The dynamical gcd problem

Thomas J. Tucker

Number fields

Let $a, b \in \mathbb{Z}$ be multiplicatively independent (i.e., there are no positive integers i, j such that $a^i = b^j$). Bugeaud, Corvaja, and Zannier showed that for any $\epsilon > 0$, we have

$$\gcd(a^n - 1, b^n - 1) \ll_\epsilon \exp(n\epsilon).$$

The proof involves deep tools from diophantine approximation, in particular an ingenious use of the Schmidt subspace theorem.

It can be viewed as Vojta's conjecture applied to the blow up of $\mathbb{P}^2$ at $[1 : 1 : 1]$ (and is one of the only cases of Vojta conjeture in dimension greater than one to be proved).

A function field analog

Ailon and Rudnick later showed that over function fields one can obtain something even stronger, and with much more elementary techniques.

Theorem

(Ailon-Rudnick) Let $f, g \in \mathbb{C}[t]$ be nonconstant, multiplicatively independent polynomials. Then there is an $h(t) \in \mathbb{C}[t]$ such that for any n, we have

$$\gcd(f^n - 1, g^n - 1) | h.$$

Note of course that such a finiteness result could never hold over number fields since p always divides $a^{p-1} - 1$ and $b^{p-1} - 1$. Similarly, one cannot obtain such a finiteness result in a function field in characteristic p (work of Silverman makes this explicit).

A function field analog, continued

The degree of h can explicitly bounded in terms of the degrees of f and g.

Observe that bounding the degree of $\gcd(f^n - 1, g^n - 1)$ in terms of $\deg f$ and $\deg g$ bounds the degree of h, since 1 is fixed under powering, so if $m|n$, then $\gcd(f^m - 1, g^m - 1)$ divides $\gcd(f^n - 1, g^n - 1)$. (This would not be true of 1 was not fixed under powering.)

Proof of Ailon-Rudnick

Let $(f,g) : \mathbb{A}^1 \longrightarrow \mathbb{A}^2$ by $(f,g)(t) = (f(t), g(t))$ and let C be the image of (f,g). If $f(\lambda)^n = g(\lambda)^n = 1$, then $f(\lambda)$ and $g(\lambda)$ are roots of unity.

By the Serre-Ihara-Tate theorem (Manin-Mumford for $\mathbb{G}_m^2$), C contains infinitely many pairs (ξ, ξ') for ξ, ξ' roots of unity exactly when the equation defining C has the form $x^i - y^j = 0$. Thus, if f, g are multiplicatively independent, there are at most finitely many such λ (the multiplicity of the roots is easy to control here), so this shows the existence of the polynomial h such that

$$\gcd(f^n - 1, g^n - 1) | h$$

for all h.

Different “target points”

What if instead one fixed $c(x), d(x)$ nonzero polynomials and asked if there is an h such that

$$\gcd(f^n - c, g^n - d) | h$$

for all n?

Note that since c and d may not be roots of unity, an upper bound on the *degree* of $\gcd(f^n - c, g^n - d)$, independent of n, is not enough, since we could be constantly getting new factors for different n.

Note also that one does not get roots of unity, so cannot apply Serre-Ihara-Tate (you might guess what we can apply instead).

Different "target points" continued

We obtain the following generalization of the Ailon-Rudnick theorem.

Theorem

Let $f, g, c, d \in \mathbb{C}[t]$. Suppose that $c, d \neq 0$, that $\deg f, \deg g > 0$ and that f and g are multiplicatively independent. Then there is an $h(t) \in \mathbb{C}[t]$ such that

$$\gcd(f^n - c, g^n - d) | h$$

for all n.

Different “target points” proof

One can attack the problem using Zhang's Bogomolov conjecture theorem that generalizes Serre-Ihara-Tate, using “small points” (families of points with height tending to zero) rather than roots of unity.

Indeed, if everything is defined over $\bar{\mathbb{Q}}$, and we have infinitely many λ_i such that $f(\lambda_i)^{n_i} = c(\lambda_i)$ and $g(\lambda_i)^{n_i} = d(\lambda_i)$, then it is easy to see, using some very simple height estimates, that

$$\lim_{i\to\infty} h(f(\lambda_i)) + h(g(\lambda_i)) = 0.$$

Different “target points” proof continued

Then, Zhang's proof of Bogomolov conjecture for $\mathbb{G}_m^2$ implies that the curve given by $(f(t), g(t))$ is of the form $x^i - y^j = 0$ so f and g are not multiplicatively independent.

To obtain the result over $\mathbb{C}$, we pass to the finitely generated field generated K by the coefficients of f, g, c, d, put suitable absolute values on K to give a height, use a suitable form of function field Bogomolov (e.g. Ghioca's) to descend to the case where everything is defined over $\bar{\mathbb{Q}}$ (one can say that points of sufficiently small height on C must lie in $\bar{\mathbb{Q}}$).

Questions and remarks

- Can the degree of h be made effective, as in the case where $c = d = 1$? As noted above, not enough to bound the degree of each individual $\gcd(f^n - c, g^n - d)$.
- Is there an "elementary proof" for this in general (as with $c = d = 1$)?
- Also note that if $b = d = 0$, the degree of gcd goes up, though its number of *distinct factors* does not. (Simple example of how ramification increases multiplicity in this context.)
- One can obtain something similar for multiplication of points on elliptic curves, using Bogomolov (notably DeMarco-Wang-Ye's "small points" generalization of Masser-Zannier's finiteness result for torsion specializations on fiber products of the Legendre elliptic curve).
- In fact, the proof gives an h such that $\gcd(f^m - c, g^n - d) | h$ for all m, n (not just $m = n$).

Compositional Ailon-Rudnick

For a polynomial f, w let $f^{\circ n}$ be the n-th compositional power of f, that is f composed with itself n times.

Can one prove anything about

$$\gcd(f^{\circ n}, g^{\circ n})$$

for f and g compositionally independent (whatever that means) of degree greater than one? This was raised by Ostafe in a recent preprint.

One also has to rule out an obvious counterexample, where some power of $f^{\circ i}$ and $g^{\circ j}$ share a common root that is periodic and ramified, so that it appears with higher and higher multiplicity in $\gcd(f^{\circ n}, g^{\circ n})$.

Compositional dependence, example

Let $f(x) = x^2 + 1$ and $g(x) = -(x^2 + 1)$. Then $f^{\circ n} = -g^{\circ n}$ for all n so clearly the degree of $\gcd(f^{\circ n}, g^{\circ n})$ goes to infinity as n goes to infinity, although $f^{\circ n} \neq g^{\circ m}$ for any m, n.

Thus, one has to allow for composition with finite linear maps at least to have an adequate notion of compositional independence. More on this later....

Compositional Ailon-Rudnick, continued

However, modulo the discussion above answer is “yes” and the proof is much the same as above, using small points. We will show that under reasonable hypotheses, there is an h such that

$$\gcd(f^{\circ n}, g^{\circ n}) | h$$

for all n.

Here, we will explicitly use equidistribution (which can be used to give Bogomolov on $\mathbb{G}_m^2$), following the model of Baker/DeMarco and Yuan/Zhang, plus some results on Julia sets due to Beardon/Schmidt/Steimetz.

Compositional Ailon-Rudnick, proof

We start with the proof over $\bar{\mathbb{Q}}$ so $f, g \in K[t]$ where K is a number field. We let

- $h_v(\alpha) = \max \log(\|\alpha\|_v, 1)$, for $\| \cdot \|_v$ a suitably normalized place extended to all of $\bar{\mathbb{Q}}$, be the Weil local height for v;
- $h_{f,v}(\alpha) = \lim_{n\to\infty} \frac{h_v(f^{\circ n}(\alpha))}{(\deg f)^n}$ be the canonical local height for f at v.
- $h_f((\alpha)) = \sum_v h_{f,v}(\alpha)$ be the global canonical height for α. Then $h_f(f(\alpha)) = (\deg f) h_f(\alpha)$.

Likewise we define canonical global and local heights for g.

Compositional Ailon-Rudnick, proof continued

If we have infinitely many λ_i with $f^{\circ n_i}(\lambda_i) = g^{\circ n_i}(\lambda_i) = 0$, then we must have

$$\lim_{i \to \infty} h_f(\lambda_i) = h_g(\lambda_i) = 0.$$

By standard equidistribution theorems (Favre/Rivera-Letelier, Baker/Rumely, Chambert-Loir, Yuan), this means that the λ_i equidistribute with respect to the canonical measures associated to the local canonical heights at each place v. This forces $h_{f,v} = h_{g,v}$ up to a constant, and thus $h_f = h_g$.

Since $h_f(\alpha) = 0$ exactly when α is preperiodic for f and $h_g(\alpha) = 0$ exactly when α is preperiodic for g.

Thus, the Julia sets of f and g are equal (the boundary of the set of points in $\mathbb{C}$ that don't escape to infinity under iteration).

Compositional Ailon-Rudnick, proof continued

Now a result of Schmidt and Steinmetz shows that either:

1. f and g are both powers of the same polynomial p, up to composition with finite order linear automorphisms.
2. f and g are conjugate to powering maps x^i, x^j.
3. f and g are conjugate to plus/minus Chebychev polynomials T_i, T_j.

The idea here is that the finite order linear automorphism is a symmetry of the Julia set. (In the case of powering and Chebychev maps, the symmetry group is infinite.)

Compositional Ailon-Rudnick, proof continued

Perhaps the powering and Chebychev cases above do not count as "compositionally dependent". But if one specifics that $\deg f = \deg g$, then they will satisfy a reasonable notion of compositional dependence.

We might say that f and g of the same degree are "compositionally dependent" if the diagonal is preperiodic under the action of $(f, g) : \mathbb{A}^2 \longrightarrow \mathbb{A}^2$ by $(f, g)(x, y) = (f(x), g(y))$.

In this case we see then that there is an h such that $\gcd(f^{\circ n}, g^{\circ n}) | h$ for all n unless f and g are compositionally dependent or share a common ramified periodic point.

Descent to $\bar{\mathbb{Q}}$

To prove over $\mathbb{C}$, we let K be the (finitely generated) field generated by the coefficients f and g. One may apply the same equidistribution theorems (still valid over K) to derive $h_f = h_g$, where now the canonical heights are with respect to a complete set of absolute values satisfying the product formula for K.

- If f and g are not isotrivial (conjugate to a polynomial over $\bar{\mathbb{Q}}$), then $h_f(\alpha) = h_g(\alpha) = 0$ exactly when α is preperiodic for both f and g, by work of Benedetto, Baker, Chatzidakis, Hrushovski, and one obtains equality of Juli sets and proceeds as above.
- If one of f or g is isotrivial over $\bar{\mathbb{Q}}$, then both are and we reduce to the $\bar{\mathbb{Q}}$ case just treated.

There are still some issues with the descent argument: over $\bar{\mathbb{Q}}$ we can treat the case of $\gcd(f^{\circ n} - c, g^{\circ n} - d)$ for any $c, d \in \bar{\mathbb{Q}}[x]$, but we need the same "target point" to descend from $\mathbb{C}$.

Possible extensions

Will all of the above work when f and g are *rational* functions of the same degree?

Are there counterexamples when f and g are Lattès maps for conjugate endomorphisms of an elliptic curve with complex multiplication?

This isn't clear, since in this case f and g have infinitely many preperiodic points in common and the diagonal is not preperiodic under (f, g), but is isn't clear that iterates of f and g have infinitely many roots in common.

So it is possible the results here avoid the counterexamples to dynamical Bogomolov conjecture. We might also ask for some effectivity here, and perhaps a more elementary proof. We might also ask for something in characteristic p.

Varying the points under nontrivial maps

Let us return now to doing one map f but applying it to different points $a(t)$ and $b(t)$ (instead of just to $a(t) = t$, $b(t) = t$ as above), we also allow f to have coefficients in $\mathbb{C}[t]$ rather than $\mathbb{C}$.

Here is the idea: in the initial powering map case, we have a map $\mathbb{P}^1_{\mathbb{C}(t)} \longrightarrow \mathbb{P}^1_{\mathbb{C}(t)}$ but the map is actually defined over $\mathbb{C}$. One should be able to treat the general case of non-isotrivial maps $\mathbb{P}^1_{\mathbb{C}(t)} \longrightarrow \mathbb{P}^1_{\mathbb{C}(t)}$.
Let $f \in \mathbb{C}[t][x]$ have degree greater than one. Let $a(t), b(t) \in \mathbb{C}[t]$. What can one say about

$$\gcd(f^{\circ n}(a(t)), f^{\circ n}(b(t))?$$

It turns out we can say something unless there is some kind of compositional dependence again.

Varying the points under nontrivial maps, continued

Theorem

(Hsia-Silverman-T) With notation as above, assume further that f is in normal form (monic and with no second-highest degree term). Then there is a finite set S such that every root of $\gcd(f^{\circ n}(a(t)), f^{\circ n}(b(t))$ lies in S unless there is a polynomial $h_t(z) \in \mathbb{C}[t][z]$ that commutes with some power of f and maps some iterate of $a(t)$ to an iterate of $b(t)$, or some iterate of $b(t)$ to $a(t)$.

(When $\deg h_t = 1$, it plays the same role as the symmetry of the Julia set above.)

Note that we do not necessarily get an h such that $\gcd(f^{\circ n}(a(t)), f^{\circ n}(b(t))|h$ for all h. There could be increasing multiplicities. We believe this can happen even without ramification, that $f^{\circ n}(a(t))$ can get closer and closer to zero at some place λ for non-obvious dynamical reasons (not just because of obvious ramification).

Varying the points under nontrivial maps, proof sketch

The proof follows the general outline used before, using points of small height to obtain equality of canonical heights, which gives equality of the preperiodic locus, after a descent to $\bar{\mathbb{Q}}$. The points in this case are specializations λ such that $f_\lambda^{\circ n}(a(\lambda)) = f_\lambda^{\circ n}(b(\lambda)) = 0$. The height functions here are $h_{f_\lambda}(a(\lambda))$ and $h_{f_\lambda}(b\lambda))$. One obtains that the heights of these specializations goes to zero, using Silverman specialization. The theorem then follows from a recent result of Baker and DeMarco.

Generalizing further

- Can f be replaced by a rational function? No obvious counterexamples, and the "Lattès" Legendre family has been treated by DeMarco-Wang-Ye, as mentioned above.
- Can $\gcd(f^{\circ n}(a(t)), f^{\circ n}(b(t)))$ be replaced more generally with $\gcd(f^{\circ n}(a(t)) - c(t), f^{\circ n}(b(t)) - d(t))$? Yes, this all works over $\bar{\mathbb{Q}}$, via Silverman specialization, and the only issue is the descent from $\mathbb{C}$.
- Can $\gcd(f^{\circ n}(a(t)), f^{\circ n}(b(t)))$ be replaced more generally with $\gcd(f^{\circ n}(a(t)), g^{\circ n}(b(t)))$? This is less clear, given some examples in other contexts (mentioned above) for rational functions coming from conjugate endomorphisms of elliptic curves with complex multiplication. But as of yet, no counterexamples have been constructed.