with additional analytic and topological arguments, Wang proved a similar result where the polynomials f_i are replaced by finite Blaschke products (these are endomorphisms of the complex unit disk).

Idea of the proof

One can show that indeed the special case of the diagonal line in $\mathbb{A}^2$ is the general case. Assume now that x_0, y_0, and all coefficients of f and g are integers. Then knowing that $f^m(x_0) = g^n(y_0)$ for infinitely many positive integers m and n yields that for *each* pair of positive integers r and s, the curve

$$f^r(X) - g^s(Y) = 0$$

contains infinitely many integral points. Using Siegel's Theorem (together with the refinements of Bilu and Tichy) one gets that f^r and g^s must have very special forms. Using the recent results of Zieve who improved on the old results of Ritt on polynomial decomposition, one derives that the only possibility is that f and g must share a common iterate.

Siegel's Theorem

Theorem

(Siegel 1929!) If C is a plane curve defined over $\mathbb{Q}$ of positive genus, or of genus 0 but with at least 3 points at infinity, then it contains finitely many points with both coordinates integers.

Note that if $f, g \in \mathbb{Q}[z]$ are *generic* polynomials of degree larger than 2, then the curve given by the equation

$$f(x) - g(y) = 0$$

has no irreducible component which is a genus 0 curve with at most 2 points at infinity.

All of the previous results suggest that the Mordell-Lang principle holds in the dynamical setting *always* for a cyclic semigroup of endomorphisms of the ambient space, and it holds only *sometimes* when S is not cyclic.

The Dynamical Mordell-Lang Conjecture

Let X be any quasiprojective variety defined over $\mathbb{C}$, and let Φ be any endomorphism of X. Then for any subvariety $V \subset X$, and for any point $\alpha \in X(\mathbb{C})$, the set of integers $n \in \mathbb{N}$ such that $\Phi^n(\alpha) \in V(\mathbb{C})$ is a union of at most finitely many arithmetic progressions.

In particular, the above conjecture predicts that only periodic subvarieties may contain a Zariski dense set of points from a given orbit. This reformulation is in line with the principle of *special points* and *special subvarieties* outlined before: this time, the special points are the points of the orbit, while the special subvarieties are the ones which are periodic.

Cases of the Skolem approach

We will consider the following case of DML.

1. $\Phi : \mathbb{A}^N \longrightarrow \mathbb{A}^N$
2. $\Phi(x_1, \ldots, x_N) = (f_1(x_1, \ldots, x_N), \ldots, f_N(x_1, \ldots, x_N))$ where each $f_i \in \overline{\mathbb{Q}}[x_1, \ldots, x_N]$.

We will say that that DML holds for a given Φ and $\alpha \in \mathbb{A}^N(\overline{\mathbb{Q}})$ if the set of integers $n \in \mathbb{N}$ such that $\Phi^n(\alpha) \in V(\overline{\mathbb{Q}})$ is a union of at most finitely many arithmetic progressions.

The ramification locus

The ramification locus R_φ is defined by the vanishing of the Jacobian determinant $\det(D\Phi)$.
We have the following: *if there is a prime p such that the orbit of α avoids the ramification locus of R_Φ modulo p, then DML holds for Φ and α.*
The reason is that in this case we can construct a p-adic parametrization of the orbit of α.

Example

Start with the case $N = 1$. Consider $\Phi(x) = x^3 + 1$ and $\alpha = 3$. Then R_Φ is defined by the vanishing of $\Phi'(x) = 3x^2$. So R_Φ is simply the point zero.
Try the primes 2, 3, 5, 7, 11, and see when the orbit of $\alpha = 3$ avoids R_Φ.
Mod 2: $1 \mapsto 0$
Mod 3: (starts at 0)
Mod 5: $3 \mapsto 3$ (avoids)
Mod 7: $3 \mapsto 0$
Mod 11: $3 \mapsto 6 \mapsto 10 \mapsto 0$

More examples of R_Φ

Example

$\Phi(x_1, x_2) = (x_1^2 + x_2, x_1^5 + x_2^3)$. Then R_Φ is defined by the vanishing of

$$\det \begin{pmatrix} 2x_1 & 1 \\ 5x_1^4 & 3x_2^2 \end{pmatrix} = 6x_1x_2^2 - 5x_1^4$$

In particular, it is a curve. More generally, we always get that R_Φ has dimension $N - 1$ (unless it is empty, which means Φ is unramified) when we are in dimension N.

Avoiding ramification, the birthday problem, and random maps

Note that $\mathbb{A}^N$ has p^N points modulo p. By the "birthday problem", it should therefore take about $\sqrt{p^N}$ iterations for an orbit to start to cycle modulo p.

The ramification locus R_Φ has codimension 1 in $\mathbb{A}^N$, so it has p^{N-1} points modulo p. So each iterate has a $1/p$ chance of hitting R_Φ modulo p (regardless of N).

Moral: when N is large and p is large, any given orbit probably goes through R_Φ. But...you can probably avoid ramification modulo some p in dimensions 1, 2, and 3.

Empirical evidence for random maps

In dimension 1, analysis predicts an average orbit length of $\sqrt{p}$. Here is the data we obtained fixing a polynomial and starting point α and looking at the first 1000 primes. (The chart is normalized by dividing orbit length by $\sqrt{2p}$.) We graph against what a random model would predict.

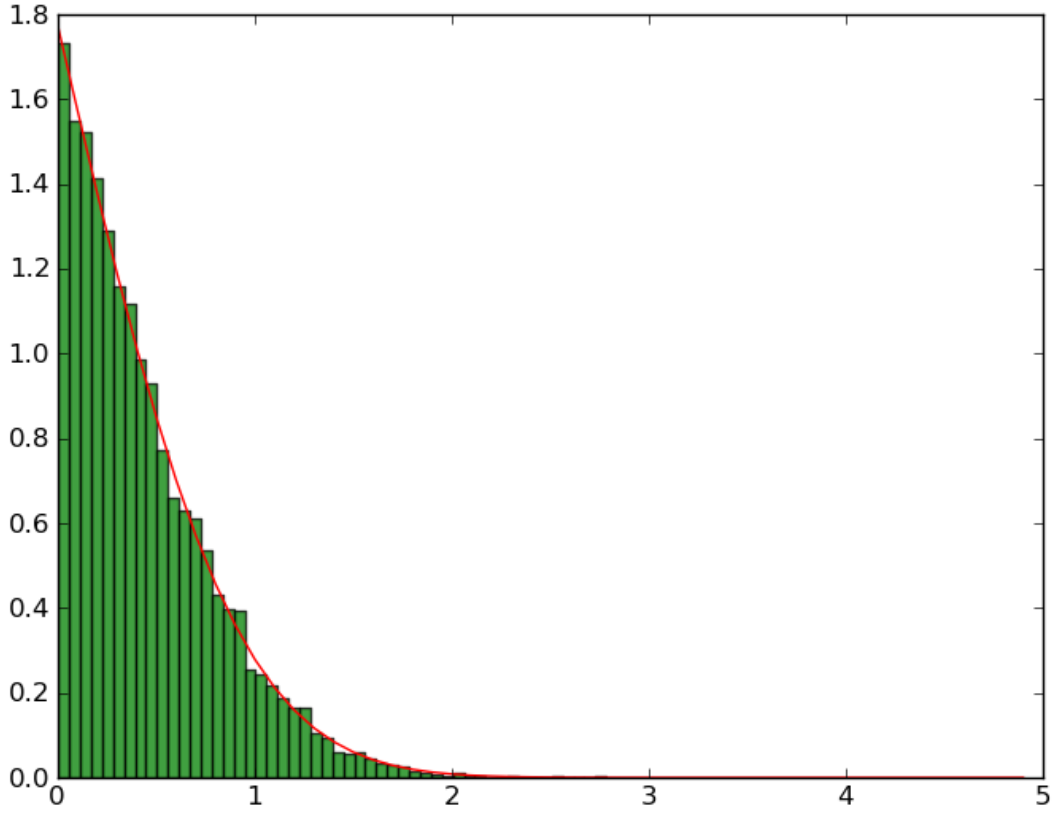

Empirical evidence for (failure of) ramification avoidance

See the chart below for the chance that a point passes through ramification modulo p (fixed starting point, varying maps of fixed degree).

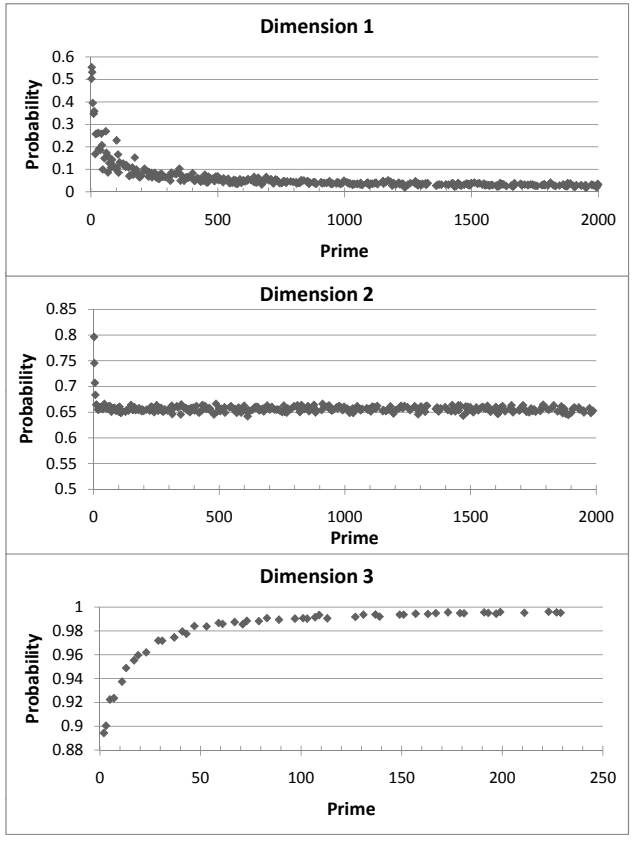

Skolemdämmerung: Twilight of the Skolem approach

Let $\mathrm{A}(p)$ be the chance that the orbit passes through R_Φ modulo p for a give p. Then, when $N = 3$, the last chart shows that

$$\lim_{p \to \infty} \mathrm{A}(p) = 1.$$

But we only need to avoid ramification at a single primes, so we're okay as long as

$$\prod_{\text{primes } p} \mathrm{A}(p) = 0$$

(which it does for $N = 3$).
When $N \geq 5$, the chance of passing through ramification at a prime p goes to 1 so quickly that

$$\prod_{\text{primes } p} \mathrm{A}(p) > 0$$

and we have empirical evidence that there are points and maps in dimension 5 where we pass through ramification at every prime p. So the Skolem approach to DML seems hopeless in general.

How bad is it when you can't avoid ramification?

Bad: When you do not avoid ramification at p, there are counterexamples if you replace *algebraic subvariety* V with *p-adic analytic subvariety* V. Note that when SLM works at a prime p, it proves a result for *p-adic analytic subvarieties*.
Good.

- We can still prove DML for algebraic subvarieties in many cases where there are p-adic analytic counterexamples.
- The counterexamples are highly transcendental, and a "dynamical transcendence theory" (suggested by Zhang) might rule these out in the algebraic case.

A special case

Consider the following special case ("split coordinates").

- polynomials $f_1, \ldots, f_N \in \overline{\mathbb{Q}}[x]$.
- $\Phi = (f_1, \ldots, f_N) : \mathbb{A}^N \longrightarrow \mathbb{A}^N$ by

$$(f_1, \ldots, f_N)(x_1, \ldots, x_N) \mapsto (f_1(x_1), \ldots, f_N(x_N))$$

Here we think we *can* avoid ramification modulo p for infinitely many p. (Note that we think we can "specialize" to get a proof over $\mathbb{C}$, via work of Medevdev-Scanlon, Silverman, Baker-Benedetto.)

Birthday revisited

Idea: in each coordinate, should have orbit length of $\sqrt{p}$, so there's only $\sqrt{p}/p = 1/\sqrt{p}$ chance of passing through a given point modulo p, so should avoid ramification mod p for a density 1 set of primes.

An intersection of g sets of primes of density 1 is still density 1, so we get lots of "ramification-avoiding" primes (so the method works) for $(f_1, \ldots, f_N)$.

A rogues gallery of counterexamples

Problem: For $f(x) = x^3 + 1$, a given α will have its orbit pass through 0 (the critical point) with density 1/4. (Empirical evidence: Roberts-Vivaldi, Towsley.)
Why?: If $p \equiv 2 \pmod 3$, then f acts as a *random permutation*, which have much longer orbits and there is a one half chance that a given orbit contains a given point. Since for 1/2 of p, we have $p \equiv 2 \pmod 3$, this gives one quarter.

Work of Schur/Fried-Guralnick-Mueller-Saxl shows that powering maps, Chebychev maps and Lattès maps are the only maps where you get permutations modulo p infinitely often (up to composition with automorphisms).

Question

Are there other rational functions with this sort of (failure of) "ramification avoidance" property?

Densities

Given a set S of primes, we define the *density* of S as

$$\lim_{M \to \infty} \frac{\#\{p \in S \text{ with } p \leq M\}}{\#\{p \leq M\}}$$

(assuming this limit exists). Here's a sample of the type of result that one gets.

Theorem

(Odoni-Stoll-Jones) Let $f(x) = x^2 + c$ where $c \neq 0, -1, -2$, and let $\alpha \in \mathbb{Z}$. Suppose that the orbit of α does not pass through 0. Let T be the set of primes such that α passes through 0 modulo p. Then T has density 0.

The main tool here is the Chebotarev density theorem, which you will treat in discussion today.

A positive result

Theorem

(BGHKST) Suppose that

- *polynomials* $f_1, \ldots, f_N \in \mathbb{Z}[x]$ *with* $f_i(x) = x^2 + c_i$.
- $\Phi = (f_1, \ldots, f_N) : \mathbb{A}^N \longrightarrow \mathbb{A}^N$ *by*

$$(f_1, \ldots, f_N)(x_1, \ldots, x_N) \mapsto (f_1(x_1), \ldots, f_N(x_N))$$

Then one can avoid ramification modulo infinitely many primes p and DML holds for any subvariety, i.e for any $\alpha \in \mathbb{A}^N(\overline{\mathbb{Q}})$ *whose orbit avoids ramification, there are infinitely many p such that the orbit of* α *avoids ramification modulo p.*

Proof.

At each c_i with $c_i \neq 0, -1, -2$, we avoid ramification with density 1. We treat the cases $0, -1, -2$ separately. □

Remark

There are similar results for other special polynomials.

A theorem of Lang

We deal now with another problem of "special points" and "special varieties". Consider the special points be points in the affine plane with both coordinates roots of unity. In order to determine the special curves in $\mathbb{A}^2$ we search for some *obvious* candidates of curves C which would contain infinitely many points with both coordinates roots of unity. It's easy to see that whenever the equation of C is

$$\alpha \cdot x^m y^n = 1,$$

for some root of unity α, then indeed C contains infinitely many points with both coordinates roots of unity. The interesting thing is that these are the only *special* ones according to the following

Theorem

If the curve $C \subset \mathbb{A}^2$ defined over $\mathbb{C}$ contains infinitely many points with both coordinates roots of unity, then the equation of C is $x^m y^n = \alpha$, for some $m, n \in \mathbb{Z}$ and some root of unity α.

Proof of Lang's Theorem in a special case

Assume C is a curve defined by an equation: $y = g(x)$ for some rational function g. Furthermore, using a multiplicative translation by a root of unity, we may assume that $g(1) = 1$.
Now, since C contains infinitely many points whose coordinates are roots of unity, we conclude that C is defined over a cyclotomic extension of $\mathbb{Q}$ (since, on one hand, it is defined over the infinite extension of $\mathbb{Q}$ generated by all roots of unity, and on the other hand, it is defined over a finitely generated extension of $\mathbb{Q}$). So, $g \in \mathbb{Q}(\zeta_m)(x)$ and let now $\{\zeta_{n_k}\}_k$ an infinite sequence of primitive roots of unity of order n_k with the property that $g(\zeta_{n_k})$ is also a root of unity.
Obviously, we may assume that $n_1 < n_2 < \cdots < n_k < \cdots$.

Proof (continued)

For each $k \in \mathbb{N}$ let σ_k be an automorphism of $\bar{\mathbb{Q}}$ such that

$$\sigma_k(\zeta_{n_k}) = e^{2\pi i/n_k}.$$

Since $\mathbb{Q}(\zeta_m)/\mathbb{Q}$ is a Galois extension, at the expense of replacing the sequence $\{n_k\}_k$ with a subsequence, we may assume that for *each* $k \in \mathbb{N}$, the restriction of σ_k on $\mathbb{Q}(\zeta_m)$ is the same automorphism σ. Furthermore, at the expense of replacing g with g^σ, we obtain that for each $k \in \mathbb{N}$ we have

$$\sigma_k(g(\zeta_{n_k})) = g\left(e^{2\pi i/n_k}\right) \text{ is a root of unity.}$$

Furthermore $g\left(e^{2\pi i/n_k}\right) \in \mathbb{Q}(\zeta_m, \zeta_{n_k}) \subset \mathbb{Q}(\zeta_{2mn_k})$, which means that there exists some $d_k \in \{0, 1, \ldots, 2mn_k - 1\}$ such that

$$g\left(e^{2\pi i/n_k}\right) = e^{2d_k\pi i/2mn_k}.$$

Clearly $e^{2\pi i/n_k} \to 1$ as $k \to \infty$, and since $g(1) = 1$, we get that also

$$2^{2d_k\pi i/2mn_k} \to 1 \text{ as } k \to \infty.$$

At the expense of replacing g by $1/g$, we may assume that

$$\frac{d_k}{mn_k} \to 0 \text{ as } k \to \infty.$$

Complex analysis

On the other hand, using basic complex analysis, there exists some positive real number C such that for all sufficiently large $k \in \mathbb{N}$ we have

$$|g\left(2^{2\pi i/n_k}\right) - g(1)| \leq C \cdot |e^{2\pi i/n_k} - 1| \leq \frac{2C\pi}{n_k}.$$

Using that $g(1) = 1$ and that $g\left(e^{2\pi i/n_k}\right) = e^{2d_k\pi i/2mn_k}$ we conclude that also

$$|e^{2d_k\pi i/2mn_k} - 1| = 2\sin(d_k\pi/2mn_k) \leq \frac{2C\pi}{n_k}.$$

Using the inequality $\sin(\alpha) > \alpha/\pi$ for any angle $\alpha \in [0, \pi/2]$ allows us to conclude that for each $k \in \mathbb{N}$,

$$d_k < 2C\pi.$$

Conclusion

By the pigeonhole principle, again at the expense of replacing $\{n_k\}$ with a subsequence, we may assume that $d_k = d$ for some $d \in \mathbb{N}$, for *all* $k \in \mathbb{N}$. So,

$$g\left(e^{2\pi i/n_k}\right) = e^{2d\pi i/n_k}$$

for infinitely many $n_k \in \mathbb{N}$.
Therefore $g(x) = x^d$ as desired.

The Manin-Mumford Conjecture

We note that roots of unity are precisely the torsion points for $\mathbb{G}_m(\mathbb{C})$. Also, note that an equation of the form $x^m y^n = \alpha$ for some root of unity α cuts a curve which is a torsion translate of a 1-dimensional algebraic subgroup of $\mathbb{G}_m^2$.
So, Lang's Theorem can be reformulated (and extended) as follows.

Conjecture

If $V \subset \mathbb{G}_m^N$ is an irreducible subvariety which contains a Zariski dense set of torsion points, then V is a torsion translate of an algebraic subgroup of $\mathbb{G}_m^N$.

The above conjecture was proven by Raynaud in a more general form which deals with *semiabelian varieties*. An alternative proof to the above conjecture can be given using the results of Pila and Wilkie.

An abelian variety is a projective, connected, group variety. The simplest example is an elliptic curve. A product of abelian varieties is also an abelian variety. For any curve of positive genus g, its Jacobian is an abelian variety of dimension g.
A semiabelian variety (over $\mathbb{C}$) is an extension G of an abelian variety A by a power of the multiplicative group, i.e., we have a short exact sequence of group varieties (over $\mathbb{C}$)

$$1 \longrightarrow G_m^g \longrightarrow G \longrightarrow A \longrightarrow 0.$$

Raynaud's Theorem yields that if an irreducible subvariety V of a semiabelian variety G defined over $\mathbb{C}$ contains a Zariski dense set of torsion points, then V is a torsion translate of an algebraic subgroup of G.

A dynamical reformulation

Here's an alternative interpretation of Raynaud's Theorem: we have an abelian variety A and the multiplication-by-2-map Φ on A. Then the torsion points of A are the preperiodic points for Φ, while the torsion translates of algebraic subgroups of A are the subvarieties which are preperiodic under the action of Φ (this part follows from a result of Hindry).

So, it is natural to ask whether it is always true for a projective variety X, and an endomorphism Φ of X that if a subvariety V of X contains a Zariski dense set of preperiodic points, then V is preperiodic itself.

The above formulation is once again an instance of the philosophy that only "special" varieties may contain a Zariski dense set of "special" points.

However there are some obvious counterexamples to the above formulation as shown by the following example.
Consider the endomorphism $\Phi : \mathbb{P}^1 \times \mathbb{P}^1 \longrightarrow \mathbb{P}^1 \times \mathbb{P}^1$ be the projectivized version of the map $\varphi : \mathbb{A}^1 \times \mathbb{A}^1 \longrightarrow \mathbb{A}^1 \times \mathbb{A}^1$ given by $\varphi(x, y) = (x^2, y^3)$. Then the diagonal subvariety contains infinitely many preperiodic points (note that both $x \mapsto x^2$ and $y \mapsto y^3$ share the same set of preperiodic points) without being itself preperiodic.
Same counterexample can be constructed in general: for any abelian variety A, the endomorphism $\Phi : A \times A \longrightarrow A \times A$ given by $\Phi(x, y) = (2x, 3y)$ allows the diagonal subvariety of $A \times A$ contain a Zariski dense set of preperiodic points (note that the torsion points are dense in A) without being itself preperiodic.

The Dynamical Manin-Mumford Conjecture (first formulation)

The above problem can be easily fixed through the following terminology.

Definition

We call the endomorphism Φ of the projective variety X polarizable if there exists some ample line bundle $\mathcal{L}$ and some integer $d \geq 2$ such that $\Phi^*(\mathcal{L}) = d \cdot \mathcal{L}$.

In the examples above that would force the map Φ act with the same "weights" on both factors. So, it's natural to formulate the following conjecture (of Zhang).

Conjecture

Let Φ be a polarizable endomorphism of a projective variety X defined over $\mathbb{C}$. If a subvariety V contains a Zariski dense set of preperiodic points, then V is preperiodic.

In positive characteristic, the above statement would fail due to the presence of varieties V defined over finite fields which would contain a Zariski dense set of torsion points (assuming V is contained in an abelian variety X defined over $\overline{\mathbb{F}_p}$ under the action of the multiplication-by-2-map Φ) without being themselves preperiodic under Φ.

Introduction to height functions

Before continuing with dynamical Manin-Mumford, a detour into heights...
Height functions play a crucial role in the study of diophantine geometry. We begin with the simplest definition. Let $\alpha = a/b$ with $a, b \in \mathbb{Z}$ in lowest terms, that is there no p that divides both a and b. Then we define the (logarithmic) *height* of a/b as

$$h(a/b) = \log \max(|a|, |b|)$$

1. $h(1) = 0$.
2. $h(99/100) = \log 100$.
3. $h(0) = 0$.
4. $h(-23/4) = \log 23$.

Basic properties of height functions

For any $\alpha, \beta \in bQ$, we have

1. $h(\alpha\beta) \leq h(\alpha) + h(\beta)$.
2. $h(\alpha + \beta) \leq h(\alpha) + h(\beta) + \log 2$.
3. For any M, there are finitely many $\alpha \in \mathbb{Q}$ such that $h(\alpha) \leq M$.

Extending to other algebraic numbers

What if we wanted to extend this to $\overline{\mathbb{Q}}$?

We could intuit $h(\sqrt{2})$ to be $\frac{\log 2}{2}$.

But what if we wanted to know (or at least define!) the height of α where α is a root of $x^3 + 9x + 3$ say? There are two ways to proceed. We will take the algebraic approach. First off, we redefine our p-adic absolute values on $\mathbb{Q}$ very carefully.

$$|p|_p = p^{-1}$$

In other words,

$$|\alpha|_p = p^{-v_p(\alpha)}$$

with the convention that $v_p(0) = \infty$ so $|0|_p = 0$.

An alternate definition of height

Why this choice of normalization?
First of all, the product formula. We let $\mathbb{M}$ be the set $\{\infty, 2, 3, 5, 7, 11, 13, ...\}$. We set $|\cdot|_\infty$ to be the usual absolute value. Then for any nonzero $\alpha \in \mathbb{Q}$, we have the famous product formula of Weil

$$\prod_{v \in \mathbb{M}} |\alpha|_v = 1.$$

And crucially,

$$h(\alpha) = \sum_{v \in \mathbb{M}} \log \max(|\alpha|_v, 1).$$

From now on, we let $\log^+ z = \log \max(z, 1)$. Then the above becomes

$$h(\alpha) = \sum_{v \in \mathbb{M}} \log^+ |\alpha|_v$$

Extending to algebraic numbers

Now, let K be a number field with $[K : \mathbb{Q}] = n$. Then for any algebraically closed field F, there are n distinct field embeddings $\sigma_1, \ldots, \sigma_n$ with $\sigma_i : K \hookrightarrow F$.

Each absolute value $|\cdot|_v$ extends to the completion $\mathbb{Q}_v$ (either $\mathbb{Q}_p$ or $\mathbb{R}$) and then to $\overline{\mathbb{Q}}_v$. For any $\alpha \in K$, we may define

$$h(\alpha) = \frac{1}{[K : \mathbb{Q}]} \sum_{v \in \mathbb{M}} \sum_{\sigma_i : K \hookrightarrow \overline{\mathbb{Q}}_v} \log^+ |\sigma_i(\alpha)|_v.$$

(Note that this does not depend on K.)

Extending to algebraic numbers, continued

$$h(\alpha) = \frac{1}{[K:\mathbb{Q}]} \sum_{v \in \mathbb{M}} \sum_{\sigma_i : K \hookrightarrow \overline{\mathbb{Q}}_v} \log^+ |\sigma_i(\alpha)|_v.$$

We will not work explicitly with this construction very much. Let's do a few examples in $K = \mathbb{Q}(\sqrt{5})$.

1. $\sqrt{5}$.
2. $\frac{\sqrt{5}}{2}$.
3. $1 + \sqrt{5}$.
4. $\frac{\sqrt{5}}{3}$.

Properties of general height functions

For any $\alpha, \beta \in \overline{\mathbb{Q}}$, we have

1. $h(\alpha\beta) \leq h(\alpha) + h(\beta)$.
2. $h(\alpha + \beta) \leq h(\alpha) + h(\beta) + \log 2$.
3. For any M and D, there are finitely many $\alpha \in \overline{\mathbb{Q}}$ such that $h(\alpha) \leq M$ and $[\mathbb{Q}(\alpha) : \mathbb{Q}] \leq D$. (Northcott's theorem.) That is, there are finitely many points of bounded height and bounded degree.

Proof of Northcott's theorem

Simple:

1. Let α have bounded degree and bounded height.
2. Let $\alpha_1, \ldots, \alpha_n$ be its conjugates.
3. Let
$$F(x) = \prod_{i=1}^{n} (x - \alpha_i).$$

Then the coefficients of F have bounded height from 1. and 2. of the previous page, since they are just symmetric polynomials in the α_i

This leaves finitely many choices for each coefficient by properties of height functions on $\mathbb{Q}$.

Points of height zero

When can a point have height zero?

Reasoning dynamically, we see that if $h(\alpha) = 0$, then $h(\alpha^m) = 0$ for all m. But all $\alpha^m \in \mathbb{Q}(\alpha)$ so they all have bounded degree. So there must be only finitely many *distinct* α^m, so we must have $\alpha^n = \alpha^m$ for some $n > m$, so either

1. $\alpha = 0$; or
2. $\alpha^{n-m} = 1$ (α is a root of unity).

Equidistribution

By Northcott, we have the following: if $\lim_{n\to\infty} h(\alpha_n) = 0$, then $\lim_{n\to\infty} \deg \alpha_n = \infty$.

Some families where $\lim_{n\to\infty} h(\alpha_n) = 0$:

1. Primitive n-th roots of unity ($\alpha^n = 1$, $\alpha^m \neq 1$ for $m < n$).
2. $\sqrt[n]{2}$.
3. Solutions to $(x^n - 1)(x - 2) + 3$ (Autissier example)

Do you see a pattern?

A theorem of Bilu

Bilu proved that points with height tending to zero equidistribute on the unit circle. More specifically we have.

Theorem

(Bilu) Let $(\alpha_n)_{n=1}^{\infty}$ *be a nonrepeating sequence of points in* $\overline{\mathbb{Q}}$ *such that* $\lim_{n\to\infty} h(\alpha_n) = 0$. *Let* f *be a continuous bounded function on* $\mathbb{C}$. *Then*

$$\lim_{n\to\infty} \frac{1}{\deg \alpha_n} \sum_{\text{conjugates } \alpha_n^{\sigma} \text{ of } \alpha_n} f(\alpha_n^{\sigma}) = \int_0^1 f(e^{2\pi i z})\, dz$$

Heights in higher dimension

We can naturally extend heights to $\mathbb{G}_m^N$ by

$$h(\alpha_1, \ldots, \alpha_N) = \sum_{n=1}^{N} h(\alpha_n)$$

(by abuse of notation).
Example

1. $h(1, 3) = 0 + \log 3 = \log 3$
2. $h(4, 11) = \log 4 + \log 11$.
3. $h(i, 1) = 0$.

A point in $\mathbb{G}_m^N$ has height zero if and only if it is a root of unity in each coordinate, i.e. is a torsion point of $\mathbb{G}_m^N x(\overline{\mathbb{Q}})$.

Bogomolov conjecture

Theorem

(Zhang) Let C be an irreducible curve in $\mathbb{G}_m^2$ such that there is an infinite nonrepeating sequence $(\alpha_n)_{n=1}^{\infty}$ with $\lim_{n\to\infty} h(\alpha_n) = 0$. Then C has the form $x^i y^j = \xi$, for some $i, j \in \mathbb{Z}$ and some root of unity ξ.

This is much harder than the Serre-Ihara-Tate theorem from Lang! There are many more points with $\lim_{n\to\infty} h(\alpha_n) = 0$ than there are roots of unity.

Sketch of proof of Bogomolov à la Bilu

The proof (a sketch) is as follows: take $(\alpha_n)_{n=1}^{\infty}$ with $\lim_{n\to\infty} h(\alpha_n) = 0$ and apply maps $f_{i,j} : (x, y) \to x^i y^j$. Then (with a bit of elementary analysis) either:

1. The $(\alpha_n)_{n=1}^{\infty}$ map to a point in C, which means that the points lie on a curve $x^i y^j = \xi$ for ξ a root of unity; or
2. The $(\alpha_n)_{n=1}^{\infty}$ actually equidistribute on the circle cross the circle in $\mathbb{G}_m^2$.

The second possibility means they do not really lie on a curve at all.

Special special case, due to Zagier

Consider the line $x + y = 1$. Then the only points of height zero are $(-\xi_3, -\xi_3^2), (-\xi_3^2, -\xi_3), (1, 0), (0, 1)$ (note that the last two are not in $\mathbb{G}_m^2$).

Zagier explicitly calculates that for all other $\alpha \in \mathbb{G}_m^2$, we have $h(\alpha) > 0.24$.

It turns out this is not the actual lim inf, though. The lim inf is still not known. (Just to show how difficult the Bogomolov problem is in general.)

An interesting counterexample aside

Does it make sense to say that any family $(\alpha_n)_{n=1}^{\infty}$ with $\lim_{n\to\infty} h(\alpha_n) = 0$ is also special?

Maybe. It turns out that if $f(x) = \log(x-2)$ (which is not continuous of course!), then

$$\lim_{n\to\infty} \frac{1}{\deg \alpha_n} \sum_{\text{conjugates } \alpha_n^\sigma \text{ of } \alpha_n} f(\alpha_n^\sigma) = \int_0^1 f(e^{2\pi i z})\, dz$$

for nonrepeating sequences $(\alpha_n)_{n=1}^{\infty}$ with $h(\alpha_n) = 0$ all n but *not necessarily* if only $\lim_{n\to\infty} h(\alpha_n) = 0$ (see problem set!).

Preperiodic points and points of small canonical height

Now, let f be a polynomial of degree d with coefficients in $\mathbb{Q}$. About how large is $h(f(\alpha))$ compared to $h(\alpha)$ Let's take a look:

$$f(\alpha) = a_d\alpha^d + a_{d-1}\alpha^{d-1} + \cdots + a_0.$$

The biggest term should dominate, so after taking logs, you should get about $dh(\alpha)$. In fact...

Height change under polynomial maps

We easily see that for all but finitely many v, we have

$$\log^+ |f(\alpha)|_v = \log^+ |\alpha|_v^d.$$

(we are abusing notation a bit here, thinking of α as already embedded in $\overline{\mathbb{Q}}_v$). At the others, we have a constant B_v such that

$$\left| \log^+ |f(\alpha)|_v - \log^+ |\alpha|_v^d \right| < B_v$$

(just like the proof in calculus that odd degree polynomials have real roots). So summing up over all v, we obtain

$$|h(f(\alpha)) - dh(\alpha)| < B$$

The canonical height

One has the Call-Silverman *canonical height* for a polynomial f given by

$$h_f(\alpha) = \lim_{n\to\infty} \frac{h(f^n(\alpha))}{d^n}.$$

Why does it converge? By Tate's famous telescoping series argument.

We rewrite this as

$$h(\alpha) + \sum_{n=1}^{\infty} \frac{h(f^n(\alpha)) - dh(f^{n-1}(\alpha))}{d^n}.$$

Then $\frac{|h(f^n(\alpha))-dh(f^{n-1}(\alpha))|}{d^n} \leq \frac{B}{d^n}$, so the series converges.

Basic properties of canonical height

1. There is a constant M such that $|h(\alpha) = h_f(\alpha)| < M$ for all $\alpha \in \overline{\mathbb{Q}}$;
2. $h_f(f(\alpha)) = dh_f(\alpha)$ for all $\alpha \in \overline{\mathbb{Q}}$.
3. $h_f(\alpha) = 0$ if and only if α is preperiodic under f.

1. gives dynamical Northcott: *there are finitely many points of bounded degree and bounded canonical height*.

Then 1. and 3. give: *f has only finitely many preperiodic points of bounded degree*.

Equidistribution

If we extend the height function to include the point at infinity, setting $h(\infty) = 0$, we can define canonical heights for rational functions $\varphi(x) = f(x)/g(x)$ of degree $d > 1$.

All the same properties apply.

In particular, if $\lim_{n\to\infty} h(\alpha_n) = 0$, then $\lim_{n\to\infty} \deg \alpha_n = \infty$...so it's not unreasonable to expect that points with height tending to zero spread out (and equidistribute) in a nice way.

Dynamical equidistribution

Any rational function has a Brolin-Lyubich measure supported on the famous *Julia set* $\mathcal{J}(f)$. These are the Julia sets for $f(x) = x^2 + c$ where c is as below.

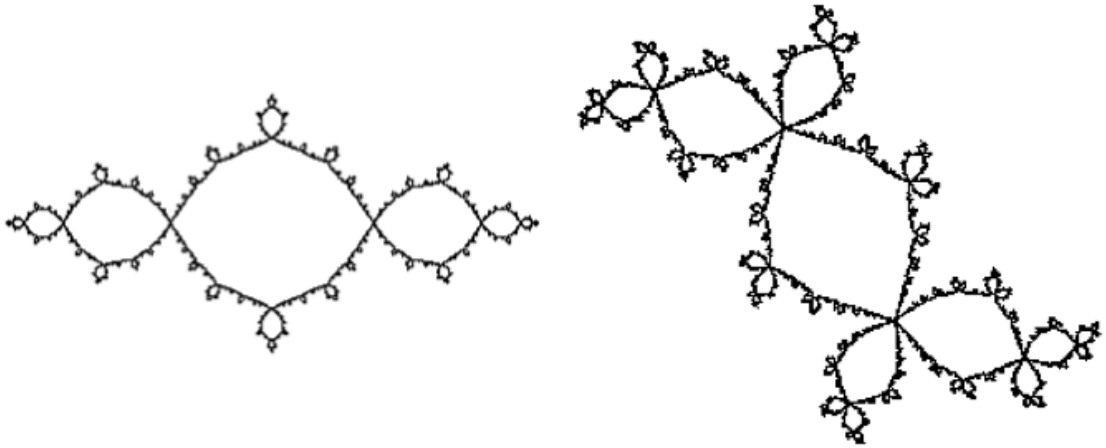

Figure 1: The basilica ($c = -1$) and the Douady rabbit ($c = -0.12256 + 0.74486i$) Julia sets.

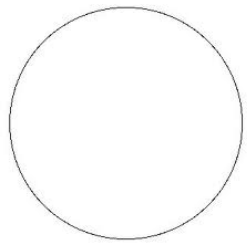

($c = 0$)

More on the Julia set

The Julia set $\mathcal{J}(f)$ is the boundary of the *filled-in Julia set* of points that do not escape to infinity:

$$\{z \in \mathbb{C} \mid \lim_{n\to\infty} |f^n(z)| \neq \infty\}$$

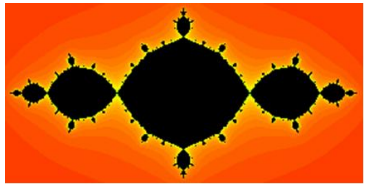

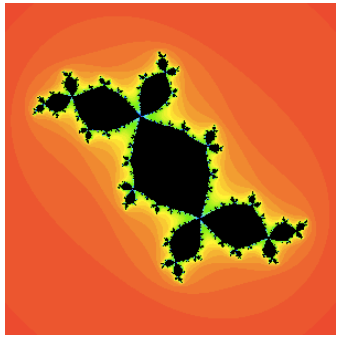

Some quick pictures

For $f(x) = x^2$.

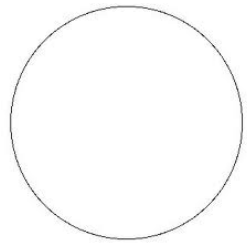

Julia set

Filled-in Julia set.

Julia sets and the Brolin-Lyubich measure

Let $\varphi \in \mathbb{C}(x)$ be a rational function of degree $d > 1$. Then the Brolin-Lyubich measure μ_φ is defined as follows. Take any z that has an infinite backwards orbit – infinitely many points such that $\varphi^n(w) = z$ for some n (such points are called nonexceptional). Then define

$$\int_{\mathbb{C}} g \, d\mu_\varphi = \lim_{n\to\infty} \frac{1}{d^n} \sum_{\varphi^n(w)=z} g(w).$$

Colloquially, it is the limit of the average of point masses on inverse images of z under φ.

The Brolin-Lyubich measure is supported on the Julia set of φ. (For us, that will be the definition of the Julia set for general rational functions.)

Equidistribution revisited

Work of Baker/Hsia, Baker/Rumely, Chambert-Loir, Favre/Rivera-Letelier, and Yuan gives the following generalization of BIlu's result.

Theorem

Let $(\alpha_n)_{n=1}^{\infty}$ be a nonrepeating sequence of algebraic points such that $\lim_{n\to\infty} h_\varphi(\alpha_n) = 0$. Then

$$\sum_{\text{conjugates } \alpha_n^\sigma \text{ of } \alpha} g(\alpha_n^\sigma) = \int_{\mathbb{C}} g \, d\mu_\varphi.$$

In other words, small points equidistribute with respect to the Brolin-Lyubich measure.

This will drive our efforts to prove dynamical Manin-Mumford.

The Dynamical Manin-Mumford Conjecture (first formulation)

Conjecture

Let Φ be a polarizable endomorphism of a projective variety X defined over $\mathbb{C}$. If a subvariety V contains a Zariski dense set of preperiodic points, then V is preperiodic.

And the Bogomolov variant.

Conjecture

Let Φ be a polarizable endomorphism of a projective variety X defined over $\overline{\mathbb{Q}}$. If a subvariety V contains a Zariski dense set of points $(\alpha_n)_{n=1}^{\infty}$ with $\lim_{n\to\infty} h_\Phi(\alpha_n) = 0$, then V is preperiodic.

Note: we haven't defined general canonical heights h_Φ yet, but for us we will be in the split case $(f_1, \ldots, f_N) : \mathbb{A}^N \longrightarrow \mathbb{A}^N$ with each $f_i \in K[x]$ so $h_\Phi(x_1, \ldots, x_N) = h_{f_1}(x_1) + \ldots h_{f_N}(x_N)$.

Bogomolov and Manin-Mumford

- Bogomolov-type results hold for small points.
- Manin-Mumford-type results hold only for preperiodic points.
- Bogomolv-type results hold only over number fields (where heights are defined).
- Manin-Mumford-type results hold over $\mathbb{C}$.

In practice, Bogomolov results seem stronger because there are various tricks for sending F to K where K is a number field and F is a finitely generated ring over which our problem is defined.

Results

The above (both Manin-Mumford and Bogomolov type) conjectures hold for lines $V \subset (\mathbb{P}^1)^N$ under the action of $\Phi := (f_1, \ldots, f_N)$ where each $f_i \in \bar{\mathbb{Q}}[z]$ (polynomials!) such that $\deg(f_1) = \cdots = \deg(f_N) \geq 2$.
In order to prove this result one easily makes the following reductions:

1. $N = 2$ (simply project on two coordinate axis at a time; each projection is still a line);
2. V projects dominantly to each axis (since otherwise it has to project on one of the axis to a preperiodic point and the result is obvious then).
3. V is the diagonal line in $\mathbb{P}^1 \times \mathbb{P}^1$ (we achieve this by replacing g with a suitable conjugate). If V is the line given by $y = \sigma(x)$ for some linear polynomial σ, then we replace g with $\sigma^{-1} \circ g \circ \sigma$ and obtain that both f and (the new) g share infinitely many preperiodic points in common.

The case N = 2

In the case $N = 2$, we have:

- we have two maps f and g of the same degree with $(f, g) : \mathbb{A}^2 \longrightarrow \mathbb{A}^2$ by $(f, g)(x, y) = (f(x), g(y))$.
- a line L in $\mathbb{A}^2$.
- A nonrepeating sequence of points $((\alpha_n, \beta_n))_{n=1}^{\infty}$ on L such that $\lim_{n\to\infty} h_f(\alpha_n) + h_g(\beta_n) = 0$.

We want to show that L is preperiodic under the action of (f, g)

The case N = 2: reduction to the diagonal

After changing coordinates, we may assume the line is the diagonal. So we have then a nonrepeating sequence of points $((\alpha_n, \alpha_n))_{n=1}^{\infty}$ such that $\lim_{n\to\infty} h_f(\alpha_n) + h_g(\alpha_n) = 0$. Since $h_f, h_g \geq 0$, this means

$$\lim_{n\to\infty} h_f(\alpha_n) = \lim_{n\to\infty} h_g(\alpha_n) = 0.$$

Since the α_n equidistribute, they determine the Brolin-Lyubich measures...so we must have $d\mu_f = d\mu_g$. So $\mathcal{J}(f) = \mathcal{J}(g)$. *If two polynomials share a sequence of small points, then they have the same Julia set.*

If $\mathcal{J}(f) = \mathcal{J}(g)$, then what this means for f and g?

Classical results in complex dynamics (proven by Beardon) yield that the Julia set essentially governs not only the dynamics of a given polynomial mapping, but even the polynomial itself. So, $\mathcal{J}(f) = \mathcal{J}(g)$ yields that there exists a linear isometry σ of $\mathcal{J}(f) = \mathcal{J}(g)$ such that $g = \sigma \circ f$; furthermore $f \circ \sigma = \sigma^d \circ f$ where $d = \deg(f) = \deg(g)$. So $(f, g)^n$ always looks like $(f^n, \tau f^n)$ for some isometry τ of the Julia set.

1. If neither f nor g are conjugate to z^d, then $\mathcal{J}(f)$ has *finitely many isometries*, and then immediately we get that the diagonal in $\mathbb{P}^1 \times \mathbb{P}^1$ is preperiodic under the action of (f, g) since it simply moves through various a finite list of automorphic images of the diagonal $(\tau x, x)$.
2. If f (and so, also g) is conjugate to z^d, we are back to Zhang's theorem and the conclusion follows immediately.

Isometries

Which of these sets has a lot of isometries?

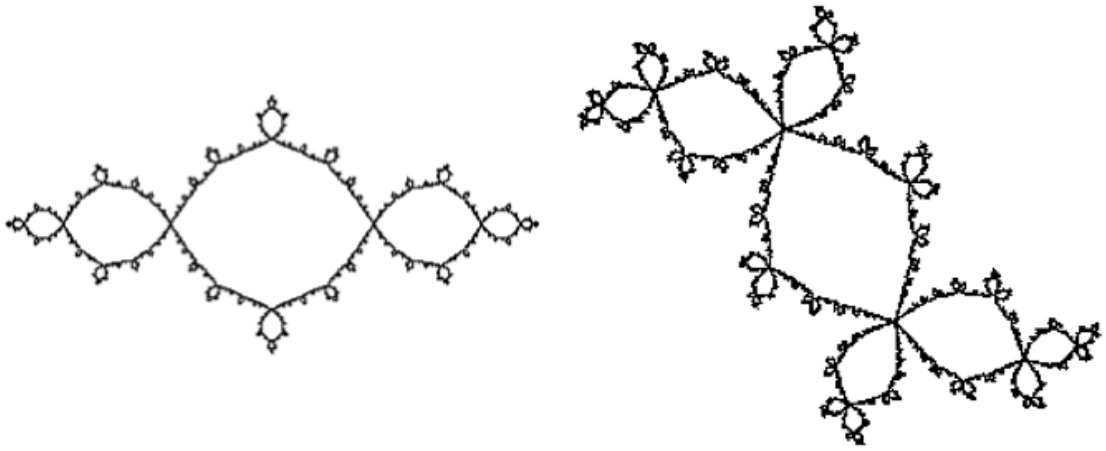

Figure 1: The basilica ($c = -1$) and the Douady rabbit ($c = -0.12256 + 0.74486i$) Julia sets.

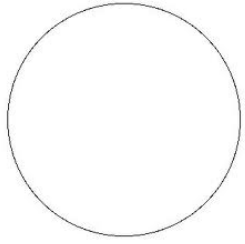

$(c = 0)$

How about rational functions?

If we try this for rational functions, ψ and φ, things work out the same.

Reducing to the diagonal, we get a sequence of points that are small for both canonical heights, so $d\mu_\varphi = d\mu_\psi$.

This means the Julia sets are the same, but that is a bit weaker (it turns out) for rational functions. Two rational functions can have the same Julia set without having the same Brolin-Lyubich measure.

Counterexample with Lattès maps

However, there are counterexamples when φ and ψ are Lattès maps corresponding to endomorphisms of same degree of the same CM elliptic curve.

(See diagram on board)

It does turn out these are the only counterexamples.

Why are Lattès maps good examples for counterexamples?

Lattès maps have a number of properties that make them good candidates for counterexamples here.

- They all have Julia sets that are equal to the whole Riemann sphere. (This has many symmetries.)
- They commute with many other maps, and these all share the same Brolin-Lyubich measure.
- They are counterexamples to many important statements on moduli of rational functions (Thurston rigidity, McMullen multipier theorem, etc.).

Counterexample with CM elliptic curves

Let E be the elliptic curve given by the equation $y^2 = x^3 - x$. The endomorphism ring of E is $\mathbb{Z}[i]$; note that

$$[i](x, y) = (-x, iy).$$

Let φ be the endomorphism of E corresponding to $[2 + i]$ and ψ be the endomorphism corresponding to $[2 - i]$. We let $A = E \times E$ and let $\Psi := (\varphi, \psi) \in \mathrm{End}(A)$. It's easy to see that Ψ is polarizable (here we use the fact that $2 + i$ and $2 - i$ have the same norm). Also, the preperiodic points for Φ are precisely pairs of torsion points for E.
So, the diagonal $\Delta \subset E \times E$ contains infinitely many preperiodic points for Φ, but it is not preperiodic (here we use the fact that for each $n \in \mathbb{N}$ we have $(2 + i)^n \neq (2 - i)^n$).

A simple explanation with exponential maps

Our old friend the exponential map – which helped so much with dynamical Mordell-Lang – is the villain here. Let $A = E \times E$. We have a map exp (over $\mathbb{C}$) that sends the tangent space at the identity to surjectively onto $\mathbb{C}$, that is

$$\exp : T_{A,0} \longrightarrow A(\mathbb{C}).$$

What does the Jacobian of $([2+i],[2-i])$ look like at the identity? $\begin{pmatrix} 2+i & 0 \\ 0 & 2-i \end{pmatrix}$. Now, exp sends the space $\begin{pmatrix} x \\ x \end{pmatrix}$ onto the diagonal in $E \times E$...but it obviously is not preperiodic under $\begin{pmatrix} 2+i & 0 \\ 0 & 2-i \end{pmatrix}$.

The Dynamical Manin-Mumford Conjecture (new version)

We believe that maybe all counterexamples to the original version of the Dynamical Manin-Mumford Conjecture descend to endomorphisms of abelian varieties (similar to the above construction). Motivated by this we have the following reformulation of the conjecture:

Conjecture

Let X be a projective variety, let $\varphi : X \longrightarrow X$ be an endomorphism defined over $\mathbb{C}$ with a polarization, and let Y be a subvariety of X which has no component included into the singular part of X. Then Y is preperiodic under φ if and only if there exists a Zariski dense subset of smooth points $x \in Y \cap \mathrm{Prep}_{\varphi}(X)$ such that the tangent subspace of Y at x is preperiodic under the induced action of φ on the Grassmanian $\mathrm{Gr}_{\dim(Y)}(T_{X,x})$. (Here we denote by $T_{X,x}$ the tangent space of X at the point x.)

Preview of coming attractions

Here is something you may be even more likely to have seen than a Julia set: Mandelbrot set.

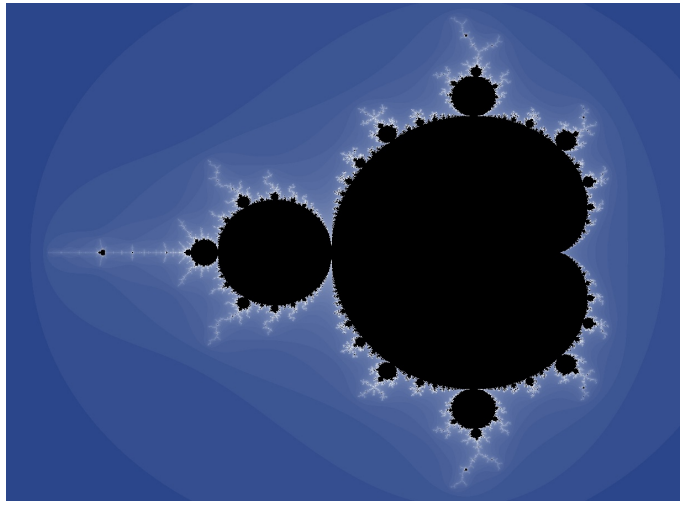

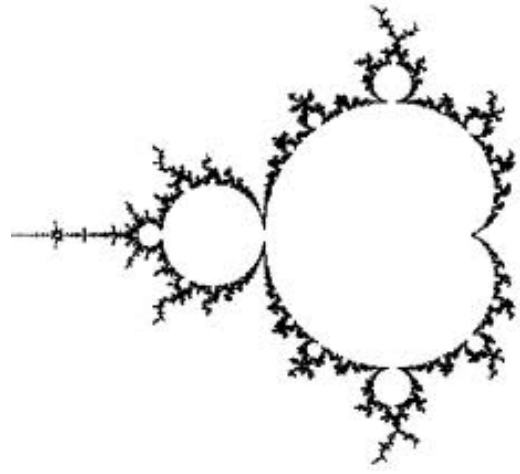

WIth the Julia set, you fix a map, vary points and see which points do not go to infinity. With a Mandelbrot set, you fix a point (namely 0), vary the map (over the family $x^2 + c$), and see for which maps that point goes to infinity. The same equidistribution results apply, and they allow us treat a whole new class of questions...next time!

The case of abelian varieties for the new conjecture

It is easy to check the above conjecture for all algebraic group endomorphisms of abelian varieties. Also, the approach we had above for the case of lines under polynomial actions works for the above reformulation for all lines under the coordinatewise action of any rational maps.
First note that since $\Phi : A \longrightarrow A$ is a polarizable algebraic group endomorphism, all its preperiodic points are the torsion points of A.

Indeed, if $\Phi^n(\alpha) - \Phi^m(\alpha)$ for some positive integers $n \neq m$, then $\alpha \in \ker(\Phi^n - \Phi^m)$. Because Φ is polarizable, Φ acts on the tangent space of A at any point with eigenvalues of absolute value strictly larger than 1. Therefore $\Phi^n - \Phi^m$ is a finite endomorphism of A, and thus it has finite kernel, which means that α is a torsion point.
If V is an irreducible subvariety of A which contains a Zariski dense set of preperiodic points for Φ, then V contains a Zariski dense set of torsion points of A and so, by Laurent's Theorem we obtain that V is a torsion translate of an algebraic subgroup H of A.
Using the fact that the tangent subspace of V at any preperiodic point for Φ is preperiodic under the induced action of Φ on the tangent space of A, we get that H itself must be preperiodic under the action of Φ.

Unlikely intersections

We return again to Lang's Theorem:

Theorem

If the curve $C \subset \mathbb{A}^2$ defined over $\mathbb{C}$ contains infinitely many points with both coordinates roots of unity, then the equation of C is $x^m y^n = \alpha$, for some $m, n \in \mathbb{Z}$ and some root of unity α.

Lang's result yields (assuming C projects dominantly to both axis) that for *each* point $(x, y) \in C$, we have that x is a root of unity *if and only if* y is a root of unity. In other words, if there is a point $(x_0, y_0) \in C$ such that x_0 is a root of unity while y_0 is not a root of unity (or vice-versa), then there exist *at most finitely many* points $(x, y) \in C$ such that both x and y are roots of unity.
Masser and Zannier proved a similar result for torsion points in fibres of the Legendre family of elliptic curves. Both their result and the above theorem of Lang fit into a more general program which is highlighted by the Pink-Zilber Conjecture.

The Legendre family of elliptic curves

For each $\lambda \notin \{0, 1\}$ we let E_λ be the elliptic curve given by the equation

$$y^2 = x(x-1)(x-\lambda)$$

and let $P_\lambda := (2, \sqrt{2(2-\lambda)})$ and $Q_\lambda := (3, \sqrt{6(3-\lambda)})$.
It is easy to check that $P_2 = (2, 0)$ is a 2-torsion point on E_2, while $Q_2 = (3, \sqrt{6})$ is a nontorsion point for E_2. Similarly, $Q_3 = (3, 0)$ is a 2-torsion point on E_3, while $P_3 = (2, i\sqrt{2})$ is a nontorsion point for E_3.
On the other hand, one can find infinitely many $\lambda \in \bar{\mathbb{Q}}$ such that P_λ is a torsion point for E_λ, and similarly one can find infinitely many $\lambda \in \bar{\mathbb{Q}}$ such that Q_λ is a torsion point for E_λ.
Can we find infinitely many $\lambda \in \mathbb{C}$ such that *both* P_λ and Q_λ are torsion points for E_λ?
No.

The dynamical interpretation of the Masser and Zannier's result

For each elliptic curve E_λ we consider [2] be the multiplication-by-2-map on E_λ. Then the preperiodic points of this map are precisely the torsion points of E_λ; hence Masser and Zannier's theorem is similar to Lang's result on roots of unity on curves (also recall that roots of unity are preperiodic points for the squaring map).
For example, one can find a parametrization of the $C \subset \mathbb{A}^2$ in Lang's theorem given by two algebraic functions $f(z)$ and $g(z)$ and then the setting is the same as in the elliptic case: can one find infinitely many $\lambda \in \mathbb{C}$ such that both $f(\lambda)$ and $g(\lambda)$ are roots of unity at the same time?

The general problem

Let $\{X_\lambda\}$ be an algebraic family of varieties parametrized by $\lambda \in \mathbb{C}$. Let $P_\lambda, Q_\lambda \in X_\lambda$ be two algebraic families of points, and let $\Phi_\lambda : X_\lambda \longrightarrow X_\lambda$ be an algebraic family of endomorphisms.
Question: Under what conditions there exist infinitely many $\lambda \in \mathbb{C}$ such that both P_λ and Q_λ are preperiodic points for Φ_λ?
One cannot expect to find *explicit* necessary and sufficient conditions satisfying the conclusion of the above question. However we note that since for *each* family of points R_λ one expects that there exist infinitely many $\lambda \in \mathbb{C}$ such that R_λ is preperiodic for Φ_λ, then we expect that the following trichotomy be the answer to the above question:

1. P_λ is preperiodic for Φ_λ, for all λ; or
2. Q_λ is preperiodic for Φ_λ, for all λ; or
3. for each λ, P_λ is preperiodic for Φ_λ if and only if Q_λ is preperiodic for Φ_λ.

A theorem of Baker and DeMarco

It is immediate to see that the conclusion in the following result fits exactly into the trichotomy for *unlikely intersections in dynamics*.

Theorem

Let $a, b \in \mathbb{C}$, let d be an integer greater than 1, and let $\Phi_\lambda(x) = x^d + \lambda$. There exist infinitely many $\lambda \in \mathbb{C}$ such that both a and b are preperiodic for Φ_λ if and only if $a^d = b^d$.

The proof uses various techniques, from an equidistribution result on Berkovich spaces (proven by Baker an Rumely, Favre and Rivera-Letelier, and Chambert-Loir) to classical theorems in complex dynamics and complex analysis (regarding univalent functions).

An example

Consider the family of polynomials $f_\lambda(x) = x^3 - \lambda x^2 + (\lambda^2 - 1)x + \lambda$ indexed by all $\lambda \in \mathbb{C}$. Let $\mathbf{a}(\lambda) = \lambda$ and $\mathbf{b}(\lambda) = \lambda^3 - 1$.

Question: Are there infinitely many $\lambda \in \mathbb{C}$ such that both $\mathbf{a}(\lambda)$ and $\mathbf{b}(\lambda)$ are preperiodic for the same f_λ?

For example, $\lambda = 0$ satisfies the above conditions since then

- $f_0(x) = x^3 - x$;
- $\mathbf{a}(0) = 0$ and $\mathbf{b}(0) = -1$,

and $f_0(0) = 0$ while $f_0(-1) = 0$.

Also $\lambda = 1$ works since then

- $f_1(x) = x^3 - x^2 + 1$;
- $\mathbf{a}(1) = 1$ and $\mathbf{b}(1) = 0$,

and $f_1(1) = 1$ while $f_1(0) = 1$.

Are there infinitely many more such λ's? Note that *individually*, there exist infinitely many $\lambda \in \mathbb{C}$ such that either $\mathbf{a}(\lambda)$ or $\mathbf{b}(\lambda)$ are preperiodic for f_λ (simply solve the equation $f_\lambda^n(\mathbf{a}(\lambda)) = \mathbf{a}(\lambda)$ for varying $n \in \mathbb{N}$).

On the other hand, $\lambda = -1$ does not work since

- $f_{-1}(x) = x^3 + x^2 - 1$;
- $\mathbf{a}(-1) = -1$ and $\mathbf{b}(-1) = -2$,

and $f_{-1}(-1) = -1$, while

$$f_{-1}(-2) = -5;\ f_{-1}^2(-2) = -101;\ \ldots\ldots$$

So, it's not true that $\mathbf{a}(\lambda)$ is preperiodic exactly when $\mathbf{b}(\lambda)$ is preperiodic, and it's not true that $\mathbf{b}(\lambda)$ is always preperiodic under f_λ. Nor it is true that $\mathbf{a}(\lambda)$ is always preperiodic, as it's shown by the case $\lambda = 2$. In that case,

- $f_2(x) = x^3 - 2x^2 + 3x + 2$ and $\mathbf{a}(2) = 2$, while
- $f_2(2) = 8$, $f_2^2(2) = 410$, $\ldots\ldots$.

The above two examples coupled with our conjecture suggest that there should only be finitely many $\lambda \in \mathbb{C}$ such that both $\mathbf{a}(\lambda)$ and $\mathbf{b}(\lambda)$ are preperiodic for f_λ since all three conditions from the trichotomy principle regarding unlikely intersections in dynamics fail in this example. This follows from the next result.

Theorem

Let $\{f_\lambda\}$ be any family of polynomials defined over $\bar{\mathbb{Q}}$ of degree larger than 1, and let $\mathbf{a}, \mathbf{b} \in \bar{\mathbb{Q}}[\lambda]$. Suppose that $\mathbf{a}$ and $\mathbf{b}$ are not preperiodic under $\{f_\lambda\}$ in $\bar{\mathbb{Q}}[\lambda]$. If there exist infinitely many $\lambda \in \bar{\mathbb{Q}}$ such that both $\mathbf{a}(\lambda)$ and $\mathbf{b}(\lambda)$ are preperiodic for f_λ, then:

- *for each $\lambda \in \bar{\mathbb{Q}}$, $\mathbf{a}(\lambda)$ is preperiodic for f_λ if and only if $\mathbf{b}(\lambda)$ is preperiodic for f_λ.*

Sketch of proof

The idea here, due to Baker and DeMarco, is a beautiful one. Instead of using a map to define a height on points, you use a point to define a height on maps. Their family is $f_\lambda(x) = x^2 + \lambda$. Fixing a point a gives us a height on the parameter space $\{f_\lambda\}_{\lambda \in \overline{\mathbb{Q}}}$ by

$$h_a(\lambda) = h_{f_\lambda}(a).$$

Then crucially: $h_a(\lambda) = 0$ *if and only if* a *is preperiodic for* f_λ.

Also: One can obtain a measure corresponding to this height in various ways and equidistribution results apply.
Let's do an example that we touched on last time.

Mandelbrot set revisited

As before, let $f_\lambda(x) = x^2 + \lambda$. Let $a = 0$. Then the Mandelbrot set below is the exact analog of the filled in Julia set. It is:

$$\{\lambda \in \mathbb{C} \mid \lim_{n \to \infty} |f_\lambda^n(0)| \neq \infty\}$$

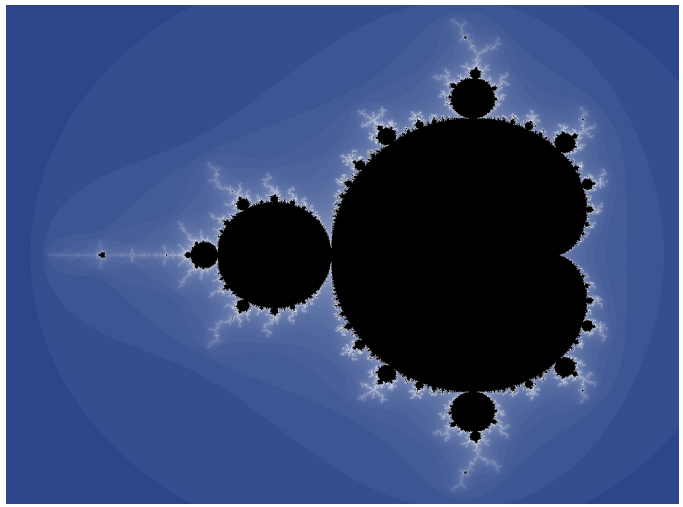

Mandelbrot set revisited

There is a natural measure one can associate to this set, the equilibrium measure (one can also obtain the measure in other ways), concentrated on the boundary. Let's call the measure μ_0.

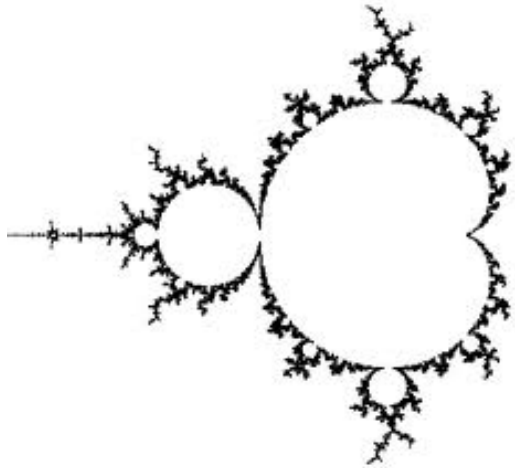

Equidistribution

Now, let's go back to our slightly more general setting. The family is $f_\lambda(x) = x^2 + \lambda$. Fixing a point a gives us a height on the parameter space $\{f_\lambda\}_{\lambda \in \overline{\mathbb{Q}}}$ by

$$h_a(\lambda) = h_{f_\lambda}(a).$$

Theorem

(Baker-DeMarco et al.) Let $(\lambda)_{n=1}^{\infty}$ be a nonrepeating sequence such that $\lim_{n\to\infty} h_a(\lambda_n) = 0$. Then the $(\lambda_n)_{n=1}^{\infty}$ equidistribute with respect to $d\mu_a$. (In the usual sense of weak convergence of average point masses.)

In particular, when $a = 0$, such a family $(\lambda_n)_{n=1}^{\infty}$ equidistributes on the boundary of the Mandelbrot set.

Examples with heights

With this family again and $a = 0$, for which $\lambda \in \mathbb{Q}$, do we have $h_0(\lambda) = 0$?

In other words, for which λ is 0 preperiodic under f_λ?

We see right off that $\lambda \in \mathbb{Z}$ because once a denominator gets in, it will never come out. Similarly if $\lambda > 1$, then 0 must go off to infinity. The same if $|\lambda| > 2$. So we are just left with $0, -1, -2$. We check that 0 is indeed preperiodic for these λ.

Let's look at the Mandelbrot set again.

Idea of proof of Baker-DeMarco

Take a and b. Let $h_a(\lambda)$ and $h_b(\lambda)$ be as above. Suppose that we have $(\lambda_n)_{n=1}^{\infty}$, be a nonrepeating sequence such that $\lim_{n\to\infty} h_a(\lambda_n) = \lim_{n\to\infty} h_b(\lambda_n) = 0$. Then

1. $(\lambda_n)_{n=1}^{\infty}$ equidistribute for both μ_a and μ_b
2. So $\mu_a = \mu_b$.
3. This forces a "special relation" between a and b (as with our proof of Manin-Mumford, using Julia sets), namely that $a^2 = b^2$.

Notes on Baker-DeMarco

Some notes on Baker-DeMarco:

1. It gives a "Bogomolov-type" result for points of small height, not just preperiodic points.
2. It extends from $\overline{\mathbb{Q}}$ to $\mathbb{C}$ via specialization.
3. The $(\lambda_n)_{n=1}^{\infty}$ equidistribute for both μ_a and μ_b, so $\mu_a = \mu_b$ idea extends vastly to other situations..
4. The "special relation" $a^2 = b^2$ is difficult to prove and may be hard to extend to more generality.

(One more chart on special)

Recall from earlier.

Theorem

Let $\{f_\lambda\}$ be any family of polynomials defined over $\bar{\mathbb{Q}}$ of degree larger than 1, and let $\mathbf{a}, \mathbf{b} \in \bar{\mathbb{Q}}[\lambda]$. Suppose that $\mathbf{a}$ and $\mathbf{b}$ are not preperiodic under $\{f_\lambda\}$ in $\bar{\mathbb{Q}}[\lambda]$. If there exist infinitely many $\lambda \in \bar{\mathbb{Q}}$ such that both $\mathbf{a}(\lambda)$ and $\mathbf{b}(\lambda)$ are preperiodic for f_λ, then:

- *for each $\lambda \in \bar{\mathbb{Q}}$, $\mathbf{a}(\lambda)$ is preperiodic for f_λ if and only if $\mathbf{b}(\lambda)$ is preperiodic for f_λ.*

In fact, it's enough to have a sequence of "small points".

Proof

The same basic trickeration works but with a weaker last step. Define

$$h_a(\lambda) = h_{f_\lambda}(a(\lambda))$$

for $\lambda \in \overline{\mathbb{Q}}$.

1. $(\lambda_n)_{n=1}^{\infty}$ equidistribute for both μ_a and μ_b.
2. So $\mu_a = \mu_b$.
3. This forces $h_a = h_b$ so the preperiodic maps are the same for a as for b.

When we can get a "special relation"

Theorem

(Ghioca, Hsia, Tucker) Let d be an integer greater than 1, let $c_d \in \mathbb{C}^*$, let $c_{d-1}, \ldots, c_0 \in \mathbb{C}[\lambda]$, and let

$$f_\lambda(x) = c_d x^d + c_{d-1}(\lambda)x^{d-1} + \cdots + c_1(\lambda)x + c_0(\lambda).$$

Let $\mathbf{a}, \mathbf{b} \in \mathbb{C}[\lambda]$ such that

- $\deg(\mathbf{a}) = \deg(\mathbf{b}) \geq d \cdot \max\{\deg(c_0), \ldots, \deg(c_{d-1})\}$;
- $\mathbf{a}$ *and* $\mathbf{b}$ *have the same leading coefficient.*

If there exist infinitely many $\lambda \in \mathbb{C}$ such that both $\mathbf{a}(\lambda)$ and $\mathbf{b}(\lambda)$ are preperiodic for f_λ, then $\mathbf{a} = \mathbf{b}$.

In particular, we get that $\mathbf{a}(\lambda)$ is preperiodic if and only if $\mathbf{b}(\lambda)$ is preperiodic.

Getting the "special relation"

The special relation trick will only ever work for polynomials because there you Bötcher's Uniformization Theorem, which allows you to choose analytic coordinate near infinity that make your polynomial look like a powering map (obviously can't do this for more general rational functions).

For any (monic) polynomial $g \in \mathbb{C}[x]$ of degree $d \geq 2$, there exists a real number $R \geq 1$ and an analytic map $\Phi : U_R \longrightarrow U_R$, where

$$U_R = \{z \in \mathbb{C} : |z| > R\}$$

satisfying the following two conditions:

(i) Φ is univalent on U_R and at ∞,

$$\Phi(z) = z + O\left(\frac{1}{z}\right);$$

(ii) for all $z \in U_R$ we have

$$\Phi(g(z)) = \Phi(z)^d.$$

All happy families are alike

You can actually get a result that compares different families of maps.

Theorem

Let $\{f_\lambda\}$ and $\{g_\lambda\}$be any family of polynomials defined over $\bar{\mathbb{Q}}$ of degree larger than 1, and let $\mathbf{a}, \mathbf{b} \in \bar{\mathbb{Q}}[\lambda]$. Suppose that $\mathbf{a}$ is not preperiodic under $\{f_\lambda\}$ in $\bar{\mathbb{Q}}[\lambda]$ and $\mathbf{b}$ is not preperiodic under $\{g_\lambda\}$ in $\bar{\mathbb{Q}}[\lambda]$. If there exist infinitely many $\lambda \in \bar{\mathbb{Q}}$ such that both $\mathbf{a}(\lambda)$ and $\mathbf{b}(\lambda)$ are preperiodic for f_λ, then:

- *for each $\lambda \in \bar{\mathbb{Q}}$, $\mathbf{a}(\lambda)$ is preperiodic for f_λ if and only if $\mathbf{b}(\lambda)$ is preperiodic for g_λ.*

More general special relations

We would like to say something like this: if for each $\lambda \in \bar{\mathbb{Q}}$, $\mathbf{a}(\lambda)$ is preperiodic for f_λ if and only if $\mathbf{b}(\lambda)$ is preperiodic for g_λ, then $(\mathbf{a}(\lambda), \mathbf{b}(\lambda))$ lie on a curve that is preperiodic under (f_λ, g_λ).

Example

When $f_\lambda = g_\lambda = x^2 + \lambda$ and a and b are constant, we get $a^2 = b^2$. So (a, b) lives on the (periodic) curve $f_\lambda(x) = f_\lambda(y)$ in the (x, y)-plane.

A counterexample revisited

But in full generality, you cannot always get a dynamically special relation.

Example

Take the constant families of maps Lattès maps f and g where f corresponds to multiplication by $2+i$ on the elliptic curve $y^2 = x^3 - x$ and g corresponds to multiplication by $2+i$. Let $a(\lambda) = b(\lambda) = \lambda$ (a nonconstant family). Then there are clearly infinitely many λ_n such that λ_n is preperiodic for both f and g, but there is no nice dynamically relation between a and b (although they are actually equal).

Both sides now

We would like a Bogomolov version of Masser-Zannier's result.

Conjecture

Let E be the elliptic curve over $\overline{\mathbb{Q}}(\lambda)$ given by

$$y^2 = x(x-1)(x-\lambda)$$

and let P, Q be two points in $E(\overline{\mathbb{Q}}(\lambda))$. Suppose that there values λ_n such that $\lim_{n\to\infty} h_{E_{\lambda_n}}(P_{\lambda_n}) = \lim_{n\to\infty} h_{E_{\lambda_n}}(P_{\lambda_n}) = 0$. Then $[m]P = [n]Q$ in $E(\overline{\mathbb{Q}}(\lambda))$ for some m, n.

Curiously, getting the "small points" form via Pila-Wilkie seems impossible.
Getting the "special relation" $[m]P = [n]Q$ via equidistribution seems impossible.
But maybe we can use them together.

Come together, right now

- Use equidistribution results on the family of Lattès maps to prove that if $\lim_{n\to\infty} h_{E_{\lambda_n}}(P_{\lambda_n}) = \lim_{n\to\infty} h_{E_{\lambda_n}}(P_{\lambda_n}) = 0$, then P_λ is torsion at λ exactly when Q_λ is.
- Apply Masser-Zannier (which uses Pila-Wilkie) to conclude that $[m]P = [n]Q$ in $E(\overline{\mathbb{Q}}(\lambda))$.