Questions about isogenies, automorphisms and bounds

B.M.

July 20, 2015

Very rough notes for my lecture at the Model Theory Seminar: MSRI April 29, 2014

Contents

1 Introductory remarks

These notes are just to record a few (closely related) Diophantine questions related to finite upper bounds, that have some connection with model theory. None of this are my own results. I'm just formulating some things of interest that one can easily read off from the literature[1]. A number of people have helped me think a bit about some of them and have made computations related to them[2].

In all the cases we'll discuss, the bounds attained come from (model-theoretic) consideration of specific pairs of algebraic varieties yoked together by the existence of a (transcendental) analytic mapping from an open domain of the one onto an open domain of the other. The prototype of this structure is, of course, the exponential function mapping the additive group to the multiplicative group, or the elliptic modular function, $z \mapsto j(z)$, that maps the upper half plane onto the entire complex plane, or finite products of these.

2 Bounds

Alexandru Buium has a very important technology for producing strikingly explicit bounds in certain situations (e.g., twenty years ago—in [1]—he found an explicit upper bound for the number of points on the intersection a curve of genus $g > 1$ embedded in an abelian variety defined over a function field with the group of division points of a subgroup of finite rank r in that abelian variety. For this, Buium had to invent special terminology to describe that huge number–but it was beautifully explicit: $N(!)^{6N+6}$ where $N \geq \max(4, g, r)$. Here $a(!)^n$ means the n-th iterated factorial of a. Explicit effective (doubly exponential) bounds were proved in the context of semi-algebraic varieties over $\mathbb{C}$ by Ehud Hrushovski and Anand Pillay [6] for an analogous question regarding transcendental points in the intersection of finite rank subgroups with subvarieties (that don't contain the 'sum' of two proper positive dimensional subvarieties).

In certain cases of interest here, Jim Freitag and Tom Scanlon will also be producing explicit bounds (strikingly smaller). See the forthcoming [4].

3 A nonmodular hyperelliptic involution

I have been intrigued for quite a while by the following example. Let X be the modular curve $X := X_0(37)$ classifying elliptic curves with a chosen 37-isogeny. That is, any (noncuspidal) point $x \in X$ represents an isomorphism class of pairs $(E_x, C_x \subset E_x)$ where E_x is an elliptic curve and C_x is a subgroup of order 37. Or equivalently one can say that the point represents an isogeny of degree 37,

$$E_x \to E_{w(x)}$$

where $w : X \to X$ is the natural involution—the *Atkin-Lehner involution*– that sends the pair (E, C) to (E', C') where $E' = E/C$ is the quotient elliptic curve obtained by dividing

[1]and that have been—for me–a useful introduction to that literature

[2]Kiren Kedlaya, David Harvey, Kirsten Wickelgren, Jacob Tsimerman, and—in this model theory semester— discussion and correspondence with Jim Freitag, Tom Scanlon and Alexandru Buium

by C, the cyclic group of order 37; and where $C' = E[37]/C$. That is, C' is the cyclic group of order 37 in E' given by the quotient of the group $E[37]$ of all 37-torsion of E, by C.

The modular curve X is of genus two, so there is also a hyperelliptic involution

$$\alpha : X \to X.$$

This automorphism α is defined over $\mathbf{Q}$ and is a natural map having a perfectly nice algebro-geometric interpretation, and commutes with the Atkin-Lehner involution $w : X \to X$. But the reason for my focusing on α is that it is important for the diophantine analysis of this curve over $\mathbf{Q}$ and yet (or I might say: because) it seems to ignore the rest of the "modular structure" (Its graph is not "weakly special" in the product $X_0(37) \times X_0(37)$.) The involution α even destroys the distinction between cuspidal and noncuspidal points. There are two cusps $0, \infty$ in X and the image of this set of cusps under the involution α consists of two noncuspidal points, interchanged by the Atkin-Lehner operator. These two points classify a certain 37-isogeny $E \to E'$ and its dual $E' \to E$, these being the only 37-isogenies defined over $\mathbf{Q}$. It is natural to ask: *what does this automorphism α do to the rest of the Hecke orbit partition of $X_0(37)$?* For example, does it have the following extremely transverse property:

Question 3.1. *Is $|\alpha(Hx) \cap H\alpha(x)| \leq 2$ for all $x \in X$?*

4 The general set-up

Let $X = \mathcal{H}/\Gamma \cup \{\text{cusps}\}$ be a modular curve with $\Gamma \subset \mathrm{GL}_2(\mathbf{Z})$ some congruence subgroup (e.g., $X = X(N)$, or $X_0(N)$, etc.). For $x \in X(\mathbb{C})$ denote by $Hx \subset X(\mathbb{C})$ its full Hecke orbit, meaning the set of points of $X(\mathbb{C})$ achievable by iterated application—starting from the point x— of the Hecke operators T_p for prime numbers p not dividing the level of Γ and the operators U_p when they exist. To the points of Hx will correspond elliptic curves with the sort of level structure classified by X that are isogenous to the elliptic curve that corresponds to the point x. The Hecke orbits partition $X(\mathbb{C})$ into equivalence classes that are each countably infinite for noncuspidal points (and finite for the cusps). In a sense, the quotient by this equivalence relation is a parameter space representing isogeny classes of elliptic curves with the relevant level structure.

Let $\alpha : X \to X$ be an automorphism that *does not preserve the cusps.* For the moment, we'll assume that α is defined over some number field. How does α behave vis a vis the structure defined by Hecke orbits? That is, how does the image of Hecke orbits of $X(\mathbf{C})$ under the automorphism α relate to the Hecke orbit partition?

Here is convenient terminology for the intersection of the Hecke orbit of $\alpha(x)$ and image under α of the Hecke orbit of x:

Definition 4.1. *The* **(Hecke) α-inter-orbit of the point x in X** *is the subset:*

$$H_X(\alpha, x) = H(\alpha, x) := \alpha(Hx) \cap H\alpha(x) \ \subset \ X(\mathbb{C}).$$

5 Basic questions

1. Is $H_X(\alpha, x)$ finite?

   The answer is yes, as can be seen from the model-theoretic literature. More about that below, but see especially [15] and also Martin Orr's [8]. I'm very grateful to Jacob Tsimerman for enlightening conversations we had about this problem and for suggesting Orr's article. In [8] a significantly more extensive, and related, theorem is proved:

   Theorem 5.1. *Let Y be any 'non-special' algebraic curve in $\mathcal{A}_g$, the moduli space of principally polarized abelian varieties. Then the intersection of Y with any isogeny class of abelian varieties represented by $\mathcal{A}_g$ is finite.*

   The fact that Hecke inter-orbits are finite tells us that the curve Y_α given by the graph of α in $X \times X$ can be thought of as parameterizing a certain family of isogeny classes of elliptic curves—*with finite ambiguity.*

   Explicit bounds of any sort—even those like the ones in section 2—expressing specific dependence on the variables α and x would be interesting. For example:

2. Is $|H_X(\alpha, x)|$ bounded, by an upper bound B_X dependent only on X and not on x or α (satisfying our hypotheses; i.e., not preserving cusps)?

   I don't know the answer.

3. If $x \in X(\mathbb{C})$ and α are transcendental, is $|H_X(\alpha, x)|$ bounded, by an upper bound B_X dependent only on X and not on x or α (satisfying our hypotheses; i.e., not preserving cusps)?

   The answer here is yes, by a recent theorem of Freitag and Scanlon. The bound B_X can be taken to be a simple utterly explicit function of the genus of X and the index of Γ.

4. Let p be a prime number (say not dividing the level of X) and consider $X(\bar{\mathbf{F}}_p)$. Is $H_X(\alpha, x)$ finite?

   I don't know the answer. But it is not inconceivable that there is an affirmative answer to this. Here's a heuristic reason why. Let $q = p^n$. Let $X_{/}\mathrm{Spec}(\mathbf{F}_q)$ be the pull-back of a modular curve (of level, say, prime to p). We'll be taking $q >> 0$. We have

   $$|X(\mathbf{F}_q)| = O(q),$$

   and noting—by Honda-Tate, that there are $O(q^{1/2})$ isogeny classes represented by the points in $X(\mathbf{F}_q)$. On the average, therefore, there are $O(q^{1/2})$ elements in each isogeny class. So, if we have an automorphism $\alpha : X \to X$ that we think of as randomly permuting the $O(q)$ points of $X(\mathbf{F}_q)$, then—by this very vague heuristic—on the average there should be $O(1)$ points in the intersection of any $\alpha(Hx)$ and $H\alpha(x)$, i.e., $|H(\alpha, x)|$ should have a finite average, independent of q.

6 General background about isogenies and Galois action

If k is a finite field, or a number field, the *isogeny-class* of an abelian variety over k is determined by the associated ℓ-adic Galois representation on ℓ-power torsion points (for ℓ different from the characteristic of k) thanks to classical results of Tate, and Faltings, respectively. In the finite field case, then, complete knowledge of the isogeny class over the algebraic closure of the finite field is given by a finite collection of Weil numbers (taken up to multiplication by roots of unity).

This is in curious contrast to the information required in order to pin down the precise *isomorphism-class* of the abelian variety over the algebraic closure of the finite field.

When is it the case that knowing the isogeny-class**es** of *two* abelian varieties that are connected by a well-defined recipe, pins down—or at least pins down with finite ambiguity—the isomorphism classes of both?

If the abelian varieties are over a number field, this type of 'knowledge' is equivalent to knowing two Galois representations plus that 'well-defined recipe.' Thinking of it this way puts it—very very vaguely!— into the perspective of general anabelian questions that ask whether curves or varieties can be pinned down in their isomorphism class by Galois representations (this is usually phrased in terms of the natural representation on the fundamental group or some quotient of it).

7 Some ingredients in the proof of finiteness of Hecke inter-orbits

Daniel Bertrand tells me that given the history of the evolution of the technique of proof here, one might call the underlying basic strategy *Pila-Zannier's Strategy* (see [11]). Here I want to recall Tom Scanlon's recent MSRI lecture on o-minimal structures related to algebraic group actions. The format, well known to people in this room, is that we have two algebraic varieties (I'll call them "upstairs" and "downstairs") with an algebraic group acting 'upstairs' and such that there are open analytic domains both upstairs and downstairs where the downstairs analytic domain is the quotient of the upstairs analytic domain by the action of a discrete subgroup of the algebraic group. BUT, the magic here is that one also has an o-minimal $\mathbf{R}_{\mathrm{an,exp}}$ structure giving powerful control of the upstairs $\mathbf{R}$-semi-algebraic geometry and its relation to both up-and-downstairs complex algebraic geometry. Here is what we have specifically:

Let $X = \mathcal{H}/\Gamma \cup \{\text{cusps}\}$ be a modular curve with $\Gamma \subset \mathrm{GL}_2(\mathbf{Z})$ as in section 4. Consider the group scheme $G := \mathrm{GL}_2 \times \mathrm{GL}_2$ (over $\mathbf{Z}$) and have it act coordinate-wise in the product $\mathbf{P}^1 \times \mathbf{P}^1$. In the underlying complex analytic manifold of $\mathbf{P}^1 \times \mathbf{P}^1$, we have:

- the open analytic domain given by the product of two copies of the upper-half-plane

$\mathcal{H} \times \mathcal{H}$,

- the restriction of the above action on $\mathbf{P}^1 \times \mathbf{P}^1$ to yield the natural action of $G(\mathbf{R})$ on $\mathcal{H} \times \mathcal{H}$—rendering $\mathcal{H} \times \mathcal{H}$ a homogenous space for $G(\mathbf{R})$. This action restricts to
- the bi-Hecke orbit' natural action of $G(\mathbf{Q})$ on $\mathcal{H} \times \mathcal{H}$, which in turn restricts to
- the discrete action of $\Gamma \times \Gamma$ on $\mathcal{H} \times \mathcal{H}$, and finally we have
- $\mathcal{F} \subset \mathcal{H} \times \mathcal{H}$, a fundamental domain for this latter action.

$$\begin{array}{ccc} \mathcal{H} \times \mathcal{H} & \xrightarrow{\supset} & \mathcal{F} \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi_{\mathcal{F}}} \\ \mathcal{H}/\Gamma \times \mathcal{H}/\Gamma & \xrightarrow{\subset} & X \times X \end{array}$$

Here $\pi_{\mathcal{F}}$ is π restricted to $\mathcal{F}$. As already mentioned, the action of $G(\mathbf{Q})$ on points $z = (z_1, z_2)$ in $\mathcal{H} \times \mathcal{H}$ relates to the Hecke orbits of the points $x_1, x_2 \in X(\mathbb{C})$, where $x = (x_1, x_2) = \pi(z_1, z_2) \in X \times X$. To make the connection more specific, for $n \geq 1$, let $G(\mathbf{Q})_n \subset G(\mathbf{Q})$ denote the elements of $G(\mathbf{Q}) \subset M_2(\mathbf{Q}) \times M_2(\mathbf{Q})$ consisting of those pairs (in $G(\mathbf{Q})$) of matrices which have the property that their $4+4$ entries, when taken projectively as a point in the bi-projective space $\mathbf{P}^1(\mathbf{Q}) \times \mathbf{P}^1(\mathbf{Q})$, has (standard, multiplicative) height equal to n.

We have:

$$G(\mathbf{Q}) = \sqcup_{n=1}^{\infty} G(\mathbf{Q})_n,$$

and

$$\pi\big(G(\mathbf{Q}) \cdot z\big) = Hx_1 \times Hx_2 \subset X(\mathbf{C}) \times X(\mathbf{C}).$$

Now let's bring α into the picture by considering its graph, Y_α, viewed as a curve in $X \times X$,

$$Y_\alpha(\mathbb{C}) := \{y = (x, \alpha(x)) \ \mid \ \text{where } x \in X(\mathbb{C}))\}.$$

Form its inverse image in $\mathcal{F}$ under the projection mapping $\pi_{\mathcal{F}}$:

$$Z = Z_\alpha := \pi_{\mathcal{F}}^{-1} Y_\alpha(\mathbb{C}),$$

so we have the diagram:

$$\begin{array}{ccccc} \mathcal{H} \times \mathcal{H} & \xrightarrow{\supset} & \mathcal{F} & \xrightarrow{\supset} & Z_\alpha \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi} \\ \mathcal{H}/\Gamma \times \mathcal{H}/\Gamma & \xrightarrow{\subset} & X \times X & \xrightarrow{\supset} & Y_\alpha(\mathbb{C}). \end{array}$$

Here Z_α is an analytic one-dimensional manifold (perhaps with boundary) in $\mathcal{F}$. Since α doesn't preserve cusps, it follows that Y_α is not 'special' and from the work of Ullmo and Yafaev ([15]) one concludes that Z_α contains *no* real semi-algebraic arc.

Consider a point $y = (x, \alpha(x)) \in Y_\alpha(\mathbb{C})$ such that neither x nor $\alpha(x)$ is a cusp. Then there is a point $z \in Z_\alpha$ such that $\pi_{\mathcal{F}}(z) = x$.

Consider

$$\pi_{\mathcal{F}}\big((G(\mathbf{Q}) \cdot z) \cap Z\big),$$

which is just

$$\{(h_1 x, h_2\alpha(x)) \mid \alpha(h_1(x)) = h_2\alpha(x)) \ \mid \ h_1, h_2 \text{ ranging through all Hecke operators}\} = H(\alpha, x),$$

i.e., we have that $\pi_{\mathcal{F}}$ sends $(G(\mathbf{Q}) \cdot z) \cap Z$ onto $H(\alpha, x)$, and therefore the partition

$$\big(G(\mathbf{Q}) \cdot z\big) \cap Z = \sqcup_{n=1}^{\infty} \big(G(\mathbf{Q})_n \cdot z\big) \cap Z,$$

induces a corresponding partition on $H(\alpha, x)$ after projection via $\pi_{\mathcal{F}}$:

$$\begin{array}{ccccc}
\big(G(\mathbf{Q})_n \cdot z\big) \cap Z & \xrightarrow{\subset} & \big(G(\mathbf{Q}) \cdot z\big) \cap Z & \xrightarrow{\subset} & \mathcal{F} \\
\downarrow{\scriptstyle \pi_{\mathcal{F}}} & & \downarrow{\scriptstyle \pi_{\mathcal{F}}} & & \downarrow{\scriptstyle \pi_{\mathcal{F}}} \\
H(\alpha, x)_n & \xrightarrow{\subset} & H(\alpha, x) & \xrightarrow{\subset} & X \times X.
\end{array}$$

Finiteness of $H(\alpha, x)$ will follow from the following upper and lower bounds for $H(\alpha, x)_n$:

- There is a positive number $d := d(\alpha, x) > 0$ such that for any n for which $|H(\alpha, x)_n| > 0$ we have the lower bound $|H(\alpha, x)_n| >> n^d$ (see 7.1, 7.2, 7.3 below)

and

- $|H(\alpha, x)_n| << n^\epsilon$ for any $\epsilon > 0$ (see 7.4 and 7.5 below).

7.1 Galois orbits $\implies$ bounds for $H(\alpha, x)_n$.

Lemma 7.1. *There exists a number $d > 0$ such that for $n >>_d 0$, either $H(\alpha, x)_n$ is empty, or else*

$$|H(\alpha, x)_n| \quad > \quad n^d.$$

In the proof of Lemma 7.1 one might distinguish between two cases:

- the point $x \in X(\mathbb{C})$ is defined over an algebraic number field K, or
- the point x is transcendental,

for they do lead to different types of (proved) bounds.

7.2 The Algebraic Case

If $x = (x_1, x_2)$ is algebraic, and noting that the index n of $H(\alpha, x)_n$ is sufficiently closely related to the degree of isogeny between the elliptic curve representing the point x_1 (respectively x_2) and and the elliptic curves representing the points ξ_1 (respectively ξ_2) for elements $(\xi_1, \xi_2) \in H(\alpha, x)_n$. Since we also assumed that the automorphism α itself is algebraic, one can use the classical study of Galois actions on CM elliptic curves, or–in the non-CM case, the theorem of Serre [14][3] about the relative fullness of action of Galois on torsion points of elliptic curves defined over number fields to get—for an appropriate positive number d—an n^d lower bound on the size of any Galois orbit of any point $(\xi_1, \xi_2) \in H(\alpha, x)_n$.

What is behind this connection between *degree of isogeny* and *Galois orbit* is the following simple computation—simple thought, really; it's not much of a computation. If E is an elliptic curve over a number field $K \subset \bar{\mathbf{Q}}$, and $E[n] \simeq \mathbf{Z}/n\mathbf{Z} \times \mathbf{Z}/n\mathbf{Z}$ denotes the finite $G_K := \mathrm{Gal}(\bar{\mathbf{Q}}/K)$ module given by the kernel of multiplication by n in $E(\bar{\mathbf{Q}})$, then the set of cyclic degree n isogenies of E is in one:one correspondence with the set of all cyclic subgroups of order n in $\mathbf{Z}/n\mathbf{Z} \times \mathbf{Z}/n\mathbf{Z} \simeq E[n]$. This latter set is of cardinality $n + o(n)$ and forms a single orbit under the natural action of $\mathrm{PGL}_2(\mathbf{Z}/n\mathbf{Z})$. So, if, as n tends to infinity, the image of G_K in $\mathrm{PGL}_2(\mathbf{Z}/n\mathbf{Z})$ is of index admitting a finite upper bound b we get that the size of a G_K-orbit of an isogeny of degree n is $\geq n/b$ for $n >> 0$.

Of course, there may be no points at all in $H(\alpha, x)_n$. This gives Lemma 7.1.

7.3 The Transcendental Case

In fact, one has a fairly similar proof if x is transcendental! This uses the standard technique of, for example, dealing with elliptic curves—say—of the form $E : y^2 = x^3 + \pi x + e$. Now assume, for fun, that π and e are independent transcendental numbers. To keep from being hypnotized by their lore, the situation is best psychologically dealt

---

[3]One also can use the Masser-Wüstholz isogeny theorem for abelian varieties [7]:

Theorem 7.2. *Let Aand B be principally polarised abelian varieties over a number field K. Suppose that there exists some isogeny $A \to B$. Then there is an isogeny $A \to B$ of degree at most $c \max([K : \mathbf{Q}], h(A))^\kappa$ where c and κ are constants depending only on the dimension of A.*

with by just renaming the transcendentals π and e (call them, say, u and v) and noting that you now have a perfect algebraic replica of your elliptic curve given by the equation $y^2 = x^3 + ux + v$ which can be comfortably viewed as a two-parameter family of elliptic curves, the parameters being (u, v). We get a member of that family for any specialization of the variables u and v leading to a cubic equation with nonvanishing discriminant $\Delta = -4u^3 - 27v^2$. This particular example, then, forms a family $E \to S$ of elliptic curves over $S := \mathrm{Spec}(\mathbf{Z}[u, v, \Delta^{-1}]$. The initial elliptic curve is now conceived as just the specialization of this family gotten by sending (u, v) to (π, e). But after all, how drastically specialized can it be, with its transcendental, and independent, coefficients?

Passing to a generic point one gets $y^2 = x^3 + ux + v$ over $\mathbf{Q}(u, v)$ and since the absolute Galois group of $\mathbf{Q}(u, v)$ is substantial, you continue to work, as in the algebraic case. But this time, you must use a version of Masser-Wüstholz 's isogeny theorem [7] valid for finitely generated fields. This is done fairly explicitly in section 5 of Orr's [8].

7.4 Consequences of the Pila-Wilkie and Pila Theorems

Here we must deal with **definable blocks** which—for this context— let us mean either *points* or *definable positive dimensional subsets of semi-algebraic sets in* $\mathbf{R}^4 \supset \mathcal{H} \times \mathcal{H}$ this latter type of 'block' I'll refer to as an *honest block*. Although I'd like to understand this, and I don't yet, using either [9], or Proposition 5.1 of [12] or Proposition 5.2 of [13], we get that the analytic manifold Z does not contain any honest blocks. This is because, if so, Z would be algebraic and by the Ullmo-Yafaev result [15] quoted earlier—plus our hypothesis on α—this is not the case.

But we also get the following corollary of Pila-Wilkie's Theorem, [10], as formulated by Pila; see especially Martin Orr's [8].

Lemma 7.3. *For any $\epsilon > 0$ and $n >>_\epsilon 0$, then*

$$\cup_{k \le n} \big(G(\mathbf{Q})_k \cdot z\big) \cap Z_\alpha$$

is contained in at most n^ϵ definable blocks.

Putting the above together, Z would contain an honest block and then be algebraic (by [9] and Prop. 5.1 of [12]) which it is not.

Corollary 7.4. *$H(\alpha, x)$ is finite.*

Proof: Comparing upper and lower bounds, it is clear that there must be at least one honest block in Z if $H(\alpha, x)_n$ were nonempty for infinitely many n.

7.5 The Transcendental Case revisited

The families $E \to S$ we produce from such transcendental points $x \in X$ give us a variation of Hodge structure over the parameter space $S \otimes \mathbb{C}$, and related to this: there are corresponding 'Picard-Fuchs equations.' One has the Hodge filtration on $H_{DR}(E/S)$ viewed as a flag of vector bundles over S and the comparison theorem with singular cohomology provides us with a connection on the vector bundle over $S \otimes \mathbb{C}$. This connects with the manner in which Freitag and Scanlon re-express Hecke orbits in terms of a differential algebraic structure on the third jet spaces related to the modular curve X, and (using [6]) view $H(\alpha, x)$ as the K-valued points of a differential algebraic variety (call it $V_{\mathrm{F-S},X}(\alpha, x)$) whose equations are given fairly explicitly and whose degree one can bound, via nothing more complicated than Bezout's Theorem. Moreover—thanks to Orr's finiteness theorem and the identification of the Hecke α-inter-orbit of x with $V = V_{\mathrm{F-S},X}(\alpha, x)(\mathbb{C})$, the algebraic variety V is consequently zero-dimensional, so one gets an actual numerical bound directly from its degree.

8 A more general question about abelian varieties

Take two Shimura varieties X, Y, classifying abelian varieties of specific types (with specified choices of additional structure) and a mapping $f : X \to Y$ (sending X onto its image in Y by a finite morphism) *that does not respect the Shimura variety structures of X and Y*. More precisely, assume that the graph of f in $X \times Y$ contains no positive dimensional 'weakly special subvarieties' of the Shimura variety $X \times Y$.

Given a point $x \in X$ that 'classifies' an abelian variety A_x let $H_{\{X\}}(x)$ be the subset of points $\{x'\}$ in the Shimura variety X that classify abelian varieties $A_{x'}$ isogenous to A_x. Consider, now the image $f(H_{\{X\}}(x)) \subset Y$ and, again, we would like theorems that guarantee that $H_{\{Y\}}(f(x') \cap f(H_{\{X\}}(x'))$ is finite for any $x' \in H_{\{X\}}(x)$. Here $H_{\{Y\}}(f(x')$ is the full Hecke orbit of $f(x')$ in the Shimura variety Y. What types of upper bounds can one expect?

This question is phrased about the graph Γ_f of the function f in $Z := X \times Y$ and its relation to the partition of $X \times Y$ into subsets that give full Hecke orbits

$$H_{\{X\}}(x) \times H_{\{Y\}}(y).$$

In the same spirit, but more generally, one might ask, for any point z in any non-weakly special irreducible algebraic curve contained in a Shimura variety, $\Gamma \subset Z$, what can be said about upper bounds for

$$|H_{\{Z\}}(z) \cap \Gamma|?$$

This intersection is shown to be finite in [8]. That is, when the Shimura variety represents abelian varieties with extra structure, we have:

Theorem 8.1. (Orr; [8]) *If the curve* Γ *is not weakly special, for any abelian variety (with extra structure)* A_z *'classified by' a point* $z \in \Gamma(\mathbb{C})$ *there are only finitely many abelian varieties* $A_{z'}$ *classified by points* $z' \in \Gamma(\mathbb{C})$ *that are isogenous to* A_z*.*

A variant would be to ask the same type of finiteness questions, with bounds, for ℓ-power isogenies for a given prime ℓ. This connects to:

9 Isogeny classes of abelian varieties constructed from elliptic curves

Here is just one specific example of such a construction.

Let E be an elliptic curve over an algebraically closed field k. Let $m, n > 1$ be numbers prime to the characteristic of k and P_n a choice of n-division polynomial for E (unique up to nonzero scalar multiple). So P_n has a pole of order $n^2 - 1$ at the origin and simple zeroes at each point of $E(\mathbb{C})$ of order n and is regular elsewhere. Let $\psi_{m,n} : C_{m,n} \to E$ be the (smooth projective) curve $C_{m,n}(E)$ obtained by extracting an m-th root of P_n, and let $A_{m,n}(E)$ denote the quotient:

$$0 \to \mathrm{Pic}^0(E) \to \mathrm{Pic}^0\big(C_{m,n}(E)\big) \to A_{m,n}(E) \to 0.$$

The rule that sends $E \mapsto A_{m,n}(E)$ can be construed to be a morphism $\phi_{m,n} : X \to Y$ where X is the moduli stack classifying elliptic curves, and Y the Shimura variety that classifies the type of abelian varieties that this $\{m, n\}$-construction provides.

Definition 9.1. *The abelian variety* $A_{m,n}(E)$ *is called the* $\{m, n\}$**-abelian variety constructed from** E*.*

What can one say about $|H_{\{Y\}}\phi_{m,n}(x) \cap \phi_{m,n}(H_{\{X\}}x)|$ for points $x \in X$?[4]

Or, what amounts to the same question, we merely consider the rule sending $E \mapsto \mathrm{Pic}^0\big(C_{m,n}(E)\big)$ viewed as a morphism $f_{m,n} : X \to \mathcal{A}_g$, the moduli space of principally polarized abelian varieties of dimension $g = g(m,n) = \frac{1}{2}\big((m-1)n^2 - d + 3\big)$, where $d := g.c.d(m, n^2 - 1)$.

Question 9.1. *For which pairs* $m, n > 1$*—if any—is the image of this morphism* $f_{m,n}$ *weakly special?*

[4] I'm thankful to Kirsten Wickelgren, David Harvey, and Kirin Kedlaya, for discussions—some years ago—about concrete aspects of this problem, and specific computations over finite fields.

References

[1] Buium, A.: On a question of B. Mazur, Duke Mathematical Journal, **75** (3) (1994) 639-644

[2] Buium, A.: Geometry of differential polynomial functions, iii:moduli spaces, American Journal of Mathematics, (1995) 1-73

[3] Cluckers, R., Compte, G., Loeser, F.: Non-archimedean Yomdin-Gromov parametrizations and points of bounded height, arXiv:1404.1952

[4] Freitag, J., Scanlon, T. To appear.

[5] Habegger, P., Pila,J.: Some unlikely intersections beyond André-Oort, Compos. Math. **148** (2012), no. 1, 1-27.

[6] Hrushovski, E., Pillay, A.: , Effective bounds for the number of transcendental points on subvarieties of semi-abelian varieties. Amer. J. Math. **122** (2000), no. 3, 439-450.

[7] Masser, D., Wüstholz, G.: Isogeny estimates for abelian varieties, and finiteness theorems, Ann. of Math. (2) **137** (1993), no. 3, 459-472.

[8] Orr, M.: Families of abelian varieties with many isogenous fibers, arXiv:1209.3653v2

[9] Pila, J.: Rational points of definable sets and results of André-Oort-Manin-Mumford type, Int. Math. Res. Not. IMRN (2009), **13**, 2476-2507.

[10] J. Pila, J., Wilkie, A.J.: The rational points of a denable set, Duke Math. J., **33** (2006), 591-616

[11] Pila, J., Zannier, U.: Rational points in periodic analytic sets and the Manin-Mumford conjecture, arXiv:0802.4016

[12] Scanlon, T.: A proof of the André-Oort conjecture via mathematical logic [after Pila, Wilkie, and Zanner] Séminaire Bourbaki, **1037** 63ème année 2010-2011

[13] Scanlon, T.: O-minimality as an approach to the André-Oort conjecture, preprint (U. Cal., Berkeley) 1-50.

[14] Serre J.-P., Propriétés galoisiennes des points d'ordre fini des courbes elliptiques, Inventiones mathematicae, **15** (1972) 259-331

[15] Ullmo, E., Yafaev, A.: The André-Oort conjecture for products of modular curves, Clay Mathematics Proceedings **8** (2009) 431-439