Geometric quadratic chabauty and other topics in number theory

Lido, G.M.

Introduction

This thesis consists of three parts. The first part is devoted to the quadratic Chabauty method, the second part to automorphisms of modular curves of Cartan type and the third to the discrete logarithm problem over finite fields whose characteristic is small compared to the cardinality.

The first two chapters are the result of a joint work with Bas Edixhoven and describe a method that, in certain cases, determines the set of rational points on a curve $C/\mathbb{Q}$ of genus at least 2. The finiteness of the set $C(\mathbb{Q})$ is a special case of a theorem proved by Faltings in [43], but computing this set for each curve C is still an unsolved problem. In [24], Chabauty proposed a method to solve this problem when $C(\mathbb{Q})$ contains at least one point b and the rank r of the Mordell-Weil group of the jacobian of C is smaller than the genus g of the curve. Denoting J the jacobian of C and $j_b\colon C \to J$ the map sending a point x to $[x-b]$, Chabauty's method is based on the following diagram, which is commutative for every choice of a prime p

$$\begin{tikzcd}
C(\mathbb{Q}) \arrow[r, "j_b"] \arrow[d] & J(\mathbb{Q}) \arrow[d] \\
C(\mathbb{Q}_p) \arrow[r, "j_b"] & J(\mathbb{Q}_p)
\end{tikzcd} .$$

The commutativity of the diagram implies that $C(\mathbb{Q})$, considered as a subset of $J(\mathbb{Q}_p)$, is contained in the intersection of $C(\mathbb{Q}_p)$ and the closure $\overline{J(\mathbb{Q})}$ of $J(\mathbb{Q})$. Up to computing generators for $J(\mathbb{Q})$, both the sets $C(\mathbb{Q}_p)$ and $\overline{J(\mathbb{Q})}$ can be computed with arbitrarily large precision inside $J(\mathbb{Q}_p)$ and their intersection is finite when r is smaller than g. Chabauty's method is to compute such an intersection, so to determine a finite subset of $C(\mathbb{Q}_p)$ containing $C(\mathbb{Q})$. Such an intersection can be larger than $C(\mathbb{Q})$ but in practice the Mordell-Weil sieve is usually enough to get rid of the undesired points.

In [62] and [63], Minhyong Kim proposes a non-abelian generalization of the Chabauty method, using the Galois cohomology of the $\mathbb{Q}_p$-pro-unipotent fundamental group of C.

The most interesting application of Kim's method is the so-called "quadratic Chabauty method", which is explicit and works when the rank ρ of the group $\mathrm{Pic}(J)/\mathrm{Pic}^0(J)$ is larger than $r-g+1$. In [10] this method is applied to the so-called cursed curve ($r=g=3$).

In chapter 1 we aim to make the quadratic Chabauty method *small* and *geometric* again: our generalization of Chabauty's method works by substituting J with a product of $\mathbb{G}_{\mathrm{m}}$-torsors over J and by extending the geometry over $\mathbb{Z}$.

Let $J^\vee$ be the dual abelian variety of J and let P be the Poincaré bundle on $J \times J^\vee$, the universal translational-invariant line bundle on J. After removing the zero-section of P we get a $\mathbb{G}_{\mathrm{m}}$-torsor $P^\times \to J \times J^\vee$, named *Poincaré torsor of* J, which is the main actor in our method. For any $\mathbb{Q}$-scheme S and any choice of points $x, x_1, x_2 \in J(S)$ and $y, y_1, y_2 \in J^\vee(S)$, the theorem of the cube implies the existence of canonical isomorphisms $(x_1,y)^*P \otimes (x_2,y)^*P = (x_1+x_2,y)^*P$ and $(x,y_1)^*P \otimes (x,y_2)^*P = (x,y_1+y_2)^*P$. This implies the existence of maps

$$\begin{aligned} +_1 &\colon (x_1,y)^*P^\times \times_S (x_2,y)^*P^\times \longrightarrow (x_1+x_2,y)^*P^\times\,, \\ +_2 &\colon (x,y_1)^*P^\times \times_S (x,y_2)^*P^\times \longrightarrow (x,y_1+y_2)^*P^\times\,. \end{aligned}$$

These partial operations $+_1, +_2$ give the Poincaré torsor a structure of *biextension*.

Moreover, the group of line bundles on J that arise as $(\mathrm{id}, g)^*P$ for some morphism $g\colon J \to J^\vee$ is a subgroup of $\mathrm{Pic}(J)$ of finite index: all the elements of $\mathrm{Pic}^0(J)$ can be obtained with g constant and, for any class $[\mathcal{L}] \in \mathrm{Pic}(J)/\mathrm{Pic}^0(J)$, the class $2[\mathcal{L}]$ can be obtained choosing $g\colon x \mapsto \mathrm{tr}_x^*\mathcal{L}\otimes\mathcal{L}^{-1}$, where tr_x is the translation by x on J. This implies the existence of maps $g_1, \dots, g_{\rho-1}\colon J \to J^\vee$ such that the line bundles $\mathcal{L}_i := (\mathrm{id}, g_i)^*P$ are linearly independent in $\mathrm{Pic}(J)$ and, for every $i \in \{1, \dots, \rho-1\}$, the line bundle $j_b^*(\mathcal{L}_i)$ on C is the trivial. Let $\mathcal{L}_i^\times$ be the $\mathbb{G}_{\mathrm{m}}$-torsor on J obtained removing the zero-section from $\mathcal{L}_i$ and let T be the $\mathbb{G}_{\mathrm{m}}^{\rho-1}$-torsor on J obtained as the product of all the $\mathcal{L}_i^\times$. Then j_b^*T is a trivial $\mathbb{G}_{\mathrm{m}}^{\rho-1}$-torsor on C, implying that the map $j_b\colon C \to J$ can be lifted to a map

$$\widetilde{j_b}\colon C \longrightarrow T\,.$$

This construction can be extended over $\mathbb{Z}$. The abelian varieties J and $J^\vee$ admit Néron models over $\mathbb{Z}$ and the Poincaré torsor uniquely extends, as a biextension, to a $\mathbb{G}_{\mathrm{m}}$-torsor over the product of the Néron model of J and the scheme $J^{\vee 0}$, defined as the fibrewise connected component of 0 in the Néron model of $J^\vee$. Up to composing g_i with a certain multiplication map on $J^\vee$, we can suppose that the image of the Néron model of J under g_i is contained in $J^{\vee 0}$. This gives the extension of $\mathcal{L}_i$ and T as torsors over the Néron model of J. The curve $C/\mathbb{Q}$ can be extended to a regular proper curve $C/\mathbb{Z}$, but to apply our method we need to restrict to certain open sub-schemes. Inside the smooth part of

C let U be an open sub-scheme obtained by removing, for each prime q of bad reduction, all but one irreducible component of the fibre at $\mathbb{F}_q$. The map j_b extends to the smooth part of C and the line bundles $j_b^*\mathcal{L}_i$ are trivial on U. Hence there exists a lift $\widetilde{j_b}\colon U \to T$ of j_b making the following diagram commutative for every prime p

$$\begin{tikzcd} U(\mathbb{Z}) \arrow[r, "\widetilde{j_b}"] \arrow[d] & T(\mathbb{Z}) \arrow[d] \\ U(\mathbb{Z}_p) \arrow[r, "\widetilde{j_b}"] & T(\mathbb{Z}_p) \end{tikzcd} .$$

For simplicity we suppose $p > 2$. Since $T(\mathbb{Z})$ is a $\mathbb{G}_\mathrm{m}(\mathbb{Z})^{\rho-1}$-torsor over $J(\mathbb{Z})$ and since $\mathbb{G}_\mathrm{m}(\mathbb{Z})^{\rho-1}$ is a finite group, we expect the closure $\overline{T(\mathbb{Z})}$ of $T(\mathbb{Z})$ inside $T(\mathbb{Z}_p)$ to be a p-adic variety of dimension at most r. This is a consequence of Theorem 1.4.10: the set of points in $\overline{T(\mathbb{Z})}$ with a given reduction modulo p, when not empty, is the image of an analytic map $\kappa\colon \mathbb{Z}_p^r \to T(\mathbb{Z}_p)$, constructed using the biextension structure on $P^\times$. Since $U(\mathbb{Z}_p)$ is 1-dimensional and $T(\mathbb{Z}_p)$ has dimension $g+\rho-1$, we expect the set $\overline{T(\mathbb{Z})} \cap U(\mathbb{Z}_p)$ to be finite when ρ is larger than $r-g+1$. This is proven in Section 1.9.2.

The geometric quadratic Chabauty method is to compute $\overline{T(\mathbb{Z})} \cap U(\mathbb{Z}_p)$, so to determine a finite subset of $U(\mathbb{Z}_p)$ containing $U(\mathbb{Z})$. Since $C(\mathbb{Q})$ is the union of the sets $U(\mathbb{Z})$ for all possible U's and since there are finitely many U's, the method can be used to prove that a certain list of points in $C(\mathbb{Q})$ is complete. In Theorem 1.4.12 we explain how, sometimes, computations in $T(\mathbb{Z}/p^2\mathbb{Z})$ imply a bound on the cardinality of $\overline{T(\mathbb{Z})} \cap U(\mathbb{Z}_p)$. In Sections 1.6 and 1.7 we explain how to make the method explicit. In Section 1.8 we apply our method to a specific example, with $g = r = \rho = 2$. Chapter 2 is devoted to an alternative proof of Theorem 1.4.10, using formal biextensions.

A motivation to study modular curves associated to Cartan and Cartan-plus subgroups of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$, as we do in chapter 3, comes from Serre's uniformity conjecture. This conjecture states that, for p prime big enough, the natural Galois representation $G_\mathbb{Q} \to \mathrm{GL}(E[p])$ is surjective for any elliptic curve $E/\mathbb{Q}$. The conjecture would be solved if we knew, for each prime p and each maximal subgroup $H < \mathrm{GL}_2(\mathbb{F}_p)$ such that $\det(H) = \mathbb{F}_p^\times$, the rational points on the modular curve associated with H. All the H's for which we do not know the answer are Cartan-plus subgroups, which are maximal for $p > 3$. This also gives motivation to study the so-called cursed curve, which is a modular curve associated to a Cartan-plus subgroup of $\mathrm{GL}_2(\mathbb{F}_{13})$.

Given a positive integer n, a Cartan subgroup of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$ is a subgroup arising as $A^\times \subset \mathrm{GL}(A) \cong \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$ for some étale $\mathbb{Z}/n\mathbb{Z}$-algebra A of rank 2. We call Cartan-plus subgroup of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$ a subgroup generated by $A^\times$ and the group of ring automorphisms of A, for some étale $\mathbb{Z}/n\mathbb{Z}$-algebra A of rank 2. For example, if n is prime there are two Cartan subgroups and two Cartan-plus subgroups up to conjugacy: the split

Cartan, respectively Cartan-plus, if $A \cong \mathbb{F}_n \times \mathbb{F}_n$ and the non-split Cartan, respectively Cartan-plus, if $A \cong \mathbb{F}_{n^2}$. We notice that the term *Cartan-plus* is not common in the literature: the most studied cases are the ones where $n > 3$ is prime and in these cases Cartan-plus subgroups are just normalizers of Cartan subgroups. We also deal with composite level and studying Cartan-plus subgroups allows us to state certain results with more uniformity than we could have done if we had studied normalizers of Cartan subgroups.

When a modular curve X is geometrically connected, the set $Y(\mathbb{C})$, made of its complex non-cuspidal points, is the quotient of $\mathbb{H} = \{z \in \mathbb{C} : \mathrm{Im}(z) > 0\}$ by the action of a subgroup $\Gamma < \mathrm{PSL}_2(\mathbb{Z})$. Every matrix $m \in \mathrm{PSL}_2(\mathbb{R})$ defines a complex automorphism of $\mathbb{H}$, that descends to an automorphism of $Y(\mathbb{C})$ if and only if the matrix m lies in the normalizer of Γ. When this happens, the automorphism extends to the whole $X(\mathbb{C})$. We call *modular* automorphism of $X_{\mathbb{C}}$ any such automorphism. We call *Cartan curve* a modular curve associated to a Cartan or to a Cartan-plus subgroup of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$. Using this terminology we state the main result of chapter 3.

Theorem. *Let n be either an integer larger than 10^{400} or a prime power such that $n > 11$ and $n \notin \{3^3, 2^4, 2^5, 2^6\}$. Then, over $\mathbb{C}$, all the automorphisms of a Cartan curve of level n are modular.*

For each Cartan or Cartan-plus subgroup $H < \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$, the group of modular automorphisms of the modular curve associated to H is easy to compute: it is either isomorphic to $N'/H' \times \mathbb{Z}/2\mathbb{Z}$ or to N'/H', where $N' < \mathrm{SL}_2(\mathbb{Z}/n\mathbb{Z})$ is the normalizer of $H' := H \cap \mathrm{SL}_2(\mathbb{Z}/n\mathbb{Z})$. This is stated more precisely in Proposition 3.6.13. Remark 3.6.16 gives N'/H' for each possible H.

In the proof of the main result of chapter 3, one of the steps is the following generalization of a result of Chen.

Theorem. *Let n be a positive integer. Then the jacobian of a Cartan curve of level n is a quotient of the jacobian of the modular curve $X_0(n^2)$.*

Using the last theorem and a result of Shimura characterizing the CM sub-abelian varieties of $J_0(n^2)$, we prove that, for all but finitely many n, a large part of the jacobian of a Cartan curve does not contain any CM sub-abelian variety. This, using a result of Ribet, implies that all the automorphisms of a Cartan curve of level n are defined over a compositum of quadratic fields for all but finitely many n.

The main result of chapter 3 then follows from Abramovich's lower bound of the gonality of modular curves and the following criterion.

Lemma. *Let n be a positive integer and let X be the base change to $\mathbb{C}$ of a modular curve associated with a subgroup $H < \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$. Suppose that H contains the scalar* *matrices, that* $\det(H)$ *is the whole* $(\mathbb{Z}/n\mathbb{Z})^\times$ *and that there are two primes* $\ell_1 < \ell_2$ *not dividing* n *such that* $5 \leq \ell_2 < \frac{1}{2}\operatorname{gon}(X) - 1$*, with* $\operatorname{gon}(X)$ *the gonality of* X*. Then every automorphism of* X *which is defined over a compositum of quadratic fields is modular.*

For an automorphism $u\colon X \to X$ to be modular it is necessary and sufficient that u preserves the set of cusps, so that u restricts to an automorphism of the non-cuspidal locus Y, and preserves the set of elliptic points, namely the branch points of the map $\mathbb{H} \to Y(\mathbb{C})$.

In Section 3.3 we see how to distinguish cusps, elliptic points and all the other points on X by looking at the action of Hecke operators T_{l_1}, T_{l_2}. More precisely, we look for multiple points in the divisors $T_{l_i}(x)$ for $x \in X(\mathbb{C})$: if x is a cusp, then $T_{l_i}(x)$ contains a point of multiplicity at least l_i; if $x = (E, \phi)$ is an elliptic point such that $j(E) = 0$, then $T_{l_i}(x)$ contains a point of multiplicity 3; if $x = (E, \phi)$ is an elliptic point such that $j(E) = 1728$, then $T_{l_i}(x)$ contains $\lfloor (l_i - 1)/2 \rfloor$ points of multiplicity 2.

These characterizations help proving the Lemma because of the following commutation rule in the group of divisors of X

$$uT_{l_i} = T_{l_i} u^{\sigma_i} \,,$$

where $\sigma_i \in G_{\mathbb{Q}}$ is a l_i-th Frobenius and $u\colon X \to X$ is supposed to be defined over a compositum of quadratic fields. The Eichler-Shimura relations imply the above equality in $\operatorname{Pic}^0(X)$ and the hypothesis on the gonality implies that the equality extends to the group of divisors of X.

In the last chapter we describe an algorithm to solve the discrete logarithm problem. Given a group G with a generator $g \in G$, solving the discrete logarithm problem means, for each element $h \in G$, computing an integer z such that $g^z = h$. The security of certain public-key cryptographic protocols depends on the hardness of this problem, depending on the choice of G. We are concerned with the cases where G is the multiplicative group of a finite field of *small characteristic*, which, for us, means a field of characteristic p and cardinality p^n for some integer $n > p$. The main result of the last chapter states that the discrete logarithm problem on finite fields of small characteristic is quasi-polynomial, hence not too hard.

Theorem. *There exists a probabilistic algorithm, described in Section 4.4, that solves the discrete logarithm problem in* $K^\times$ *for all finite fields* K *of small characteristic (namely the fields* $\mathbb{F}_{p^n}$ *with* $n > p$*) in expected time*

$$(\log \#K)^{O(\log\log \#K)} \,.$$

Our algorithm uses some ideas of the algorithm in [19], whose running time is only heuristic, and adapts them to finite fields with a different type of presentation. Let $\mathbb{F}_q$ be a finite field with $q > 2$ elements, let $E/\mathbb{F}_q$ be an elliptic curve and let P_1 be a point on E such that $\phi(P_1) - P_1 \in E(\mathbb{F}_q)$, where $\phi \colon E \to E$ is the q-th Frobenius. If $K = \mathbb{F}_q(P_1)$, then the coordinates of P_1 are generators of the extension $\mathbb{F}_q \subset K$ on which the q-th Frobenius acts "simply". If this happens and if, moreover $[K : \mathbb{F}_q] > 2$, the elliptic curve E and the point P_1 give an *elliptic presentation* of K. Given the abundance of elliptic curves over $\mathbb{F}_q$, for q big enough, it is easy to prove that every finite field of small characteristic can be embedded in a slightly larger field admitting an elliptic presentation such that q is small compared to $\#K$. A more precise statement is given in Proposition 4.1.5.

Given a finite field K with an elliptic presentation, we represent elements in $K^\times$ as $f(P_1)$ with f varying among the rational functions in $\mathbb{F}_q(E)$ that are regular and non-vanishing on P_1. Hence, we extend the discrete logarithm to these rational functions and, in a weak sense, to divisors on E. Notice that each divisor defined over $\mathbb{F}_q$ is a linear combination of irreducible divisors, namely those divisors that are the sum, with multiplicity 1, of all the $G_{\mathbb{F}_q}$-conjugates of a point in $E(\overline{\mathbb{F}_q})$.

Our algorithm is an index calculus using divisors: the idea is looking for linear relations among the discrete logarithm of h and the "discrete logarithms" of irreducible divisors of small degree; when many relations are found, we compute the discrete logarithm of h by solving a linear system.

We find relations using a descent procedure, which, given an irreducible divisor D of degree $4d \geq 320$, computes irreducible divisors D_i of degree dividing $2d$ such that the "discrete logarithm" of D is a linear combination of the "discrete logarithms" of the D_i's. Most of the last chapter is devoted to the description and the proof of the correctness of this descent procedure. It mainly uses the following equalities

$$f(P_1)^q = f^\phi(\phi(P_1)) = f^\phi(P_1 + P_0) = f^\phi \circ \tau_{P_0}(P_1)\,,$$

where $f \in \overline{\mathbb{F}_q}(E)$ is a function regular and non vanishing in P_1, the point $P_0 \in E(\mathbb{F}_q)$ is equal to $\phi(P_1) - P_1$, the map $f \to f^\phi$ is the automorphism on $\overline{\mathbb{F}_q}(E)$ that acts as the q-th Frobenius on $\overline{\mathbb{F}_q}$ and sends x and y to themselves. In Section 4.5 we see that, in order to compute the divisors D_i, it is sufficient to find a rational function f and a matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ satisfying certain conditions. After parametrizing the possible f's, this problem boils down to finding points in $\mathcal{C}(k)$, where $k \subset \overline{\mathbb{F}_q}$ is the extension of $\mathbb{F}_q$ of degree d and $\mathcal{C}$ is a variety of dimension at most 2 whose definition depends on D. We prove that $\mathcal{C}(k)$ is large using Weil's estimates. To prove that the geometrically irreducible components of $\mathcal{C}$ are defined over k, we use a little bit of Galois theory, condensed in Proposition 4.6.1, and some tedious computations, mostly contained in Proposition 4.6.3 and in the Claims 4.8.2.3, 4.8.2.6, 4.8.3.2.