

Heights of pre-special points of Shimura varieties

Article

Accepted Version

Daw, C. and Orr, M. (2016) Heights of pre-special points of Shimura varieties. *Mathematische Annalen*, 365 (3-4). pp. 1305-1357. ISSN 0025-5831 doi:
<https://doi.org/10.1007/s00208-015-1328-3> Available at
<https://centaur.reading.ac.uk/70364/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

Published version at: <http://dx.doi.org/10.1007/s00208-015-1328-3>

To link to this article DOI: <http://dx.doi.org/10.1007/s00208-015-1328-3>

Publisher: Springer

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Heights of pre-special points of Shimura varieties

Christopher Daw Martin Orr

May 15, 2017

Abstract

Let s be a special point on a Shimura variety, and x a pre-image of s in a fixed fundamental set of the associated Hermitian symmetric domain. We prove that the height of x is polynomially bounded with respect to the discriminant of the centre of the endomorphism ring of the corresponding \mathbb{Z} -Hodge structure. Our bound is the final step needed to complete a proof of the André–Oort conjecture under the conjectural lower bounds for the sizes of Galois orbits of special points, using a strategy of Pila and Zannier.

1 Introduction

Our aim in this paper is to prove a bound for the height of a pre-special point in a fundamental set of a Hermitian symmetric domain covering a Shimura variety. This generalises a theorem of Pila and Tsimerman ([PT13] Theorem 3.1) concerning the heights of pre-images of CM points in a fundamental set of the Siegel upper half-space. Our motivation for considering this bound is that it completes a strategy, originating in the work [PZ08] of Pila and Zannier, for a new proof of the André–Oort conjecture under the conjectural lower bounds for the sizes of Galois orbits of special points, known to hold under the Generalised Riemann Hypothesis (GRH).

The following is the precise statement of our primary bound, comparing the height of a pre-special point x with the discriminant of the centre of the endomorphism ring of the \mathbb{Z} -Hodge structure associated with x . Note that the association of \mathbb{Z} -Hodge structures with points of X depends on the choice of a representation

Christopher Daw, Institut des Hautes Études Scientifiques, Bures-sur-Yvette, France
Martin Orr, Dept. of Mathematics, Imperial College, South Kensington, London, UK
2010 Mathematics Subject Classification: 11G18

of the group \mathbf{G} and of a lattice in this representation – this is the only purpose of ρ and $E_{\mathbb{Z}}$ in the theorem.

Theorem 1.1. *Let (\mathbf{G}, X) be a Shimura datum with \mathbf{G} being an adjoint group. Let $\rho: \mathbf{G} \rightarrow \mathbf{GL}(E)$ be a faithful self-dual \mathbb{Q} -representation, and fix a lattice $E_{\mathbb{Z}} \subset E$.*

Let $\Gamma \subset \mathbf{G}(\mathbb{Q})$ be a congruence subgroup and let $\mathcal{F} \subset X$ be a fundamental set for Γ as in Théorème 13.1 of [Bor69], with respect to a pre-special base point x_0 .

Choose a realisation of X such that the action of $\mathbf{G}(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ on X is semialgebraic and defined over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$. Let $H(x)$ denote the multiplicative Weil height of a point in $X(\mathbb{Q}^{\text{alg}})$ with respect to this realisation.

There are constants C_1, C_2 (depending on $\mathbf{G}, X, \mathcal{F}, \rho$ and the choice of a realisation for X) such that for all pre-special points $x \in \mathcal{F}$,

$$H(x) \leq C_1 |\text{disc } R_x|^{C_2},$$

where R_x is the centre of the endomorphism ring of the \mathbb{Z} -Hodge structure $\rho \circ x$.

Our motivation for proving this theorem was to apply it to the André–Oort conjecture. In order to do this, we prove an additional bound (Theorem 4.1) comparing the discriminant of R_x with certain invariants associated with the Mumford–Tate group of the pre-special point. These invariants are the same invariants which appear in lower bounds for the sizes of Galois orbits of special points.

The upshot of Theorems 1.1 and 4.1 is that the absence of lower bounds for Galois orbits of special points becomes the only remaining obstacle to an unconditional proof of the André–Oort conjecture. Specifically, the combination of Theorems 1.1 and 4.1 together with [PW06], [PT14], [Ull14], [KUY] and the strategy of [PZ08] implies the following.

Theorem 1.2. *Let (\mathbf{G}, X) be a Shimura datum and K a compact open subgroup of $\mathbf{G}(\mathbb{A}_f)$.*

Assume that there exist positive constants C_3, C_4, C_5, C_6 (depending only on \mathbf{G}, X and K) such that, for each pre-special point $x \in X$, its image $[x, 1]$ in the Shimura variety $\text{Sh}_K(\mathbf{G}, X)$ satisfies

$$\left| \text{Gal}(\mathbb{Q}^{\text{alg}}/L_x) \cdot [x, 1] \right| \geq C_3 C_4^{i(\mathbf{M})} [K_{\mathbf{M}}^m : K_{\mathbf{M}}]^{C_5} |\text{disc } L_x|^{C_6}$$

(where we use notations (a)–(d) from Theorem 1.4).

Then the André–Oort conjecture (Conjecture 1.3) holds for $\text{Sh}_K(\mathbf{G}, X)$.

It is known that the Generalised Riemann Hypothesis for CM fields implies the Galois bounds assumed in Theorem 1.2 (see [UY15]). Thus Theorem 1.2 gives a new proof of the André–Oort conjecture assuming the GRH for CM fields.

The André–Oort conjecture is already known under the GRH due to the work of Klingler, Ullmo and Yafaev (see [KY14] and [UY14a]).

However, Theorem 1.2 is stronger than simply “GRH implies André–Oort” because it is conceivable that the necessary Galois bounds could be proved without using the GRH. Tsimerman has recently proved these bounds for the case of the moduli space of principally polarised abelian varieties \mathcal{A}_g (see [Tsi]). This represents an advantage of Theorem 1.2 over the proof of the André–Oort conjecture under the GRH by Klingler, Ullmo and Yafaev, as their proof depends much more heavily on the GRH.

The \mathcal{A}_g case of Theorem 1.1 was proved by Pila and Tsimerman (see [PT13] Theorem 3.1). Earlier known cases of Theorem 1.1 included modular curves (see [Pil09]) and Hilbert modular surfaces (see [DY11]). The analogous bound for abelian varieties in place of Shimura varieties, used in Pila and Zanner’s proof of the Manin–Mumford conjecture, is trivial.

It is not difficult to deduce Theorem 1.1 for all Shimura varieties of abelian type from the case of \mathcal{A}_g , so the new contribution of this paper is that it holds for Shimura varieties which are not of abelian type. Our proof works uniformly for all Shimura varieties, but as we explain later in the introduction, moving beyond Shimura varieties of abelian type introduced substantial new difficulties.

1.1 The Pila–Zannier strategy for proving André–Oort

We will outline the various ingredients used in Pila and Zannier’s strategy for proving the André–Oort conjecture, and how Theorem 1.1 fits into this. There are several expositions available on how this strategy is implemented once one has the necessary ingredients (see, for example, [PT13] for the case of \mathcal{A}_2 , the moduli space of principally polarised abelian surfaces, [Ull14] for the case of \mathcal{A}_g^n and [Daw15] for the case of a general Shimura variety).

We begin by recalling the statement of the André–Oort Conjecture.

Conjecture 1.3. *Let S be a Shimura variety and let Σ be a set of special points in S . Every irreducible component of the Zariski closure of Σ in S is a special subvariety.*

The most appealing feature of the Pila–Zannier strategy is the manner in which it combines a number of independent ingredients to deliver a relatively simple proof of the conjecture. The ingredients themselves are substantially more complicated and belong to various branches of mathematics. The primordial result is the so-called Pila–Wilkie counting theorem (see [PW06]), yielding strong upper bounds on the number of algebraic points of bounded height and degree away from the algebraic part of a set definable in an o-minimal structure.

The fact that Shimura varieties are amenable to tools from o-minimality is due to the second ingredient, which states that, if we write $S := \Gamma \backslash X^+$, the restriction of the uniformisation map $X^+ \rightarrow S$ to \mathcal{F} is definable in the o-minimal structure $\mathbb{R}_{\text{an,exp}}$. For the moduli space \mathcal{A}_g of principally polarised abelian varieties this is a theorem of Peterzil and Starchenko (see [PS10]). It has since been demonstrated for any Shimura variety in the work [KUY] of Klingler, Ullmo and Yafaev.

The third ingredient is the so-called hyperbolic Ax-Lindemann-Weierstrass conjecture, which is a statement in functional transcendence regarding the uniformisation map. This was first proved for products of modular curves by Pila in his seminal work [Pil11], then for compact Shimura varieties by Ullmo and Yafaev in [UY14b] and for \mathcal{A}_g by Pila and Tsimerman in [PT14]. The case of a general Shimura variety has recently been demonstrated by Klingler, Ullmo and Yafaev (see [KUY]).

The final two ingredients are arithmetic in nature and serve as opposing forces. The goal of these two ingredients is to show that there are constants C_7 and C_8 , such that for every pre-special point $x \in \mathcal{F}$ with image $[x, 1]$ in the Shimura variety,

$$H(x) \leq C_7 \left| \text{Gal}(\mathbb{Q}^{\text{alg}}/L_x) \cdot [x, 1] \right|^{C_8}.$$

This is broken down into two parts. The first is the lower bound for the sizes of Galois orbits of special points. As we discussed previously, this is known unconditionally for \mathcal{A}_g and under the GRH for all Shimura varieties. Without the GRH, this bound remains an open problem in general. The second of the final ingredients is the upper bound for the heights of pre-special points in the fundamental domain \mathcal{F} , which is the subject of this paper. We have already discussed previous work on special cases of these final two ingredients.

Gao has generalised the Pila–Zannier strategy to the mixed André–Oort conjecture (see [Gaob]). In the note [Gaoa], he shows that Theorems 1.1 and 4.1 imply the mixed André–Oort conjecture for all mixed Shimura varieties, conditional on the same Galois bounds for pure Shimura varieties as Theorem 1.2.

1.2 Pre-special points and realisations

We emphasise that we are talking about the heights of *pre-special* points, namely points belonging to X^+ , rather than *special points*, which in our terminology are precisely the images of pre-special points in S . Recall that X is a $\mathbf{G}(\mathbb{R})$ -conjugacy class of morphisms $\mathbb{S} \rightarrow G_{\mathbb{R}}$, where \mathbb{S} is the Deligne torus. An element $x \in X$ is said to be **pre-special** if it factors through a subtorus of \mathbf{G} defined over \mathbb{Q} . In the classical setting, where X^+ is *realised* as the upper half-plane or the Siegel upper half-space \mathcal{H}_g , points of X^+ are often referred to as **period matrices**. The period matrices corresponding to CM abelian varieties always have algebraic entries. The

theorem of Pila and Tsimerman, which we are generalising, bounds the height of the period matrices of CM abelian varieties.

In general, in order to define the height $H(x)$ of a point $x \in X$, we must choose a **realisation** \mathcal{X} of X . By this we mean an analytic subset of a quasi-projective complex variety, equipped with a transitive action of $\mathbf{G}(\mathbb{R})$ by holomorphic automorphisms of \mathcal{X} and with an isomorphism of $\mathbf{G}(\mathbb{R})$ -homogeneous spaces $\mathcal{X} \rightarrow X$ such that, for every $x_0 \in \mathcal{X}$, the map

$$\mathbf{G}(\mathbb{R}) \rightarrow \mathcal{X} : g \mapsto g \cdot x_0$$

is semi-algebraic (regarding \mathcal{X} as a subset of a real algebraic variety by taking real and imaginary parts of the coordinates). A morphism of realisations is defined to be a $\mathbf{G}(\mathbb{R})$ -equivariant biholomorphism commuting with the respective isomorphisms with X . By [Ull14] Lemme 2.1, every realisation is a semialgebraic set and every morphism of realisations is semialgebraic.

We restrict ourselves to realisations \mathcal{X} for which the the action of $\mathbf{G}(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ is **semialgebraic over** $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$. By this, we mean that:

- (i) \mathcal{X} is an analytic subset of the complex points of a quasi-projective variety defined over \mathbb{Q}^{alg} , and
- (ii) for each \mathbb{Q}^{alg} -point $x_0 \in \mathcal{X}$, the map $g \mapsto g \cdot x_0$ is definable in the language of ordered rings with constants $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$. (In other words, the graph of this map is a semialgebraic set which can be defined by equations and inequalities whose coefficients are in $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$.)

Note that throughout this paper, \mathbb{Q}^{alg} denotes the subfield of algebraic numbers in \mathbb{C} , so $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ is unambiguously defined. The Borel realisation is an example of a realisation which is semialgebraic over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ (see [UY11]) and any two such realisations are related by an isomorphism which is semialgebraic over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$.

By [UY11] Proposition 3.7, pre-special points in such a realisation have coordinates in \mathbb{Q}^{alg} . Alternatively, this fact follows from the fact that each pre-special point in X is the unique fixed point of some element of $\mathbf{G}(\mathbb{Q})$ (see [DH] Theorem 2.3). Therefore, we can write $H(x)$ for a pre-special point $x \in X$ to mean its multiplicative Weil height in a chosen realisation \mathcal{X} . Henceforth, for any Shimura datum (\mathbf{G}, X) , we tacitly assume the choice of a realisation for X .

In order for the Pila–Zannier strategy to work, it is necessary that pre-special points in \mathcal{X} are not only algebraic but are defined over number fields of uniformly bounded degree. In the case of the Borel realisation, it follows from the proof of [UY11] Lemma 3.8 that, given a faithful representation $\rho: \mathbf{G} \rightarrow \mathbf{GL}_n$, a pre-special point $x \in \mathcal{X}$ is defined over the splitting field L of a maximal torus in \mathbf{GL}_n containing the Mumford-Tate group of x . The rank d of this torus is of course

bounded by n , and the Galois action on its group of characters is given by an embedding

$$\mathrm{Gal}(L/\mathbb{Q}) \hookrightarrow \mathbf{GL}_d(\mathbb{Z}).$$

Therefore, it follows from a classical result of Minkowski on finite subgroups of $\mathbf{GL}_d(\mathbb{Z})$ that the degree $[L : \mathbb{Q}]$ is bounded by a constant depending only on n . The fact that any two realisations semialgebraic over $\mathbb{Q}^{\mathrm{alg}} \cap \mathbb{R}$ are related by an isomorphism which is semialgebraic over $\mathbb{Q}^{\mathrm{alg}} \cap \mathbb{R}$ implies that this holds for all such realisations.

1.3 Precise statement of our final bound and minor remarks

For convenience, we include here a precise statement of our final bound for heights of pre-special points (i.e. the combination of Theorem 1.1 and Theorem 4.1). The form of the bound in Theorem 1.4 matches the form of the conjectural Galois bounds (as mentioned in Theorem 1.2). Hence the invariants which appear seem to be the natural invariants of a (pre-)special point of a general Shimura variety.

Theorem 1.4. *Let (\mathbf{G}, X) be a Shimura datum in which \mathbf{G} is an adjoint group.*

Choose a realisation of X such that the action of $\mathbf{G}(\mathbb{Q}^{\mathrm{alg}} \cap \mathbb{R})$ on X is semialgebraic and defined over $\mathbb{Q}^{\mathrm{alg}} \cap \mathbb{R}$. Let $H(x)$ denote the multiplicative Weil height of a point in $X(\mathbb{Q}^{\mathrm{alg}})$ with respect to this realisation.

Let K be a compact open subgroup of $\mathbf{G}(\mathbb{A}_f)$ and let \mathcal{F} be a fundamental set in X for $K \cap \mathbf{G}(\mathbb{Q})$, as in Théorème 13.1 of [Bor69], with respect to a pre-special base point x_0 .

There exist constants $C_9, C_{10} > 0$ such that for all $C_{11} > 0$, there exists $C_{12} > 0$ (where C_9 and C_{10} depend only on $\mathbf{G}, X, \mathcal{F}$ and the realisation of X , while C_{12} depends on these data and also C_{11}) such that:

For each pre-special point $x \in \mathcal{F}$:

- (a) *Let \mathbf{M} denote the Mumford–Tate group of x (which is a torus because x is pre-special).*
- (b) *Let $K_{\mathbf{M}} = K \cap \mathbf{M}(\mathbb{A}_f)$ and let $K_{\mathbf{M}}^m$ be the maximal compact subgroup of $\mathbf{M}(\mathbb{A}_f)$.*
- (c) *Let $i(\mathbf{M})$ be the number of primes p for which $K \cap \mathbf{M}(\mathbb{Q}_p)$ is strictly contained in the maximal compact subgroup of $\mathbf{M}(\mathbb{Q}_p)$.*
- (d) *Let L_x be the splitting field of \mathbf{M} .*

Then

$$H(x) \leq C_{12} C_{11}^{i(\mathbf{M})} [K_{\mathbf{M}}^m : K_{\mathbf{M}}]^{C_9} |\mathrm{disc} L_x|^{C_{10}}.$$

The quantifiers associated with the constants in Theorem 1.4 are complicated. In particular, why is C_{11} universally quantified while the others are existentially quantified? The reason for this is the need to compare Theorem 1.4 with the Galois bound

$$C_3 C_4^{i(\mathbf{M})} [K_{\mathbf{M}}^m : K_{\mathbf{M}}]^{C_5} |\text{disc } L_x|^{C_6} \leq |\text{Gal}(\mathbb{Q}^{\text{alg}}/L_x) \cdot [x, 1]|$$

from the assumption of Theorem 1.2 and end up with a conclusion

$$H(x) \leq C_7 |\text{Gal}(\mathbb{Q}^{\text{alg}}/L_x) \cdot [x, 1]|^{C_8}. \quad (*)$$

The difficulty in comparing these two bounds is that C_4 might be less than 1. We will need to choose a large value for C_8 in order for the $[K_{\mathbf{M}}^m : K_{\mathbf{M}}]$ and $|\text{disc } L_x|$ factors on the right hand side of (*) to be larger than the corresponding factors on the left. In order for the $i(\mathbf{M})$ factors in (*) to work out, we need

$$C_{11} \leq C_4^{C_8}. \quad (\dagger)$$

But if $C_4 < 1$, then we cannot achieve (\dagger) by making C_8 large. Instead we need the freedom to choose C_{11} in Theorem 1.4. Furthermore, our choice of C_{11} depends on C_8 which in turn depends on C_9 and C_{10} so the quantifiers of these constants must be ordered as in the statement of the theorem. Meanwhile, “there exists C_{12} ” has to come last due to the proof of Theorem 1.4.

In the statements of Theorems 1.1 and 1.4 we assume that \mathbf{G} is an adjoint group i.e. it has trivial centre. In the context of the André–Oort conjecture this is entirely inconsequential (see [EY03] §2). The fact that \mathbf{G} is adjoint ensures that the Hodge structures induced by any representation of \mathbf{G} are of pure weight, which is essential to our proof (indeed, it ensures that these Hodge structures are pure of weight 0, but this is not essential). We have taken advantage of the hypothesis that \mathbf{G} is adjoint to make some other, non-essential, simplifications.

Note that it is a trivial observation that Theorems 1.1 and 1.4 fail if one does not restrict to a fundamental set in X .

1.4 Comparison with Theorem 3.1 of [PT13]

In essence, the proof of [PT13] Theorem 3.1 has three steps: studying the relationship between polarisations of a CM abelian variety and its endomorphism ring, choosing a suitable symplectic basis and reduction theory for matrices in the Siegel upper half-space \mathcal{H}_g . (The division between the cases of simple and non-simple abelian varieties which appears in [PT13] affects only the first of these steps. It is true that our paper could be greatly simplified if we could ignore the case of non-irreducible isotypic Hodge structures, but given the final structure of our proof,

we cannot easily separate it into parts dealing with irreducible and non-irreducible Hodge structures.)

We use polarisations in the same way as [PT13] (Proposition 2.9). We also use reduction theory (section 2.6), although in dealing with general Shimura varieties, our calculations are necessarily less explicit.

The main additional difficulty for general Shimura varieties concerns the second of these steps. Instead of a symplectic basis, we must find a basis having the property we call **G**-admissibility, which we can describe using a multilinear form Φ on our Hodge structure whose stabiliser is equal to **G** (the existence of Φ is guaranteed by Chevalley's theorem). In the case of \mathcal{A}_g , $\mathbf{G} = \mathbf{GSp}_{2g}$ and this multilinear form is the standard symplectic form. (Of course, \mathbf{GSp}_{2g} is the group of similitudes of the standard symplectic form, rather than its stabiliser, but this is an unimportant technical difference.) It is much easier to manipulate symplectic forms than general multilinear forms, and this leads to the most important difficulties in this paper (sections 2.5 and 3).

It is perhaps worth saying a word about why this issue of **G**-admissibility is so important. Throughout the paper, we work with a faithful representation $\mathbf{G} \hookrightarrow \mathbf{GL}_n$ and we embed $X \subset \mathrm{Hom}(\mathbb{S}, \mathbf{G})$ in a $\mathbf{GL}_n(\mathbb{C})$ -conjugacy class $X_{\mathbf{GL}_n}$ of morphisms $\mathbb{S}_{\mathbb{C}} \rightarrow \mathbf{GL}_{n,\mathbb{C}}$. A basis is said to be **G**-admissible if it lies in the $\mathbf{G}(\mathbb{R})$ -orbit of some fixed reference basis. Thus **G**-admissibility has two parts: the relevant matrix must be both real and in **G**. The requirement that this matrix be real is the purpose of section 3 while the requirement that it be in **G** is the core of the difficulties in section 2.

In earlier attempts to prove our theorem, we ignored **G**-admissibility and attempted to use reduction theory directly in $\mathbf{GL}_n(\mathbb{R})$ rather than in $\mathbf{G}(\mathbb{R})$. This fails because reduction theory in $\mathbf{GL}_n(\mathbb{R})$ works with the symmetric space $\mathbf{GL}_n(\mathbb{R})/\mathbb{R}^\times \mathbf{O}_n(\mathbb{R})$, which is not the same as $X_{\mathbf{GL}_n}$.

1.5 Outline of paper

The proof of Theorem 1.1 is found in section 2. In order to attach Hodge structures to points of X , we have to choose a representation of the group **G**. Of course, the constants we get depend on which representation we choose but the fact that there always exists some representation satisfying the conditions of Theorem 1.1 means that this is sufficient for proving Theorem 1.4.

In section 3 we prove the following theorem, saying that an affine variety defined over $\mathbb{Q}^{\mathrm{alg}} \cap \mathbb{R}$ has a $(\mathbb{Q}^{\mathrm{alg}} \cap \mathbb{R})$ -point whose height is polynomially bounded relative to the coefficients of polynomials defining the variety. A version of this with $\mathbb{Q}^{\mathrm{alg}} \cap \mathbb{R}$ replaced by $\mathbb{Q}^{\mathrm{alg}}$ is straightforward, but we need the version with $\mathbb{Q}^{\mathrm{alg}} \cap \mathbb{R}$ in section 2.

Theorem 1.5. *For all positive integers m, n, D , there are constants C_{13} and C_{14} depending on m, n and D such that:*

For every affine algebraic set $V \subset \mathbb{A}^n$ defined over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ by polynomials $f_1, \dots, f_m \in (\mathbb{Q}^{\text{alg}} \cap \mathbb{R})[X_1, \dots, X_n]$ of degree at most D and height at most H , if $V(\mathbb{R})$ is non-empty, then $V(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ contains a point of height at most $C_{13}H^{C_{14}}$.

Most of this section is elementary, based on the proof of the Noether normalisation lemma. However in some cases, in order to prove the $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ version of Theorem 1.5, we need to show that if we have bounds for the heights of polynomials defining a variety V , then we can find a proper algebraic subset $V' \subset V$ such that $\text{Sing } V \subset V'$ and the heights of polynomials defining V' are also bounded. The proof of this uses Philippon's arithmetic Bézout theorem ([Phi95]), Nesterenko's study of the Chow form ([Nes77]) and an idea of Bombieri, Masser and Zannier ([BMZ07]) to use the Chow form in studying the singular locus.

Finally in section 4 we relate the bound in terms of the discriminant of the centre of the endomorphism ring which appears in Theorem 1.1 to the bound in terms of invariants of the Mumford–Tate torus which appears in Theorem 1.4. This generalises arguments of Tsimerman for the \mathcal{A}_g case (see [Tsi12] section 7.2).

Let us say a little more about the proof of Theorem 1.1. This proof is inspired by the proof of [PT13] Theorem 3.1 but turns out to be significantly more difficult, as explained above (and we have found it convenient to write it in a very different way). We fix a pre-special base point x_0 for X , and we aim to find an element (which we will call g_4) in $\mathbf{G}(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ such that

$$g_4 x_0 = x \text{ and the height of } g_4 \text{ is polynomially bounded.}$$

In order to construct g_4 , we first construct several different elements of $\mathbf{GL}(E_{\mathbb{C}})$ which map x_0 to x (by conjugation in $\text{Hom}(\mathbb{S}_{\mathbb{C}}, \mathbf{GL}(E_{\mathbb{C}}))$) but which do not satisfy all the conditions we want for g_4 : they are not always in $\mathbf{G}(\mathbb{R})$ and they satisfy weaker bounds than a straightforward height bound. (Note: in section 2, we talk about bases for $E_{\mathbb{C}}$ rather than elements of $\mathbf{GL}(E_{\mathbb{C}})$. These are equivalent once we have fixed a reference basis.) These constructions use Minkowski's theorem, the theory of maximal tori and some calculations with Hermitian forms.

We then use Theorem 1.5 to construct g_4 itself. Initially we do not get a height bound for g_4 , only for the matrix relating g_3 and g_4 . However up to this point we have not used the fact that x is in the fundamental set \mathcal{F} . We use the definition of the fundamental set, together with our various other bounds for g_4 , to conclude that the height of g_4 is bounded.

Some explanatory paragraphs of this preprint are omitted in the published version.

2 Height bound in terms of the discriminant of the endomorphism ring

In this section we show that the heights of pre-special points in suitable fundamental sets are polynomially bounded with respect to the discriminant of the centre of the endomorphism ring of the associated Hodge structure. In other words, we prove Theorem 1.1.

To prove Theorem 1.1, we construct an element $g_4 \in \mathbf{G}(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ such that

$$g_4 x_0 = x$$

and g_4 has polynomially bounded height. (Throughout this section, **polynomially bounded** means bounded above by an expression of the form $C |\text{disc } R_x|^{C'}$.) This suffices to prove the theorem because the action of $\mathbf{G}(\mathbb{R})$ on X is semialgebraic.

On the way to constructing g_4 , we will need to talk about gx_0 not just for $g \in \mathbf{G}(\mathbb{R})$ but also for $g \in \mathbf{GL}(E_{\mathbb{C}})$. In order for this to make sense, we enlarge

$$X = \text{a } \mathbf{G}(\mathbb{R})\text{-conjugacy class in } \text{Hom}(\mathbb{S}, \mathbf{G}_{\mathbb{R}})$$

to a $\mathbf{GL}(E_{\mathbb{C}})$ -conjugacy class in $\text{Hom}(\mathbb{S}_{\mathbb{C}}, \mathbf{GL}(E_{\mathbb{C}}))$. Whenever we write gx_0 for $g \in \mathbf{GL}(E_{\mathbb{C}})$ we always refer to the action of $\mathbf{GL}(E_{\mathbb{C}})$ on $\text{Hom}(\mathbb{S}_{\mathbb{C}}, \mathbf{GL}(E_{\mathbb{C}}))$ by conjugation.

Rather than talking about elements of $\mathbf{GL}(E_{\mathbb{C}})$, we will often talk about bases for $E_{\mathbb{C}}$. Since $\mathbf{GL}(E_{\mathbb{C}})$ acts simply transitively on the set of bases, this is purely a matter of language. It is more convenient to use the language of bases because, while we are ultimately interested in how a general basis \mathcal{B} is related to $\mathcal{B}^{(0)}$ (a fixed basis linked to the base point x_0), for our computations we shall want to consider the coordinates of \mathcal{B} relative to $\mathcal{B}_{\mathbb{Z}}$ (a fixed basis for the lattice $\mathcal{B}_{\mathbb{Z}}$).

We shall define a number of adjectives to describe special types of (ordered) bases for $E_{\mathbb{C}}$. For now, we omit certain technical complications from these definitions; the full details will appear in section 2.1.4.

1. A basis \mathcal{B} for $E_{\mathbb{C}}$ is a **Hodge basis** for $x \in X$ if every vector in \mathcal{B} is an eigenvector for the Hodge parameter x (and \mathcal{B} is ordered in a correct way).

The purpose of this definition is to relate bases to elements of $x \in X$: \mathcal{B} will be a Hodge basis for x if and only if the element $g \in \mathbf{GL}(E_{\mathbb{C}})$ mapping $\mathcal{B}^{(0)}$ to \mathcal{B} satisfies $gx_0 = x$.

2. A basis \mathcal{B} for $E_{\mathbb{C}}$ is a **diagonal basis** for E_x if each vector in \mathcal{B} is an eigenvector for the centre of $\text{End}_{\mathbb{Q}\text{-HS}} E_x$, and \mathcal{B} is a Hodge basis. (If x is pre-special, then the condition that \mathcal{B} is a Hodge basis serves only to control the ordering of \mathcal{B}).

3. A basis \mathcal{B} for $E_{\mathbb{C}}$ is **G-admissible** if the element $g \in \mathbf{GL}(E_{\mathbb{C}})$ which maps $\mathcal{B}^{(0)}$ to \mathcal{B} is in $\mathbf{G}(\mathbb{R})$.
4. A diagonal basis \mathcal{B} for E_x is **Galois-compatible** if the action of $\text{Aut}(\mathbb{C})$ on coordinates permutes the vectors of \mathcal{B} in a suitable way.
The benefit of Galois-compatibility is that, if we have a Galois-compatible basis and we can bound the absolute values and the denominators of its coordinates, then we can deduce a bound for the heights of the coordinates.
5. A basis \mathcal{B} for $E_{\mathbb{C}}$ is **weakly bounded** if all its coordinates are algebraic numbers with polynomially bounded denominators and its covolume is polynomially bounded.

Using these adjectives, we can translate our search for g_4 into the following: we construct a **G-admissible** Hodge basis $\mathcal{B}^{(4)}$ for the pre-special point x whose coordinates are algebraic numbers with polynomially bounded height. In order to do this, we will first construct a series of other bases satisfying some but not all of the properties we want:

1. $\mathcal{B}^{(1)}$, a **G-admissible** diagonal basis with no bounds at all;
2. $\mathcal{B}^{(2)}$, a Galois-compatible weakly bounded diagonal basis which need not be **G-admissible**;
3. $\mathcal{B}^{(3)}$, a Galois-compatible weakly bounded diagonal basis on which the values of a polarisation are bounded;
4. $\mathcal{B}^{(4)}$, a weakly bounded **G-admissible** diagonal basis, such that the coordinates of $\mathcal{B}^{(4)}$ with respect to $\mathcal{B}^{(3)}$ are algebraic numbers of polynomially bounded height.

Once we have obtained $\mathcal{B}^{(4)}$ as above, we will then prove that $\mathcal{B}^{(4)}$ has coordinates of polynomially bounded height (relative to $\mathcal{B}_{\mathbb{Z}}$). The point of bounding the height of the coordinates of $\mathcal{B}^{(4)}$ relative to $\mathcal{B}^{(3)}$ is that it allows us to pass bounds back and forth between $\mathcal{B}^{(3)}$ and $\mathcal{B}^{(4)}$. In particular this means that we can exploit the Galois-compatibility of $\mathcal{B}^{(3)}$ while bounding $\mathcal{B}^{(4)}$ (which is not Galois-compatible).

The constructions of $\mathcal{B}^{(1)}$, $\mathcal{B}^{(2)}$ and $\mathcal{B}^{(3)}$ are fairly straightforward. $\mathcal{B}^{(1)}$ is constructed geometrically using maximal tori, while $\mathcal{B}^{(2)}$ is constructed arithmetically using the Minkowski bound. $\mathcal{B}^{(3)}$ is obtained from $\mathcal{B}^{(2)}$ by some calculations with positive-definite Hermitian forms.

The hardest part of the proof is to obtain $\mathcal{B}^{(4)}$ from $\mathcal{B}^{(3)}$. To do this we will use Chevalley's theorem: we can choose a multilinear form $\Phi: E_{\mathbb{Z}}^{\otimes k} \rightarrow \mathbb{Z}$ such that

$$\mathbf{G} = \{g \in \mathbf{GL}(E) \mid \Phi(gv_1, \dots, gv_k) = \Phi(v_1, \dots, v_k) \text{ for all } v_j \in E\}.$$

Then a basis \mathcal{B} is \mathbf{G} -admissible if and only if the values of Φ on \mathcal{B} are the same as its values on $\mathcal{B}^{(0)}$.

We can use the bound for the polarisation on $\mathcal{B}^{(3)}$ together with the Galois-compatibility of $\mathcal{B}^{(3)}$ to show that the values of the multilinear form Φ on $\mathcal{B}^{(3)}$ are algebraic numbers of polynomially bounded height. This allows us to conclude that the following algebraic set is defined by equations of polynomially bounded height:

$$V = \{h \in \mathbf{GL}_n(\mathbb{C}) \mid \mathcal{B}^{(3)}h \text{ is diagonal and } \mathbf{G}\text{-admissible}\}.$$

Here we write $\mathcal{B}^{(3)}h$ with h on the right to mean the basis whose coordinates relative to $\mathcal{B}^{(3)}$ are given by the columns of h . This is different from $g\mathcal{B}^{(3)}$ with a linear map g on the left, which means the basis obtained by applying g to each element of $\mathcal{B}^{(3)}$.

We can apply Theorem 3.1 to V to obtain $\mathcal{B}^{(4)}$ (the existence of $\mathcal{B}^{(1)}$ ensures that $V(\mathbb{R})$ is non-empty).

Finally we have to bound the height of $\mathcal{B}^{(4)}$. In order to do this we will use the fact that x is in the fundamental domain \mathcal{F} – this fact has not been used so far in the construction of $\mathcal{B}^{(4)}$. In particular, this means that the element $g_4 \in \mathbf{G}(\mathbb{R})$ mapping $\mathcal{B}^{(0)}$ to $\mathcal{B}^{(4)}$ is contained in a Siegel set. Using the definition of Siegel sets together with bounds for $\mathcal{B}^{(4)}$ deduced from the fact that $\mathcal{B}^{(3)}$ is weakly bounded and Galois-compatible, we prove that the coordinates of $\mathcal{B}^{(4)}$ have polynomially bounded absolute values. Using again the height bound for the relation between $\mathcal{B}^{(3)}$ and $\mathcal{B}^{(4)}$ and the fact that $\mathcal{B}^{(3)}$ is Galois-compatible, we deduce that the coordinates of $\mathcal{B}^{(4)}$ have polynomially bounded height.

2.1 Notation

In the first three parts of this section, we define notation for data which depend only on the Shimura datum (\mathbf{G}, X) and the representation ρ (when we say “choose” something, we choose it once depending on (\mathbf{G}, X, ρ) and thereafter regard it as fixed). In particular these data do not depend on the special point x . In section 2.1.3, we define notation for data which does depend on x .

Choose a basis $\mathcal{B}_{\mathbb{Z}}$ for $E_{\mathbb{Z}}$. Whenever we talk about the coordinates of a basis in $E_{\mathbb{C}}$, we mean with respect to $\mathcal{B}_{\mathbb{Z}}$.

Let $n = \dim E$.

If \mathcal{B} and \mathcal{B}' are two bases of the vector space E , we define

$$\text{Vol}(\mathcal{B} : \mathcal{B}') = \text{covol}(\mathcal{B}') / \text{covol}(\mathcal{B}).$$

2.1.1 Multilinear forms

By Chevalley's theorem, we can choose a multilinear form $\Phi: E_{\mathbb{Z}}^{\otimes k} \rightarrow \mathbb{Z}$ such that

$$\mathbf{G} = \{g \in \mathbf{GL}(E) \mid \Phi(gv_1, \dots, gv_k) = \Phi(v_1, \dots, v_k) \text{ for all } v_j \in E\}.$$

(By Deligne's extension of Chevalley's theorem, because \mathbf{G} is reductive, we can require that g exactly preserves Φ rather than just up to a scalar. Because of the hypothesis that ρ is self-dual, we can say that Φ is in $E_{\mathbb{Z}}^{\vee \otimes k}$ rather than a mixture of $E_{\mathbb{Z}}^{\vee}$ and $E_{\mathbb{Z}}$.)

Let $\Psi: E_{\mathbb{Q}} \times E_{\mathbb{Q}} \rightarrow \mathbb{Q}$ denote a \mathbf{G} -invariant C -polarisation as in [Del79] 1.1.15 and 1.1.18(b).

By the definition of a C -polarisation, for each $x \in X$, the bilinear form on $E_{\mathbb{C}}$ defined by

$$\Theta_x(u, v) = \Psi_{\mathbb{C}}(\bar{u}, x(i)v)$$

is Hermitian and positive definite, where $\Psi_{\mathbb{C}}$ denotes the \mathbb{C} -bilinear extension of Ψ .

2.1.2 Base point

Choose a pre-special point $x_0 \in \mathcal{F}$ such that the fundamental set \mathcal{F} is a finite union of $\mathbf{G}(\mathbb{Q})$ -translates of $\mathfrak{S}.x_0$, where \mathfrak{S} is a Siegel set for \mathbf{G} over \mathbb{Q} .

Let \mathbf{T}_0 be a maximal torus of \mathbf{G} defined over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ which contains the image of the Hodge parameter x_0 .

Choose a basis $\mathcal{B}^{(0)}$ for $E_{\mathbb{C}}$ which satisfies the following conditions:

- (a) every vector in $\mathcal{B}^{(0)}$ is an eigenvector for \mathbf{T}_0 ;
- (b) for each $v \in \mathcal{B}^{(0)}$, the complex conjugate \bar{v} is also in $\mathcal{B}^{(0)}$;
- (c) each vector of $\mathcal{B}^{(0)}$ has coordinates in \mathbb{Q}^{alg} (relative to $\mathcal{B}_{\mathbb{Z}}$).

Note that each eigenspace of \mathbf{T}_0 is contained in one of the Hodge components $E_{x_0}^{p,-p}$, so condition (a) implies that $\mathcal{B}^{(0)}$ consists of eigenvectors for the Hodge parameter x_0 . In condition (b), \bar{v} may be equal to v .

To construct such a basis $\mathcal{B}^{(0)}$, observe that each eigenspace of \mathbf{T}_0 is defined over \mathbb{Q}^{alg} , and that the complex conjugate of an eigenspace of \mathbf{T}_0 is also eigenspace of \mathbf{T}_0 (perhaps the same one). So we simply take the union of the following:

- (i) for each pair of distinct complex conjugate eigenspaces $E_{\chi}, E_{\bar{\chi}}$ of \mathbf{T}_0 , choose one of these eigenspaces E_{χ} and choose a basis for E_{χ} defined over \mathbb{Q}^{alg} ;
- (ii) the complex conjugates of (i);
- (iii) for each real eigenspace of \mathbf{T}_0 , choose any basis for the eigenspace which is defined over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$.

2.1.3 Variable data

Let F_x denote the centre of $\text{End}_{\mathbb{Q}\text{-HS}} E_x$ and R_x the centre of $\text{End}_{\mathbb{Z}\text{-HS}} E_x$. We can decompose F_x as a direct product of fields:

$$F_x = \prod_{i=1}^s F_i$$

where each F_i is either a CM field or \mathbb{Q} (\mathbb{Q} appears if $E_x^{0,0} \neq \{0\}$).

Let $\Sigma = \text{Hom}(F_x, \mathbb{C})$, and let $\Sigma_i = \text{Hom}(F_i, \mathbb{C})$ for $1 \leq i \leq s$. Then Σ is the disjoint union of the Σ_i .

For each $\sigma \in \Sigma$, let r_σ denote the dimension of the eigenspace in E_x on which F_x acts via σ . Then r_σ is the same for all σ in the same Σ_i , and we also denote this value r_i .

2.1.4 Types of basis

We will define a number of adjectives which we will use to describe bases of $E_{\mathbb{C}}$. Implicitly, our bases will be ordered, so that given two bases \mathcal{B} and \mathcal{B}' , the “linear map g such that $g\mathcal{B}' = \mathcal{B}$ ” is well-defined. However we will never explicitly write down the ordering of a basis except during the definition of a Hodge basis – indeed, most of our bases will be diagonal bases labelled as explained in the definition of a diagonal basis.

Choose and fix an ordering of the basis $\mathcal{B}^{(0)}$ which we have already chosen.

1. An ordered basis $\mathcal{B} = \{v_1, \dots, v_n\}$ of $E_{\mathbb{C}}$ is a **Hodge basis** for E_x if:
 - (a) every vector in \mathcal{B} is an eigenvector for the Hodge parameter x ; and
 - (b) the Hodge type (in E_x) of each v_j is the same as the Hodge type (in E_{x_0}) of the corresponding $v_j^{(0)} \in \mathcal{B}^{(0)}$.

Equivalently, a basis \mathcal{B} is a Hodge basis for E_x if and only if the linear map $g \in \mathbf{GL}(E_{\mathbb{C}})$ defined by $g\mathcal{B}^{(0)} = \mathcal{B}$ satisfies

$$gx_0 = x \text{ in } \text{Hom}(\mathbb{S}_{\mathbb{C}}, \mathbf{GL}(E_{\mathbb{C}})).$$

2. An ordered basis \mathcal{B} of $E_{\mathbb{C}}$ is a **diagonal basis** for V_x if it is a Hodge basis and every vector in \mathcal{B} is an eigenvector for the centre of $\text{End}_{\mathbb{Q}\text{-HS}} E_x$.

If x is pre-special (which is the only case we care about), every eigenvector of the centre of $\text{End}_{\mathbb{Q}\text{-HS}} E_x$ is automatically an eigenvector for x , so requiring that \mathcal{B} be a Hodge basis only adds conditions (b), that \mathcal{B} is ordered so that its Hodge types match those of $\mathcal{B}^{(0)}$.

We shall label the elements of a diagonal basis as

$$\{v_{\sigma j} \mid \sigma \in \Sigma, 1 \leq j \leq r_\sigma\}$$

such that F_x acts on $v_{\sigma j}$ via the character σ .

3. A basis \mathcal{B} for $E_{\mathbb{C}}$ is **G-admissible** if the linear map $g \in \mathbf{GL}(E_{\mathbb{C}})$ defined by $g\mathcal{B}^{(0)} = \mathcal{B}$ satisfies

$$g \in \rho(\mathbf{G}(\mathbb{R})).$$

Note that this imposes two conditions on g : that it is in the image of \mathbf{G} , and that it is real.

4. A diagonal basis $\mathcal{B} = \{v_{\sigma j}\}$ for E_x is **Galois-compatible** if

$$\tau(v_{\sigma j}) = v_{(\tau\sigma)j}$$

for all $\tau \in \text{Aut}(\mathbb{C})$ (where the LHS means: apply τ to each coordinate of $v_{\sigma j}$ with respect to $\mathcal{B}_{\mathbb{Z}}$).

Equivalently, for each $i \in \{1, \dots, s\}$, there are vectors

$$w_{i1}, \dots, w_{ir_i} \in F_i^n$$

i such that

$$v_{\sigma j} = \sigma(w_{ij})$$

for all $\sigma \in \Sigma_i$ and $j \in \{1, \dots, r_i\}$.

We may also apply the term Galois-compatible to other collections indexed by Σ , with the obvious meaning.

5. A basis \mathcal{B} for $E_{\mathbb{C}}$ is **weakly bounded** if all its coordinates (with respect to $\mathcal{B}_{\mathbb{Z}}$) are algebraic numbers with polynomially bounded denominators and its covolume is polynomially bounded.

2.1.5 The Siegel set

The notation we define here for the Siegel set will be used only in section 2.6.

The definition of Siegel set is taken from [Bor69] 12.3, but we have reversed the order of multiplication in \mathbf{G} so that it has a left action on X instead of a right action. Consequently we also reverse the inequality in the definition of A_t .

Note also that for Borel, $\mathbf{G}_{\mathbb{R}}^0$ means the identity component of $\mathbf{G}(\mathbb{R})$ in the real topology, which we denote by $\mathbf{G}(\mathbb{R})^+$, while Borel uses \mathbf{G}^0 to mean the identity component of the algebraic group \mathbf{G} in the Zariski topology, which we denote \mathbf{G}° .

By definition, the Siegel set \mathfrak{S} has the form

$$\mathfrak{S} = \omega A_t K$$

where

- (i) K is the stabiliser of x_0 in $\mathbf{G}(\mathbb{R})$,
- (ii) $A_t = \{a \in \mathbf{S}(\mathbb{R})^+ \mid \varphi(a) \geq t \text{ for all simple } \mathbb{Q}\text{-roots } \varphi \text{ of } \mathbf{G} \text{ w.r.t. } \mathbf{U}\}$ for some $t \in \mathbb{R}_{>0}$,
- (iii) ω is a compact neighbourhood of the identity in $\mathbf{U}(\mathbb{R})\mathbf{M}(\mathbb{R})^+$,
- (iv) \mathbf{S} is a maximal \mathbb{Q} -split torus in \mathbf{G} ,
- (v) \mathbf{M} is the maximal \mathbb{Q} -anisotropic subgroup of $Z_{\mathbf{G}}(\mathbf{S})^\circ$, and
- (vi) \mathbf{U} is the unipotent radical of a minimal \mathbb{Q} -parabolic subgroup of \mathbf{G} containing \mathbf{S} .

Without loss of generality, we may assume that $\rho(\mathbf{S})$ is contained in the set of diagonal matrices and $\rho(\mathbf{U})$ in the set of upper triangular matrices of $\mathbf{GL}(E)$ with respect to $\mathcal{B}_{\mathbb{Z}}$ – this simply requires us to replace $\mathcal{B}_{\mathbb{Z}}$ by a fixed $\mathbf{GL}(E_{\mathbb{Q}})$ -conjugate, which makes a bounded change to $|\text{disc } R_x|$.

By [Bor69] 13.1, our fundamental set \mathcal{F} has the form

$$\mathcal{F} = C.\mathfrak{S}.x_0$$

for some finite set $C \subset \mathbf{G}(\mathbb{Q})$.

2.2 A \mathbf{G} -admissible diagonal basis

We prove that there exists a \mathbf{G} -admissible diagonal basis for E_x . First we construct $g_1 \in \mathbf{G}(\mathbb{R})$ which maps x_0 to x and which conjugates the fixed maximal torus \mathbf{T}_0 to a maximal torus containing the Mumford–Tate group of x . Then we let $\mathcal{B}^{(1)} = g_1 \mathcal{B}^{(0)}$. This is automatically a \mathbf{G} -admissible Hodge basis which diagonalises the Mumford–Tate group of x , and we prove that this implies that it is a diagonal basis.

Lemma 2.1. *There exists $g_1 \in \mathbf{G}(\mathbb{R})$ such that $g_1 x_0 = x$ and $g_1 \mathbf{T}_0 g_1^{-1}$ contains the Mumford–Tate group of E_x .*

Proof. First choose any $g \in \mathbf{G}(\mathbb{R})$ such that $gx_0 = x$. Then $g\mathbf{T}_0g^{-1}$ is a maximal \mathbb{R} -torus of \mathbf{G} containing the image of x , but it might not contain the Mumford–Tate group of E_x .

Choose any maximal torus $\mathbf{T} \subset \mathbf{G}$ defined over \mathbb{R} which contains the Mumford–Tate group of E_x . Then $g\mathbf{T}_0g^{-1}$ and \mathbf{T} are both maximal \mathbb{R} -tori in K_x , the stabiliser of x in $\mathbf{G}(\mathbb{R})$. Since K_x is compact, there is some $g' \in K_x$ which conjugates $g\mathbf{T}_0g^{-1}$ into \mathbf{T} .

Taking $g_1 = g'g$ proves the lemma ($g_1x_0 = x$ because g' stabilises x). \square

Lemma 2.2. *Let \mathcal{B} be a Hodge basis for E_x . If \mathcal{B} diagonalises the Mumford–Tate group of E_x , then \mathcal{B} is a diagonal basis for E_x .*

Proof. Let $\mathbf{M} \subset \mathbf{G}$ be the Mumford–Tate group of E_x .

We write $\text{End } \rho|_{\mathbf{M}}$ for the endomorphism algebra of the \mathbb{Q} -representation $\rho|_{\mathbf{M}}: \mathbf{M} \rightarrow \mathbf{GL}(E)$, considered as a subalgebra of $\text{End}(E_{\mathbb{Q}})$. It is a standard property of Mumford–Tate groups that

$$\text{End } \rho|_{\mathbf{M}} = \text{End}_{\mathbb{Q}\text{-HS}} E_x$$

and hence the centre of $\text{End } \rho|_{\mathbf{M}}$ is F_x .

Let \mathbf{Z} denote the centre of the centraliser of \mathbf{M} in $\mathbf{GL}(E)$. Since the \mathbb{Q} -points of the centraliser of \mathbf{M} in $\mathbf{GL}(E)$ are the same as the invertible elements of $\text{End } \rho|_{\mathbf{M}}$, we deduce that $\mathbf{Z}(\mathbb{Q}) = F_x^{\times}$ (as subgroups of $\mathbf{GL}(E_{\mathbb{Q}})$).

Since \mathbf{M} is commutative, $\mathbf{M} \subset \mathbf{Z}$ and so each eigenspace for \mathbf{Z} in $E_{\mathbb{C}}$ is contained in an eigenspace for \mathbf{M} in $E_{\mathbb{C}}$. We need to establish the inclusion of eigenspaces in the opposite direction.

The isotypic decomposition of the \mathbb{Q} -Hodge structure E_x is

$$E_x = \bigoplus_i E_{x,i} \tag{*}$$

where $E_{x,i}$ is isomorphic to the r_i -th power of an irreducible \mathbb{Q} -Hodge structure with endomorphism algebra F_i . Because sub- \mathbb{Q} -Hodge structures are the same as subrepresentations of the Mumford–Tate group, this is the same as the isotypic decomposition of the \mathbb{Q} -representation $\rho|_{\mathbf{M}}$. Since \mathbf{M} is a torus, it follows that each eigenspace of \mathbf{M} in $E_{\mathbb{C}}$ is contained in one of the $E_{x,i}$ and has dimension r_i .

By looking at the action of $\text{End } \rho|_{\mathbf{M}}$, we see that (*) is also the isotypic decomposition of the \mathbb{Q} -representation $\rho|_{\mathbf{Z}}$, and hence that every eigenspace of \mathbf{Z} in $E_{\mathbb{C}}$ is contained in one of the $E_{x,i}$ and has dimension r_i .

Thus each eigenspace of \mathbf{Z} in $E_{\mathbb{C}}$ has the same dimension r_i as the eigenspace of \mathbf{M} which contains it, so the eigenspaces of \mathbf{Z} are equal to the eigenspaces of \mathbf{M} . Thus the fact that \mathcal{B} diagonalises \mathbf{M} implies that it also diagonalises \mathbf{Z} (or in other words, it diagonalises the action of F_x on E_x). \square

Proposition 2.3. *There exists a \mathbf{G} -admissible diagonal basis $\mathcal{B}^{(1)} = \{v_{\sigma_j}^{(1)}\}$ for E_x .*

Proof. Choose g_1 as in Lemma 2.1, and let $\mathcal{B}^{(1)} = g_1\mathcal{B}^{(0)}$. Since $g_1x_0 = x$, this gives a Hodge basis for E_x .

Since $\mathcal{B}^{(0)}$ diagonalises \mathbf{T}_0 , $\mathcal{B}^{(1)}$ diagonalises $g_1\mathbf{T}_0g_1^{-1}$ and hence a fortiori \mathbf{M} . Lemma 2.2 implies that $\mathcal{B}^{(1)}$ is a diagonal basis for E_x .

Finally $\mathcal{B}^{(1)}$ is \mathbf{G} -admissible because $g_1 \in \mathbf{G}(\mathbb{R})$. □

2.3 A weakly bounded diagonal basis

In order to prove the existence of a weakly bounded diagonal basis, we need a generalisation of Minkowski's bound to arbitrary lattices in a module over a product of number fields. This is essentially Claim 3.1 from [PT13] (the generalisation of Minkowski's bound from ideals in a ring of integers to lattices), but we have included all the generalisations we will need in a single statement: working in a vector space over a number field rather than in the field itself, and working with a product of number fields rather than just a single field.

The deduction of Proposition 2.5 from Lemma 2.4 is essentially just bookkeeping. The construction gives us Galois-compatibility for free.

Lemma 2.4. *For all positive integers d, s and r_1, \dots, r_s , there are constants C_{15} and C_{16} depending on d, s and r_1, \dots, r_s such that:*

For all number fields F_1, \dots, F_s of degree at most d and all \mathbb{Z} -lattices

$$L \subset \prod_i F_i^{r_i},$$

if we let $F = \prod_i F_i$, $R = \text{Stab}_F L$ and $L_0 = \prod_i \mathcal{O}_{F_i}^{r_i}$, then there exists $\nu \in \prod_i \mathbf{GL}_{r_i}(F_i)$ such that

$$\nu L \subset L_0 \text{ and } [L_0 : \nu L] \leq C_{15} |\text{disc } R|^{C_{16}}.$$

Proof. Let $\mathcal{O}_F = \prod_i \mathcal{O}_{F_i} \subset F$, and let $e_R = [\mathcal{O}_F : R]$. Then

$$\text{disc } R = e_R^2 \prod_i \text{disc } F_i.$$

Let $L' = \mathcal{O}_F.L$. Then L' is an \mathcal{O}_F -module such that

$$e_R L' \subset L \subset L'.$$

This implies that $[L' : L] \leq e_R^{\text{rk}_{\mathbb{Z}} L}$.

Since L' is an \mathcal{O}_F -module, it splits as a direct sum

$$L' = \bigoplus_i L'_i$$

where each $L'_i \subset F_i^{r_i}$ is an \mathcal{O}_{F_i} -module of rank r_i . Furthermore since \mathcal{O}_{F_i} is a Dedekind domain, L'_i is isomorphic (as an \mathcal{O}_{F_i} -module) to $\mathcal{O}_{F_i}^{r_i-1} \oplus J_i$ for some ideal $J_i \subset \mathcal{O}_{F_i}$. By Minkowski's first theorem, we can choose J_i such that $[\mathcal{O}_{F_i} : J_i]$ is bounded above by $|\text{disc } F_i|^{1/2}$ times a constant depending only on d .

In other words, there exists $\nu_i \in \mathbf{GL}_{r_i}(F_i)$ such that $\nu_i L'_i = \mathcal{O}_{F_i}^{r_i-1} \oplus J_i$, which is contained in $\mathcal{O}_{F_i}^{r_i}$ with index bounded by a constant times $|\text{disc } F_i|^{1/2}$.

Letting $\nu = (\nu_1, \dots, \nu_s) \in \prod_i \mathbf{GL}_{r_i}(F_i)$, we get that

$$\nu L \subset \nu L' \subset L_0 \text{ and } [L_0 : \nu L] = [L' : L] \prod_i [\mathcal{O}_{F_i}^{r_i} : \nu L'_i]$$

and these indices are bounded relative to $|\text{disc } R|$ as required. \square

Proposition 2.5. *There exists a Galois-compatible weakly bounded diagonal basis $\mathcal{B}^{(2)} = \{v_{\sigma j}^{(2)}\}$ for E_x .*

Proof. Choose an isomorphism of F_x -modules $\iota: \prod_i F_i^{r_i} \rightarrow E_{\mathbb{Q}}$. Choose ν as in Lemma 2.4 applied to $L = \iota^{-1}(E_{\mathbb{Z}})$, and replace ι by $\iota \circ \nu^{-1}$. Then

$$E_{\mathbb{Z}} \subset \iota(L_0) \text{ and } [\iota(L_0) : E_{\mathbb{Z}}] \text{ is polynomially bounded.}$$

There is a natural isomorphism of \mathbb{C} -vector spaces

$$\prod_i F_i^{r_i} \otimes_{\mathbb{Q}} \mathbb{C} \rightarrow \mathbb{C}^n$$

built using all the embeddings $F_i \rightarrow \mathbb{C}$ on each copy of F_i . Let \mathcal{B}_F be the basis for $\prod_i F_i^{r_i}$ which maps to the standard basis of \mathbb{C}^n under this isomorphism, and let

$$\mathcal{B}^{(2)} = \iota(\mathcal{B}_F).$$

This consists of eigenvectors for the action of F_x on E_x . If we order the vectors of $\mathcal{B}^{(2)}$ appropriately, then it will satisfy the condition on Hodge types to be a Hodge basis for E_x , giving us a diagonal basis for E_x .

Let P denote the matrix giving the coordinates of $\mathcal{B}_{\mathbb{Z}}$ with respect to $\mathcal{B}^{(2)}$, in other words

$$e_{j'} = \sum_{\sigma \in \Sigma} \sum_{j=1}^{r_{\sigma}} P_{\sigma j j'} v_{\sigma j}^{(2)}$$

where the rows of P are indexed by $(\sigma, j) \in \bigcup_{i=1}^s \Sigma_i \times \{1, \dots, r_i\}$ and the columns are indexed by $j' \in \{1, \dots, n\}$.

Because $\mathcal{B}_{\mathbb{Z}} \subset \iota(L_0)$, the entries of P are algebraic integers and $\text{Aut}(\mathbb{C})$ permutes the rows of P in the natural way i.e.

$$\tau(P_{\sigma j j'}) = P_{(\tau\sigma) j j'}.$$

Furthermore

$$\det P = \text{Vol}(\mathcal{B}^{(2)} : \mathcal{B}_{\mathbb{Z}}) = \text{Vol}(\mathcal{B}_F : L_0) \text{Vol}(L_0 : E_{\mathbb{Z}}) = \left(\prod_i |\text{disc } F_i|^{r_i/2} \right) [L_0 : E_{\mathbb{Z}}]$$

so $(\det P)^2$ is a polynomially bounded rational integer.

Hence P^{-1} has entries which are algebraic numbers with polynomially bounded denominators and its determinant is polynomially bounded (indeed $\det P^{-1} \leq 1$). Since P^{-1} gives the coordinates of $\mathcal{B}^{(2)}$ with respect to $\mathcal{B}_{\mathbb{Z}}$, this says that $\mathcal{B}^{(2)}$ is weakly bounded.

Since $\text{Aut}(\mathbb{C})$ permutes the rows of P , it permutes the columns of P^{-1} in the natural way, and thus $\mathcal{B}^{(2)}$ is Galois-compatible. \square

2.4 A weakly bounded diagonal basis with polynomially bounded polarisation

In order to construct $\mathcal{B}^{(3)}$, we perform a long calculation with Hermitian forms. The key step is Lemma 3.5 from [PT13]. Essentially we want to apply this lemma to the values of the positive definite Hermitian form

$$\Theta_x(u, v) = \Psi_{\mathbb{C}}(\bar{u}, x(i)v)$$

on $\mathcal{B}^{(2)}$, but we have to tweak these values slightly (by using ζ constructed using Lemma 2.6) to get a set of values which are Galois-compatible.

Lemma 2.6. *For all positive integers d , there exist constants C_{17} and C_{18} such that:*

For every totally real number field F of degree d and every collection of signs $s_{\sigma} \in \{\pm 1\}$ indexed by the embeddings $\sigma: F \rightarrow \mathbb{R}$, there exists an element $\zeta \in F^{\times}$ such that

$$H(\zeta) \leq C_{17} |\text{disc } F|^{C_{18}}$$

and for each $\sigma: F \rightarrow \mathbb{R}$, the sign of $\sigma(\zeta)$ is s_{σ} .

Proof. Consider the lattice of integers \mathcal{O}_F in $F \otimes_{\mathbb{Q}} \mathbb{R} \cong \mathbb{R}^d$. Minkowski's second theorem implies that the covering radius μ of this lattice is bounded by a polynomial in $|\text{disc } F|$ (with the polynomial depending only on d).

Consider a ball B in $F \otimes_{\mathbb{Q}} \mathbb{R}$ which has radius greater than μ and which is contained in the hypercube

$$\{z \in F \otimes_{\mathbb{Q}} \mathbb{R} \mid 0 < s_{\sigma} \sigma(z) < 3\mu \text{ for all } \sigma: F \rightarrow \mathbb{R}\}.$$

By the definition of the covering radius, B contains an element ζ of the lattice \mathcal{O}_F . By our choice of B , $\sigma(\zeta)$ has the correct sign for each $\sigma: F \rightarrow \mathbb{R}$. Since ζ is an algebraic integer, its height satisfies

$$H(\zeta) \leq \prod_{\sigma: F \rightarrow \mathbb{R}} |\sigma(\zeta)| < (3\mu)^d$$

which is bounded by a polynomial in $|\text{disc } F|$ as required. \square

Lemma 2.7. *The determinant of the Hermitian matrix*

$$\Theta_x(\mathcal{B}^{(2)}, \mathcal{B}^{(2)})$$

is polynomially bounded.

Proof. We have

$$\det \Theta_x(\mathcal{B}^{(2)}, \mathcal{B}^{(2)}) = \left| \text{Vol}(\mathcal{B}^{(1)} : \mathcal{B}^{(2)}) \right|^2 \det \Theta_x(\mathcal{B}^{(1)}, \mathcal{B}^{(1)}).$$

Now

$$\text{Vol}(\mathcal{B}^{(0)} : \mathcal{B}^{(1)}) = 1$$

because $\mathcal{B}^{(1)}$ is \mathbf{G} -admissible and \mathbf{G} is semisimple (so $\det \rho(\mathbf{G})$ must be trivial).

Hence

$$\text{Vol}(\mathcal{B}^{(1)} : \mathcal{B}^{(2)}) = \text{Vol}(\mathcal{B}^{(0)} : \mathcal{B}^{(2)})$$

is polynomially bounded because $\mathcal{B}^{(2)}$ is weakly bounded. Also

$$\Theta_x(\mathcal{B}^{(1)}, \mathcal{B}^{(1)}) = \Theta_{x_0}(\mathcal{B}^{(0)}, \mathcal{B}^{(0)})$$

is constant, so the lemma is proved. \square

The following lemma is proved in the same way as Lemma 3.5 of [PT13], but for convenience we have written it down as a general statement with the precise list of conditions required.

Lemma 2.8. *For all positive integers d and r , there are constants C_{19} and C_{20} depending on d and r such that:*

For every number field F of degree at most d and every matrix $A \in M_r(F)$, if

- (i) the entries of A have denominator at most D ,*
- (ii) $\text{Nm}_{F/\mathbb{Q}} \det A \leq D$, and*
- (iii) for every $\sigma: F \rightarrow \mathbb{C}$, $\sigma(A)$ is Hermitian and positive definite,*

then there exists $Q \in \mathbf{GL}_r(F)$ such that the entries of Q are algebraic integers and the entries x_{ij} of $Q^\dagger A Q$ satisfy

$$|\sigma(x_{ij})| \leq C_{19}(D |\text{disc } F|)^{C_{20}}$$

for all $\sigma: F \rightarrow \mathbb{C}$.

Proposition 2.9. *There exists a Galois-compatible weakly bounded diagonal basis $\mathcal{B}^{(3)} = \{v_{\sigma j}^{(3)}\}$ for E_x such that the complex numbers*

$$\Theta_x(v_{\sigma j}^{(3)}, v_{\sigma j'}^{(3)})$$

are polynomially bounded for all $\sigma \in \Sigma$, $1 \leq j, j' \leq r_\sigma$.

Proof. Since the Hodge structure E_x has weight 0, $\rho x(i)$ has eigenvalues ± 1 . By Lemma 2.6, we can choose an element $\zeta \in F_x^\times$ which is totally real, such that $\sigma(\zeta)$ has the same sign as the eigenvalue of $\rho x(i)$ on $v_{\sigma 1}^{(2)}$ for all $\sigma \in \Sigma$, and whose height is polynomially bounded.

Define

$$A_{\sigma j j'} = \sigma(\zeta) \Psi_{\mathbb{C}}(\bar{v}_{\sigma j}^{(2)}, v_{\sigma j'}^{(2)})$$

for each $\sigma \in \Sigma$ and $1 \leq j, j' \leq r_\sigma$.

The way we chose the signs of $\sigma(\zeta)$ implies that

$$A_{\sigma j j'} = |\sigma(\zeta)| \Theta_x(v_{\sigma j}^{(2)}, v_{\sigma j'}^{(2)}).$$

Since Θ_x is a positive definite Hermitian form, we deduce that the square matrix A_σ is Hermitian and positive definite for each $\sigma \in \Sigma$.

Observe that, for every $\tau \in \text{Aut}(\mathbb{C})$ and $\sigma \in \Sigma$, we have $\tau \bar{\sigma} = \overline{\tau \sigma}$ because F_x is a product of CM fields and \mathbb{Q} . Since Ψ is defined over \mathbb{Q} , we deduce that

$$\tau(A_{\sigma j j'}) = A_{(\tau \sigma) j j'}$$

or in other words there are matrices $A_i \in M_{r_i}(F_i)$ such that

$$A_{\sigma j j'} = \sigma(A_{i j j'})$$

for all $1 \leq i \leq s$, $\sigma \in \Sigma_i$, $1 \leq j, j' \leq r_i$.

Because $\mathcal{B}^{(2)}$ has polynomially bounded denominators and ζ has polynomially bounded height, the entries of the A_i have polynomially bounded denominators.

Furthermore

$$\prod_{i=1}^s \text{Nm}_{F_i/\mathbb{Q}} \det A_i = \prod_{\sigma \in \Sigma} \det A_\sigma = \left(\prod_{\sigma \in \Sigma} |\sigma(\zeta)| \right) \det \Theta_x(\mathcal{B}^{(2)}, \mathcal{B}^{(2)}).$$

(We are using the fact that the eigenspaces of F_x are orthogonal to each other with respect to Θ_x , so the matrix $\Theta_x(\mathcal{B}^{(2)}, \mathcal{B}^{(2)})$ is block diagonal with one block for each σ .) The values $|\sigma(\zeta)|$ are polynomially bounded because ζ has polynomially bounded height. Using Lemma 2.7, we conclude that

$$\prod_{i=1}^s \text{Nm}_{F_i/\mathbb{Q}} \det A_i$$

is polynomially bounded.

Each value $(\text{Nm}_{F_i/\mathbb{Q}} \det A_i)^{-1}$ is polynomially bounded because the denominators of A_i are polynomially bounded. Hence

$$\text{Nm}_{F_i/\mathbb{Q}} \det A_i$$

is polynomially bounded for each i .

Hence $A_i \in \text{M}_{r_i}(F_i)$ satisfies the conditions of Lemma 2.8 (with D being polynomial in $|\text{disc}(R)|$). So we can choose $Q_i \in \mathbf{GL}_{r_i}(F_i)$ as in the conclusion of that lemma.

Let

$$v_{\sigma j}^{(3)} = \sum_{j'=1}^{r_\sigma} \sigma(Q_{ij'j}) v_{\sigma j'}^{(2)}$$

for all $1 \leq i \leq s$, $\sigma \in \Sigma_i$ and $1 \leq j \leq r_\sigma$.

Now $\mathcal{B}^{(3)} = \{v_{\sigma j}^{(3)}\}$ is a diagonal basis because each $v_{\sigma j}^{(3)}$ is a linear combination of σ -eigenvectors for F_x . Clearly it is also Galois-compatible.

Because the archimedean absolute values of entries of $Q_i^\dagger A_i Q_i$ are polynomially bounded,

$$\text{Nm}_{F_i/\mathbb{Q}} \det Q_i^\dagger A_i Q_i$$

is polynomially bounded. As remarked above, $(\text{Nm}_{F_i/\mathbb{Q}} \det A_i)^{-1}$ is polynomially bounded. Hence

$$\left| \text{Nm}_{F_i/\mathbb{Q}} \det Q_i \right|^2 = (\text{Nm}_{F_i/\mathbb{Q}} \det Q_i^\dagger A_i Q_i) (\text{Nm}_{F_i/\mathbb{Q}} \det A_i)^{-1}$$

is polynomially bounded. We deduce that the covolume of $\mathcal{B}^{(3)}$, namely

$$\left(\prod_{i=1}^s \text{Nm}_{F_i/\mathbb{Q}} \det Q_i \right) \text{Vol}(\mathcal{B}_Z : \mathcal{B}^{(2)}),$$

is polynomially bounded.

Furthermore $\mathcal{B}^{(3)}$ has bounded denominators because $\mathcal{B}^{(2)}$ had bounded denominators and the entries of Q_i are algebraic integers. Hence $\mathcal{B}^{(3)}$ is weakly bounded.

Finally the values

$$|\sigma(\zeta)| \Theta_x(v_{\sigma j}^{(3)}, v_{\sigma j'}^{(3)})$$

are the entries of $\sigma(Q_i^\dagger A_i Q_i)$ and so are polynomially bounded. Since ζ has polynomially bounded height, we conclude that

$$\Theta_x(v_{\sigma j}^{(3)}, v_{\sigma j'}^{(3)})$$

are polynomially bounded. □

2.5 A \mathbf{G} -admissible diagonal basis whose height is bounded relative to $\mathcal{B}^{(3)}$

In this section we will construct the basis $\mathcal{B}^{(4)}$ which is \mathbf{G} -admissible, diagonal, and whose coordinates with respect to $\mathcal{B}^{(3)}$ have polynomially bounded height.

We begin by considering the coordinates of $\mathcal{B}^{(3)}$ with respect to $\mathcal{B}^{(1)}$. Using the fact that Θ_x is polynomially bounded on $\mathcal{B}^{(3)}$, we show that these coordinates are polynomially bounded, and deduce that the multilinear form Φ is polynomially bounded on $\mathcal{B}^{(3)}$. Because $\mathcal{B}^{(3)}$ is Galois-compatible and has polynomially bounded denominators, the bound on (the standard absolute value of) the values of Φ on $\mathcal{B}^{(3)}$ implies a bound for the heights of these values.

The set of coordinates (with respect to $\mathcal{B}^{(3)}$) of \mathbf{G} -admissible diagonal bases forms an affine algebraic set V_0 defined over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$. Because \mathbf{G} is precisely the stabiliser of Φ , we can use Φ to write down equations for $V_{0, \mathbb{Q}^{\text{alg}}}$, and then describe a suitable $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -form of this algebraic set. The bound for the heights of values of Φ on $\mathcal{B}^{(3)}$ implies that these equations have coefficients of polynomially bounded height, and the existence of a \mathbf{G} -admissible diagonal basis $\mathcal{B}^{(1)}$ implies that $V_0(\mathbb{R})$ is non-empty. We can therefore use Theorem 3.1 to obtain an element of $V_0(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ of polynomially bounded height, and thus $\mathcal{B}^{(4)}$.

Lemma 2.10. *For each $\sigma \in \Sigma$, let R_σ be the matrix in $\mathbf{GL}_{r_\sigma}(\mathbb{C})$ such that*

$$v_{\sigma j}^{(3)} = \sum_{j'=1}^{r_\sigma} R_{\sigma j' j} v_{\sigma j'}^{(1)}.$$

(Such matrices exist because $\mathcal{B}^{(1)}$ and $\mathcal{B}^{(3)}$ are both diagonal bases for E_x .)

The entries of R_σ are polynomially bounded.

Proof. Consider the $r_\sigma \times r_\sigma$ matrices

$$B_{\sigma j j'} = \Theta_x(v_{\sigma j}^{(1)}, v_{\sigma j'}^{(1)})$$

and

$$B'_{\sigma j j'} = \Theta_x(v_{\sigma j}^{(3)}, v_{\sigma j'}^{(3)}).$$

These matrices are Hermitian positive definite, so they can be diagonalised by unitary matrices:

$$B_\sigma = U_\sigma^\dagger D_\sigma U_\sigma \text{ and } B'_\sigma = U_\sigma'^\dagger D'_\sigma U_\sigma'$$

where U_σ, U_σ' are unitary matrices and D_σ, D'_σ are diagonal with positive real entries. We can then choose further diagonal matrices Λ_σ with positive real entries such that

$$\Lambda_\sigma^2 D_\sigma = D'_\sigma.$$

Observe that

$$B_{\sigma jj'} = \Theta_{x_0}(v_{\sigma j}^{(0)}, v_{\sigma j'}^{(0)})$$

and hence does not depend on x . Thus D_σ and U_σ are constant.

The entries of U_σ' are bounded by 1 because it is a unitary matrix. The entries of B'_σ are polynomially bounded by the construction of $\mathcal{B}^{(3)}$. Hence the entries of D'_σ , and therefore also Λ_σ , are polynomially bounded.

Let

$$R'_\sigma = U_\sigma'^\dagger \Lambda_\sigma U_\sigma' \in \mathbf{GL}_{r_\sigma}(\mathbb{C}).$$

This has polynomially bounded entries and satisfies

$$B'_\sigma = R_\sigma'^\dagger B_\sigma R'_\sigma.$$

But by the definition of R ,

$$B'_\sigma = R_\sigma^\dagger B_\sigma R_\sigma.$$

Hence $R_\sigma R_\sigma'^{-1}$ lies in the compact group of B_σ -unitary matrices. Because B_σ is constant, there is a uniform bound for entries of the entries of B_σ -unitary matrices.

Combining this bound for $R_\sigma R_\sigma'^{-1}$ with the polynomial bound for R'_σ , we deduce that R_σ has polynomially bounded entries. \square

Lemma 2.11. *The values*

$$\Phi(v_{\sigma_1 j_1}^{(3)}, \dots, v_{\sigma_k j_k}^{(3)})$$

are algebraic numbers of polynomially bounded height for all

$$((\sigma_1, j_1), \dots, (\sigma_k, j_k)) \in \left(\bigcup_i \Sigma_i \times \{1, \dots, r_i\} \right)^k.$$

Proof. The definition of R_σ implies that we can express the values of Φ on $\mathcal{B}^{(3)}$ in terms of its values on $\mathcal{B}^{(1)}$ as follows:

$$\Phi(v_{\sigma_1 j_1}^{(3)}, \dots, v_{\sigma_k j_k}^{(3)}) = \sum_{j'_1=1}^{r_{\sigma_1}} \cdots \sum_{j'_k=1}^{r_{\sigma_k}} R_{\sigma_1 j'_1 j_1} \cdots R_{\sigma_k j'_k j_k} \Phi(v_{\sigma_1 j'_1}^{(1)}, \dots, v_{\sigma_k j'_k}^{(1)}).$$

The facts that $\mathcal{B}^{(1)}$ is \mathbf{G} -admissible and Φ is \mathbf{G} -invariant imply that the values of Φ on $\mathcal{B}^{(1)}$ are equal to its values on $\mathcal{B}^{(0)}$ and hence constant. Thus Lemma 2.10 implies that the values

$$\Phi(v_{\sigma_1 j_1}^{(3)}, \dots, v_{\sigma_k j_k}^{(3)})$$

are polynomially bounded.

The facts that $\mathcal{B}^{(3)}$ is Galois-compatible and that Φ is defined over \mathbb{Q} imply that

$$\tau(\Phi(v_{\sigma_1 j_1}^{(3)}, \dots, v_{\sigma_k j_k}^{(3)})) = \Phi(v_{(\tau\sigma_1)j_1}^{(3)}, \dots, v_{(\tau\sigma_k)j_k}^{(3)})$$

for all $\tau \in \text{Aut}(\mathbb{C})$. Hence the values of Φ on $\mathcal{B}^{(3)}$ are algebraic numbers.

Because each $\Phi(v_{\sigma_1 j_1}^{(3)}, \dots, v_{\sigma_k j_k}^{(3)})$ is polynomially bounded and because these values are Galois-compatible, all Galois conjugates of $\Phi(v_{\sigma_1 j_1}^{(3)}, \dots, v_{\sigma_k j_k}^{(3)})$ are polynomially bounded. Furthermore the denominators of $\Phi(v_{\sigma_1 j_1}^{(3)}, \dots, v_{\sigma_k j_k}^{(3)})$ are polynomially bounded because $\mathcal{B}^{(3)}$ is weakly bounded. The combination of bounded denominators and all Galois conjugates being bounded implies that the heights are polynomially bounded. \square

Proposition 2.12. *There exist matrices $S_\sigma \in \mathbf{GL}_{r_i}(\mathbb{C})$ whose entries are algebraic numbers of polynomially bounded height and such that the diagonal basis $\mathcal{B}^{(4)} = \{v_{\sigma_j}^{(4)}\}$ defined by*

$$v_{\sigma_j}^{(4)} = \sum_{j'=1}^{r_\sigma} S_{\sigma j' j} v_{\sigma j'}^{(3)}$$

is \mathbf{G} -admissible.

Proof. Let V denote the \mathbb{Q}^{alg} -algebraic set consisting of those elements

$$S := (S_\sigma)_{\sigma \in \Sigma} \in \prod_{\sigma \in \Sigma} \mathbf{GL}_{r_\sigma}$$

which satisfy the equations

$$\sum_{j_1=1}^{r_1} \cdots \sum_{j_k=1}^{r_k} S_{\sigma_1 j'_1 j_1} \cdots S_{\sigma_k j'_k j_k} \Phi(v_{\sigma_1 j'_1}^{(3)}, \dots, v_{\sigma_k j'_k}^{(3)}) = \Phi(v_{\sigma_1 j_1}^{(1)}, \dots, v_{\sigma_k j_k}^{(1)}) \quad (*)$$

for all tuples

$$((\sigma_1, j_1), \dots, (\sigma_k, j_k)) \in \left(\bigcup_i \Sigma_i \times \{1, \dots, r_i\} \right)^k.$$

Observe that the basis

$$\mathcal{B}^{(3)} S = \left\{ \sum_{j'=1}^{r_\sigma} S_{\sigma j' j} v_{\sigma j'}^{(3)} \mid \sigma \in \Sigma, 1 \leq j \leq r_\sigma \right\}$$

is always diagonal. The purpose of the algebraic set V is that the basis $\mathcal{B}^{(3)}S$ is of the form $g\mathcal{B}^{(0)}$ for some $g \in \mathbf{G}(\mathbb{C})$ if and only if $S \in V(\mathbb{C})$.

We will now construct a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -form V_0 of V such that $\mathcal{B}^{(3)}S$ is \mathbf{G} -admissible if and only if $S \in V_0(\mathbb{R})$.

Recall that, when choosing $\mathcal{B}^{(0)}$, we required that for each $v \in \mathcal{B}^{(0)}$, its complex conjugate is also in $\mathcal{B}^{(0)}$. Since $\mathcal{B}^{(1)} = g_1\mathcal{B}^{(0)}$ for some real linear map g_1 , it follows that $\mathcal{B}^{(1)}$ has the same property. Clearly, if F_x acts on v via the character σ , then it acts on \bar{v} via the character $\bar{\sigma}$. Hence for each $\sigma \in \Sigma$, there is a permutation $\beta(\sigma, -)$ of $\{1, \dots, r_\sigma\}$ defined by

$$\bar{v}_{\sigma j}^{(1)} = v_{\bar{\sigma}\beta(\sigma, j)}^{(1)}.$$

(The subscript on the right hand side of that equation might be hard to read: it is $\bar{\sigma}\beta(\sigma, j)$.) On the other hand, because $\mathcal{B}^{(3)}$ is Galois-compatible, we have

$$\bar{v}_{\sigma j}^{(3)} = v_{\bar{\sigma} j}^{(3)}.$$

Define a semilinear involution θ of $\prod_{\sigma \in \Sigma} \mathbf{GL}_{r_\sigma, \mathbb{Q}^{\text{alg}}}$ by

$$\theta(S)_{\sigma j' j} = \bar{S}_{\bar{\sigma} j' \beta(\sigma, j)}.$$

By construction, a tuple of matrices $S \in \prod_{\sigma} \mathbf{GL}_{r_\sigma}(\mathbb{C})$ satisfies $\theta(S) = S$ if and only if the basis $\mathcal{B}^{(3)}S$ is permuted by complex conjugation in the same fashion as $\mathcal{B}^{(1)}$, or in other words, if and only if the linear map transforming $\mathcal{B}^{(1)}$ into this basis is defined over \mathbb{R} . Thus $\mathcal{B}^{(3)}S$ is \mathbf{G} -admissible if and only if $S \in V(\mathbb{C})$ and $\theta(S) = S$.

Because Φ is defined over \mathbb{Q} (so *a fortiori* over \mathbb{R}), composing one of the equations (*) with θ transforms it into another of the equations (*). Hence there is a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -form V_0 of V on which the action of complex conjugation is given by θ . In particular

$$V_0(\mathbb{R}) = \{S \in V(\mathbb{C}) \mid \theta(S) = S\}$$

and thus $\mathcal{B}^{(3)}S$ is \mathbf{G} -admissible if and only if $S \in V_0(\mathbb{R})$.

The $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -variety V_0 is defined by the equations

$$f + \theta(f), \quad i(f - \theta(f))$$

as f runs over the equations (*). The number of equations (*) depends only on the combinatorial data s, r_1, \dots, r_s (for which there are finitely many possibilities), and their degrees are fixed. By Lemma 2.11, their coefficients are algebraic numbers of polynomially bounded height. Hence V_0 is defined by a finite set of polynomials such that the number of polynomials in the set is uniformly bounded, the degree

of each of the polynomials is uniformly bounded and their coefficients are elements of $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ of polynomially bounded height.

Furthermore $V_0(\mathbb{R})$ is non-empty because $R^{-1} \in V_0(\mathbb{R})$ (recall that R was defined in Lemma 2.10 as the tuple of matrices such that $\mathcal{B}^{(3)} = \mathcal{B}^{(1)}R$).

Thus we can apply Theorem 3.1 to V_0 , to deduce that $V_0(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ contains a point S of polynomially bounded height. This is the S whose existence is asserted by the proposition. \square

2.6 Bounding the height of $\mathcal{B}^{(4)}$

The main work in this section is showing that the coordinates of $\mathcal{B}^{(4)}$ (with respect to $\mathcal{B}_{\mathbb{Z}}$) are polynomially bounded in absolute value. We will then (by going to $\mathcal{B}^{(3)}$ and back) deduce that these coordinates have polynomially bounded height.

Notation-wise, we shall consider $E_{\mathbb{C}}^{\vee}$ as a space of row vectors and write $\|\cdot\|$ for the standard Hermitian norm on $E_{\mathbb{C}}^{\vee}$ with respect to the basis dual to $\mathcal{B}_{\mathbb{Z}}$. For any matrix $M \in \mathbf{GL}_n(\mathbb{C})$, we shall write $\lambda_{\min}(M)$ for the minimum of the absolute values of eigenvalues of M .

Let $g_0, g_3, g_4 \in \mathbf{GL}(E_{\mathbb{C}})$ be the linear maps such that

$$\mathcal{B}^{(0)} = g_0\mathcal{B}_{\mathbb{Z}}, \mathcal{B}^{(3)} = g_3\mathcal{B}^{(0)} \text{ and } \mathcal{B}^{(4)} = g_4\mathcal{B}^{(0)}.$$

We will interpret these linear maps as matrices in $\mathbf{GL}_n(\mathbb{C})$ with respect to the basis $\mathcal{B}_{\mathbb{Z}}$. Then the coordinates of $\mathcal{B}^{(3)}$ and $\mathcal{B}^{(4)}$ (with respect to $\mathcal{B}_{\mathbb{Z}}$) are given by the columns of g_3g_0 and g_4g_0 respectively.

Let S' denote the matrix in $\mathbf{GL}(E_{\mathbb{C}})$ which expresses the coordinates of $\mathcal{B}^{(4)}$ with respect to $\mathcal{B}^{(3)}$. In other words,

$$g_4g_0 = g_3g_0S'.$$

By the construction of $\mathcal{B}^{(4)}$, S' can be obtained by taking the block diagonal matrix with blocks S_{σ} from Proposition 2.12, and permuting the rows and columns to match the orderings of the bases $\mathcal{B}^{(3)}$ and $\mathcal{B}^{(4)}$.

Since $\mathcal{B}^{(4)}$ is \mathbf{G} -admissible, we have $g_4 \in \mathbf{G}(\mathbb{R})$. Since $\mathcal{B}^{(4)}$ is diagonal for E_x , we have $g_4x_0 = x$ in the symmetric space X . Since x_4 is in the fundamental set $\mathcal{F} = C.\mathfrak{S}.x_0$, we have $g_4 \in C.\mathfrak{S}$. We can therefore write

$$g_4 = \gamma\nu\alpha\kappa$$

with $\gamma \in C$, $\nu \in \omega$, $\alpha \in A_t$ and $\kappa \in K$.

Lemma 2.13. *For all $e^{\vee} \in E_{\mathbb{Z}}^{\vee} - \{0\}$,*

$$\|e^{\vee}g_4\|^{-1}$$

is polynomially bounded.

Proof. We have

$$\|e^\vee g_4\| = \|e^\vee g_3 g_0 S' g_0^{-1}\| \geq \lambda_{\min}(S' g_0^{-1}) \|e^\vee g_3 g_0\|.$$

By Proposition 2.12, the entries of S' are algebraic numbers with polynomially bounded height, while g_0 is a constant matrix whose entries are algebraic numbers. Hence the entries of $S' g_0^{-1}$ are algebraic numbers with polynomially bounded height, so the same is true of its eigenvalues. Hence $\lambda_{\min}(S' g_0^{-1})^{-1}$ is polynomially bounded.

The row vector $e^\vee g_3 g_0$ is a \mathbb{Z} -linear combination of rows of $g_3 g_0$. Since $\mathcal{B}^{(3)}$ is Galois-compatible and weakly bounded, the coordinates of $e^\vee g_3 g_0$ are a set of algebraic numbers which are closed under Galois conjugation and have polynomially bounded denominators. Hence $\|e^\vee g_3 g_0\|^2$ is a rational number with polynomially bounded denominator.

Furthermore $\|e^\vee g_3 g_0\|^2 \neq 0$ because $g_3 g_0$ is invertible. A non-zero rational number with polynomially bounded denominator has polynomially bounded inverse, so $\|e^\vee g_3 g_0\|^{-1}$ is polynomially bounded. \square

Lemma 2.14. *For $1 \leq j \leq n$,*

$$\|e_j^\vee \nu \alpha\|^{-1}$$

is polynomially bounded.

Proof. By definition, $\nu \alpha = \gamma^{-1} g_4 \kappa^{-1}$.

Since κ is contained in the constant compact set K , for any $v^\vee \in E_{\mathbb{C}}^\vee$, $\|v^\vee \kappa^{-1}\|$ is bounded between constant multiples of $\|v^\vee\|$. Hence it will suffice to prove that $\|e_j^\vee \gamma^{-1} g_4\|^{-1}$ is polynomially bounded.

Since γ lies in the constant finite set $C \subset \mathbf{G}(\mathbb{Q})$, there is a uniform bound for the denominator of γ^{-1} . Thus $e_j^\vee \gamma^{-1}$ lies in a constant rational multiple of $E_{\mathbb{Z}}^\vee$, and so Lemma 2.13 implies that $\|e_j^\vee \gamma^{-1} g_4\|^{-1}$ is polynomially bounded. \square

Lemma 2.15. *There exists a positive integer ℓ depending only on \mathbf{G} such that:*

For any $\nu \in \mathbf{U}(\mathbb{R})\mathbf{M}(\mathbb{R})^+$ and $\alpha \in A_t$, if $\nu_{jj'} \neq 0$, then

$$\alpha_{jj} \geq \min(1, t^\ell) \alpha_{j'j'}.$$

Proof. Let $\nu = um$ with $u \in \mathbf{U}(\mathbb{R})$ and $m \in \mathbf{M}(\mathbb{R})^+$.

Since $\nu_{jj'} \neq 0$, there exists $i \in \{1, \dots, n\}$ such that $u_{ji} m_{ij'} \neq 0$.

Since $\mathbf{U}(\mathbb{R})$ is unipotent, its Lie algebra is the vector subspace of $\mathbf{GL}_n(\mathbb{R})$ spanned by $\{u - 1 \mid u \in \mathbf{U}(\mathbb{R})\}$. In particular, since $u_{ji} \neq 0$, either $i = j$ or there is some $U \in \text{Lie } \mathbf{U}(\mathbb{R})$ such that $U_{ji} \neq 0$.

Suppose that there is $U \in \text{Lie } \mathbf{U}(\mathbb{R})$ such that $U_{ji} \neq 0$. Since the conjugation action of the \mathbb{Q} -split torus \mathbf{S} on \mathbf{U} is diagonalisable over \mathbb{Q} , we may assume that U is an eigenvector for \mathbf{S} , while still satisfying $U_{ji} \neq 0$.

Let φ be the \mathbb{Q} -root of \mathbf{G} specifying how \mathbf{S} acts on U . Because \mathbf{S} is contained in the diagonal torus of \mathbf{GL}_n , for every $s \in \mathbf{S}$,

$$((\text{Ad } s)U)_{ji} = (s_{jj}/s_{ii})U_{ji}$$

and so $\varphi(s) = s_{jj}/s_{ii}$.

By definition, φ is a positive root with respect to \mathbf{U} , so it can be expressed as a sum of at most ℓ positive simple \mathbb{Q} -roots of \mathbf{G} , for some constant ℓ depending only on \mathbf{G} . Since $\alpha \in A_t$, we conclude that $\varphi(\alpha) \geq t^\ell$.

Thus we get $\alpha_{jj}/\alpha_{ii} \geq t^\ell$ whenever $i \neq j$. Combining with the trivial case $i = j$, we get that

$$\alpha_{jj} \geq \min(1, t^\ell) \alpha_{ii}.$$

Since \mathbf{M} centralises \mathbf{S} and \mathbf{S} is diagonal,

$$\alpha_{ii}m_{ij'} = m_{ij'}\alpha_{j'j'}.$$

Since $m_{ij'} \neq 0$, we deduce that $\alpha_{ii} = \alpha_{j'j'}$ so the lemma is proved. \square

Lemma 2.16. *The values α_{jj}^{-1} are polynomially bounded.*

Proof. We have

$$\|e_j^\vee \nu \alpha\|^2 = \sum_{j'=1}^n (\nu_{jj'} \alpha_{j'j'})^2.$$

Using Lemma 2.15, we deduce that

$$\|e_j^\vee \nu \alpha\|^2 \leq \sum_{j'=1}^n \nu_{jj'}^2 \min(1, t^\ell)^{-2} \alpha_{jj}^2.$$

Since ν is in the fixed compact set ω , there is a uniform upper bound for $|\nu_{jj'}|$.

Hence α_{jj} is bounded below by a uniform constant multiple of $\|e_j^\vee \nu \alpha\|$. Combining with Lemma 2.14, we deduce that α_{jj}^{-1} is polynomially bounded. \square

Proposition 2.17. *The coordinates of $\mathcal{B}^{(4)}$ with respect to $\mathcal{B}_{\mathbb{Z}}$ are polynomially bounded in absolute value.*

Proof. Since S' has determinant of polynomially bounded height and $\mathcal{B}^{(3)}$ is weakly bounded, g_4 has polynomially bounded determinant. Now

$$\prod_{j=1}^n \alpha_{jj} = \det \alpha = |\det \gamma|^{-1} |\det g_4|$$

because $\det \nu = |\det \kappa| = 1$. Since γ is in the fixed finite set C , $\prod \alpha_{jj}$ is polynomially bounded.

Combining this bound for $\prod \alpha_{jj}$ with the bound for all α_{jj}^{-1} from Lemma 2.16, we deduce that each α_{jj} is polynomially bounded.

Since each of γ , ν and κ is contained in a fixed finite set, this implies that the entries of $g_4 = \gamma\nu\alpha\kappa$ are polynomially bounded.

Since the coordinates of $\mathcal{B}^{(4)}$ are given by g_4g_0 , they are also polynomially bounded. \square

Proposition 2.18. *The coordinates of $\mathcal{B}^{(4)}$ with respect to $\mathcal{B}_{\mathbb{Z}}$ are algebraic numbers of polynomially bounded height.*

Proof. Since S' gives the coordinates of $\mathcal{B}^{(4)}$ with respect to $\mathcal{B}^{(3)}$ and has entries of polynomially bounded height, Proposition 2.17 implies that the coordinates of $\mathcal{B}^{(3)}$ with respect to $\mathcal{B}_{\mathbb{Z}}$ are polynomially bounded in absolute value.

Since $\mathcal{B}^{(3)}$ is Galois-compatible, this tells us that each coordinate of a vector in $\mathcal{B}^{(3)}$ in fact has all its Galois conjugates polynomially bounded. Furthermore its denominator is polynomially bounded, because $\mathcal{B}^{(3)}$ is weakly bounded. Hence the heights of the coordinates of $\mathcal{B}^{(3)}$ are polynomially bounded.

Again using the fact that S' has polynomially bounded height, we conclude that the heights of the coordinates of $\mathcal{B}^{(4)}$ are polynomially bounded. \square

Proof of Theorem 1.1. Proposition 2.18 implies that the entries of g_4 are algebraic numbers of polynomially bounded height. Since $x = g_4x_0$ and the action of $G(\mathbb{R})$ on X is semialgebraic over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$, it follows that the coordinates of $x \in X$ also are algebraic with polynomially bounded height. \square

3 Points of small height on real affine varieties

Let V be an affine algebraic set defined over \mathbb{Q}^{alg} . Suppose that V can be defined by a fixed number of polynomials in a fixed number of variables and of fixed degrees, such that the coefficients of these polynomials have height at most H . Then the arithmetic Bézout theorem and Zhang's theorem on the essential minimum imply that V has a \mathbb{Q}^{alg} -point whose height is bounded by a polynomial in H .

For our application to pre-special points, we need a variant of this result: we assume that V is defined over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ and we need to show that $V(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ contains a point of polynomially bounded height. (For us, \mathbb{Q}^{alg} means the set of algebraic numbers in \mathbb{C} rather than an abstract algebraic closure of \mathbb{Q} , so $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ is well-defined.) To avoid trivial counter-examples, we must also assume that $V(\mathbb{R})$ is non-empty.

Theorem 3.1. *For all positive integers m , n , D , there are constants C_{13} and C_{14} depending on m , n and D such that:*

For every affine algebraic set $V \subset \mathbb{A}^n$ defined over $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ by polynomials $f_1, \dots, f_m \in (\mathbb{Q}^{\text{alg}} \cap \mathbb{R})[X_1, \dots, X_n]$ of degree at most D and height at most H , if $V(\mathbb{R})$ is non-empty, then $V(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ contains a point of height at most $C_{13}H^{C_{14}}$.

We will prove the theorem by an elementary argument, using a quantitative version of the Noether normalisation lemma. The proof is by induction on n .

Given an algebraic set in \mathbb{A}^n , the following lemma constructs an algebraic set in \mathbb{A}^{n-1} satisfying appropriate bounds, which we can then study inductively. Part (i) of the lemma is simply the inductive step of the Noether normalisation lemma, without any height bound. Parts (ii) to (iv) allow us to find a point of small height in V , once we have got such a point in $\pi\phi(V)$. Part (v) ensures that as we carry out the induction, replacing $V \subset \mathbb{A}^n$ by $\pi\phi(V) \subset \mathbb{A}^{n-1}$, the number, degrees and heights of the defining polynomials remain controlled. (Parts (i) to (iii) work when k is any field of characteristic zero. For (iv) and (v) the field must be contained in \mathbb{Q}^{alg} so that we have a notion of height.)

Lemma 3.2. *Let $V \subset \mathbb{A}^n$ be an affine algebraic set defined over a field $k \subset \mathbb{Q}^{\text{alg}}$. Suppose that V can be defined by m polynomials with coefficients in k , each of degree at most D and height at most H .*

If $V \neq \mathbb{A}^n$, then there exists a linear automorphism ϕ of k^n such that:

- (i) $(\pi\phi)|_V: V \rightarrow \mathbb{A}^{n-1}$ is a finite morphism (where $\pi: \mathbb{A}^n \rightarrow \mathbb{A}^{n-1}$ means projection onto the first $n-1$ coordinates);
- (ii) the matrix representing ϕ w.r.t. the standard basis has entries in \mathbb{Z} with absolute values at most D ;
- (iii) there exists a polynomial $g \in k[X_1, \dots, X_{n-1}][X_n]$ which vanishes on $\phi(V)$ and which is monic of degree at most D in X_n (we do not require that $\phi(V) = \{x \mid \pi(x) \in \pi\phi(V) \text{ and } g(x) = 0\}$);
- (iv) the height of g is bounded by a polynomial in H ;
- (v) the image $\pi\phi(V)$ can be defined by polynomials in $k[X_1, \dots, X_{n-1}]$ such that the number and degree of these polynomials are bounded by bounds depending on m and D only, and their heights are bounded by a polynomial in H .

Here “bounded by a polynomial in H ” means bounded by an expression of the form $C_{21}H^{C_{22}}$ where C_{21} and C_{22} are constants which depend on n , m and D only.

Using Lemma 3.2, we can directly prove the simplified version of Theorem 3.1 which only concerns \mathbb{Q}^{alg} -points. Indeed, by induction on n we get a point

$$x = (x_1, \dots, x_{n-1}) \in \pi\phi(V)(\mathbb{Q}^{\text{alg}})$$

of polynomially bounded height. Parts (iii) and (iv) of the lemma imply that any point in the fibre $\pi^{-1}(x) \cap \phi(V)(\mathbb{Q}^{\text{alg}})$ also has polynomially bounded height, because its x_n -coordinate must be a root of $g(x_1, \dots, x_{n-1})(X_n)$. Then part (ii) allows us to transform this into a point of polynomially bounded height in V .

This argument is not sufficient to obtain a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -point of small height on V . This is because the fibre $\pi^{-1}(x) \cap \phi(V)$ at our chosen x is not guaranteed to contain any real points.

In order to deal with this, we will strengthen the inductive hypothesis slightly as follows:

Theorem 3.3. *In the setting of Theorem 3.1, each connected component Z of $V(\mathbb{R})$ (in the real topology) contains a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -point of height at most $C_{13}H^{C_{14}}$.*

We then consider the boundary of $\pi\phi(Z)$, viewed as a subset of $\pi\phi(V)(\mathbb{R})$. If $\pi\phi(Z)$ is an entire connected component of $\pi\phi(V)(\mathbb{R})$, then by the inductive hypothesis it contains a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -point of polynomially bounded height and we can apply the same argument as above to get a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -point of small height in Z .

So the interesting case is when $\pi\phi(Z)$ has non-empty boundary in $\pi\phi(V)(\mathbb{R})$. At a point (x_1, \dots, x_{n-1}) on the boundary of $\pi\phi(Z)$, an analytic argument implies that either $g(x_1, \dots, x_{n-1})(X_n)$ has a double root, or $\pi\phi(V)$ is singular. In either case we can find an algebraic subset $V' \subset V$, of smaller dimension and defined by polynomials satisfying the appropriate bounds on number, degree and height, such that $Z \cap V'(\mathbb{R})$ is non-empty, so we can replace V by V' and repeat the argument.

If $g(x_1, \dots, x_{n-1})(X_n)$ has a double root at some point of $\pi\phi(Z)$, then it is straightforward to construct V' . In order to construct V' when $\pi\phi(V)$ has a singular point, we use Proposition 3.11, which we prove using the Chow form and an idea of Bombieri, Masser and Zannier.

3.1 Definitions

Throughout this section, k denotes a field of characteristic 0.

We will denote by π the linear map $\mathbb{A}^n \rightarrow \mathbb{A}^{n-1}$ given by projection onto the first $n-1$ coordinates.

We identify an algebraic set over k with its set of k^{alg} -points. When we say that an algebraic set V is defined by polynomials f_1, \dots, f_m , we mean simply that

$$V(k^{\text{alg}}) = \{(x_1, \dots, x_n) \in \mathbb{A}^n(k^{\text{alg}}) \mid f_i(x_1, \dots, x_n) = 0 \text{ for all } i\}.$$

In particular, we do not require the ideal $(f_1, \dots, f_m) \subset k[X_1, \dots, X_m]$ to be a radical ideal.

Symbols C_i will denote constants; the first time we introduce each constant we will specify the parameters on which it depends. Usually C_i depends on degrees but not heights.

We define the **height** of a point $(x_1, \dots, x_n) \in \mathbb{A}^n(\mathbb{Q}^{\text{alg}})$ to be the multiplicative absolute Weil height of $[x_1 : \dots : x_n : 1] \in \mathbb{P}^n(\mathbb{Q}^{\text{alg}})$, that is

$$H(x_1, \dots, x_n) = \prod_v \max(|x_1|_v, \dots, |x_n|_v, 1)^{d_v/[K:\mathbb{Q}]}$$

where K is any number field containing x_1, \dots, x_n , the product is over all places of K and $d_v = [K_v : \mathbb{Q}_p]$ or $[K_v : \mathbb{R}]$ as appropriate, with the absolute values normalised so that they agree with the standard absolute value when restricted to \mathbb{Q}_p or \mathbb{R} .

The **height** of a polynomial in $\mathbb{Q}^{\text{alg}}[X_1, \dots, X_n]$ means the multiplicative absolute Weil height of the point in projective space whose homogeneous coordinates are the coefficients of the polynomial.

Lemma 3.4. *If $f \in \mathbb{Q}^{\text{alg}}[X]$ is a monic polynomial in one variable of degree d and x is a root of f , then*

$$H(x) \leq dH(f).$$

Proof. Let

$$f(X) = X^d + a_{d-1}X^{d-1} + \dots + a_0.$$

and let K be a number field containing a_0, \dots, a_{d-1}, x .

For each absolute value $|\cdot|$ of K ,

$$\begin{aligned} \max(|x|, 1)^d &= \max(|a_{d-1}x^{d-1} + \dots + a_0|, 1) \\ &\leq d \max(|a_{d-1}x^{d-1}|, \dots, |a_0|, 1) \\ &\leq d \max(|a_{d-1}|, \dots, |a_0|, 1) \max(|x|, 1)^{d-1} \end{aligned}$$

so that

$$\max(|x|, 1) \leq d \max(|a_{d-1}|, \dots, |a_0|, 1)$$

and the factor of d (coming from the triangle inequality) can be replaced by 1 for non-archimedean absolute values.

Taking the product over all places of K gives the result. \square

3.2 Bounded normalisation lemma

We will now prove parts (i) to (iv) of Lemma 3.2. Our proof is based on the proof of the Noether normalisation lemma in exercise 5.16 of [AM69].

Let us remark that the bound in part (ii) is stronger than most of our bounds: the height of ϕ is bounded only in terms of the degrees of the defining equations for V , independent of the height of these equations. We do not need this extra strength. In order to obtain this bound in part (ii), we will use the following lemma.

Lemma 3.5. *Let $f \in k[X_1, \dots, X_n]$ be a non-zero polynomial of degree at most D (over any field k of characteristic zero). There exist integers $\lambda_1, \dots, \lambda_n$, such that $|\lambda_i| \leq D$ for all i , $\lambda_i \neq 0$ for all i and*

$$f(\lambda_1, \dots, \lambda_n) \neq 0.$$

Proof. The proof is by induction on n . Write

$$f(X_1, \dots, X_n) = \sum_{i=0}^D g_i(X_1, \dots, X_{n-1})X_n^i.$$

The g_i are polynomials in $k[X_1, \dots, X_{n-1}]$ of degree at most D .

Choose some i such that $g_i \neq 0$. By induction, there are non-zero integers $\lambda_1, \dots, \lambda_{n-1}$, with absolute value at most D , such that

$$g_i(\lambda_1, \dots, \lambda_{n-1}) \neq 0.$$

Then $f(\lambda_1, \dots, \lambda_{n-1}, X_n)$ is a non-zero polynomial in one variable X_n . It has degree at most D , so it has at most D roots.

Hence the set of $D + 1$ values $\{-1, 1, 2, \dots, D\}$ contains at least one non-root of $f(\lambda_1, \dots, \lambda_{n-1}, X_n)$, as required. \square

Proof of Lemma 3.2 (i)–(iv). Let $f = f_1$. This is a polynomial in $k[X_1, \dots, X_n]$ of degree at most D which vanishes on V , and without loss of generality $f \neq 0$.

Let F denote the homogeneous polynomial consisting of the terms of f with highest total degree. By Lemma 3.5, we can find non-zero integers $\lambda_1, \dots, \lambda_n$, with absolute value at most D , such that

$$F(\lambda_1, \dots, \lambda_n) \neq 0.$$

Let $\phi: \mathbb{A}^n \rightarrow \mathbb{A}^n$ be the k -linear map with matrix

$$\begin{pmatrix} \lambda_n & & & -\lambda_1 \\ & \ddots & & \vdots \\ & & \lambda_n & -\lambda_{n-1} \\ & & & 1 \end{pmatrix}$$

This is invertible because $\lambda_n \neq 0$. It satisfies (ii) by construction.

Let $g \in k[X_1, \dots, X_n]$ be the polynomial

$$g(X_1, \dots, X_n) = f(\phi^{-1}(X_1, \dots, X_n)).$$

Observe that g vanishes on $\phi(V)$.

Let G be the homogeneous polynomial consisting of the terms of g with highest total degree. Then

$$G(0, \dots, 0, \lambda_n) = F(\phi^{-1}(0, \dots, 0, \lambda_n)) = F(\lambda_1, \dots, \lambda_n) \neq 0.$$

Since G is homogeneous and $\lambda_n \neq 0$, we deduce that

$$G(0, \dots, 0, 1) \neq 0.$$

Let d denote the total degree of G . Then $G(0, \dots, 0, 1)$ is the coefficient of the X_n^d term of g , and all other terms in g have a lower power of X_n . Hence after scaling by $G(0, \dots, 0, 1)^{-1}$, $g(X_1, \dots, X_n)$ is monic of degree $d \leq D$ as a polynomial in the one variable X_n with coefficients in $k[X_1, \dots, X_{n-1}]$. Thus (iii) is satisfied.

The coordinate ring $k[\phi(V)]$ is generated as a $k[X_1, \dots, X_{n-1}]$ -algebra by X_n . Furthermore, since g is monic in X_n and vanishes on $\phi(V)$, we deduce that X_n (viewed as an element of $k[\phi(V)]$) is integral over $k[X_1, \dots, X_{n-1}]$. Hence $k[\phi(V)]$ is finite as a $k[X_1, \dots, X_{n-1}]$ -algebra. In other words, the morphism $\pi|_{\phi(V)}$ is finite. Since ϕ is invertible, (i) is satisfied.

Since f , ϕ and $G(0, \dots, 0, 1)$ all have height bounded by a polynomial in H , the same is true of g , so (iv) holds. \square

3.3 Bounding the image of normalisation

We wish to prove part (v) of Lemma 3.2, i.e. that the image $\pi\phi(V)$ has polynomially bounded height-complexity. Note first that $\pi\phi(V)$ is indeed Zariski-closed because $\pi\phi|_V$ is finite.

Take the polynomials $f_1 \circ \phi^{-1}, \dots, f_m \circ \phi^{-1}$ which define $V' = \phi(V)$, and label them $g_1, \dots, g_r, h_1, \dots, h_s$ such that the g_i all contain X_n while the h_i do not contain X_n . A point $(x_1, \dots, x_{n-1}) \in \mathbb{A}^{n-1}(k^{\text{alg}})$ is in $\pi\phi(V)$ if and only if all the h_i vanish at (x_1, \dots, x_{n-1}) and all the one-variable polynomials $g_i(x_1, \dots, x_{n-1})(X_n)$ have a common root x_n . So the problem is to write down algebraic conditions on the coefficients which recognise when a set of one-variable polynomials with varying coefficients has a common root.

Given two one-variable polynomials, we can recognise when they have a common root using the resultant. More precisely: for each pair of integers (d, e) , there is a universal polynomial $\text{Res}_{d,e} \in \mathbb{Z}[a_0, \dots, a_d, b_0, \dots, b_e]$ such that for all one-variable polynomials u of degree d and v of degree e , $\text{Res}_{d,e}$ vanishes on the coefficients of u and v if and only if u and v have a common root in k^{alg} .

We must take care about what happens when we plug polynomials into $\text{Res}_{d,e}$ whose degrees are smaller than the expected degrees d and e . It turns out that, if both $\deg u < d$ and $\deg v < e$, then $\text{Res}_{d,e}(u, v)$ is always zero. If $\deg u = d$ and $\deg v < e$ (or vice versa), then as usual, $\text{Res}_{d,e}(u, v) = 0$ if and only if u and v have a common root. Essentially, in the case where both $\deg u < d$ and $\deg v < e$, the fact that $\text{Res}_{d,e}(u, v) = 0$ is telling us that the degree- d homogenisation of u and the degree- e homogenisation of v have a common root at $\infty \in \mathbb{P}^1(k^{\text{alg}})$.

To determine whether a set of more than two polynomials have a common root, we just have to look at the resultants of sufficiently many linear combinations of them.

Lemma 3.6. *Let $g_1, \dots, g_r \in k[X]$ be polynomials in one variable, with degree at most D . Let $d = \deg g_1$.*

The polynomials g_1, \dots, g_r share a common root in k^{alg} if and only if

$$\text{Res}_{d,D}(g_1, \lambda_2 g_2 + \lambda_3 g_3 + \dots + \lambda_{r-1} g_{r-1} + g_r) = 0$$

for all $(\lambda_2, \dots, \lambda_{r-1}) \in \{0, \dots, d\}^{r-2}$.

Proof. The “only if” direction is clear: if g_1, \dots, g_r all share a common root, then every linear combination of them also vanishes at that point.

For the “if” direction, the proof is by induction on r . For $r = 2$, this is just the definition of the resultant. We know that g_1 has exactly degree d , so even if $\deg g_2 < D$ the resultant still tells us whether they have a common root.

For $r > 2$: Suppose that the given resultants are all zero. By induction, for every $\lambda \in \{0, \dots, d\}$, the polynomials

$$g_1, g_2, \dots, g_{r-2}, \lambda g_{r-1} + g_r$$

share a common root α_λ . Since each α_λ is a root of g_1 , and g_1 has at most d roots, there must be two distinct values $\lambda, \mu \in \{0, \dots, d\}$ such that

$$\alpha_\lambda = \alpha_\mu.$$

Then by construction

$$\lambda g_{r-1}(\alpha_\lambda) + g_r(\alpha_\lambda) = \mu g_{r-1}(\alpha_\lambda) + g_r(\alpha_\lambda) = 0$$

which implies that

$$g_{r-1}(\alpha_\lambda) = g_r(\alpha_\lambda) = 0.$$

Thus α_λ is the required common root of g_1, \dots, g_r . □

The following lemma establishes part (v) of Lemma 3.2. For defining equations of $V' = \phi(V)$, we take $f_1 \circ \phi^{-1}, \dots, f_m \circ \phi^{-1}$ as well as g from part (iii) of Lemma 3.2. Since g vanishes on V' by construction, adding it to the set of defining polynomials does not change V' ; g ensures that the condition that one of the polynomials should be monic in X_n is satisfied.

Lemma 3.7. *For all positive integers m, n, D , there are constants C_{23}, C_{24} and C_{25} depending on m, n and D such that:*

For every affine algebraic set $V' \subset \mathbb{A}^n$ defined over $k \subset \mathbb{Q}^{\text{alg}}$ by at most m polynomials of degree at most D and height at most H , such that one of these polynomials is monic in X_n , the image under projection

$$\pi(V') \subset \mathbb{A}^{n-1}$$

is defined by at most C_{23} polynomials over k , each of degree at most C_{23} and height at most $C_{24}H^{C_{25}}$.

Proof. Take the polynomials defining V' , and label them

$$g_1, \dots, g_r, h_1, \dots, h_s$$

such that the g_i all contain X_n while the h_i do not contain X_n . Choose the labelling such that g_1 is monic in X_n .

A point $(x_1, \dots, x_{n-1}) \in \mathbb{A}^{n-1}(k^{\text{alg}})$ is in $\pi(V')$ if and only if all the h_i vanish at (x_1, \dots, x_{n-1}) and the one-variable polynomials $g_i(x_1, \dots, x_{n-1})(X_n)$ have a common root.

Let d be the degree of g_1 with respect to X_n . Because g_1 is monic in X_n , d is equal to $\deg g_1(x_1, \dots, x_{n-1})(X_n)$ for every $(x_1, \dots, x_{n-1}) \in \mathbb{A}^{n-1}(k^{\text{alg}})$.

Hence we can apply Lemma 3.6 at each point of $\mathbb{A}^{n-1}(k^{\text{alg}})$, the polynomials $g_i(x_1, \dots, x_{n-1})(X_n)$ have a common root if and only if

$$\text{Res}_{d,D}(g_1, \lambda_2 g_2 + \lambda_3 g_3 + \dots + \lambda_{r-1} g_{r-1} + g_r) = 0$$

at (x_1, \dots, x_{n-1}) for all $(\lambda_2, \dots, \lambda_{r-1}) \in \{0, \dots, d\}^{r-2}$.

These resultants give us a set of $(d+1)^{r-2}$ polynomials in $k[X_1, \dots, X_{n-1}]$ which, together with h_1, \dots, h_s , define $\pi(V')$. Because these polynomials are constructed as universal polynomials in the coefficients of g_1, \dots, g_r , their degrees are bounded by a constant depending on n and D and their heights by a polynomial in H . \square

We have completed the proof of Lemma 3.2 and, as remarked below the statement of Lemma 3.2, it is straightforward to deduce a version of Theorem 3.1 which only concerns \mathbb{Q}^{alg} -points.

3.4 Behaviour of real points under projection

In order to obtain $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -points of small height, and not just \mathbb{Q}^{alg} -points, we look at the image of a single connected component of $V(\mathbb{R})$ under $\pi\phi$. As explained in the introduction to the section the difficult case is when this image is strictly contained in a connected component of $\pi\phi(V)(\mathbb{R})$.

The principle we use to deal with this is that, in a smoothly varying family of monic one-variable polynomials such as $g(x_1, \dots, x_{n-1})(X_n)$, real roots cannot disappear spontaneously: whenever a real root disappears, there must be either a repeated root or a singular point in the base. The condition that the polynomials are monic ensures that a root cannot disappear off to infinity (such as occurs in $x_1 X_2 - 1 = 0$ when $x_1 \rightarrow 0$).

The set of points of $\pi\phi(V)$ at which $g(x_1, \dots, x_{n-1})(X_n)$ has a repeated root is itself an algebraic set, which is (usually) of smaller dimension than $\pi\phi(V)$ itself, defined by adding $\partial g / \partial X_n$ to the set of polynomials defining $\pi\phi(V)$. Thus clearly it is still defined by polynomials of bounded height, and we can apply the induction hypothesis to this set.

The singular locus of $\pi\phi(V)$ is likewise an algebraic set of smaller dimension. Controlling the height of polynomials defining the singular locus is more difficult. Indeed, instead of using the singular locus directly, we use Proposition 3.11 to obtain an algebraic set which sits in between $\pi\phi(V)$ and its singular locus, and whose height we can control.

It is convenient to restrict attention to irreducible algebraic sets in the following lemma. Lemmas 3.9 and 3.10 ensure that we can take irreducible components of V while retaining control of the heights of defining polynomials.

Proposition 3.8. *Let $W \subset \mathbb{A}^n$ be an irreducible affine algebraic set defined over \mathbb{R} such that $\pi|_W$ is finite. Let $g \in \mathbb{R}[X_1, \dots, X_{n-1}][X_n]$ be a polynomial which is monic in X_n and which vanishes on W .*

Let Z be a connected component of $W(\mathbb{R})$ (in the real topology).

Then at least one of the following holds:

- (i) $\pi(Z)$ is a connected component of $\pi(W)(\mathbb{R})$;
- (ii) there is a point $(x_1, \dots, x_n) \in Z$ such that x_n is a repeated root of the one-variable polynomial $g(x_1, \dots, x_{n-1})(X_n)$;
- (iii) $\pi(Z)$ contains a point at which $\pi(W)$ is singular.

Proof. Assume for contradiction that none of the statements of the lemma hold.

Since $\pi|_W$ is proper and Z is a closed subset of $W(\mathbb{R})$, $\pi(Z)$ is a closed subset of $\pi(W)(\mathbb{R})$. Furthermore $\pi(Z)$ is connected because Z is connected. Because (i) does not hold, $\pi(Z)$ is not open in $\pi(W)(\mathbb{R})$. Hence there exists some point

(x_1, \dots, x_{n-1}) which is contained in the intersection of $\pi(Z)$ with the closure of $\pi(W)(\mathbb{R}) - \pi(Z)$.

Consider the algebraic set

$$\Sigma = \{(z_1, \dots, z_n) \in \mathbb{A}^n \mid (z_1, \dots, z_{n-1}) \in \pi(W) \text{ and } g(z_1, \dots, z_n) = 0\}.$$

W is an irreducible Zariski closed subset of Σ of the same dimension, so is an irreducible component of Σ .

Since (iii) does not hold, $\pi(W)$ is smooth at (x_1, \dots, x_{n-1}) . Hence there is an open neighbourhood U of (x_1, \dots, x_{n-1}) in $\pi(W)(\mathbb{C})$ which is isomorphic (as a complex manifold) to a complex ball.

Since $(x_1, \dots, x_{n-1}) \in \pi(Z)$ we can find some $x_n \in \mathbb{R}$ such that

$$(x_1, \dots, x_n) \in Z.$$

Since (ii) does not hold, x_n is a simple root of $g(x_1, \dots, x_{n-1})(X_n)$. Thus the partial derivative $\partial g / \partial X_n$ is non-zero at (x_1, \dots, x_n) . We can therefore apply the implicit function theorem to get a holomorphic function $\alpha: U \rightarrow \Sigma(\mathbb{C})$ which induces a bijection between U and some open neighbourhood \tilde{U} of (x_1, \dots, x_n) in $\Sigma(\mathbb{C})$, such that α is inverse to $\pi|_{\tilde{U}}$.

It follows that Σ is smooth at (x_1, \dots, x_n) and so (x_1, \dots, x_n) is contained in only one irreducible component of Σ , namely W . Thus $\tilde{U} \subset W(\mathbb{C})$.

Furthermore, α is the only continuously differentiable function which induces a bijection $U \rightarrow \tilde{U}$ with inverse π . But since Σ and (x_1, \dots, x_n) are defined over \mathbb{R} , the continuously differentiable function $\bar{\alpha} = c \circ \alpha \circ c$ (where c denotes complex conjugation) also induces a bijection $U \rightarrow \tilde{U}$ with inverse π . So the uniqueness of α implies that

$$\alpha = \bar{\alpha}.$$

Hence α maps $U \cap \pi(W)(\mathbb{R})$ into $W(\mathbb{R})$, and indeed induces a homeomorphism

$$U \cap \pi(W)(\mathbb{R}) \rightarrow \tilde{U} \cap W(\mathbb{R}).$$

Since $W(\mathbb{R})$ is a real algebraic set, it is locally connected and Z is open in $W(\mathbb{R})$. Hence $\tilde{U} \cap Z$ is open in $W(\mathbb{R})$. Thanks to the above homeomorphism, $\pi(\tilde{U} \cap Z)$ is open in $\pi(W)(\mathbb{R})$.

Hence $\pi(\tilde{U} \cap Z)$ is an open neighbourhood of (x_1, \dots, x_{n-1}) in $\pi(W)(\mathbb{R})$ which is contained in $\pi(Z)$. But the existence of such a neighbourhood contradicts the fact that (x_1, \dots, x_{n-1}) is in the closure of $\pi(W)(\mathbb{R}) - \pi(Z)$. \square

We are now ready to finish the proof of Theorem 3.3, and hence that of Theorem 3.1.

Proof of Theorem 3.3. We prove Theorem 3.3 by induction on $n + \dim V$.

Choose a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -irreducible component V_1 of V whose real points intersect Z .

There are two cases, depending on whether V_1 is geometrically irreducible or not. First consider the case in which it is not geometrically irreducible. Then $V_1 = V_2 \cup \overline{V_2}$ for some geometrically irreducible algebraic set V_2 defined over \mathbb{Q}^{alg} , and $V_1(\mathbb{R})$ is contained in the intersection $V_2 \cap \overline{V_2}$. By Lemmas 3.9 and 3.10, V_2 can be defined by polynomials P_1, \dots, P_t with coefficients in \mathbb{Q}^{alg} whose number and degree are bounded and whose heights are polynomially bounded. Then $V_2 \cap \overline{V_2}$ is defined by the real and imaginary parts

$$\operatorname{Re} P_1, \operatorname{Im} P_1, \dots, \operatorname{Re} P_t, \operatorname{Im} P_t \in (\mathbb{Q}^{\text{alg}} \cap \mathbb{R})[X_1, \dots, X_n]$$

which again are bounded in number and degree and polynomially bounded in height. Since $\dim(V_2 \cap \overline{V_2}) < \dim V$, we can apply the induction hypothesis to $V_2 \cap \overline{V_2}$ and a connected component of $(V_2 \cap \overline{V_2})(\mathbb{R}) \cap Z$, proving the theorem in this case.

Otherwise V_1 is geometrically irreducible. According to Lemmas 3.9 and 3.10, V_1 can be defined by polynomials with coefficients in $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ whose number and degree are bounded and whose heights are polynomially bounded, so we can replace V by V_1 and Z by a connected component of $V_1(\mathbb{R}) \cap Z$.

Choose ϕ and g as in Lemma 3.2 (with $k = \mathbb{Q}^{\text{alg}} \cap \mathbb{R}$) and let $W = \phi(V)$. We will need to assume that $\partial g / \partial X_n$ does not vanish identically on W : to achieve this, simply replace g by $\partial^r g / \partial X_n^r$ for the largest $r < \deg g$ such that this vanishes identically on W (scaled to remain monic).

If we are in case (i) of Proposition 3.8, then $\pi\phi(Z)$ is equal to a connected component Z' of $\pi(W)(\mathbb{R})$. By part (v) of Lemma 3.2, $\pi(W)$ is defined by polynomials over $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ which are bounded in number and degree and polynomially bounded in height. Hence by induction Z' contains a $\mathbb{Q}^{\text{alg}} \cap \mathbb{R}$ -point (x_1, \dots, x_{n-1}) of polynomially-bounded height. Choose x_n such that $(x_1, \dots, x_n) \in Z$; then x_n must also have bounded height because it is a root of $g(x_1, \dots, x_{n-1})(X_n)$ and the theorem is proved for this case.

If we are in case (ii) of Proposition 3.8, let

$$W' = \{x \in W \mid (\partial g / \partial X_n)(x) = 0\}.$$

Because we have assumed that $\partial g / \partial X_n$ does not vanish identically on W , and because W is irreducible, W' is a subset of W of smaller dimension. Clearly W' is defined by equations of polynomially bounded height. Because we are in case (ii) of Proposition 3.8, $W'(\mathbb{R}) \cap Z$ is non-empty. Hence by induction $W'(\mathbb{R}) \cap Z$ contains a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -point of small height.

If we are in case (iii) of Proposition 3.8, choose $W' \subset W$ as in Proposition 3.11. By construction, W' is defined by polynomials of bounded height and has smaller

dimension than W . By case (iii) of Proposition 3.8, $W'(\mathbb{R}) \cap Z$ is non-empty. Hence by induction $W'(\mathbb{R}) \cap Z$ contains a $(\mathbb{Q}^{\text{alg}} \cap \mathbb{R})$ -point of small height. \square

3.5 Height bounds for varieties and the singular locus

We will use Chow polynomials, together with an idea of Bombieri, Masser and Zannier, to prove that the singular locus of an irreducible affine algebraic variety is contained in an algebraic subset of positive codimension and which is defined by equations of bounded height. Along the way, we also get for free a bound for the height of polynomials defining an irreducible component of an affine algebraic set, a necessary reduction in the proof of Theorem 3.1.

Let W be an irreducible algebraic variety defined over \mathbb{Q}^{alg} . We define a certain polynomial associated with W , the **Chow form** of W , and define the **height** $H(W)$ of the variety W to be the height of its Chow form. (Because we are working with multiplicative heights, $H(W)$ is $\exp h(W)$ where $h(W)$ is the height defined by Philippon.) From the Chow form, we can obtain a special set of polynomials defining W , which we call the **Chow polynomials** of W .

Our first two lemmas show that, given any set of equations defining W , the height of the Chow form is polynomially bounded with respect to the height of these equations, and in the other direction, the heights of the Chow polynomials are polynomially bounded with respect to the height of the Chow form. Combining these two lemmas gives us a bound for the heights of equations defining an irreducible component of an affine algebraic set. Note that some care is required when we consider irreducible components and when we allow arbitrary affine algebraic sets because Chow forms are only defined for irreducible varieties, and the arithmetic Bézout theorem only talks about geometrically irreducible components.

The Chow polynomials do not necessarily generate a radical ideal, and so the algebraic set W' defined by applying the Jacobian criterion to the Chow polynomials may be bigger than $\text{Sing } W$. Following an idea of Bombieri, Masser and Zannier, we show that this W' is strictly contained in W , and this is enough for us to deal with case (iii) of Proposition 3.8.

We believe that it should be possible to show that $\text{Sing } W$ itself can be defined by equations of bounded number and degree and polynomially bounded height by using the theory of Gröbner bases.

Philippon and Nesterenko work with projective rather than affine algebraic sets, but we can turn our affine sets into projective ones simply by embedding $\mathbb{A}^n \hookrightarrow \mathbb{P}^n$ and replacing W by its Zariski closure (i.e. homogenising the defining polynomials).

Lemma 3.9. *For all positive integers m, n, D , there are constants C_{26}, C_{27} and C_{28} depending on m, n and D such that:*

For every affine algebraic set $V \subset \mathbb{A}^n$ defined over \mathbb{Q}^{alg} by at most m polynomials of degree at most D and height at most H , each irreducible component of V has degree at most C_{26} and height at most $C_{27}H^{C_{28}}$.

Proof. Note that V is equal to the intersection $V(f_1) \cap \cdots \cap V(f_m)$, where $V(f_i)$ is the variety defined by

$$V(f_i) = \{(x_1, \dots, x_n) \in \mathbb{A}^n \mid f_i(x_1, \dots, x_n) = 0\}.$$

For an irreducible component W of V we may choose an irreducible factor g_i of each f_i such that W is contained in $V(g_1) \cap \cdots \cap V(g_m)$.

By [Hab08] §2, there exists a constant $C_{29}(D)$ such that the height of $V(g_i)$ is at most $C_{29}H$. Of course, the degree of $V(g_i)$ is at most D . Therefore, by the arithmetic Bézout theorem (see [Hab08] Theorem 3), there exist constants $C_{30}(D)$ and $C_{31}(D)$ such that the multiplicative height of any irreducible component of the intersection of any two of the $V(g_i)$ is at most $C_{30}H^{C_{31}}$. Similarly, the standard Bézout theorem (see [RU09] Lemme 3.4) implies that there exists a constant $C_{32}(D)$ which bounds the degree of such a component.

Reiterating this procedure we eventually reach W itself (after at most m steps). \square

Lemma 3.10. *For all positive integers n and D , there are constants C_{33} , C_{34} and C_{35} depending on n and D such that:*

For every irreducible affine variety $W \subset \mathbb{A}^n$ defined over a field $k \subset \mathbb{Q}^{\text{alg}}$ such that W has degree at most D and height at most H , the Chow polynomials of W are a set of at most C_{33} polynomials in $k[X_1, \dots, X_n]$ defining W , each of degree at most C_{33} and height at most $C_{34}H^{C_{35}}$.

Proof. We follow [Nes77] §1. Let $A = k[\mathbf{X}]$ be the ring of polynomials over k in the variables X_0, \dots, X_n , let $r = \dim(W) + 1$ and let $U_{i,j}$ for $i \in \{1, \dots, r\}$ and $j \in \{0, \dots, n\}$ be algebraically independent variables over A . We introduce the linear forms

$$L_i(\mathbf{X}) = \sum_{j=0}^n U_{i,j} X_j$$

and let I be the homogeneous prime ideal of A corresponding to the Zariski closure W' of W in \mathbb{P}^n . We let χ denote the ideal in $A[\mathbf{U}] = A[U_{1,0}, \dots, U_{r,n}]$ generated by X_0, \dots, X_n and define the ideals

$$\begin{aligned} \tilde{I} &= \bigcup_{k \geq 0} ((I, L_1, \dots, L_r) : \chi^k) \subset A[\mathbf{U}], \\ \bar{I} &= \tilde{I} \cap k[\mathbf{U}]. \end{aligned}$$

By [Nes77] Lemma 5 (3), \bar{I} is a principal ideal and we choose a generator $F \in k[\mathbf{U}]$ for \bar{I} . We refer to this polynomial as a **Chow form** of W' (it is unique up to scalar). The height of W is, by definition, $H(F)$. In particular, this does not depend on the choice of F .

We consider the variables

$$S_{j,k}^{(i)}, \quad j, k \in \{0, \dots, n\}, \quad i \in \{1, \dots, r\},$$

algebraically independent over A except for the skew symmetry

$$S_{j,k}^{(i)} + S_{k,j}^{(i)} = 0.$$

We define an A -algebra homomorphism

$$\theta: A[\mathbf{U}] \rightarrow A[\mathbf{S}] := A[S_{0,0}^{(1)}, \dots, S_{n,n}^{(r)}]$$

by sending $U_{i,j}$ to $\sum_{k=0}^n S_{j,k}^{(i)} X_k$. The **Chow ideal** J of W' is the ideal of $k[\mathbf{X}]$ generated by the coefficients P'_1, \dots, P'_m of $\theta(F)$ as a polynomial in the $S_{j,k}^{(i)}$ ($j < k$). We refer to these coefficients as Chow polynomials of W' and, by [Nes77] Lemma 11, they define W' . Therefore, if we set $X_0 = 1$, we obtain polynomials P_1, \dots, P_M defining W and we refer to these as **Chow polynomials**. The number M of Chow polynomials is clearly bounded by a constant depending only on n and D .

Consider the diagram

$$\Gamma \subset \mathbb{P}^n \times (\mathbb{P}^n)^r \xrightarrow{\varphi} (\mathbb{P}^n)^r,$$

where Γ is the variety defined by the ideal generated by I and the L_i . By definition, this is a projective and hence proper morphism. By the proof of [Nes77] Lemma 4, the image $\varphi(\Gamma)$ is precisely the hypersurface \mathcal{H} defined by F . Therefore, we have a proper, surjective morphism

$$\Gamma \subset \mathbb{P}^n \times (\mathbb{P}^n)^r \xrightarrow{\varphi} \mathcal{H},$$

which is clearly generically finite. By the projection formula, the degree of \mathcal{H} is at most the degree of Γ and, therefore, so is the degree of the Chow form F .

Note that the degree of the subvariety Γ_W of $\mathbb{P}^n \times (\mathbb{P}^n)^r$ defined by the ideal generated by I in $A[\mathbf{U}]$ is precisely the degree of W . Therefore, we can calculate the degree of Γ by applying the standard Bézout theorem to the intersection of Γ_W with the subvarieties defined by the L_i . We conclude that the degree of F and, hence, the degrees of the Chow polynomials are bounded by a constant $C_{33}(n, D)$. Furthermore, given the explicit nature of θ , there exist constants $C_{34}(n, D)$ and $C_{35}(n, D)$ such that the heights of the Chow polynomials are at most $C_{34}H^{C_{35}}$. \square

Proposition 3.11. *For all positive integers m, n, D , there are constants C_{36}, C_{37} and C_{38} depending on m, n and D such that:*

For every geometrically irreducible affine algebraic set $W \subset \mathbb{A}^n$ defined over a field $k \subset \mathbb{Q}^{\text{alg}}$ by at most m polynomials of degree at most D and height at most H , if $\text{Sing } W(\mathbb{Q}^{\text{alg}}) \neq \emptyset$, then there exists an algebraic set $W' \subset \mathbb{A}^n$ satisfying:

- (i) $\text{Sing } W \subset W' \subset W$;
- (ii) $W' \neq W$;
- (iii) W' can be defined by at most C_{36} polynomials, each of degree at most C_{36} and having coefficients in k of height at most $C_{37}H^{C_{38}}$.

Proof. We follow the argument of [BMZ07] pp. 9–10.

The hypothesis that W is geometrically irreducible allows us to apply Lemma 3.9. We can then also apply Lemma 3.10 to deduce that the number and degree of the Chow polynomials of W are bounded by constants depending only on m, n and D , and that the heights of the Chow polynomials are polynomially bounded with respect to H .

Let J be the Chow ideal of W , which is generated by the Chow polynomials P_1, \dots, P_M . Let I denote the prime ideal of W , and let $Q_1, \dots, Q_{M'}$ be generators for I . We shall study the rank of the Jacobian matrix

$$\text{Jac}(P_1, \dots, P_M) = \left(\frac{\partial P_i}{\partial X_j} \right)_{i,j} \in M_{M,n}(k[W])$$

and compare it with the rank of the analogous matrix $\text{Jac}(Q_1, \dots, Q_{M'})$.

According to [Nes77] Lemma 11, I is an isolated primary component of J (in fact it is the only isolated primary component). Hence we can find a minimal primary decomposition

$$J = I \cap I_2 \cap \dots \cap I_t.$$

Because this decomposition is minimal, $I_2 \cap \dots \cap I_t \not\subset I$ so we can choose some $f \in I_2 \cap \dots \cap I_t$ which is not in I . We thus have $f \notin I$ but $fI \subset J$.

Because $fI \subset J \subset I$, and because all polynomials in I vanish on W , a calculation shows that

$$\text{rk Jac}(fQ_1, \dots, fQ_{M'}) \leq \text{rk Jac}(P_1, \dots, P_M) \leq \text{rk Jac}(Q_1, \dots, Q_{M'})$$

at each point of W .

A similar calculation shows that

$$\text{Jac}(fQ_1, \dots, fQ_{M'}) = f \text{Jac}(Q_1, \dots, Q_{M'}) \text{ in } M_{M',n}(k[W]).$$

Since $f \notin I$, it is invertible in the function field $k(W)$. We conclude that

$$\mathrm{rk} \mathrm{Jac}(fQ_1, \dots, fQ_{M'}) = \mathrm{rk} \mathrm{Jac}(Q_1, \dots, Q_{M'})$$

at the generic point of W . It is well-known that

$$\mathrm{rk} \mathrm{Jac}(Q_1, \dots, Q_{M'}) = n - \dim(W)$$

at the generic point, and we deduce that the same is true of $\mathrm{rk} \mathrm{Jac}(P_1, \dots, P_M)$.

It follows that the inequality

$$\mathrm{rk} \mathrm{Jac}(P_1, \dots, P_M) < n - \dim(W)$$

defines a proper subset W' of W . This subset can be defined as the subset on which all $(n - \dim(W)) \times (n - \dim(W))$ minors of $\mathrm{Jac}(P_1, \dots, P_M)$ vanish, and so it is an algebraic subset defined by polynomials satisfying the bounds of condition (iii).

On the other hand, at any point in $\mathrm{Sing} W$, the Jacobian criterion for singular points tells us that

$$\mathrm{rk} \mathrm{Jac}(P_1, \dots, P_M) \leq \mathrm{rk} \mathrm{Jac}(Q_1, \dots, Q_{M'}) < n - \dim(W)$$

and so $\mathrm{Sing} W \subset W'$. □

4 Discriminants and tori

The bound for the height of a pre-special point x in Theorem 1.1 is relative to the discriminant of the centre of the endomorphism ring of the Hodge structure E_x , while the bound in Theorem 1.4 (as required to obtain Theorem 1.2) is relative to certain invariants of the Mumford–Tate torus of x . In this section, we prove the following theorem which links these bounds together.

Theorem 4.1. *Let (\mathbf{G}, X) be a Shimura datum, $\rho: \mathbf{G} \rightarrow \mathbf{GL}(E)$ a faithful \mathbb{Q} -representation and $E_{\mathbb{Z}}$ a lattice in $E_{\mathbb{Q}}$. Let K be a compact open subgroup of $\mathbf{G}(\mathbb{A}_f)$.*

There exist constants $C_{39}, C_{40} > 0$ such that for all $C_{41} > 0$, there exists $C_{42} > 0$ (where C_{39} and C_{40} depend only on $\mathbf{G}, X, \mathcal{F}$ and the realisation of X , while C_{42} depends on these data and also C_{41}) such that:

For every pre-special point $x \in X$:

- (a) *Let \mathbf{M} denote the Mumford–Tate group of x (which is a torus because x is pre-special).*
- (b) *Let $K_{\mathbf{M}} = K \cap \mathbf{M}(\mathbb{A}_f)$ and let $K_{\mathbf{M}}^m$ be the maximal compact subgroup of $\mathbf{M}(\mathbb{A}_f)$.*

- (c) Let $i(\mathbf{M})$ be the number of primes p for which $K_{\mathbf{M},p} := K \cap \mathbf{M}(\mathbb{Q}_p)$ is strictly contained in the maximal compact subgroup $K_{\mathbf{M},p}^m$ of $\mathbf{M}(\mathbb{Q}_p)$.
- (d) Let L be the splitting field of \mathbf{M} .
- (e) Let F_x be the centre of $\text{End}_{\mathbb{Q}\text{-HS}} E_x$, let R_x the centre of $\text{End}_{\mathbb{Z}\text{-HS}} E_x$, and let \mathcal{O}_x be the maximal order in F_x .

Then

$$[\mathcal{O}_x : R_x] \leq C_{42} C_{41}^{i(\mathbf{M})} [K_{\mathbf{M}}^m : K_{\mathbf{M}}]^{C_{39}} |\text{disc } L|^{C_{40}}.$$

For an explanation of the quantifiers associated with the constants in the above theorem, see section 1. Throughout this section, constants which depend only on \mathbf{G} , X , ρ and K will be referred to as **uniform**.

Our proof is based on Tsimerman's proof of this theorem for the special case of \mathcal{A}_g ([Tsi12] Lemma 7.2). Our only changes to Tsimerman's argument are that in section 4.2, we consider general CM Hodge structures instead of CM Hodge structures of type $(-1, 0) + (0, -1)$ (which is only a minor change), and that we give a proof for Proposition 4.11 which is used tacitly by Tsimerman but which appears to us to be a non-trivial statement.

We will reduce to the case in which the compact open subgroup $K \subset \mathbf{G}(\mathbb{A}_f)$ is a direct product of compact open subgroups $K_p \subset \mathbf{G}(\mathbb{Q}_p)$, so that

$$[K_{\mathbf{M}}^m : K_{\mathbf{M}}] = \prod_p [K_{\mathbf{M},p}^m : K_{\mathbf{M},p}].$$

We then study the p -part of $[\mathcal{O}_x : R_x]$, one prime at a time. We use the fact that tori splitting over unramified local fields have nice integral models to show that, outside a fixed finite set of bad primes, if p divides $[\mathcal{O}_x : R_x]$ then either p is ramified in L (i.e. p divides $\text{disc } L$) or $K_{\mathbf{M},p}^m \neq K_{\mathbf{M},p}$.

Note that, even if $K_{\mathbf{M},p}^m \neq K_{\mathbf{M},p}$, it need not be true that p divides $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$. We will deal with this by using [UY14a] Proposition 3.15 which says that, outside the bad primes and those ramified in L , if $K_{\mathbf{M},p}^m \neq K_{\mathbf{M},p}$ then $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}] \gg p$.

The hardest parts of the proof (section 4.2 and Proposition 4.12) control which power of p divides $[\mathcal{O}_x : R_x]$, relative to the power of p which divides $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$.

4.1 Generating \mathbb{Q} -algebras

Let $x \in X$ be a pre-special point, \mathbf{M} its Mumford–Tate group and L the splitting field of \mathbf{M} . We can naturally embed $\mathbf{M}(\mathbb{Q})$ as a subgroup of F_x^\times . Hence the reciprocity morphism

$$r : \text{Res}_{L/\mathbb{Q}} \mathbb{G}_m \rightarrow \mathbf{M}$$

induces a homomorphism

$$L^\times = \text{Res}_{L/\mathbb{Q}} \mathbb{G}_m(\mathbb{Q}) \rightarrow F_x^\times.$$

We will use Goursat's lemma for algebras to show that the image of L^\times generates F_x^\times as a \mathbb{Q} -algebra.

We will reuse the notation from the proof of Lemma 2.2: the isotypic decomposition of the \mathbb{Q} -Hodge structure E is

$$E_x = \bigoplus_{i=1}^s E_{x,i}$$

where $E_{x,i}$ is isomorphic to a power of an irreducible \mathbb{Q} -Hodge structure with endomorphism algebra F_i , the centre of $\text{End}_{\mathbb{Q}\text{-HS}} E_x$ is

$$F_x = \prod_{i=1}^s F_i,$$

and the centre of the centraliser of the Mumford–Tate group \mathbf{M} in $\mathbf{GL}(E)$ is denoted \mathbf{Z} . We recall that $\mathbf{Z}(\mathbb{Q})$ is the image of $F_x^\times \hookrightarrow \mathbf{GL}(E_{\mathbb{Q}})$, and so

$$\mathbf{Z} \cong \prod_{i=1}^s \text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m,$$

and that $\mathbf{M} \subset \mathbf{Z}$ because \mathbf{M} is commutative. Hence we get an injective homomorphism $\mathbf{M}(\mathbb{Q}) \rightarrow F_x^\times$.

Now the image of $\mathbf{M}(\mathbb{Q})$ generates F_x as a \mathbb{Q} -algebra. This is an instance of the general fact that for any semisimple representation of a commutative group, the group generates the centre of its endomorphism algebra, but we will give more details of the proof below.

Lemma 4.2. *Suppose that the image of \mathbf{M} under the projection*

$$\mathbf{Z} \rightarrow \text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m$$

is contained in a subtorus $\text{Res}_{F'/\mathbb{Q}} \mathbb{G}_m \subset \text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m$ for some subfield F' of F_i . Then $F' = F_i$.

Proof. Consider an irreducible \mathbf{M} -subrepresentation $E' \subset E_{x,i}$. This is an F_i -vector space of dimension 1, and hence an F' -vector space of dimension $[F_i : F']$. Since $\mathbf{M}(\mathbb{Q}) \subset F'$, every F' -subspace of E' is an \mathbf{M} -subrepresentation. But since E' is irreducible, we must have $[F_i : F'] = 1$. \square

Lemma 4.3. *Suppose the map from \mathbf{M} to $\text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m \times \text{Res}_{F_j/\mathbb{Q}} \mathbb{G}_m$ factors through the map*

$$\text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m \rightarrow \text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m \times \text{Res}_{F_j/\mathbb{Q}} \mathbb{G}_m : x \mapsto (x, \varphi(x)),$$

induced by a field isomorphism $\varphi : F_i \rightarrow F_j$. Then $i = j$.

Proof. Let $E'_{x,i}$ and $E'_{x,j}$ denote irreducible \mathbf{M} -subrepresentations of $E_{x,i}$ and $E_{x,j}$ respectively. The field isomorphism φ induces an isomorphism of \mathbf{M} -representations $E'_{x,i} \rightarrow E'_{x,j}$ and hence these are both in the same isotypic component i.e. $i = j$. \square

These properties yield the following conclusion:

Proposition 4.4. *The image of $\mathbf{M}(\mathbb{Q})$ generates the \mathbb{Q} -algebra $F_x = \prod_i F_i$.*

Proof. Consider the image of $\mathbf{M}(\mathbb{Q})$ in $\text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m(\mathbb{Q})$ for each i . The \mathbb{Q} -algebra generated by these elements constitutes a subfield F' of F_i . However, by Lemma 4.2, $F' = F_i$. Therefore, the image of $\mathbf{M}(\mathbb{Q})$ in F_x generates a \mathbb{Q} -algebra that surjects on to each factor F_i . Furthermore, however, by Lemma 4.3 the projection of $\mathbf{M}(\mathbb{Q})$ to $F_i \times F_j$ for any distinct i, j is not contained in the graph of an isomorphism $F_i \cong F_j$. Therefore, the proposition follows from the following lemma. \square

Lemma 4.5. *Let $A = \bigoplus_{i=1}^r K_i$ be a direct sum of finite field extensions K_i of a field k . Let Θ be a subset of A . Then Θ generates A as a k -algebra if and only if*

- (1) *for each projection $\pi_i : A \rightarrow K_i$, the image $\pi_i(\Theta)$ is not contained in a proper subfield of K_i and*
- (2) *for each projection $\pi_{i,j} : A \rightarrow K_i \oplus K_j$, the image $\pi_{i,j}(\Theta)$ is not contained in the graph of a k -isomorphism $K_i \cong K_j$.*

Proof. This is an immediate consequence of Goursat's lemma for algebras. \square

We will now consider the reciprocity morphism $r : \text{Res}_{L/\mathbb{Q}} \mathbb{G}_m \rightarrow \mathbf{M}$. Let us briefly recall its definition. Extending $x : \mathbb{S} \rightarrow \mathbf{G}_{\mathbb{R}}$ to a morphism over \mathbb{C} yields

$$\mathbb{G}_{m,\mathbb{C}}^2 \cong \mathbb{S}_{\mathbb{C}} \xrightarrow{hc} \mathbf{M}_{\mathbb{C}}$$

and we denote by μ the cocharacter obtained by pre-composing with the embedding $\mathbb{G}_{m,\mathbb{C}} \hookrightarrow \mathbb{G}_{m,\mathbb{C}}^2$ sending z to $(z, 1)$. As a cocharacter of \mathbf{M} , μ is necessarily defined over L i.e. $\mu : \mathbb{G}_{m,L} \rightarrow \mathbf{M}_L$ and restricting scalars to \mathbb{Q} yields $\text{Res}_{L/\mathbb{Q}} \mathbb{G}_m \rightarrow \text{Res}_{L/\mathbb{Q}} \mathbf{M}_L$. The morphism r is obtained by composing with the norm map $\text{Res}_{L/\mathbb{Q}} \mathbf{M}_L \rightarrow \mathbf{M}$. Since r is surjective (as a morphism of algebraic groups) we deduce the following corollary:

Corollary 4.6. *The image of $\mathrm{Res}_{L/\mathbb{Q}} \mathbb{G}_m(\mathbb{Q})$ generates the \mathbb{Q} -algebra F_x .*

Proof. If the corollary did not hold then the algebra generated by the image of $\mathrm{Res}_{L/\mathbb{Q}} \mathbb{G}_m(\mathbb{Q})$ in F_x would violate one of the conditions of Lemma 4.5. This would imply that the image of $\mathrm{Res}_{L/\mathbb{Q}} \mathbb{G}_m$ in the product of the $\mathrm{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m$ would factor non-trivially in one of the two ways referred to in Lemmas 4.2 and 4.3. However, since r is surjective (as a morphism of algebraic groups) and the \mathbb{Q} -points of $\mathrm{Res}_{L/\mathbb{Q}} \mathbb{G}_m(\mathbb{Q})$ are Zariski dense in $\mathrm{Res}_{L/\mathbb{Q}} \mathbb{G}_m$, this would be a contradiction. \square

4.2 Generators of small index

We have seen that the reciprocity map induces a homomorphism $L^\times \rightarrow F_x^\times$ whose image generates F_x as a \mathbb{Q} -algebra. We will construct a \mathbb{Q} -linear map $S: L \rightarrow F_x$ which can be interpreted as the logarithm of this homomorphism (this interpretation can be made precise p -adically). We then show that, for each prime p , there is an element $z_p \in \mathcal{O}_L \otimes \mathbb{Z}_p$ such that $S(z_p)$ generates a \mathbb{Z}_p -algebra whose index in the maximal order of $F_x \otimes \mathbb{Q}_p$ is a bounded power of p . We will later use $S(z_p)$ to bound the power of p dividing $[\mathcal{O}_x : R_x]$ relative to the power of p dividing $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$.

The morphism $L^\times \rightarrow \mathbf{M}(\mathbb{Q})$ defined by r is given explicitly by

$$z \mapsto \prod_{\sigma \in \mathrm{Gal}(L/\mathbb{Q})} \mu^\sigma(\sigma(z)) = \prod_{\sigma \in \mathrm{Gal}(L/\mathbb{Q})} \sigma(\mu(z)).$$

For $i = 1, \dots, s$, let μ_i denote the composition of $\mu: \mathbb{G}_{m,L} \rightarrow \mathbf{M}_L$ with the projection to $(\mathrm{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m) \times_{\mathbb{Q}} L$. Since \mathbf{M} is split over L , the fields F_i are all isomorphic to subfields of the Galois field L and so

$$(\mathrm{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m) \times_{\mathbb{Q}} L \cong \prod_{\tau: F_i \rightarrow L} \mathbb{G}_{m,L}.$$

Let $(-n_{i,\tau}, -n'_{i,\tau})$ denote the Hodge type of the F_i -eigenspace in $E_{x,L}$ on which F_i acts via τ . Then the definition of the cocharacter μ of \mathbf{M} implies that

$$\mu_i(z) = (z^{n_{i,\tau}})_{\tau: F_i \rightarrow L} \text{ in } \prod_{\tau: F_i \rightarrow L} \mathbb{G}_{m,L}.$$

Therefore, the projection r_i of r to $\mathrm{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m$ sends $z \in L^\times$ to the element

$$\prod_{\sigma \in \mathrm{Gal}(L/\mathbb{Q})} \sigma(\mu_i(z)) = \left(\prod_{\sigma \in \mathrm{Gal}(L/\mathbb{Q})} \sigma(z)^{n_{i,\sigma^{-1}\tau}} \right)_{\tau: F_i \rightarrow L}$$

in $\mathrm{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m(L) \cong \prod_{\tau: F_i \rightarrow L} L^\times$, for uniform integers $n_{i,\tau}$. Observe that this value is $\mathrm{Gal}(L/\mathbb{Q})$ -invariant, so in fact it is an element of $\mathrm{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m(\mathbb{Q}) = F_i^\times$.

We define a \mathbb{Q} -linear map $S: L \rightarrow F_x$ by

$$S(z) \otimes 1 = \left(\sum_{\sigma \in \text{Gal}(L/\mathbb{Q})} n_{i, \sigma^{-1}\tau} \sigma(z) \right)_{i, \tau} \quad \text{in } F_x \otimes_{\mathbb{Q}} L \cong \prod_{i=1}^s \prod_{\tau: F_i \rightarrow L} L^\times.$$

Observe that the right hand side of the above expression is invariant under the action of $\text{Gal}(L/\mathbb{Q})$ on $F_x \otimes_{\mathbb{Q}} L$, so it does indeed give us values in $F_x \otimes 1$. This map S is the same as the map from the tangent space of $\text{Res}_{L/\mathbb{Q}} \mathbb{G}_m$ to the tangent space of \mathbf{Z} induced by r .

Lemma 4.7. *The image of S generates the full \mathbb{Q} -algebra F_x .*

Proof. By Corollary 4.6, we know that $r_i: \text{Res}_{L/\mathbb{Q}} \mathbb{G}_m \rightarrow \text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m$ does not factor through $\text{Res}_{F'/\mathbb{Q}} \mathbb{G}_m$ for any proper subfield $F' \subset F_i$. Since r_i is a morphism of algebraic groups over a field of characteristic zero, it follows that the induced map of tangent spaces does not factor through the tangent space of the algebraic subgroup $\text{Res}_{F'/\mathbb{Q}} \mathbb{G}_m$. In other words, the projection of S onto F_i does not factor through any proper subfield $F' \subset F_i$.

A similar argument shows that the projection of S onto $F_i \times F_j$ does not factor through the graph of any isomorphism $F_i \rightarrow F_j$ (when $i \neq j$).

By Lemma 4.5, these properties imply the statement of the lemma. \square

For each prime p , the elements of $\text{Gal}(L/\mathbb{Q})$ extend \mathbb{Q}_p -linearly to the \mathbb{Q}_p -algebra $L \otimes_{\mathbb{Q}} \mathbb{Q}_p$ and so the map S also extends \mathbb{Q}_p -linearly to $L \otimes \mathbb{Q}_p$. Provided $n > 1$ (to ensure convergence), we have

$$r(\exp p^n z) = \exp p^n S(z)$$

as maps $\mathcal{O}_L \otimes \mathbb{Z}_p \rightarrow (\mathcal{O}_x \otimes \mathbb{Z}_p)^\times$.

Proposition 4.8. *There is a uniform constant C_{43} such that, for every prime p , there exists $z \in \mathcal{O}_L \otimes \mathbb{Z}_p$ such that the index of $\mathbb{Z}_p[S(z)]$ in*

$$\mathcal{O}_{x,p} := \mathcal{O}_x \otimes_{\mathbb{Z}} \mathbb{Z}_p = \prod_{i=1}^s \mathcal{O}_{F_i} \otimes_{\mathbb{Z}} \mathbb{Z}_p$$

is at most $p^{C_{43}}$.

Proof. Enumerate the complete set of ring homomorphisms ψ_i from $\mathcal{O}_{x,p}$ to a fixed p -adic completion L_v of L . Let v_p be the p -adic valuation on \mathbb{Q} and its extension to L_v . For $z \in \mathcal{O}_L \otimes \mathbb{Z}_p$ the index of $\mathbb{Z}_p[S(z)]$ in $\mathcal{O}_{x,p}$ is given by p raised to the value of v_p on the Vandermonde determinant

$$\iota(z) := \prod_{i \neq j} (\psi_i(S(z)) - \psi_j(S(z))),$$

divided by the discriminant of $\mathcal{O}_{x,p}$.

To find a suitable z we first choose a \mathbb{Z}_p -basis $\{\beta_j\}$ of $\mathcal{O}_L \otimes \mathbb{Z}_p$ such that $v_p(\beta_j)$ is bounded by a uniform positive integer. To see that this is possible, first choose any \mathbb{Z}_p -basis $\{\beta_j\}$ of $\mathcal{O}_L \otimes \mathbb{Z}_p$ and recall that, if we enumerate the automorphisms σ_i of $L \otimes \mathbb{Q}_p$, then the discriminant of $L \otimes \mathbb{Q}_p$ is p raised to the value of v_p on the determinant of the matrix $(\sigma_i(\beta_j))$. This discriminant is uniformly bounded and so we conclude that $v_p(\beta_j)$ is uniformly bounded for one of the β_j . However, by the properties of v_p we can therefore obtain our desired basis by, if necessary, taking sums of the form $\beta_j + \beta_{j'}$. Since the value of v_p is preserved under automorphism, $v_p(\iota(\beta_j))$ is uniformly bounded for any β_j .

Writing

$$z = \sum_j a_j \beta_j \text{ (for } a_j \in \mathbb{Z}_p\text{),}$$

we can express ι as a polynomial over \mathbb{Z}_p of uniformly bounded degree in the a_j . By Lemma 4.7, this polynomial is not identically zero. Hence by Lemma 3.5, we can find a linear combination z of the β_j whose coefficients have uniformly bounded p -adic valuation such that $\iota(z) \neq 0$. \square

4.3 Generating the maximal order

In this section we show that, for all primes p outside of a uniformly fixed finite set Σ , if p divides $[\mathcal{O}_x : R_x]$ then either p is ramified in L (and hence divides $\text{disc } L$) or $K_{\mathbf{M},p}$ is strictly contained in $K_{\mathbf{M},p}^m$. Note that, once we know that $K_{\mathbf{M},p} \neq K_{\mathbf{M},p}^m$, [UY14a] Proposition 3.15 implies that

$$[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}] \gg p.$$

The uniform finite set which we have to omit is the following: let Σ be the set of primes p for which the compact open subgroup $K_p \subset \mathbf{G}(\mathbb{Q}_p)$ is not contained in $\mathbf{GL}(E_{\mathbb{Z}_p})$.

Lemma 4.9. *If p is unramified in L , the group $K_{\mathbf{M},p}^m$ generates the ring $\mathcal{O}_{x,p}$.*

Proof. Recall that $\text{Res}_{L/\mathbb{Q}} \mathbf{G}_m$, \mathbf{M} and $\text{Res}_{F_i/\mathbb{Q}} \mathbf{G}_m$ possess smooth models \mathcal{R}_L , \mathcal{T} and \mathcal{R}_{F_i} over \mathbb{Z}_p whose generic fibres are $(\text{Res}_{L/\mathbb{Q}} \mathbf{G}_m)_{\mathbb{Q}_p}$, $\mathbf{M}_{\mathbb{Q}_p}$ and $(\text{Res}_{F_i/\mathbb{Q}} \mathbf{G}_m)_{\mathbb{Q}_p}$, respectively. Furthermore, since p is unramified in L , we may insist that the special fibres of these models are tori over \mathbb{F}_p and that their \mathbb{Z}_p -points are the respective unique hyperspecial (maximal) compact open subgroups. The extension of r over \mathbb{Q}_p then extends uniquely to a morphism

$$\mathcal{R}_L \rightarrow \mathcal{T} \rightarrow \prod_{i=1}^s \mathcal{R}_{F_i}$$

over \mathbb{Z}_p and, reducing modulo p , we obtain a map

$$\mathcal{R}_{L, \mathbb{F}_p} \rightarrow \mathcal{T}_{\mathbb{F}_p} \rightarrow \prod_{i=1}^s \mathcal{R}_{F_i, \mathbb{F}_p}.$$

Since $\mathbf{M}(\mathbb{Q})$ generates F_x , $\mathbf{M}(\mathbb{Q}_p)$ generates

$$F_x \otimes \mathbb{Q}_p = \prod_{i=1}^s \prod_{v|p} F_{i,v}$$

where the second product is over places of F_i dividing p .

Therefore, by Lemma 4.5, the embedding of $\mathbf{M}_{\mathbb{Q}_p}$ in

$$\prod_{i=1}^s (\text{Res}_{F_i/\mathbb{Q}} \mathbb{G}_m)_{\mathbb{Q}_p} = \prod_{i=1}^s \prod_{v|p} \text{Res}_{F_{i,v}/\mathbb{Q}_p} \mathbb{G}_m$$

has the analogous properties to the embedding over \mathbb{Q} , namely that the projection to any factor $\text{Res}_{F_{i,v}/\mathbb{Q}_p} \mathbb{G}_m$ does not factor through $\text{Res}_{F'/\mathbb{Q}_p} \mathbb{G}_m$ for any proper p -adic subfield F' of $F_{i,v}$ and the projection to the product of any two factors $\text{Res}_{F_{i,v}/\mathbb{Q}_p} \mathbb{G}_m$ and $\text{Res}_{F_{j,w}/\mathbb{Q}_p} \mathbb{G}_m$ does not factor through the graph of an isomorphism. By [DG63] Exp. X Lemme 4.1, these two properties immediately transfer to the analogous conditions on the embedding of $\mathcal{T}_{\mathbb{F}_p}$ in

$$\prod_{i=1}^s \mathcal{R}_{F_i, \mathbb{F}_p} = \prod_{i=1}^s \prod_{v|p} \text{Res}_{k_{i,v}/\mathbb{F}_p} \mathbb{G}_m,$$

where $k_{i,v}$ denotes the residue field of $F_{i,v}$. By Lemma 4.5, we conclude that $\mathcal{T}_{\mathbb{F}_p}(\mathbb{F}_p)$ generates the full \mathbb{F}_p -algebra

$$\mathcal{O}_x \otimes_{\mathbb{Z}} \mathbb{F}_p = \prod_{i=1}^s \prod_{v|p} k_{i,v}$$

The lemma is now a direct consequence of Nakayama's lemma. □

Lemma 4.10. *If $p \notin \Sigma$, then $K_{\mathbf{M},p}$ is contained in $R_{x,p}$.*

Proof. Since $p \notin \Sigma$,

$$K_{\mathbf{M},p} \subset \mathbf{M}(\mathbb{Q}_p) \cap \mathbf{GL}(E_{\mathbb{Z}_p}).$$

On the other hand, we have already seen that

$$\mathbf{M}(\mathbb{Q}_p) \subset \mathbf{Z}(\mathbb{Q}_p) = (F_x \otimes_{\mathbb{Q}} \mathbb{Q}_p)^\times.$$

Hence

$$K_{\mathbf{M},p} \subset (F_x \otimes_{\mathbb{Q}} \mathbb{Q}_p) \cap \text{End}_{\mathbb{Z}_p} E_{\mathbb{Z}_p} = R_{x,p}. \quad \square$$

Proposition 4.11. *If $p \notin \Sigma$ is unramified in L and p divides $[\mathcal{O}_x : R_x]$ then $K_{\mathbf{M},p}$ is strictly contained in $K_{\mathbf{M},p}^m$.*

Proof. Suppose $K_{\mathbf{M},p} = K_{\mathbf{M},p}^m$. By Lemma 4.9, $K_{\mathbf{M},p}^m$ generates the ring $\mathcal{O}_{x,p}$. However, by Lemma 4.10, $K_{\mathbf{M},p}^m$ is contained in the ring $R_{x,p}$. Hence $R_{x,p} = \mathcal{O}_{x,p}$. \square

4.4 Proof of Theorem 4.1

The following proposition is the final ingredient in the proof of Theorem 4.1. We use it in two different ways: if p divides $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$, then the proposition tells us that the p -part of $[\mathcal{O}_x : R_x]$ is bounded by a suitable power of the p -part of $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$, while if p does not divide $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$, then the proposition tells us that the power of p dividing $[\mathcal{O}_x : R_x]$ is uniformly bounded.

Proposition 4.12. *There is a uniform integer C_{44} such that, if p^n is the maximum power of p dividing $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$, then*

$$p^{C_{44}(n+1)} \nmid [\mathcal{O}_x : R_x].$$

Proof. Choose $z \in \mathcal{O}_L \otimes \mathbb{Z}_p$ as in Proposition 4.8 so that the index of $\mathbb{Z}_p[S(z)]$ in $\mathcal{O}_{x,p}$ is at most $p^{C_{43}}$. Let h be the dimension of $\mathcal{O}_{x,p}$ as a free module over \mathbb{Z}_p and define $y_j := \exp(jp^N S(z))$ for $j \in \{0, \dots, h-1\}$ and $N > 1$ an integer. Note that $y_j = r(\exp(jz))^{p^N}$ is an element of $K_{\mathbf{M},p}^m$. However, it also lies in the image I of the p -adic exponential on $p^2\mathcal{O}_{x,p}$, which is a pro- p group. Therefore, the quotient of $K_{\mathbf{M},p}^m \cap I$ by $K_{\mathbf{M},p} \cap I$ is a p -group contained in the finite abelian group $K_{\mathbf{M},p}^m / K_{\mathbf{M},p}$. In particular, its order divides p^n and, if n divides N , we conclude that $y_j \in K_{\mathbf{M},p}$. If $p \notin \Sigma$ we conclude that $y_j \in R_{x,p}$. If $p \in \Sigma$, we can replace N by a uniformly bounded multiple of n to obtain the same result.

By definition, modulo p^{Nh-h} ,

$$y_j \equiv 1 + jp^N S(z) + \frac{j^2 p^{2N} S(z)^2}{2!} + \dots + \frac{j^{h-1} p^{N(h-1)} S(z)^{h-1}}{(h-1)!}.$$

Therefore, the span Λ of the elements y_j in $\mathbb{Z}_p[p^N S(z)]$ modulo p^{Nh-h} is of index bounded by the maximum power p^{m_h} of p dividing the (non-zero) determinant of the matrix

$$\left(\frac{j^k}{k!} \right)_{j,k \in \{0, \dots, h-1\}}.$$

On the other hand, $\mathbb{Z}_p[p^N S(z)]$ certainly contains $p^{N(h-1)+C_{43}}\mathcal{O}_{x,p}$, so we conclude that Λ contains

$$p^{N(h-1)+C_{43}+m_h} \mathcal{O}_{x,p} / p^{Nh-h} \mathcal{O}_{x,p},$$

which is non-trivial if $N > h + C_{43} + m_h$. In particular, if all of the above conditions on N hold then Λ contains the quotient

$$p^{Nh-h-1}\mathcal{O}_{x,p}/p^{Nh-h}\mathcal{O}_{x,p}$$

and so, by Nakayama's lemma, the span of the y_j in $\mathcal{O}_{x,p}$ contain $p^{Nh-h-1}\mathcal{O}_{x,p}$. Since the $y_j \in R_{x,p}$, we conclude

$$[\mathcal{O}_{x,p} : R_{x,p}] \leq p^{Nh^2-h^2-h},$$

which finishes the proof. \square

Now we are ready to finish the proofs of Theorems 1.4 and 4.1.

Proof of Theorem 4.1. We begin by reducing to the case in which $K \subset \mathbf{G}(\mathbb{A}_f)$ is a direct product of compact open subgroups $K_p \subset \mathbf{G}(\mathbb{Q}_p)$. To see that this is sufficient, first note that an arbitrary compact open subgroup $K \subset \mathbf{G}(\mathbb{A}_f)$ always contains such a subgroup $K' = \prod_p K'_p$. Replacing K by K' changes $i(\mathbf{M})$ by at most the (finite) number of primes for which

$$K \cap \mathbf{G}(\mathbb{Q}_p) \neq K' \cap \mathbf{G}(\mathbb{Q}_p),$$

and multiplies $[K_{\mathbf{M}}^m : K_{\mathbf{M}}]$ by $[K_{\mathbf{M}} : K'_{\mathbf{M}}]$, which is at most $[K : K']$. Hence replacing K by K' changes the right hand side of the bound in Theorem 4.1 by a uniformly bounded quantity, so it suffices to prove the theorem with K' in place of K .

We may therefore assume that $K = \prod_p K_p$, so

$$[K_{\mathbf{M}}^m : K_{\mathbf{M}}] = \prod_p [K_{\mathbf{M},p}^m : K_{\mathbf{M},p}].$$

For any integer x and any prime p , let $[x]_p$ denote the largest power of p which divides x i.e.

$$[x]_p = p^{v_p(x)}.$$

We will show that we can choose constants C_{39} and C_{40} uniformly and a constant C_{45} depending on \mathbf{G} , X , ρ , K , C_{41} such that for every prime p , either

$$p \text{ and } [\mathcal{O}_x : R_x]_p \text{ are both less than } C_{45}$$

or

$$[\mathcal{O}_x : R_x]_p \leq C_{41}^{i(\mathbf{M},p)} [K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]^{C_{39}} [\text{disc } L]_p^{C_{40}} \quad (*)$$

where $i(\mathbf{M},p) = 1$ if $K_{\mathbf{M},p}^m \neq K_{\mathbf{M},p}$ and 0 if $K_{\mathbf{M},p}^m = K_{\mathbf{M},p}$. The bound in Theorem 4.1 follows by taking the product over all primes p (the first case leads to the constant C_{42}).

- If p divides $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$ then Proposition 4.12 implies that

$$[\mathcal{O}_x : R_x]_p \leq [K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]^{2C_{44}}.$$

- If we suppose further that $p^{C_{44}} \geq C_{41}^{-1}$, then

$$[\mathcal{O}_x : R_x]_p \leq C_{41} [K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]^{3C_{44}}$$

and so taking $C_{39} \geq 3C_{44}$ suffices.

- If p does not divide $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]$ then by Proposition 4.12,

$$[\mathcal{O}_x : R_x]_p \leq p^{C_{44}}.$$

- If $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}] > 1$, then by [UY14a] Proposition 3.15 there exists a uniform constant $C_{46} \leq 1$ such that

$$C_{46}p \leq [K_{\mathbf{M},p}^m : K_{\mathbf{M},p}].$$

If also p is large enough that $p^{C_{44}} \geq C_{46}^{-2C_{44}} C_{41}^{-1}$, we deduce that

$$[\mathcal{O}_x : R_x]_p \leq C_{41} [K_{\mathbf{M},p}^m : K_{\mathbf{M},p}]^{2C_{44}}.$$

Thus in this case, $C_{39} \geq 2C_{44}$ suffices.

- If $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}] = 1$ and p is ramified in L then p divides disc L and $i(\mathbf{M}, p) = 0$. Thus

$$[\mathcal{O}_x : R_x]_p \leq C_{41}^{i(\mathbf{M},p)} [\text{disc } L]_p^{C_{44}}$$

and so $C_{40} \geq C_{44}$ suffices.

- If $[K_{\mathbf{M},p}^m : K_{\mathbf{M},p}] = 1$, p is unramified in L and $p \notin \Sigma$, then Proposition 4.11 implies that $[\mathcal{O}_x : R_x] = 1$. Again $i(\mathbf{M}, p) = 0$ so both sides of (*) are equal to 1.

- We are left with the cases

$$p \in \Sigma \text{ or } p^{C_{44}} < C_{46}^{-2C_{44}} C_{41}^{-1}$$

(note that this includes $p^{C_{44}} < C_{41}^{-1}$ because $C_{46} \leq 1$). In these cases we have an upper bound for p depending on uniform data and on C_{41} (recall that Σ is finite and uniform). The fact that $[\mathcal{O}_x : R_x]_p \leq p^{C_{44}}$ implies a similar bound for $[\mathcal{O}_x : R_x]_p$. \square

Proof of Theorem 1.4. By Theorem 1.1, there are constants C_1 and C_2 depending only on \mathbf{G} , X , \mathcal{F} and ρ such that

$$H(x) \leq C_1 |\text{disc } R_x|^{C_2}.$$

and

$$\text{disc } R_x = [\mathcal{O}_x : R_x]^2 \text{disc } \mathcal{O}_x = [\mathcal{O}_x : R_x]^2 \prod_{i=1}^s \text{disc } F_i.$$

Since each F_i is isomorphic to a subfield of L , we know that $|\text{disc } F_i|$ is at most $|\text{disc } L|$. Clearly, the number of isotypic components s is at most the dimension of E . Therefore, Theorem 1.4 follows from Theorem 4.1. \square

Acknowledgements

Both authors would like to thank Emmanuel Ullmo and Andrei Yafaev for suggesting that they work together on this problem and for numerous conversations regarding the subject of this paper. The authors also owe a special thank you to Philipp Habegger who suggested the use of Chow polynomials to prove Proposition 3.11. They are grateful to Ziyang Gao for pointing out the issue with quantifiers in Theorem 1.4 which is needed for it to imply Theorem 1.2. Both authors would like to thank the referee for their reading of the manuscript and their helpful comments.

The first author is indebted to the Engineering and Physical Sciences Research Council and the Institut des Hautes Études Scientifiques for their financial support. The second author was funded by European Research Council grant 307364 “Some problems in Geometry of Shimura varieties.”

References

- [AM69] M. F. Atiyah and I. G. Macdonald. *Introduction to commutative algebra*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1969.
- [BMZ07] E. Bombieri, D. Masser, and U. Zannier. Anomalous subvarieties—structure theorems and applications. *International Mathematics Research Notices*, 2007, 2007.
- [Bor69] 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.

- [Daw15] C. Daw. The André-Oort conjecture via o-minimality. In G. O. Jones and A. J. Wilkie, editors, *O-minimality and Diophantine Geometry*, London Mathematical Society Lecture Note Series 421. Cambridge University Press, 2015.
- [Del79] 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., XXXIII, pages 247–289. Amer. Math. Soc., Providence, R.I., 1979.
- [DG63] M. Demazure and A. Grothendieck. *Schémas en groupes, SGA 3*. IHES, 1963. Exp. VIII–XIV, Fasc. 4.
- [DH] C. Daw and A. Harris. Categoricity of modular and Shimura curves. To appear in *J. Inst. Math. Jussieu*.
- [DY11] C. Daw and A. Yafaev. An unconditional proof of the André–Oort conjecture for Hilbert modular surfaces. *Manuscripta Math.*, 135:263–271, 2011.
- [EY03] B. Edixhoven and A. Yafaev. Subvarieties of Shimura varieties. *Ann. of Math.*, 157(2):621–645, 2003.
- [Gaoa] Z. Gao. About the mixed André–Oort conjecture: reduction to a lower bound for the pure case. Unpublished note, available at www.math.u-psud.fr/~gao.
- [Gaob] Z. Gao. Towards the André–Oort conjecture for mixed Shimura varieties: The Ax–Lindemann theorem and lower bounds for Galois orbits of special points. To appear in *J. Reine Angew. Math.*.
- [Hab08] P. Habegger. Intersecting subvarieties of \mathbb{G}_m^n with algebraic subgroups. *Math. Ann.*, 342(2):449–466, 2008.
- [KUY] B. Klingler, E. Ullmo, and A. Yafaev. The hyperbolic Ax–Lindemann–Weierstrass conjecture. Preprint, available at arxiv.org/abs/1307.3965.
- [KY14] B. Klingler and A. Yafaev. The André–Oort conjecture. *Ann. of Math. (2)*, 180(3):867–925, 2014.
- [Nes77] Ju V. Nesterenko. Estimates for the orders of zeros of functions of a certain class and applications in the theory of transcendental numbers. *Math. USSR Izv.*, 11(2):239–270, 1977.

- [Phi95] P. Philippon. Sur des hauteurs alternatives. III. *J. Math. Pures Appl.* (9), 74(4):345–365, 1995.
- [Pil09] J. Pila. Algebraic points of definable sets and results of André–Oort–Manin–Mumford type. *IMRN*, 13:2476–2507, 2009.
- [Pil11] J. Pila. O-minimality and the André–Oort conjecture for \mathbb{C}^n . *Ann. of Math.*, 173(3):1779–1840, 2011.
- [PS10] Y. Peterzil and S. Starchenko. Tame complex analysis and o-minimality. In *Proceedings of the ICM, Hyderabad*, 2010.
- [PT13] J. Pila and J. Tsimerman. The André–Oort conjecture for the moduli space of abelian surfaces. *Compositio Math.*, 149(2):204–216, 2013.
- [PT14] J. Pila and J. Tsimerman. Ax–Lindemann for \mathcal{A}_g . *Ann. of Math. (2)*, 179(2):659–681, 2014.
- [PW06] J. Pila and A. Wilkie. The rational points of a definable set. *Duke Math. J.*, 133:591–616, 2006.
- [PZ08] J. Pila and U. Zannier. Rational points in periodic analytic sets and the Manin–Mumford conjecture. *Rend. Mat. Acc. Lincei*, 19:149–162, 2008.
- [RU09] N. Ratazzi and E. Ullmo. Galois+equidistribution=Manin–Mumford. In *Ecole d’été arithmetic geometry*, Clay Math. Proc. Amer. Math. Soc., 2009.
- [Tsi] J. Tsimerman. A proof of the André–Oort conjecture for \mathcal{A}_g . Preprint, available at arxiv.org/abs/1506.01466.
- [Tsi12] J. Tsimerman. Brauer–Siegel for arithmetic tori and lower bounds for Galois orbits of special points. *J. Amer. Math. Soc.*, 25:1091–1117, 2012.
- [Ull14] E. Ullmo. Applications du théorème d’Ax–Lindemann hyperbolique. *Compos. Math.*, 150(2):175–190, 2014.
- [UY11] E. Ullmo and A. Yafaev. A characterisation of special subvarieties. *Mathematika*, 57(2):263–273, 2011.
- [UY14a] E. Ullmo and A. Yafaev. Galois orbits and equidistribution of special subvarieties: towards the André–Oort conjecture. *Ann. of Math. (2)*, 180(3):823–865, 2014.
- [UY14b] E. Ullmo and A. Yafaev. The hyperbolic Ax–Lindemann theorem in the compact case. *Duke Math. J.*, 163(2):267–463, 2014.

- [UY15] E. Ullmo and A. Yafaev. Nombre de classes des tores de multiplication complexe et bornes inférieures pour les orbites galoisiennes de points spéciaux. *Bull. Soc. Math. France*, 143(1):197–228, 2015.