J. reine angew. Math. **705** (2015), 211–231
DOI 10.1515/crelle-2013-0058
Journal für die reine und angewandte Mathematik
© De Gruyter 2015

Families of abelian varieties with many isogenous fibres

By *Martin Orr* at Orsay

Abstract. Let Z be a subvariety of the moduli space of principally polarised abelian varieties of dimension g over the complex numbers. Suppose that Z contains a Zariski dense set of points which correspond to abelian varieties from a single isogeny class. A generalisation of a conjecture of André and Pink predicts that Z is a weakly special subvariety. We prove this when $\dim Z = 1$ using the Pila–Zannier method and the Masser–Wüstholz isogeny theorem. This generalises results of Edixhoven and Yafaev when the Hecke orbit consists of CM points and of Pink when it consists of Galois generic points.

1. Introduction

Let $\mathcal{A}_g$ denote the Siegel moduli space of principally polarised abelian varieties of dimension g. We consider the following conjecture. In particular we prove the conjecture when Z is a curve, and make some progress on higher-dimensional cases.

Conjecture 1.1. *Let Λ be the isogeny class of a point $s \in \mathcal{A}_g(\mathbb{C})$. Let Z be an irreducible closed subvariety of $\mathcal{A}_g$ such that $Z \cap \Lambda$ is Zariski dense in Z. Then Z is a weakly special subvariety of $\mathcal{A}_g$.*

Theorem 1.2. *Conjecture* 1.1 *holds when Z is a curve.*

Theorem 1.3. *Let Λ be the isogeny class of a point $s \in \mathcal{A}_g(\mathbb{C})$. Let Z be an irreducible closed subvariety of $\mathcal{A}_g$ such that $Z \cap \Lambda$ is Zariski dense in Z. Then there is a special subvariety $S \subset \mathcal{A}_g$ which is isomorphic to a product of Shimura varieties $S_1 \times S_2$ with $\dim S_1 > 0$, and such that*

$$Z = S_1 \times Z' \subset S$$

for some irreducible closed subvariety $Z' \subset S_2$.

Theorem 1.3, but not Theorem 1.2, depends on results concerning the hyperbolic Ax–Lindemann conjecture from recent preprints of Pila and Tsimerman [19] and of Ullmo [27].

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

Conjecture 1.1 is a consequence of the Zilber–Pink conjecture on subvarieties of Shimura varieties [23]. For a statement of the Zilber–Pink conjecture and proof that it implies Conjecture 1.1, see Section 2.

Conjecture 1.1 is slightly more general than the $S = \mathcal{A}_g$ case of the following conjecture of André and Pink, because the isogeny class of $s \in \mathcal{A}_g(\mathbb{C})$ is sometimes bigger than the Hecke orbit: by *isogeny class* we mean the set of points $t \in \mathcal{A}_g(\mathbb{C})$ such that the corresponding abelian variety A_t is isogenous to A_s, with no condition of compatibility between isogeny and polarisations. On the other hand the *Hecke orbit* consists of those points for which there is a *polarised isogeny* between the principally polarised abelian varieties, that is, an isogeny $\phi\colon A_s \to A_t$ satisfying $\phi^*\lambda_t \in \mathbb{Z}\lambda_s$, where λ_s and λ_t are the polarisations. In the case of $\mathcal{A}_g$ there is no difference between Hecke orbits and Pink's generalised Hecke orbits.

Conjecture 1.4 ([1, Chapter X, Problem 3], [22, Conjecture 1.6]). *Let S be a mixed Shimura variety over $\mathbb{C}$ and $\Lambda \subset S$ the generalised Hecke orbit of a point $s \in S$. Let $Z \subset S$ be an irreducible closed algebraic subvariety such that $Z \cap \Lambda$ is Zariski dense in Z. Then Z is a weakly special subvariety of S.*

Some cases of Conjecture 1.1 are already known: If the point s is Galois generic, then the isogeny class and the Hecke orbit coincide, and Conjecture 1.1 follows from equidistribution results of Clozel, Oh and Ullmo, as was shown by Pink [22]. When s is a special point, Theorem 1.2 was proved by Edixhoven and Yafaev [8] by exploiting the fact that Galois orbits of special points are contained in $Z \cap T_g Z$ for suitable Hecke operators T_g and these Galois orbits tend to be large compared to the degree of T_g. When s corresponds to a product of elliptic curves, Habegger and Pila [10] proved the theorem using the method we extend here.

The terminology "weakly special subvariety" was introduced by Pink [22], although the concept was first studied by Moonen [14]. Moonen showed that a subvariety of a Shimura variety is totally geodesic (in the sense of differential geometry) if and only if it satisfies the following definition. An algebraic subvariety Z of $\mathcal{A}_g$ is called *weakly special* if there exist a sub-Shimura datum (H, X_H) of $(\mathrm{GSp}_{2g}, \mathcal{H}_g^{\pm})$, a decomposition

$$(H^{\mathrm{ad}}, X_H^{\mathrm{ad}}) = (H_1, X_1) \times (H_2, X_2)$$

and a point $x_2 \in X_2$ such that Z is the image in $\mathcal{A}_g$ of $X_1 \times \{x_2\}$. In other words, to say that Z is weakly special means that we can choose S, S_1, S_2 in the conclusion of Theorem 1.3 such that Z' is a single point in S_2. For more details, see Section 2.4.

In this article we will use a characterisation of weakly special subvarieties due to Ullmo and Yafaev [28]: Z is weakly special if and only if an irreducible component of $\pi^{-1}(Z)$ is algebraic, where π is the quotient map $\mathcal{H}_g \to \mathcal{A}_g$ and $\mathcal{H}_g \subset \mathrm{M}_{2g\times 2g}(\mathbb{C})$ is the Siegel upper half space. Here we call a subvariety of $\mathcal{H}_g$ *algebraic* if it is a connected component of $W_0 \cap \mathcal{H}_g$ for some algebraic variety $W_0 \subset \mathrm{M}_{2g\times 2g}(\mathbb{C})$. In order to prove Theorem 1.3 we require a strengthening of this characterisation called the hyperbolic Ax–Lindemann conjecture: if W is a maximal algebraic subvariety of $\pi^{-1}(Z)$, then $\pi(W)$ is algebraic. A proof of the Ax–Lindemann conjecture for $\mathcal{A}_g$ was recently announced by Pila and Tsimerman [19].

Our proof of Theorems 1.2 and 1.3 follows the method proposed by Pila and Zannier for proving the Manin–Mumford and André–Oort conjectures [21]. This is based upon counting rational points of bounded height in certain analytic subsets of $\mathcal{H}_g$, and applying the Pila–Wilkie counting theorem on sets definable in o-minimal structures.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

The central part of the proof of Theorems 1.2 and 1.3 is in Section 3. This uses a strong version of the Pila–Wilkie counting theorem involving definable blocks. The other ingredients are an upper bound for the heights of matrices in $\mathrm{GL}_{2g}(\mathbb{Q})$ relating isogenous points, proved in Section 4, and a lower bound for the Galois degrees of principally polarised abelian varieties in an isogeny class, derived from the Masser–Wüstholz isogeny theorem [11].

In Section 5 we use a specialisation argument to prove a version of the Masser–Wüstholz isogeny theorem for finitely generated fields of characteristic 0, generalising the original theorem which was valid only over number fields. This is necessary in order to prove Theorems 1.2 and 1.3 for points s and subvarieties Z defined over $\mathbb{C}$ and not only over $\bar{\mathbb{Q}}$.

Now we consider some generalisations of Theorem 1.2. Theorem 1.2 immediately implies Conjecture 1.4 for curves Z in Shimura varieties S of Hodge type, if we restrict to usual Hecke orbits. This is because, by the definition of a Shimura variety of Hodge type, there is a finite morphism $f: S \to \mathcal{A}_g$ for some g such that the image of each Hecke orbit in S is contained in a Hecke orbit of $\mathcal{A}_g$, and $Z \subset S$ is weakly special if and only if $f(Z) \subset \mathcal{A}_g$ is weakly special. However this does not imply Conjecture 1.4 for generalised Hecke orbits in Shimura varieties of Hodge type, as a generalised Hecke orbit in S may map into infinitely many isogeny classes in $\mathcal{A}_g$.

Because of the use of the Masser–Wüstholz theorem, our method applies only to Shimura varieties parameterising abelian varieties, i.e. those of Hodge type. In particular, let us compare with [8, Theorem 1.2]. Take any Shimura datum (G, X). Edixhoven and Yafaev generalise Hecke orbits by choosing a representation of G and considering a set of points where the induced $\mathbb{Q}$-Hodge structures are isomorphic. The Masser–Wüstholz theorem can be used only when these Hodge structures have type $(-1, 0) + (0, -1)$. Hence our method lacks a key advantage of Edixhoven and Yafaev's formulation, namely that they can replace G by a subgroup so that Z is Hodge generic, or by its adjoint group.

This restriction to isogeny classes of abelian varieties rather than generalised Hecke orbits is related to our inability to prove the full Conjecture 1.1. In the case of the André–Oort conjecture, a conclusion as in Theorem 1.3 implies the full conjecture by induction on $\dim Z$ (see [27]). This is because, when Z is of the form $S_1 \times Z'$, special points in Z project to special points in Z'.

This does not work for Conjecture 1.1 because the hypothesis that $Z = S_1 \times Z'$ contains a dense set of points from a single isogeny class does not imply the same thing for $\{x_1\} \times Z'$, where we fix a point $x_1 \in S_1$ in order to realise Z' as a subvariety of $\mathcal{A}_g$. The problem is that the decomposition $S = S_1 \times S_2$ need not have an interpretation in terms of moduli of abelian varieties. For example this happens in André and Borovoi's example of a subvariety $S \subset \mathcal{A}_8$ which decomposes as a product of Shimura varieties but where the generic abelian variety in the family parameterised by S is simple.

Acknowledgement. I am grateful to Emmanuel Ullmo for suggesting to me the problem treated in this paper and for regular conversations during its preparation. I would also like to thank Barinder Banwait for his comments on an earlier version of the manuscript, and Gaël Rémond for remarks on Theorem 5.1. I am grateful to the referee for their detailed comments.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

2. Shimura varieties and the Zilber–Pink conjecture

In this article we consider only the moduli space of abelian varieties and its subvarieties, however to place Conjecture 1.1 in its proper context we need to consider Shimura varieties. For the convenience of the reader, we briefly summarise the theory of Shimura varieties and the Zilber–Pink conjecture in this section. This contains no original material: the primary sources are [5] for Shimura varieties, [14] for special and weakly special subvarieties (special subvarieties are what Moonen calls subvarieties of Hodge type) and [23] for the Zilber–Pink conjecture. We also prove that the Zilber–Pink conjecture implies Conjecture 1.1 by an argument which is essentially due to Pink.

2.1. Shimura varieties. A *Shimura datum* is defined to be a pair (G, X) where G is a reductive algebraic group over $\mathbb{Q}$ and X is a $G(\mathbb{R})$-conjugacy class of homomorphisms

$$h\colon \mathrm{Res}_{\mathbb{C}/\mathbb{R}}\,\mathbb{G}_m \to G_{\mathbb{R}}$$

satisfying the following conditions:

(SV1) the Hodge structure on the adjoint representation of G induced by h has type contained in $\{(-1,1),(0,0),(1,-1)\}$;

(SV2) $\mathrm{int}\,h(i)$ induces a Cartan involution of the adjoint group of $G(\mathbb{R})$;

(SV3) G^{ad} has no factor defined over $\mathbb{Q}$ on which the projection of h is trivial.

Under these conditions, each connected component of X is a Hermitian symmetric domain on which the identity component $G(\mathbb{R})^+$ of $G(\mathbb{R})$ acts holomorphically.

Let K be a compact open subgroup of $G(\mathbb{A}_f)$, where $\mathbb{A}_f$ is the ring of finite adeles. We define

$$\mathrm{Sh}_K(G,X) = G(\mathbb{Q})\backslash X \times G(\mathbb{A}_f)/K,$$

where $G(\mathbb{Q})$ acts diagonally on $X \times G(\mathbb{A}_f)$ on the left and K acts on $G(\mathbb{A}_f)$ only on the right. Deligne [5] showed that $\mathrm{Sh}_K(G,X)$ can be given the structure of an algebraic variety over a number field, called a *Shimura variety*.

Choose a connected component X^+ of X. Let $G(\mathbb{R})_+$ be the preimage of the identity component (in the analytic topology) of $G^{\mathrm{ad}}(\mathbb{R})$ in $G(\mathbb{R})$; this is the stabiliser of X^+. Let $G(\mathbb{Q})_+ = G(\mathbb{R})_+ \cap G(\mathbb{Q})$.

The image of $X^+ \times \{1\} \subset X \times G(\mathbb{A}_f)$ in $\mathrm{Sh}_K(G,X)_{\mathbb{C}}$ is called the *neutral component* of $\mathrm{Sh}_K(G,X)_{\mathbb{C}}$. As a complex manifold it is canonically isomorphic to

$$\Gamma\backslash X^+$$

where $\Gamma = K \cap G(\mathbb{Q})_+$ is a congruence subgroup of $G(\mathbb{Q})_+$.

Let (G,X) be a Shimura datum and let $G^{\mathrm{ad}} = G/Z(G)$ be the adjoint group of G, where $Z(G)$ is the centre of G. We get a new Shimura datum $(G^{\mathrm{ad}}, X^{\mathrm{ad}})$ by letting X^{ad} be the $G^{\mathrm{ad}}(\mathbb{R})$-conjugacy class of morphisms $\mathrm{Res}_{\mathbb{C}/\mathbb{R}}\,\mathbb{G}_m \to G^{\mathrm{ad}}_{\mathbb{R}}$ containing $\pi \circ x$ for $x \in X$, where π is the quotient map $G \to G^{\mathrm{ad}}$. The map $X \to X^{\mathrm{ad}}$ is an injection whose image is a union of connected components of X^{ad} (see [14, Section 2.1]).

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

2.2. Moduli of abelian varieties. The fundamental example of a Shimura variety is the moduli space of principally polarised abelian varieties of dimension g. We recall briefly that a *polarisation* of an abelian variety of A is an isogeny $A \to A^\vee$ to the dual variety satisfying a certain positivity condition. Any polarisation induces a symplectic form

$$H_1(A, \mathbb{Z}) \times H_1(A, \mathbb{Z}) \to \mathbb{Z}.$$

A polarisation is *principal* if it has degree 1, that is, if it is an isomorphism $A \to A^\vee$.

The moduli space of principally polarised abelian varieties of dimension g is associated with the Shimura datum $(\mathrm{GSp}_{2g}, \mathcal{H}_g^{\pm})$ where GSp_{2g} is the group of symplectic similitudes and the conjugacy class of Hodge parameters can be identified with the union of the Siegel upper and lower half spaces

$$\mathcal{H}_g = \{Z \in \mathrm{M}_g(\mathbb{C}) \mid Z \text{ is symmetric and } \mathrm{Im}\, Z \text{ is positive definite}\},$$
$$\mathcal{H}_g^- = \{Z \in \mathrm{M}_g(\mathbb{C}) \mid Z \text{ is symmetric and } \mathrm{Im}\, Z \text{ is negative definite}\}.$$

The action of $\mathrm{GSp}_{2g}(\mathbb{R})$ on $\mathcal{H}_g^{\pm}$ is given by

$$\left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right) Z = (AZ + B)(CZ + D)^{-1}, \quad A, B, C, D \in \mathrm{M}_g(\mathbb{R}).$$

Taking $K = \mathrm{GSp}_{2g}(\hat{\mathbb{Z}})$, we get the Shimura variety $\mathcal{A}_g$ whose $\mathbb{C}$-points are in bijection with isomorphism classes of principally polarised abelian varieties of dimension g over $\mathbb{C}$.

2.3. Hecke correspondences. Let (G, X) be a Shimura datum and $K \subset G(\mathbb{A}_f)$ a compact open subgroup. Choose $g \in G(\mathbb{A}_f)$ and let $K_g = K \cap gKg^{-1}$. The inclusion $K_g \hookrightarrow K$ induces a finite morphism

$$\pi \colon \mathrm{Sh}_{K_g}(G, X) \to \mathrm{Sh}_K(G, X).$$

Because K_g is normalised by g, we also have an automorphism

$$\tau_g \colon \mathrm{Sh}_{K_g}(G, X) \to \mathrm{Sh}_{K_g}(G, X)$$

which sends the double coset $[x, \theta K_g]$ to $[x, \theta g K_g]$. We define a finite correspondence T_g on $\mathrm{Sh}_K(G, X)$ by

$$\begin{array}{ccc} & \mathrm{Sh}_{K_g}(G, X) & \\ {\scriptstyle \pi}\swarrow & & \searrow{\scriptstyle \pi \circ \tau_g} \\ \mathrm{Sh}_K(G, X) & \xrightarrow{\quad T_g \quad} & \mathrm{Sh}_K(G, X). \end{array}$$

Such correspondences are called *Hecke correspondences* and, for any point $s \in \mathrm{Sh}_K(G, X)$, the *Hecke orbit* of s is the union

$$\bigcup_{g \in G(\mathbb{A}_f)} T_g.s.$$

The Hecke correspondence T_g depends only on the double coset KgK. If we consider only Hecke correspondences which preserve the neutral component of $\mathrm{Sh}_K(G, X)$, then each double coset contains an element of $G(\mathbb{Q})_+$.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

In the case of the moduli space of principally polarised abelian varieties, whenever two points $s, t \in \mathcal{A}_g$ lie in the same Hecke orbit, the associated abelian varieties are isogenous. The converse is not true: s and t are in the same Hecke orbit if and only if there is a *polarised isogeny* between the associated principally polarised abelian varieties, that is, an isogeny $f: A_s \to A_t$ such that

$$f^*\lambda_t = n.\lambda_t$$

for some $n \in \mathbb{Z}$.

2.4. Special and weakly special subvarieties. Let (G, X) be a Shimura datum and $K \subset G(\mathbb{A}_f)$ a compact open subgroup. A Shimura datum (H, X_H) is a *Shimura sub-datum* of (G, X) if $H \subset G$ and $X_H \subset X$. Letting $K_H = K \cap H(\mathbb{A}_f)$, we get a finite morphism $\iota_H: \mathrm{Sh}_{K_H}(H, X_H) \to \mathrm{Sh}_K(G, X)$.

An irreducible subvariety of $S \subset \mathrm{Sh}_K(G, X)$ is said to be a *special subvariety* if there exist a Shimura sub-datum $(H, X_H) \subset (G, X)$ and an element $g \in G(\mathbb{A}_f)$ such that S is an irreducible component of the image of

$$T_g \circ \iota_H: \mathrm{Sh}_{K_H}(H, X_H) \to \mathrm{Sh}_K(G, X).$$

The Hecke correspondence T_g is needed in this definition only because the image of ι_H might not intersect every geometrically connected component of $\mathrm{Sh}_K(G, X)$. In the case of a geometrically connected Shimura variety such as $\mathcal{A}_g$, every special subvariety is simply an irreducible component of the image of ι_H for some Shimura sub-datum.

A *Shimura morphism* of Shimura varieties is a morphism induced by a homomorphism of the underlying algebraic groups. Let $S = \mathrm{Sh}_K(G, X)$ and $S_1 = \mathrm{Sh}_{K_1}(H_1, X_1)$, and let $f: S \to S_1$ be a surjective Shimura morphism. Then the adjoint Shimura datum $(G^{\mathrm{ad}}, X^{\mathrm{ad}})$ splits as a direct product $(H_1^{\mathrm{ad}} \times H_2, X_1^{\mathrm{ad}} \times X_2)$ for some adjoint semisimple group H_2. If the compact open subgroup $K^{\mathrm{ad}} \subset G^{\mathrm{ad}}(\mathbb{A}_f)$ splits as a direct product $K_1^{\mathrm{ad}} \times K_2$, then S is the union of some connected components of $S_1^{\mathrm{ad}} \times \mathrm{Sh}_{K_2}(H_2, X_2)$. (We can always obtain such a decomposition of K^{ad} by replacing K by a subgroup of finite index.)

An irreducible subvariety $Z \subset \mathrm{Sh}_K(G, X)$ is a *weakly special subvariety* if there exist a sub-Shimura datum $(H, X_H) \subset (G, X)$, an element $g \in G(\mathbb{A}_f)$, some other Shimura datum (H', X') and a Shimura morphism

$$f: \mathrm{Sh}_{K_H}(H, X_H) \to \mathrm{Sh}_{K_1}(H_1, X_1)$$

such that Z is an irreducible component of the image under $T_g \circ \iota_H$ of the fibre $f^{-1}(s)$ for some point $s \in \mathrm{Sh}_{K_1}(H_1, X_1)$.

An example of a weakly special subvariety is the subvariety of $\mathcal{A}_2$ parameterising principally polarised abelian surfaces of the form $E_0 \times E$, where E_0 is a fixed elliptic curve and E a varying elliptic curve. This is a special subvariety if and only if E_0 has complex multiplication. Here the Shimura variety $\mathrm{Sh}_{K_H}(H, X_H)$ is $\mathcal{A}_1 \times \mathcal{A}_1$ and the Shimura morphism f is the projection onto the first factor. Note that in this example, $\iota_H: \mathcal{A}_1 \times \mathcal{A}_1 \to \mathcal{A}_2$ is not injective but is the quotient by the action of $\mathbb{Z}/2$ exchanging the two factors $\mathcal{A}_1$.

By the above discussion of the structure of Shimura morphisms, every weakly special subvariety arises from a decomposition of $(H^{\mathrm{ad}}, X_H^{\mathrm{ad}})$ as a direct product. In particular, the definition of weakly special subvarieties in terms of fibres of Shimura morphisms is equivalent to the definition in the introduction in terms of a decomposition of an adjoint Shimura datum.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

However, not every weakly special subvariety of $\mathcal{A}_g$ arises from a product decomposition of the associated abelian varieties as in the example. This is important as it prevents us from proving Conjecture 1.1. The smallest example of a weakly special subvariety of $\mathcal{A}_g$ which does not come from a product decomposition of abelian varieties is for $g = 8$. It is due to André [2] and Borovoi and is described by Moonen [14].

2.5. The Zilber–Pink conjecture. Let S be a Shimura variety. If Z is a subvariety of S, let S_Z be the smallest special subvariety of S containing Z. This exists because every connected component of an intersection of special subvarieties is special. We define the *defect* of Z to be $\dim S_Z - \dim Z$.

Conjecture 2.1 ([23, Conjecture 1.1]). *Let S be a Shimura variety and $Z \subset S$ an irreducible subvariety over $\mathbb{C}$. Suppose that Z contains a Zariski dense set of points of defect at most d. Then the defect of Z itself is at most d.*

If we take $d = 0$ then this becomes the André–Oort conjecture: an irreducible subvariety of a Shimura variety containing a Zariski dense set of special points is a special subvariety.

Lemma 2.2. *Conjecture* 2.1 *implies Conjecture* 1.1.

Proof. This proof is based on [23, Theorem 3.3]. Since Conjecture 1.1 concerns isogeny classes rather than Hecke orbits, this lemma is slightly stronger than the $\mathcal{A}_g$ case of Pink's theorem.

The only addition we need to make to Pink's proof is to note that if $s, t \in \mathcal{A}_g$ correspond to isogenous abelian varieties, then $s \in \mathcal{A}_g$ and $(s, t) \in \mathcal{A}_g \times \mathcal{A}_g$ each have the same defect. This follows from the fact that the abelian varieties A_s and $A_s \times A_t$ have isomorphic Mumford–Tate groups. Specifically, if $M \subset \mathrm{GSp}_{2g}(\mathbb{Q})$ is the Mumford–Tate group of A_s and $g \in \mathrm{GL}_{2g}(\mathbb{Q})$ is the rational representation of some isogeny $A_t \to A_s$ (as in Section 4.2), then the Mumford–Tate group of $A_s \times A_t$ is the image of the embedding $M \to \mathrm{GL}_{2g} \times \mathrm{GL}_{2g}$ given by $x \mapsto (x, gxg^{-1})$.

Let S be the smallest special subvariety of $\mathcal{A}_g$ containing the point s, and let $d = \dim S$. Let $S' \subset \mathcal{A}_g \times \mathcal{A}_g$ be the smallest special subvariety of $\mathcal{A}_g \times \mathcal{A}_g$ containing $\{s\} \times Z$.

As we noted above, each point of $\{s\} \times \Lambda$ has defect d so by Conjecture 2.1,

$$\dim S' - \dim Z \leq d.$$

The projection of S' onto the first factor is a special subvariety of $\mathcal{A}_g$ containing $\{s\}$, so it must contain S. This projection is a Shimura morphism so comes from a product decomposition of the adjoint Shimura datum associated with S'. Hence all fibres of this projection have the same dimension, which must be at least $\dim Z$. So

$$\dim S' - \dim Z \geq \dim S = d.$$

So we must have equality in both inequalities. Equality in the latter implies that $\{s\} \times Z$ is a fibre of the Shimura morphism $S' \to S$ and so is weakly special. This implies that Z is weakly special in $\mathcal{A}_g$. □

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

3. Proof of main theorem

In this section we will deduce our main theorems (Theorems 1.2 and 1.3) from the matrix height bounds of Section 4 and the isogeny bound of Section 5. Accordingly fix a point $s \in \mathcal{A}_g(\mathbb{C})$ and let Λ be its isogeny class. Let $Z \subset \mathcal{A}_g$ be an irreducible closed algebraic subvariety such that $Z \cap \Lambda$ is Zariski dense in Z.

We begin with some definitions and notation. Let $\pi\colon \mathcal{H}_g \to \mathcal{A}_g$ denote the quotient map and $\mathcal{F}_g \subset \mathcal{H}_g$ the Siegel fundamental domain. Let

$$\tilde{Z} = \pi^{-1}(Z) \cap \mathcal{F}_g \quad \text{and} \quad \tilde{\Lambda} = \pi^{-1}(\Lambda) \cap \mathcal{F}_g.$$

Fix a point $\tilde{s} \in \mathcal{H}_g$ such that $\pi(\tilde{s}) = s$.

We define the *complexity* of a point $t \in \Lambda$ to be the minimum degree of an isogeny $A_s \to A_t$ between the abelian varieties corresponding to the points s and t of $\mathcal{A}_g$. We may also talk about the complexity of a point in $\tilde{\Lambda}$, meaning the complexity of its image in Λ.

For a matrix $\gamma \in \mathrm{M}_{n\times n}(\mathbb{Q})$, the *height* $H(\gamma)$ will mean the maximum of the standard multiplicative heights of the entries of γ. A straightforward calculation shows that if $\gamma_1, \gamma_2 \in \mathrm{M}_{n\times n}(\mathbb{Q})$ then

$$H(\gamma_1\gamma_2) \leq nH(\gamma_1)H(\gamma_2).$$

3.1. O-minimality and definability. A key role in the proof is played by the Pila–Wilkie theorem on definable sets in o-minimal structures. For an introduction to o-minimality and the Pila–Wilkie theorem, see [25]. Here we recall some of the definitions and a strengthened version of the Pila–Wilkie theorem, due to Pila.

A *structure $\mathcal{S}$ (over $\mathbb{R}$)* is a sequence $\mathcal{S}_n$ of collections of subsets of $\mathbb{R}^n$ for each natural number n such that

(1) $\mathcal{S}_n$ is closed under finite unions, intersections and complements;

(2) $\mathcal{S}_n$ contains all semialgebraic subsets of $\mathbb{R}^n$ (that is, those sets definable by polynomial inequalities);

(3) if $A \in \mathcal{S}_m$ and $B \in \mathcal{S}_n$ then $A \times B \in \mathcal{S}_{m+n}$;

(4) if $m \geq n$ and $A \in \mathcal{S}_m$ then $\pi(A) \in \mathcal{S}_n$, where $\pi\colon \mathbb{R}^m \to \mathbb{R}^n$ is the projection onto the first n coordinates.

The sets in $\mathcal{S}_n$ are called the *definable sets* of the structure $\mathcal{S}$. A function $f\colon A \to B$ for $A \subset \mathbb{R}^m$, $B \subset \mathbb{R}^n$ is said to be *definable* if its graph is a definable subset of $\mathbb{R}^{m+n}$.

A structure is *o-minimal* if every set in $\mathcal{S}_1$ is a finite union of points and intervals. The basic example of an o-minimal structure is the structure of semialgebraic sets. The o-minimality condition implies that all definable sets in the structure are topologically well-behaved: for example they have finitely many connected components, a finite cell decomposition and can be stratified and triangulated (see [6]).

For the purposes of this article, we will only use the structure $\mathbb{R}_{\mathrm{an,exp}}$ generated by the graphs of restricted analytic functions and the real exponential function. Accordingly, *definable* will henceforth mean definable in $\mathbb{R}_{\mathrm{an,exp}}$. A restricted analytic function is a function $f\colon [0,1]^n \to \mathbb{R}$ which extends to a real analytic function on some open neighbourhood of $[0,1]^n$. This structure was shown to be o-minimal by van den Dries and Miller [7].

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

According to [3], there is a map $\pi\colon \mathcal{H}_g \to \mathbb{P}^N(\mathbb{C})$ which induces an embedding of $\mathcal{A}_g$ as a quasi-projective variety. Siegel [26] used Minkowski's reduction theory to construct a semialgebraic fundamental domain for the action of $\mathrm{Sp}_{2g}(\mathbb{Z})$ on $\mathcal{H}_g$; we will call this domain $\mathcal{F}_g$. Crucially for us, Peterzil and Starchenko [16] have shown that the restriction of π to $\mathcal{F}_g$ is definable in $\mathbb{R}_{\mathrm{an,exp}}$. This implies that, if $Z \subset \mathcal{A}_g$ is an algebraic subvariety, then $\pi^{-1}(Z) \cap \mathcal{F}_g$ is definable.

The principal theorem we shall use about o-minimal structures is a strong version of the Pila–Wilkie theorem, which uses definable blocks. A definable block is a definable set which is connected and almost semialgebraic. More precisely, a *(definable) block* of dimension w in $\mathbb{R}^n$ is a connected definable subset $W \subseteq \mathbb{R}^n$ of dimension w, regular at every point, such that there is a semialgebraic set $A \subseteq \mathbb{R}^n$ of dimension w, regular at every point, with $W \subseteq A$.

A *definable block family* is a definable subset W of $\mathbb{R}^n \times \mathbb{R}^m$ such that for each $\eta \in \mathbb{R}^m$, $W_\eta = \{x \in \mathbb{R}^n \mid (x, \eta) \in W\}$ is a definable block.

The statement of the following theorem has been simplified by referring only to rational points and requiring Z to be a single definable set. Pila's theorem works for points over number fields of some chosen degree and allows Z itself to be a definable family but these are not necessary for our purposes.

Theorem 3.1 ([18, Theorem 3.6]). *Let $Z \subset \mathbb{R}^n$ be a definable set and $\epsilon > 0$. There are a finite number $J = J(Z, \epsilon)$ of definable block families*

$$\mathcal{W}^{(j)} \subset \mathbb{R}^n \times \mathbb{R}^m, \quad j = 1, \dots, J$$

and a constant $c = c(Z, \epsilon)$ such that

(1) *for all $\eta \in \mathbb{R}^m$,*

$$\mathcal{W}_\eta^{(j)} \subset Z;$$

(2) *for all $T \geq 1$, the rational points of Z of height at most T are contained in the union of at most cT^ϵ definable blocks of the form $\mathcal{W}_\eta^{(j)}$ (for some $j \in \{1, \dots, J\}$ and some $\eta \in \mathbb{R}^m$).*

3.2. Outline of proof. The key step in the proof of the main theorems is Proposition 3.2: the points of $\tilde{Z} \cap \tilde{\Lambda}$ of a given complexity are contained in subpolynomially many definable blocks, these blocks themselves contained in $\tilde{Z}$. This is proved using Theorem 3.1 (Pila's theorem) and the matrix height bounds of Section 4.

Proposition 3.2. *Let Z be a subvariety of $\mathcal{A}_g$ and $\tilde{s}$ a point in $\mathcal{H}_g$. Let $\epsilon > 0$. There is a constant $c = c(Z, \tilde{s}, \epsilon)$ such that for every $n \geq 1$, there is a collection of at most cn^ϵ definable blocks $W_i \subset \tilde{Z}$ such that the union $\bigcup W_i$ contains all points of $\tilde{Z} \cap \tilde{\Lambda}$ of complexity n.*

On the other hand, the Masser–Wüstholz isogeny theorem gives a polynomial lower bound for the Galois degree of points in Λ in terms of their complexity. Combining these two bounds, once the complexity gets large enough there are more points in $\tilde{Z} \cap \tilde{\Lambda}$ than there are blocks to contain them. Hence most points of $\tilde{Z} \cap \tilde{\Lambda}$ are contained in blocks of positive dimension. In particular the union of positive-dimensional blocks contained in $\tilde{Z}$ has Zariski dense image in Z.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

In the case $\dim Z = 1$ this implies that $\tilde{Z}$ has an algebraic irreducible component and so we can conclude using the Ullmo–Yafaev characterisation of weakly special subvarieties. When $\dim Z > 1$, we use the Ax–Lindemann theorem for $\mathcal{A}_g$ to deduce that positive-dimensional weakly special subvarieties are dense in Z and then a result of Ullmo to complete the proof of Theorem 1.3.

Let us outline the proof of Proposition 3.2. We cannot apply the counting theorem to $\tilde{\Lambda} \subset \tilde{Z}$ directly, because the points of $\tilde{\Lambda}$ are transcendental. Instead we construct a definable set Y and a semialgebraic map $\sigma: Y \to \tilde{Z}$ such that points of $\tilde{Z} \cap \tilde{\Lambda}$ have rational preimages in Y, with heights polynomially bounded in terms of their complexity. This idea is due to Habegger and Pila [10].

Consider first the case $\mathrm{End}\, A_s = \mathbb{Z}$. This case is easier because all isogenies between A_s and any abelian variety are polarised. In this case we let

$$Y = \{\gamma \in \mathrm{GSp}_{2g}(\mathbb{R})^+ \mid \gamma.\tilde{s} \in \tilde{Z}\},$$

and let $\sigma: Y \to \tilde{Z}$ be the map $\sigma(\gamma) = \gamma.\tilde{s}$.

Let $\tilde{t} \in \tilde{Z} \cap \tilde{\Lambda}$ and $t = \pi(\tilde{t})$. Then there is an isogeny $f: A_t \to A_s$ whose degree is equal to the complexity of t. By the hypothesis $\mathrm{End}\, A_s = \mathbb{Z}$ this isogeny is polarised. Hence the rational representation of f (explained in Section 4) gives a matrix $\gamma \in \mathrm{GSp}_{2g}(\mathbb{Q})^+$ such that $\pi(\gamma.\tilde{s}) = t$ and whose height is polynomially bounded with respect to the complexity. We can also find $\gamma_1 \in \mathrm{Sp}_{2g}(\mathbb{Z})$ of polynomially bounded height such that $\gamma_1\gamma.\tilde{s} = \tilde{t}$. Hence every point in $\tilde{Z} \cap \tilde{\Lambda}$ has a rational preimage in Y of polynomially bounded height. This is precisely what we need to apply Theorem 3.1 to Y.

If we drop the assumption $\mathrm{End}\, A_s = \mathbb{Z}$ then this no longer works, because the rational representation of a non-polarised isogeny is not in $\mathrm{GSp}_{2g}(\mathbb{Q})^+$. Note that even if we assume that t is in the Hecke orbit of s, so that there is some polarised isogeny $A_s \to A_t$, the isogeny of minimum degree need not be polarised. Thus we do not get an element of $\mathrm{GSp}_{2g}(\mathbb{Q})^+$ whose height is polynomially bounded in terms of the complexity.

To avoid this problem we will take Y to be a subset of $\mathrm{GL}_{2g}(\mathbb{R})$ instead of $\mathrm{GSp}_{2g}(\mathbb{R})$. This will allow us to carry out the same proof using the rational representation of a not-necessarily-polarised isogeny. Of course $\mathrm{GL}_{2g}(\mathbb{R})$ does not act on $\mathcal{H}_g$ but this does not matter: the map

$$\sigma\left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right) = (A\tilde{s} + B)(C\tilde{s} + D)^{-1} \quad \text{for } A, B, C, D \in \mathrm{M}_{g\times g}(\mathbb{R})$$

is defined on a Zariski open subset of $\mathrm{GL}_{2g}(\mathbb{R})$, and we will only consider matrices in $\mathrm{GL}_{2g}(\mathbb{R})$ where σ is defined and has image in $\mathcal{H}_g$. In particular let

$$Y = \sigma^{-1}(\tilde{Z}).$$

3.3. Proof of Proposition 3.2. Before proving Proposition 3.2, we need to check that every element of $\tilde{Z} \cap \tilde{\Lambda}$ has a rational preimage in Y whose height is polynomially bounded with respect to the complexity. Proposition 4.1 says that there is some preimage in $\mathrm{GL}_{2g}(\mathbb{R})$ with this property, and all we need to do is move it into preimage of the fundamental domain $\mathcal{F}_g$.

Lemma 3.3. *There exist constants c, k depending only on g and $\tilde{s}$ such that for any $\tilde{t} \in \tilde{Z} \cap \tilde{\Lambda}$ of complexity n, there is a rational matrix $\gamma \in Y$ such that $\sigma(\gamma) = \tilde{t}$ and $H(\gamma) \leq cn^k$.*

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

Proof. Let $t = \pi(\tilde{t})$. Let $\mathcal{B}$ be a symplectic basis for $H_1(A_s, \mathbb{Z})$ with period matrix $\tilde{s}$.

By Proposition 4.1 there is an isogeny $f: A_t \to A_s$ and a symplectic basis $\mathcal{B}'$ for $H_1(A_t, \mathbb{Z})$ such that the rational representation γ_1 of f has polynomially bounded height. As remarked in the introduction to Section 4, $\sigma({}^t\gamma_1)$ is the period matrix for (A_t, λ_t) with respect to the basis $\mathcal{B}'$. In particular,

$$\pi\sigma({}^t\gamma_1) = t.$$

Hence there is $\gamma_2 \in \mathrm{Sp}_{2g}(\mathbb{Z})$ such that $\gamma_2.\sigma({}^t\gamma_1) = \tilde{t}$. By [20, Lemma 3.2], $H(\gamma_2)$ is polynomially bounded. Then $\gamma = \gamma_2{}^t\gamma_1$ has height at most $2gH(\gamma_1)H(\gamma_2)$ and so satisfies the required conditions. □

Now we are ready to prove Proposition 3.2. We simply apply Theorem 3.1 to Y, using Lemma 3.3 to relate heights of rational points in Y to complexities of points in $\tilde{Z} \cap \tilde{\Lambda}$. We then use the fact that σ is semialgebraic, and that the blocks in Y can be chosen uniformly from finitely many definable families, to go from Y to $\tilde{Z}$.

Proof of Proposition 3.2. The set

$$Y = \sigma^{-1}(\pi^{-1}(Z) \cap \mathcal{F}_g)$$

is definable because σ is semialgebraic and $\pi_{|\mathcal{F}_g}$ is definable by a theorem of Peterzil and Starchenko [16].

Hence we can apply Theorem 3.1 to Y: for every $\epsilon > 0$, there are finitely many definable block families $\mathcal{W}^{(j)}(\epsilon) \subset Y \times \mathbb{R}^m$ and a constant $c_1(Y, \epsilon)$ such that for every $T \geq 1$, the rational points of Y of height at most T are contained in the union of at most $c_1 T^\epsilon$ definable blocks $W_i(T, \epsilon)$, taken from the families $\mathcal{W}^{(j)}(\epsilon)$.

Since σ is semialgebraic, the image under σ of a definable block in Y is a finite union of definable blocks in $\tilde{Z}$. Furthermore the number of blocks in the image is uniformly bounded in each definable block family $\mathcal{W}^{(j)}(\epsilon)$. Hence $\sigma(\bigcup W_i(T, \epsilon))$ is the union of at most $c_2 T^\epsilon$ blocks in $\tilde{Z}$, for some new constant $c_2(Z, \tilde{s}, \epsilon)$.

But by Lemma 3.3, for suitable constants c, k, every point of $\tilde{Z} \cap \tilde{\Lambda}$ of complexity n is in $\sigma(\bigcup W_i(cn^k, \epsilon))$. □

3.4. End of proof of Theorems 1.2 and 1.3. Proposition 3.2 tells us that the points of $\tilde{Z} \cap \tilde{\Lambda}$ of complexity n are contained in fewer than $c(\epsilon)n^\epsilon$ blocks for every $\epsilon > 0$. On the other hand, the Masser–Wüstholz isogeny theorem implies that the number of such points grows at least as fast as $n^{1/k}$ for some constant k. Hence most points of $\tilde{Z} \cap \tilde{\Lambda}$ are contained in a block of positive dimension. We check that this is sufficient, with the hypothesis that $Z \cap \Lambda$ is Zariski dense in Z, to deduce that the union of positive-dimensional blocks in $\tilde{Z}$ has Zariski dense image in Z.

Proposition 3.4. *Let Λ_1 be the set of points $t \in Z \cap \Lambda$ for which there is a positive-dimensional block $W \subset \tilde{Z}$ such that $t \in \pi(W)$. If $Z \cap \Lambda$ is Zariski dense in Z, then Λ_1 is Zariski dense in Z.*

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

Proof. Let Z_1 denote the Zariski closure of Λ_1 (we do not yet know that this is non-empty).

Let (A_s, λ_s) be a polarised abelian variety corresponding to the point $s \in \mathcal{A}_g(\mathbb{C})$, defined over a finitely generated field K. We choose K large enough that the varieties Z and Z_1 are also defined over K.

Let t be a point in $Z \cap \Lambda$ of complexity n. The polarised abelian variety corresponding to t might not have a model over the field of moduli $K(t)$, but it has a model (A_t, λ_t) over an extension L of $K(t)$ of uniformly bounded degree. This follows from the fact that a polarised abelian variety with full level-3 structure has no non-trivial automorphisms [13, Proposition 17.5], so is defined over its field of moduli; and the field of moduli of a full level-3 structure on the polarised abelian variety corresponding to t is an extension of $K(t)$ of degree at most $|\mathrm{Sp}_{2g}(\mathbb{Z}/3)|$.

By Theorem 5.1, the complexity n is bounded above by a polynomial $c[L:K]^k$ in $[L:K]$, with c and k depending only on A_s and K. Hence for a different constant c_1, we have

$$[K(t):K] \geq c_1 n^{1/k}.$$

But all $\mathrm{Gal}(\bar{K}/K)$-conjugates of t are contained in $Z \cap \Lambda$ and have complexity n. By Proposition 3.2, the preimages in $\mathcal{F}_g$ of these points are contained in the union of at most $c_2(Z, \tilde{s}, 1/2k)n^{1/2k}$ definable blocks, each of these blocks being contained in $\tilde{Z}$.

For large enough n, we have

$$c_1 n^{1/k} > c_2 n^{1/2k}.$$

For such n, by the pigeonhole principle there is a definable block $W \subset \tilde{Z}$ such that $\pi(W)$ contains at least two Galois conjugates of t. Since blocks are connected by definition, $\dim W > 0$. So those conjugates of t in $\pi(W)$ are in Λ_1. Since Z_1 is defined over K, it follows that t itself is also in Z_1.

In other words all points of $Z \cap \Lambda$ of large enough complexity are in Z_1. But this excludes only finitely many points of $Z \cap \Lambda$. So as $Z \cap \Lambda$ is Zariski dense in Z, we conclude that $Z_1 = Z$. □

Call a subset $W \subset \mathcal{H}_g$ *complex algebraic* if it is a connected component of $W_0 \cap \mathcal{H}_g$ for some irreducible complex algebraic variety $W_0 \subset \mathrm{M}_{2g\times 2g}(\mathbb{C})$. Let $\tilde{Z}^{\mathrm{ca}}$ denote the complex algebraic part of $\tilde{Z}$, that is, the union of positive-dimensional complex algebraic subsets of $\mathcal{H}_g$ contained in $\tilde{Z}$.

By [17, Lemma 2.1], $\tilde{Z}^{\mathrm{ca}}$ is the same as the union of the definable blocks contained in $\tilde{Z}$. So Proposition 3.4 tells us that $\pi(\tilde{Z}^{\mathrm{ca}})$ is Zariski dense in Z.

If $\dim Z = 1$, then the fact that $\tilde{Z}^{\mathrm{ca}}$ is non-empty implies that some irreducible component of $\tilde{Z}$ is complex algebraic. By [28] this implies that Z is weakly special, proving Theorem 1.2.

For $\dim Z > 1$, we use the Ax–Lindemann theorem for $\mathcal{A}_g$ proved by Pila and Tsimerman [19]: if W is a maximal complex algebraic subset of $\tilde{Z}$ then $\pi(W)$ is weakly special. Hence $\pi(\tilde{Z}^{\mathrm{ca}})$ is a union of positive-dimensional weakly special subvarieties, so these are dense in Z. Let S be the smallest special subvariety of $\mathcal{A}_g$ containing Z. By [27, Théorème 1.3], we deduce that $S = S_1 \times S_2$ for some Shimura varieties S_1 and S_2, and $Z = S_1 \times Z'$ for some subvariety $Z \subset S_2$, proving Theorem 1.3.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

4. Heights of rational representations of isogenies

Let (A, λ) and (A', λ') be principally polarised abelian varieties over $\mathbb{C}$ related by an isogeny of degree n (not necessarily compatible with the polarisations). In this section we show that, for suitable choices of bases for $H_1(A, \mathbb{Z})$ and $H_1(A', \mathbb{Z})$ and of isogeny $f: A' \to A$, the height of the rational representation of f is polynomially bounded in n. In particular we prove the following proposition, used to prove Lemma 3.3.

The notation $H(f, \mathcal{B}', \mathcal{B})$ in the proposition refers to the height of the rational representation of the isogeny f with respect to bases $\mathcal{B}'$, $\mathcal{B}$ of the period lattices. This is defined below in Section 4.2.

Proposition 4.1. *Let (A, λ) be a principally polarised abelian variety over $\mathbb{C}$ and fix a symplectic basis $\mathcal{B}$ for $H_1(A, \mathbb{Z})$. There exist constants c, k depending only on (A, λ) such that: If (A', λ') is any principally polarised abelian variety for which there exists an isogeny $A \to A'$ of degree n, then there are an isogeny $f: A' \to A$ and a symplectic basis $\mathcal{B}'$ for $H_1(A', \mathbb{Z})$ such that*

$$H(f, \mathcal{B}', \mathcal{B}) \leq cn^k.$$

In this proposition, the isogeny whose existence is assumed and the isogeny whose existence is asserted in the conclusion go in opposite directions. This is the most convenient formulation for our application, but it is not important since any isogeny $A \to A'$ of degree n gives rise to an isogeny in the opposite direction of degree n^{2g-1}.

4.1. Polarisations. Let A be an abelian variety and $A^\vee$ its dual variety. A *polarisation* of A is an isogeny $A \to A^\vee$ satisfying a certain positivity condition (being associated with an ample sheaf). Any polarisation induces a symplectic form $H_1(A, \mathbb{Z}) \times H_1(A, \mathbb{Z}) \to \mathbb{Z}$. A polarisation is *principal* if it is an isomorphism $A \to A^\vee$.

A polarisation $\lambda: A \to A^\vee$ induces an involution, called the *Rosati involution*, of $\operatorname{End} A \otimes_{\mathbb{Z}} \mathbb{Q}$ defined by

$$a^\dagger = \lambda^{-1} \circ a^\vee \circ \lambda$$

where $a^\vee$ means the morphism dual to a. This involution reverses the order of multiplication in $\operatorname{End} A \otimes_{\mathbb{Z}} \mathbb{Q}$. It gives an involution of $\operatorname{End} A$ itself if λ is principal.

Having fixed a principal polarisation λ of A, every other polarisation has the form $\lambda \circ q$ for some $q \in \operatorname{End} A$ which is *symmetric*, i.e. $q^\dagger = q$, and *positive definite*, i.e. each component of q in

$$\operatorname{End} A \otimes_{\mathbb{Z}} \mathbb{R} \cong \prod \mathrm{M}_{l_i}(\mathbb{R}) \times \prod \mathrm{M}_{m_i}(\mathbb{C}) \times \prod \mathrm{M}_{n_i}(\mathbb{H})$$

has eigenvalues which are positive real numbers.

4.2. Rational representations. We define the *rational representation* of an isogeny $f: A' \to A$ (with respect to bases $\mathcal{B}, \mathcal{B}'$ for $H_1(A, \mathbb{Z})$ and $H_1(A', \mathbb{Z})$) to be the matrix of the induced morphism

$$f_*: H_1(A', \mathbb{Z}) \to H_1(A, \mathbb{Z})$$

in terms of the chosen bases. This gives a $2g \times 2g$ integer matrix. We write

$$H(f, \mathcal{B}', \mathcal{B})$$

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

for the *height* of the rational representation of f, meaning simply the maximum of the absolute values of the entries of the matrix.

Rational representations of isogenies are particularly interesting in the case that the bases $\mathcal{B}, \mathcal{B}'$ are symplectic with respect to the polarisations λ, λ'. In this case, if $\tilde{s}, \tilde{t} \in \mathcal{H}_g$ are the period matrices of (A, λ) and (A', λ') with respect to the chosen bases and γ is the rational representation of an isogeny $A' \to A$, then

$$\tilde{t} = (A\tilde{s} + B)(C\tilde{s} + D)^{-1} \text{ where } {}^t\gamma = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right).$$

We remark also that for symplectic bases, an isogeny is polarised if and only if its rational representation is in $\mathrm{GSp}_{2g}(\mathbb{Q})$.

4.3. Outline of proof. In the situation of Proposition 4.1, let $h\colon A \to A'$ be an isogeny of degree n. Then $h^*\lambda'$ is a polarisation of A, so there is a symmetric positive definite endomorphism $q \in \mathrm{End}\, A$ such that

$$h^*\lambda' = \lambda \circ q.$$

We can identify $H_1(A', \mathbb{Z})$ with a submodule of $H_1(A, \mathbb{Z})$ of index n^{2g-1}, and so find a basis for $H_1(A', \mathbb{Z})$ whose height is at most n^{2g-1}. However this need not be a symplectic basis. We apply the standard algorithm for finding a symplectic basis: the height of this new basis is controlled by $h^*\lambda'$, in other words by q.

So we would like to bound the height of the rational representation of q in terms of $\deg h$. However this is not possible: let A be an abelian variety whose endomorphism ring is the ring of integers $\mathfrak{o}$ of a real quadratic field. In particular $\mathfrak{o}$ has infinitely many units. Let h be a unit in $\mathfrak{o}$, in other words, an isomorphism $A \to A$. If we take the same polarisation on each copy of A, then $q = h^2$ and the rational representation of this can have arbitrarily large height.

We can avoid this by replacing h by $h \circ u$ for some automorphism u of A; recall that all we have supposed about h is that it is an isogeny $A \to A'$ of degree n. This replaces q by $u^\dagger q u$. We will show that we can choose u so that the height of the rational representation of $u^\dagger q u$ is bounded by a multiple of $\deg q = n^2$.

4.4. Heights in the endomorphism ring. The following proposition is motivated by the fact that the symmetric elements of $\mathrm{End}\, A$ of a given norm fall into finitely many orbits under the action of $(\mathrm{End}\, A)^\times$ given by $(u, q) \mapsto uqu^\dagger$ (see [13, Proposition 18.2]). In geometric terms, this says that if we fix A and $\deg \mu$ then there are finitely many isomorphism classes of polarised abelian varieties (A, μ). Proposition 4.2 strengthens this by saying that each orbit contains an element whose height is bounded by a multiple of the norm. Milne's theorem is proved using the reduction theory of arithmetic groups. We also use reduction theory, but in order to get height bounds we have to go deeper into the structure of $\mathrm{End}\, A \otimes_{\mathbb{Z}} \mathbb{R}$.

The representation ρ appears in Proposition 4.2 solely to give us a convenient definition of heights and norms of elements of R. Specifically, $H(x)$ means the height of $\rho(x)$ and $\mathrm{N}(x) = \det \rho(x)$ for $x \in R$.

Proposition 4.2. *Let $(E, \dagger)$ be a semisimple $\mathbb{Q}$-algebra with a positive involution, let R be a $\dagger$-stable order in E and let $\rho : R \to \mathrm{M}_N(\mathbb{Z})$ be a faithful representation of R. There is a constant c depending only on $(R, \dagger, \rho)$ such that for any symmetric positive definite $q \in R$, there is some $u \in R^\times$ such that*

$$H(u^\dagger q u) \leq c\, \mathrm{N}(q).$$

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

Proof. We begin by checking that it suffices to prove the proposition for simple algebras E. In general, $E = \prod E_i$ for some simple $\mathbb{Q}$-algebras E_i. Let $R_i = R \cap E_i$. Then $R' = \prod R_i$ is an order of E contained in R. Let $m = [R : R']$. Given $q \in R$, we look at $mq \in R'$. Suppose that the proposition holds for each R_i; then clearly it holds for R', so there is $u \in R'^{\times}$ (a fortiori $u \in R^{\times}$) such that

$$H(umqu^{\dagger}) \leq c\,\mathrm{N}(mu).$$

Hence the proposition holds for R with constant $c\,\mathrm{N}(m)/m$.

So we suppose that E is simple. Then $E = \mathrm{M}_n(D)$ for some division algebra D, and the involution $\dagger$ is matrix transposition composed with some involution of D. We may also suppose that R is contained in the maximal order $\mathrm{M}_n(\mathfrak{o})$, where $\mathfrak{o}$ is a maximal order in D.

By the Albert classification of division algebras with positive involution, $E \otimes_{\mathbb{Q}} \mathbb{R}$ is isomorphic to one of $\mathrm{M}_{nd}(\mathbb{R})^r$, $\mathrm{M}_{nd}(\mathbb{C})^r$ or $\mathrm{M}_{nd}(\mathbb{H})^r$. Because q is symmetric, its projection onto each simple factor of $E \otimes_{\mathbb{Q}} \mathbb{R}$ is a Hermitian matrix. By the theory of Hermitian forms over $\mathbb{R}$, $\mathbb{C}$ and $\mathbb{H}$, there exist $x, d \in E \otimes_{\mathbb{Q}} \mathbb{R}$ such that d is diagonal with real entries in each factor and

$$q = x^{\dagger} d x.$$

Since q is positive definite, all the diagonal entries of d are positive so we can multiply each row of x by the square root of the corresponding entry of d to suppose that $d = 1$. We then have $q = x^{\dagger} x$.

Let G be the $\mathbb{Z}$-group scheme representing the functor $\mathbb{Z}$-**Alg** $\to$ **Grp** given by

$$G(A) = (R \otimes_{\mathbb{Z}} A)^{\times}.$$

Over $\mathbb{Q}$ this is the reductive group $\mathrm{Res}_{D/\mathbb{Q}} \mathrm{GL}_n$. We will use the following notations for subgroups of G:

(i) S is the maximal $\mathbb{Q}$-split torus of G whose $\mathbb{Q}$-points are the diagonal matrices of $\mathrm{GL}_n(D)$ with entries in $\mathbb{Q}$;

(ii) P is the minimal parabolic $\mathbb{Q}$-subgroup of G consisting of upper triangular matrices;

(iii) $U = R_u(P)$ is the group of upper triangular matrices with ones on the diagonal;

(iv) Z is the centraliser of S in G; that is, $Z(\mathbb{Q})$ consists of the diagonal matrices in $\mathrm{GL}_n(D)$;

(v) M is the maximal $\mathbb{Q}$-anisotropic subgroup of Z; that is, $M(\mathbb{Q})$ consists of the diagonal matrices in $\mathrm{GL}_n(D)$ whose diagonal entries have reduced norm ± 1;

(vi) $K = \{g \in G(\mathbb{R}) \mid g^{\dagger} g = 1\}$ is a maximal compact subgroup of $G(\mathbb{R})$.

By [4, Proposition 13.1], there exist a positive real number t, a finite set $C \subset G(\mathbb{Q})$ and a compact neighbourhood ω of 1 in $M^0(\mathbb{R})U(\mathbb{R})$ such that

$$G(\mathbb{R}) = K A_t \omega C G(\mathbb{Z})$$

where

$$A_t = \{a \in S(\mathbb{R}) \mid a_i > 0,\ a_i/a_{i+1} \leq t \text{ for all } i\}.$$

We note that $M^0(\mathbb{R})U(\mathbb{R})$ is the group of upper triangular matrices in $\mathrm{M}_n(D \otimes_{\mathbb{Q}} \mathbb{R})$ whose diagonal entries have reduced norm 1.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

Hence we can write

$$x = kazv\gamma$$

where $k \in K$, $a \in A_t$, $z \in \omega$, $v \in C$ and $\gamma \in G(\mathbb{Z}) = R^\times$.

Let $u = \gamma^{-1}$ and

$$q' = u^\dagger q u.$$

In order to prove the proposition, it will suffice to show that $H(q') \leq c\,\mathrm{N}(q)$.

Since $k^\dagger k = 1$, and using the decomposition of x, we get that

$$q' = v^\dagger z^\dagger a^\dagger a z v.$$

Fix some $\mathbb{Z}$-basis of R. We will show below that the (real) coordinates of $a^\dagger a$ are bounded above by a constant multiple of $\mathrm{N}(q)$. The coordinates of z and v are uniformly bounded because z is in the compact set ω and v is in the finite set C. Hence the coordinates of q' in this basis are bounded by a multiple of $\mathrm{N}(q)$, so $H(q')$ is likewise linearly bounded.

Let $a^\dagger a = \mathrm{diag}(a_1, \ldots, a_n)$ with $a_i \in \mathbb{R}$. In order to show that the coordinates of $a^\dagger a$ in the chosen basis are bounded, it will suffice to show that the a_i are bounded by a multiple of $\mathrm{N}(q)$. We shall show that the a_i are bounded *below* by a constant, and that their product $\prod a_i$ is bounded above by a multiple of $\mathrm{N}(q)$. These two facts together imply that the a_i are bounded above by a multiple of $\mathrm{N}(q)$.

Choose an integer m such that $mv^{-1} \in R$ for all $v \in C$. Then

$$m^2 z^\dagger a^\dagger a z = (mv^{\dagger -1}) q' (mv^{-1}) \in R$$

so every entry of $m^2 z^\dagger a^\dagger a z$, viewed as a matrix in $\mathrm{M}_n(D)$, is in $\mathfrak{o}$.

Let z_{11} denote the upper left entry of $z \in \mathrm{M}_n(D \otimes_{\mathbb{Q}} \mathbb{R})$. Because z is upper triangular, the upper left entry of $m^2 z^\dagger a^\dagger a z$ is $m^2 z_{11}^\dagger a_1 z_{11}$. So $m^2 z_{11}^\dagger a_1 z_{11} \in \mathfrak{o}$ and

$$\left|\mathrm{Nrd}_{D/\mathbb{Q}}(m^2 z_{11}^\dagger a_1 z_{11})\right| \geq 1.$$

But $\mathrm{Nrd}(z_{11}) = 1$ because $z \in \omega$, so

$$\left|\mathrm{Nrd}_{D\otimes_{\mathbb{Q}}\mathbb{R}/\mathbb{R}}(m^2 a_1)\right| \geq 1.$$

Since $m^2 a_1$ is a positive real number, $\mathrm{Nrd}_{D\otimes_{\mathbb{Q}}\mathbb{R}/\mathbb{R}}(m^2 a_1)$ is just some fixed positive power of $m^2 a_1$ so we conclude that

$$m^2 a_1 \geq 1.$$

From the definition of A_t, it follows that $a_i \geq m^{-2} t^{2-2i}$ for all i and we have established that the a_i are uniformly bounded below.

Hence there is a constant c_1 such that for every j,

$$a_j \leq c_1 \prod a_i.$$

Since ρ is faithful, we have $\dim \rho \geq n$. Together with the fact that $\prod a_i$ is bounded below this implies that

$$\prod a_i \leq c_2 \Big(\prod a_i\Big)^{\dim \rho / n} = c_2\,\mathrm{N}(a^\dagger a).$$

Now $\mathrm{N}(z) = \mathrm{N}(u) = 1$ and $\mathrm{N}(v)$ is bounded because v comes from a finite set, so $\mathrm{N}(a^\dagger a)$ is bounded above by a constant multiple of $\mathrm{N}(q)$. Combining all this we have proved that each a_i is bounded above by a constant multiple of $\mathrm{N}(q)$, and as remarked above this suffices to establish the proposition. □

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

4.5. Height of a symplectic basis. We will need the following bound for the height of a symplectic basis for a symplectic free $\mathbb{Z}$-module in terms of the values of the symplectic pairing on the standard basis. The proof is simply to apply the standard recursive algorithm for finding a symplectic basis, verifying that the new vectors introduced always have polynomially bounded heights.

Lemma 4.3. *Let $L = \mathbb{Z}^{2g}$ and let $\{e_1, \dots, e_{2g}\}$ be a basis for L. There exist constants c, k depending only on g such that for any perfect symplectic pairing $\psi : L \times L \to \mathbb{Z}$ with*

$$N = \max_{i,j} |\psi(e_i, e_j)|,$$

there exists a symplectic basis for (L, ψ) whose coordinates with respect to the original basis $\{e_1, \dots, e_{2g}\}$ are at most cN^k.

Proof. For any $x \in L$, we write $H(x)$ for the maximum of the absolute values of the coordinates of x with respect to the basis $\{e_1, \dots, e_{2g}\}$.

First let $e'_1 = e_1$ and choose e'_2 such that $\psi(e'_1, e'_2) = 1$ and $H(e'_2) \leq N$. We can do this because ψ is perfect, so that $\gcd_{i=2}^{n}(\psi(e_1, e_i)) = 1$. Hence there are integers a_i such that $|a_i| \leq N$ and

$$\sum a_i \psi(e_1, e_i) = 1.$$

We let $e'_2 = \sum a_i e_i$.

Then find $e'_3, \dots, e'_{2g}$ orthogonal to e'_1 and to e'_2 such that $\{e'_1, \dots, e'_{2g}\}$ is a basis for L and $H(e'_i) \leq 2gN^2$. We can do this by setting

$$e'_i = e_i + \psi(e'_2, e_i)e'_1 + \psi(e'_1, e_i)e'_2.$$

Here we have

$$|\psi(e'_2, e_i)| \leq \sum_{j=2}^{n} |a_j \psi(e_j, e_i)| \leq (2g-1)N^2$$

and $\psi(e'_1, e_i)e'_2$ has height at most N^2 so $H(e'_i) \leq 2gN^2$.

Finally apply the algorithm recursively to $L' = \mathbb{Z}\langle e'_3, \dots, e'_{2g}\rangle$. We have

$$|\psi(e'_i, e'_j)| \leq gNH(e'_i)H(e'_j) \leq 4g^3N^5.$$

Hence by induction L' has a symplectic basis whose coordinates with respect to $\{e'_3, \dots, e'_{2g}\}$ are bounded by a constant multiple of $N^{5k'}$, where k' is the exponent in the lemma for $\mathbb{Z}^{2(g-1)}$. Converting these into coordinates with respect to $\{e_1, \dots, e_{2g}\}$, we get that the elements of this symplectic basis for L' have height bounded by a constant multiple of $N^{2+5k'}$. This proves the lemma.

We remark that the recurrence $k(g) = 2 + 5k(g-1)$, $k(0) = 0$ is satisfied by $k(g) = (5^g - 1)/2$, so this provides a suitable choice of exponent for the lemma. □

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

4.6. Proof of Proposition 4.1. Let $h\colon A \to A'$ be an isogeny of degree n. There is $q \in \mathrm{End}\, A$ such that

$$h^*\lambda' = \lambda \circ q.$$

Apply Proposition 4.2 to get $u \in (\mathrm{End}\, A)^\times$ such that

$$H(u^\dagger q u) \leq c\, \mathrm{N}(q).$$

Then hu is an isogeny $A \to A'$ of degree n, so there is also an isogeny $f\colon A' \to A$ of degree n^{2g-1} such that

$$hu \circ f = [n]_A.$$

The image of $f_*\colon H_1(A', \mathbb{Z}) \to H_1(A, \mathbb{Z})$ is a submodule of index n^{2g-1}. By the structure theory of finitely generated $\mathbb{Z}$-modules there is a basis $\{e'_1, \ldots, e'_{2g}\}$ for $H_1(A', \mathbb{Z})$ with respect to which the rational representation of f is upper triangular and has height at most n^{2g-1}. But this need not be a symplectic basis.

Let ψ, ψ' be the symplectic forms on $H_1(A, \mathbb{Z})$ and $H_1(A', \mathbb{Z})$ induced by λ, λ' respectively. Let $q' = u^\dagger q u$. Then

$$n^2\lambda' = [n]^*_{A'}\lambda' = f^*u^*h^*\lambda' = f^*(\lambda \circ q').$$

In terms of symplectic forms this says that

$$n^2\psi'(x, y) = \psi(f_*x, q'_*f_*y).$$

In particular, since the coordinates (with respect to $\mathcal{B}$, a symplectic basis for ψ) of $\{f_*e'_1, \ldots, f_*e'_{2g}\}$ and the entries of the matrix q'_* are bounded by a polynomial in n, the same is true for

$$|\psi'(e'_i, e'_j)|.$$

Hence by Lemma 4.3 there is a symplectic basis $\mathcal{B}'$ for $H_1(A', \mathbb{Z})$ whose coordinates with respect to $\{e'_1, \ldots, e'_{2g}\}$ are polynomially bounded. Using again that the coordinates with respect to $\mathcal{B}$ of $\{f_*e'_1, \ldots, f_*e'_{2g}\}$ are polynomially bounded, we deduce that $H(f, \mathcal{B}', \mathcal{B})$ is also polynomially bounded. □

5. Isogeny theorem over finitely generated fields

The Masser–Wüstholz isogeny theorem [11] gives a bound for the minimum degree of an isogeny between two abelian varieties over number fields, as a function of one of the varieties and the degree of their joint field of definition. In order to prove Theorems 1.2 and 1.3 for points $s \in \mathcal{A}_g$ defined over $\mathbb{C}$ and not merely over $\bar{\mathbb{Q}}$, we need to extend the isogeny theorem to abelian varieties defined over finitely generated fields of characteristic 0. We will do this by a specialisation argument, using the fact that any abelian scheme has a closed fibre in which the specialisation map of endomorphism rings is surjective. The proof is based on Raynaud's proof [24] that the Manin–Mumford conjecture over $\bar{\mathbb{Q}}$ implies the conjecture over $\mathbb{C}$.

A key feature of the theorem of Masser and Wüstholz is the explicit dependence of the bound on the abelian variety A, via the Faltings height. Our theorem does not make this explicit, and it is not apparent that there is any analogy of the Faltings height over a finitely generated field which would enable it to be made explicit. Instead what matters to us is the dependence on the field of definition of B.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

Theorem 5.1. *Let K be a finitely generated field of characteristic 0 and A an abelian variety defined over K. There exist constants $c(A, K)$ and κ (κ depending only on $\dim A$) such that: If B is any abelian variety defined over a finite extension L of K and isogenous over $\bar{K}$ to A, then there exists an isogeny $A \to B$ defined over $\bar{K}$ of degree at most*

$$c(A, K)[L : K]^{\kappa}.$$

In Masser and Wüstholz's theorem, the constant c depended also on the degrees of polarisations of A and B. This dependence has been eliminated by Gaudron and Rémond [9] who also showed that we can take $\kappa = 2^{10}(\dim A)^3 + \epsilon$ for the exponent.

Proof. Let R be a finitely generated normal $\mathbb{Q}$-algebra whose field of fractions is K, and let $S = \operatorname{Spec} R$. There is an abelian scheme $\mathcal{A}$ over some open subset $U \subset S$ whose generic fibre is isomorphic to A. (Note that R is a finitely generated $\mathbb{Q}$-algebra, not a finitely generated $\mathbb{Z}$-algebra, because unlike in [24] we do not need to reduce modulo p, while Noot's specialisation result requires the base to be a variety over $\mathbb{Q}$.)

By replacing L by a larger extension of bounded degree (the bound depending only on $\dim A$), we may assume that all homomorphisms $A \to B$ are defined over L (see [12, Lemma 2.1]). Let R' be the integral closure of R in L and $S' = \operatorname{Spec} R'$. Let $\pi : S' \to S$ be the obvious finite morphism and let $U' = \pi^{-1}(U)$.

Because A and B are isogenous, there is an abelian scheme $\mathcal{B}$ over U' with generic fibre isomorphic to B, and such that $\mathcal{B}$ is isogenous to $\mathcal{A}$. We can construct this as follows: let N be the kernel of an isogeny $A \to B$. We can extend N to a finite flat subgroup scheme $\mathcal{N} \subset \mathcal{A}$. Then let $\mathcal{B}$ be the quotient $\mathcal{A}/\mathcal{N}$.

For any closed points $s' \in U'$ and $s = \pi(s') \in U$, the fibres $\mathcal{A}_s$ and $\mathcal{B}_{s'}$ are abelian varieties over the number fields k_s and $k_{s'}$, isogenous over $k_{s'}$. We can apply the Masser–Wüstholz theorem to deduce that there are constants $c(\mathcal{A}_s, k_s)$ and $\kappa(\dim A)$ and an isogeny $\mathcal{A}_s \to \mathcal{B}_{s'}$ of degree at most

$$c(\mathcal{A}_s, k_s)[k_{s'} : k_s]^{\kappa}.$$

Observe that $[k_{s'} : k_s] \leq [L : K]$.

In order to prove the theorem, all we have to do is show that this isogeny $\mathcal{A}_s \to \mathcal{B}_{s'}$ lifts to an isogeny $A \to B$ (which will have the same degree). Hence it will suffice to show that there is some closed point s such that the specialisation map

$$(*) \qquad \operatorname{Hom}_{\bar{K}}(A, B) \to \operatorname{Hom}_{\bar{k_s}}(\mathcal{A}_s, \mathcal{B}_{s'})$$

is surjective. Because we want a bound which depends only on A and not on B, we have to show that there is a single point $s \in U$ which will work for all B.

We choose a closed point $s \in U$ such that $\operatorname{End}_{\bar{K}} A \to \operatorname{End}_{\bar{k_s}} \mathcal{A}_s$ is surjective. Such an s exists by [15, Corollary 1.5] (this is proved using the Hilbert irreducibility theorem).

Let f_s be a $\bar{k_s}$-homomorphism $\mathcal{A}_s \to \mathcal{B}_{s'}$. To prove that $(*)$ is surjective, we have to show that f_s lifts to $\bar{K}$-homomorphism $A \to B$.

We are assuming that A and B are isogenous. Choose any isogeny $g_\eta : A \to B$ and let g_s be its specialisation at s. Let

$$\alpha_s = g_s^{-1} \circ f_s \in \operatorname{End}_{\bar{k_s}} \mathcal{A}_s \otimes_{\mathbb{Z}} \mathbb{Q}.$$

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

By our choice of s, this lifts to some $\alpha_\eta \in \mathrm{End}_{\bar{K}} A \otimes_{\mathbb{Z}} \mathbb{Q}$. Then $f_\eta = g_\eta \circ \alpha_\eta$ is a quasi-isogeny $A \to B$ specialising to f_s.

All we have to do is check that f_η is an isogeny and not just a quasi-isogeny. Choose an integer m such that mf_η is an isogeny. The kernel of mf_s contains $\mathcal{A}_s[m]$ so lifting to the generic fibre, the kernel of mf_η contains $A[m]$. Hence mf_η factorises as $f'_\eta \circ [m]$ for an isogeny $f'_\eta : A \to B$, and we must have $f'_\eta = f_\eta$. □

References

[1] *Y. André*, G-functions and geometry, Aspects Math. **E13**, Friedr. Vieweg & Sohn, Braunschweig 1989.

[2] *Y. André*, Mumford-Tate groups of mixed Hodge structures and the theorem of the fixed part, Compositio Math. **82** (1992), no. 1, 1–24.

[3] *W. L. Baily, Jr.*, Satake's compactification of V_n, Amer. J. Math. **80** (1958), 348–364.

[4] *A. Borel*, Introduction aux groupes arithmétiques, Publications de l'Institut de Mathématique de l'Université de Strasbourg, XV, Actualités Scientifiques et Industrielles, No. 1341, Hermann, Paris 1969.

[5] *P. Deligne*, Variétés de Shimura: interprétation modulaire, et techniques de construction de modèles canoniques, in: Automorphic forms, representations and L-functions (Part 2), Proc. Sympos. Pure Math. **33**, American Mathematical Society, Providence (1979), 247–289.

[6] *L. van den Dries*, Tame topology and o-minimal structures, London Math. Soc. Lecture Note Ser. **248**, Cambridge University Press, Cambridge 1998.

[7] *L. van den Dries* and *C. Miller*, On the real exponential field with restricted analytic functions, Israel J. Math. **85** (1994), no. 1–3, 19–56.

[8] *B. Edixhoven* and *A. Yafaev*, Subvarieties of Shimura varieties, Ann. of Math. (2) **157** (2003), no. 2, 621–645.

[9] *É. Gaudron* and *G. Rémond*, Polarisations et isogénies, preprint 2012, `http://math.univ-bpclermont.fr/~gaudron/`.

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

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

[12] *D. Masser* and *G. Wüstholz*, Periods and minimal abelian subvarieties, Ann. of Math. (2) **137** (1993), no. 2, 407–458.

[13] *J. S. Milne*, Abelian varieties, in: Arithmetic geometry (Storrs 1984), Springer, New York (1986), 103–150.

[14] *B. Moonen*, Linearity properties of Shimura varieties, I, J. Algebraic Geom. **7** (1998), no. 3, 539–567.

[15] *R. Noot*, Abelian varieties – Galois representation and properties of ordinary reduction, Compositio Math. **97** (1995), no. 1-2, 161–171.

[16] *Y. Peterzil* and *S. Starchenko*, Definability of restricted theta functions and families of abelian varieties, Duke Math. J. **162** (2013), no. 4, 731–765.

[17] *J. Pila*, Rational points of definable sets and results of André–Oort–Manin–Mumford type, Int. Math. Res. Not. IMRN (2009), no. 13, 2476–2507.

[18] *J. Pila*, O-minimality and the André–Oort conjecture for $\mathbb{C}^n$, Ann. of Math. (2) **173** (2011), no. 3, 1779–1840.

[19] *J. Pila* and *J. Tsimerman*, Ax–Lindemann for $\mathcal{A}_g$, preprint 2012, `http://arxiv.org/abs/1206.2663`.

[20] *J. Pila* and *J. Tsimerman*, The André–Oort conjecture for the moduli space of abelian surfaces, Compos. Math. **149** (2013), no. 2, 204–216.

[21] *J. Pila* and *U. Zannier*, Rational points in periodic analytic sets and the Manin–Mumford conjecture, Rend. Lincei (9) Mat. Appl. **19** (2008), no. 2, 149–162.

[22] *R. Pink*, A combination of the conjectures of Mordell-Lang and André-Oort, in: Geometric methods in algebra and number theory, Progr. Math. **235**, Birkhäuser, Boston (2005), 251–282.

[23] *R. Pink*, A common generalization of the conjectures of André–Oort, Manin-Mumford, and Mordell–Lang, preprint 2005, `www.math.ethz.ch/~pink/ftp/AOMMML.pdf`.

[24] *M. Raynaud*, Courbes sur une variété abélienne et points de torsion, Invent. Math. **71** (1983), no. 1, 207–233.

[25] *T. Scanlon*, Counting special points: Logic, diophantine geometry, and transcendence theory, Bull. Amer. Math. Soc. (N.S.) **49** (2012), no. 1, 51–71.

[26] *C. L. Siegel*, Symplectic geometry, Amer. J. Math. **65** (1943), 1–86.

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM

[27] *E. Ullmo*, Quelques applications du théorème d'Ax–Lindemann hyperbolique, preprint 2012, `www.math.u-psud.fr/~ullmo/Prebublications/cordAxLind.pdf`.
[28] *E. Ullmo* and *A. Yafaev*, A characterization of special subvarieties, Mathematika **57** (2011), no. 2, 263–273.

Martin Orr, Université Paris-Sud 11, Bat. 425, 91400 Orsay, France
e-mail: martin.orr@math.u-psud.fr

Eingegangen 28. Januar 2013, in revidierter Fassung 29. Mai 2013

Brought to you by | Imperial College London
Authenticated
Download Date | 5/13/16 4:14 PM