




Polynesian Journal of Mathematics

Volume 4, Issue 8

Special volume AĴAX in honor of René Schoof

**Computing torsion for plane quartics
without using height bounds**

Raymond van Bommel

Received 31 Jan 2026

Revised 3 May 2026

Accepted 17 May 2026

Published 28 Sep 2026

Communicated by Samuele Anni

DOI: 10.69763/polyjmath.4.8

Computing torsion for plane quartics without using height bounds

Raymond van Bommel

Massachusetts Institute of Technology, United States
University of Bristol, United Kingdom

Abstract

We describe an algorithm that provably computes the rational torsion subgroup of the Jacobian of a curve without relying on height bounds. Instead, the strategy is to find upper bounds for the torsion subgroup using reduction modulo primes and to search for torsion points, not just over $\mathbb{Q}$ but also over small number fields, until the two bounds meet. Both complex analytic and Chinese remainder theorem based methods are used to find such torsion points. The method has been implemented in Magma for plane quartic curves over $\mathbb{Q}$ with a rational point and used to provably compute the rational torsion subgroup for more than 98% of Jacobians of curves in a data set due to Sutherland consisting of 82 240 plane quartic curves.

Keywords: Abelian varieties, Galois representations, Birch–Swinnerton-Dyer conjecture, Jacobians, curves.

1 Introduction

In the 1920s, Mordell and Weil proved that, for abelian varieties over a number field K , the group of rational points is finitely generated [20, 36]. In particular, the rational torsion subgroup is finite. The torsion conjecture asserts that there are only finitely many possible torsion groups, when the dimension of the abelian variety and the degree $[K : \mathbb{Q}]$ are fixed [25]. The conjecture has been proved in the case of elliptic curves, first over $\mathbb{Q}$ by Mazur [18] and finally for all number fields by Merel [19]. The exact determination of which groups can occur, when $K \neq \mathbb{Q}$, and which of them occur infinitely often, is still an active area of research. In higher dimensions, even in the case of abelian surfaces, no upper bound is known for the rational torsion subgroup of an abelian variety over $\mathbb{Q}$.

The torsion subgroup is also of importance for the Birch and Swinnerton-Dyer conjecture [2]. The order of the rational torsion subgroup is not only one of the terms appearing in the formula, the action of the absolute Galois group on the torsion is essential for the L -function, the central object in the conjecture. Indeed, the torsion is

used to define the Tate module and subsequently the L -function of an abelian variety over a number field is defined using the action of Frobenius on this Tate module.

In this present paper, we consider the question of explicitly computing the rational torsion subgroup in case the abelian variety is the Jacobian of a curve defined over $\mathbb{Q}$. Recently, such a computation has been done by Müller and Reitsma for hyperelliptic curves of genus 3 [21]. For small genus, the torsion is typically computed by doing some form of an exhaustive search. The Néron–Tate height $\widehat{h}$, or canonical height, of torsion points is known to be 0. Then one chooses a naïve height h , finds a bound $|h - \widehat{h}| < c$, and subsequently uses the fact that the naïve height of a torsion point is at most c to find all of them [30, 31]. Such height bounds are known and relatively small for genus 1, 2, and 3, but especially in genus 3 and higher the enumeration of all rational points up to these height bounds can still be challenging. Even though methods to compute canonical heights are available, even for non-hyperelliptic curve [4], these methods do not give rise to a way to efficiently enumerate points of bounded height. Therefore, we would like to advocate an alternative approach which has been used in an ad hoc fashion in the past and has also been used to acquire information on the Galois representation. The approach could be compared with that in [17] which is used to compute the whole Galois representation.

The approach uses the fact that, in practice, the torsion points actually seem to have a way smaller naïve height than the bound given by the height bounds, and are often not too hard to find. However, the problem then still remains to prove that one has found the complete rational torsion group. For this purpose, one could use the following lemma.

Lemma 1.1. *Let K be a number field, let $\mathfrak{p}$ be a prime of $\mathcal{O}_K$ over the prime number p , and let A be an abelian variety over K . Suppose that A has good reduction at $\mathfrak{p}$ and that $A_{\mathfrak{p}}$ is its reduction. Moreover, suppose that $p > e(\mathfrak{p}/p) + 1$, where $e(\mathfrak{p}/p)$ is the ramification index of $\mathfrak{p}$ over p . Then the natural reduction map*

$$A(K)[n] \longrightarrow A_{\mathfrak{p}}(\mathcal{O}_K/\mathfrak{p})$$

is injective for any integer n such that $p \nmid n$.

Proof. In [12, Appendix], there is a proof using formal groups. Alternatively, one can also use [24, Theorem 3.4.3] to show that the closure of any subgroup scheme $(\mathbb{Z}/p\mathbb{Z})_K \hookrightarrow A[p] \subset A$ inside a Néron model $\mathcal{A}/\mathcal{O}_{K,\mathfrak{p}}$ of A must be isomorphic to $(\mathbb{Z}/p\mathbb{Z})_{\mathcal{O}_{K,\mathfrak{p}}}$, which also proves the injectivity. $\square$

In many cases, it does not suffice to only use $K = \mathbb{Q}$ in this lemma to find sharp upper bounds for the torsion subgroup. For elliptic curves, just considering the greatest common divisor of the different orders $|A_p(\mathbb{F}_p)|$, one can only deduce the existence of an elliptic curve isogenous to A with that many torsion points, see for example [12]. One could refine the method by not just considering the orders $|A_p(\mathbb{F}_p)|$, but actually the group structure of $A_p(\mathbb{F}_p)$, see for example [21, Example 1.3] but even that is not enough. For example, it is easy to construct a hyperelliptic curve with nonrational 2-torsion points P_2, P_3 , and P_6 such that P_j is defined over $\mathbb{Q}(\sqrt{j})$ for $j = 2, 3, 6$. In this case, for any

prime p of good reduction, at least one of the three points will reduce to a point defined over $\mathbb{F}_p$, causing the reduction map to never be surjective on the $\mathbb{Q}$ -rational 2-torsion points. In this case, we say that the abelian variety has a *fake torsion point*.

The solution that we propose is to also search for torsion points defined over number fields of small degree to account for the nonsurjectivity of the reduction map. Later in this paper, we will introduce an algebraic and an analytic method to find such nonrational torsion points. The algebraic method, Algorithm 4.3, constructs nonrational torsion points from their reductions modulo different medium sized auxiliary primes, without assuming anything about the field of definition of the torsion points, using a technique that we call algebraic reconstruction (see Subsection 4.1). The analytic method, Algorithm 4.6, numerically inverts the Abel–Jacobi map to find new torsion points. Altogether, this leads to the following algorithm to provably compute the torsion subgroup.

Algorithm 1.2.

Input: a curve C defined over $\mathbb{Q}$.

Output: the rational torsion subgroup of $J = \text{Jac}(C)$.

1. Find a set of small primes $p_1, \dots, p_s$ satisfying the conditions of Lemma 1.1. For each p_i compute $N_i := |J_{p_i}(\mathbb{F}_{p_i})|$ (see Steps 2 and 3 of Algorithm 4.3), where J_{p_i} is the reduction of J modulo p_i .
2. For each divisor $\ell \mid \gcd\{N_i\}$, start with $T_\ell = \{0_J\}$, and keep applying Step 3 until the condition in Step 4 is satisfied.
3. Use Algorithms 4.3 and 4.6 with input $K = T_\ell \cap J(\mathbb{Q})$ to add new possibly nonrational torsion points to T_ℓ .
4. Check if all ℓ -power torsion can be explained by T_ℓ , i.e., if for each ℓ -power torsion point $P \in T_\ell \cap J(\mathbb{Q})$ there exists a prime p_i such that every $Q_{p_i} \in J_{p_i}(\mathbb{F}_{p_i})$ with $\ell \cdot Q_{p_i} = P \bmod p_i$ is the reduction of some element of T_ℓ at some prime lying above p_i whose ramification index is at most $p_i - 2$.
5. When the conditions in Step 4 are satisfied for each ℓ , output the group T generated by the rational torsion points of all the T_ℓ .

Theorem 1.3. *The output of Algorithm 1.2 is correct.*

Proof. Suppose there exists a torsion point $Q \in J(\mathbb{Q})$ that is not contained in T . Without loss of generality, we may assume that the order of Q is a power of a prime ℓ . By Lemma 1.1, the prime ℓ must divide all N_i and hence $\gcd\{N_i\}$. Without loss of generality, we may assume that $P := \ell \cdot Q$ lies in $T_\ell \cap J(\mathbb{Q})$. Now, by the condition in Step 4, there is another nonrational torsion point Q' in T_ℓ whose reduction modulo a prime lying above p_i equals that of Q . This contradicts the injectivity of the reduction map, as shown in Lemma 1.1. $\square$

Outline. In Section 2, we discuss the background needed for this approach: the Weil pairing, Weil polynomials, the Newton–Raphson method, and different ways to do Jacobian arithmetic. In Section 3, we study the phenomenon of fake torsion points. The core section of this paper in which we explain our actual methods to find torsion points over number fields is Section 4. In the final Section 5, we talk about the computation of the torsion groups in a data set consisting of 82 240 plane quartic curves [32]. The implementation of our method in Magma [5] for plane quartics can be found at [3].

Notation. Throughout this text, C denotes a smooth projective plane quartic curve over $\mathbb{Q}$, the symbol p denotes a prime number, J is the Jacobian of C , and C_p and J_p are the reduction of C and J , respectively, modulo p , when p is a prime of good reduction.

Acknowledgements. The author has been supported by the Simons Collaboration on Arithmetic Geometry, Number Theory, and Computation (Simons Foundation grant 550033) and by Céline Maistret’s Royal Society Dorothy Hodgkin Fellowship.

The author wishes to thank Edgar Costa, Maarten Derickx, Bjorn Poonen, David Roe, and Andrew Sutherland for useful discussions that helped to improve the method and article, and for their help running the parallel computation on the servers of the Simons Collaboration at the Massachusetts Institute of Technology. The author also wishes to thank the anonymous referees for the valuable suggestions which also led to improvements of the article.

2 Preliminaries

2.1 Weil pairing

We will recall the definition and properties of the Weil pairing. For more background we refer the reader to [29, Section III.8] or [22, Chapter 8].

For abelian varieties A over a field K , the *Weil pairing* is a bilinear pairing

$$w: A(\bar{K})[n] \times A^\vee(\bar{K})[n] \rightarrow \mu_n(\bar{K}),$$

where n is an integer not divisible by the characteristic of K and where $\bar{K}$ is an algebraic closure of K . In particular, in the case of a Jacobian J of a curve, when the theta divisor on J induces a principal polarisation $J \rightarrow J^\vee$, we get a pairing

$$w: J(\bar{K})[n] \times J(\bar{K})[n] \rightarrow \mu_n(\bar{K}).$$

This pairing, which we will also call the Weil pairing, is symplectic, i.e., alternating and nondegenerate.

Now consider the case where K is a number field and where $\mathfrak{p}$ is a prime of residue characteristic $p \nmid n$, together with J satisfying the conditions of Lemma 1.1. Then the Weil pairing is compatible with the reduction map modulo $\mathfrak{p}$, i.e., the following diagram is commutative.

$$\begin{array}{ccc}
J(\overline{K})[n] \times J(\overline{K})[n] & \longrightarrow & \mu_n(\overline{K}) \\
\downarrow & & \downarrow \\
J_p(\overline{\mathbb{F}}_p)[n] \times J_p(\overline{\mathbb{F}}_p)[n] & \longrightarrow & \mu_n(\overline{\mathbb{F}}_p)
\end{array}$$

Moreover, the action of the absolute Galois group $G_K := \text{Gal}(\overline{K}/K)$ respects the Weil pairing, i.e.,

$$\sigma(w(x, y)) = w(\sigma(x), \sigma(y))$$

for all $x, y \in J(\overline{K})[n]$ and all $\sigma \in G_K$. In particular the action of G_K must factor through the general symplectic group $\text{GSp}(J(\overline{K})[n], w)$ through elements with similitude character the n^{th} cyclotomic character over K . In the case $n = \ell$ is prime, the group $\text{GSp}(J(\overline{K})[n], w)$ can be identified with the classical general symplectic group $\text{GSp}(2g, \mathbb{F}_\ell)$, where g is the dimension of J .

2.2 Weil polynomials

We will recall the definition and properties of the Weil polynomial. For more background, we refer the reader to [29, Sections III.7 and V.2] and [33].

Let A be an abelian variety over $\mathbb{Q}$ and let p be an odd prime of good reduction. Then its reduction A_p is an abelian variety over $\mathbb{F}_p$ and for any prime $\ell \neq p$, we can consider the Tate module $V_\ell := \mathbb{Q}_\ell \otimes_{\mathbb{Z}_\ell} \varprojlim_n A_p[\ell^n](\overline{\mathbb{F}}_p)$. The Frobenius element in $\text{Gal}(\overline{\mathbb{F}}_p/\mathbb{F}_p)$ acts on this $\mathbb{Q}_\ell$ -vector space of dimension $2 \dim(A)$. Its characteristic polynomial P_{A_p} , the Weil polynomial, has coefficients in $\mathbb{Z}$, is independent of the choice of ℓ , and has the property that $\#A_p(\mathbb{F}_p) = P_{A_p}(1)$.

When $J = \text{Jac}(C)$ is the Jacobian of a curve C of which p is a prime of good reduction, the polynomial P_{J_p} can be computed by computing the characteristic polynomial of Frobenius on the étale cohomology group $H_{\text{ét}}^1(C, \mathbb{Q}_\ell)$. It is feasible to compute $P_{J_p} \bmod p$ for p of size about 10^6 in several minutes using algorithms described in [8], which generalises point counting algorithm for elliptic curves (see for example [26, 27]) although not in polynomial time in $\log p$, but rather in time $O(p)$, see [8]. Therefore, we will restrict ourselves to primes of the aforementioned size.

2.3 Newton–Raphson method and precision

For part of our computation, we will use complex valued numerical computations in order to try to find algebraic torsion points in $J(\overline{\mathbb{Q}})$. For this reason, we will briefly recall the numerical methods that we use, their stability, speed of convergence, and the loss of precision that might occur. More background on the methods can be found in many of the textbooks available on numerical analysis.

The main numerical method that we use is the Newton–Raphson method. It attempts to find a zero for a holomorphic function $f: \mathbb{C} \rightarrow \mathbb{C}$ by starting with some initial guess $x_0 \in \mathbb{C}$ for the root and iteratively computing the next approximation $x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$.

If x_0 is close enough to some simple root r of f , then the sequence (x_n) will converge quadratically to r .

However, when r is a root of multiplicity at least 2 (or in practice also when f has two simple roots very close to one another) problems may arise. As x_n gets closer to r , the denominator $f'(x_n)$ will get close to 0, requiring us to use a lot more precision to reliably compute the fraction $\frac{f(x_n)}{f'(x_n)}$. In principle, the method would still converge, but the following example explains why it is inevitable that we lose precision.

Example 2.1. Suppose we are looking for a root of $x^2 - 2x + 1$ close to the starting point $x_0 = 1.1$ and we are doing computations with 1000 digits of precision. Then we might end up finding the approximate root $\tilde{r} = 1 - 10^{-500}$. As $\tilde{r}^2 - 2\tilde{r} + 1 = 10^{-1000}$, this is a root within the precision of our computation, even though the actual root $r = 1$ lies at distance 10^{-500} . We lost half of our digits of precision.

If we now continue using $\tilde{r}$ instead of r and say that we have to find a root of $x^2 - 2x + \tilde{r}$, then we could end up with the approximate root $1 + 10^{-250}$ instead of the intended solution $r = 1$. So the loss of digits in the process can accumulate, as the output of a root finding algorithm can only be expected to have precision $\sim \sqrt[n]{\varepsilon}$ given input with precision $\sim \varepsilon$ around a root of multiplicity n .

The Newton–Raphson can also be used for multivariate functions $\mathbb{C}^n \rightarrow \mathbb{C}^n$. One has to replace $f'(x_n)$ by the Jacobian matrix J of the function evaluated at x_n . Here, problems arise when $J(r)$ has an eigenvalue of zero, in which case the inverse $J(x_n)^{-1}$ either does not exist, if $J(x_n)$ also has zero as an eigenvalue, or the computation of this inverse becomes numerically unstable as there is an eigenvalue very close to zero.

2.4 Jacobian arithmetic

For hyperelliptic curves, points on the Jacobian are typically represented using Mumford coordinates [7], which gives a unique way to represent each curve. For general curves, Khuri-Makdisi [13, 14, 15] developed ways to represent points on the Jacobian and do arithmetic. Flon, Oyono, and Ritzenthaler [9] describe a method specifically tailored to nonhyperelliptic genus 3 curves.

Many of these methods have been designed with the goal of implementing fast arithmetic over finite fields, which has potential applications in cryptography. Often these methods make assumptions on the curve that are easy to satisfy over finite fields, but not over number fields. For the purpose of our computation, we used two different methods to representing points on the Jacobian and do arithmetic. Even though these methods are most likely not state of the art in terms of their efficiency, we will describe them to inform the reader about the representations we used.

From now on C is a smooth plane quartic curve over $\mathbb{Q}$ and J is its Jacobian.

2.4.1 Over exact fields

Our goal will be to reconstruct points on the Jacobian from their reductions modulo different primes p . For this reason, we would like to represent points on the Jacobian in such a way that it has the following properties.

Definition 2.2. For this representation of points on the Jacobian, we define the two properties (α) and (β) as follows:

- (α) the representation of each point is unique;
- (β) for a prime p of good reduction and any point $P \in J(\mathbb{Q})$, there is a way to reduce the representation of P modulo p and, for all but finitely many of the primes p , this reduced representation is the unique representation for the reduction $\bar{P} \in J_p(\mathbb{F}_p)$ of P modulo p .

We will assume throughout that C has a $\mathbb{Q}$ -rational point. This is satisfied by at least 99% of the curves in [32]. The method would work essentially the same for a $\mathbb{Q}$ -rational divisor of degree 1. For example, consider the curve $x^4 + y^4 + z^4 - 5xz^3$ which has no points over $\mathbb{F}_5$ and hence no points over $\mathbb{Q}$ but does have the $\mathbb{Q}$ -rational divisor $\sum(\sqrt[4]{-1} : 1 : 0) - \sum(\sqrt[4]{5} : 0 : 1)$, where the sum is taken over the conjugates. In case C does not have $\mathbb{Q}$ -rational divisors of odd degree, there will be theoretical problems as not every $\mathbb{Q}$ -rational divisor class can be represented by a $\mathbb{Q}$ -rational divisor, which poses serious challenges. One would probably want to look at the linear algebra methods as described in [13]. We will not touch this subject any further and from now on assume the existence of a $\mathbb{Q}$ -rational point P on C .

Let D be any degree 0 divisor D on C . Then we know there is an integer $m \geq 0$ such that $h^0(D + mP) = 1$. Indeed for $m = 0$, this space has dimension at most 1 by Riemann–Roch, for $m \gg 0$, the space becomes higher dimensional, and each time we increase m the dimension increases at most by 1. If there are multiple m such that $h^0(D + mP) = 1$, we take the smallest one. Then looking at the divisor of any nonzero function $f \in H^0(D + mP)$, we find a way to represent D as $E - mP$, where E is an effective divisor.

Lemma 2.3. This representation, consisting of the integer $m \geq 0$ and an effective divisor E of degree m has property (α) (as in Definition 2.2), i.e., it is unique.

Proof. Indeed, the existence of a different effective divisor E' with the property that $E - mP \sim E' - mP \sim D$ would imply that $h^0(D + mP) \geq 2$. $\square$

Lemma 2.4. This representation has property (β) (as in Definition 2.2).

Proof. Let D be a divisor of degree 0 represented as $E - mP$ with $m \geq 0$ and E as above. Let N be the product of the primes of bad reduction of C . Then C has a smooth model $\bar{C}$ over $\text{Spec}(\mathbb{Z}[1/N])$ and we can take the closures $\bar{D}$ and $\bar{P}$ inside $\bar{C}$ of D and P , respectively. We consider the sheaves $\mathcal{F} := \mathcal{O}_{\bar{C}}(\bar{D} + m\bar{P})$ and $\mathcal{G} := \mathcal{O}_{\bar{C}}(\bar{D} + (m-1)\bar{P})$. We have that $\mathcal{F}(C_{\mathbb{Q}})$ and $\mathcal{G}(C_{\mathbb{Q}})$ are $\mathbb{Q}$ -vector spaces of dimensions 1 and 0, respectively. In particular, this implies that for all but finitely many p the $\mathbb{F}_p$ -vector spaces $\mathcal{F}(C_{\mathbb{F}_p})$ and $\mathcal{G}(C_{\mathbb{F}_p})$ have dimensions 1 and 0, respectively, see [11, Theorem III.12.8, page 288] for example. For these primes p , the representation of $D \bmod p$ has the same m and uses $E \bmod p$. $\square$

Remark 2.5. Note that the proof above does not provide a way to determine which primes p actually satisfy the conditions in property (β) . This does not pose a problem

in practice for our application as the primes chosen in Step 1 of Algorithm 1.2 are so large that they (almost) always seem to satisfy these conditions, at least for the curves in our dataset. We did not make any attempt to resolve this problem theoretically.

To add two points $E_1 - m_1P$ and $E_2 - m_2P$, one can use standard methods to compute a basis for the Riemann–Roch spaces $H^0(E_1 + E_2 + mP)$, finding the smallest m such that this space has dimension 1. Then we compute the divisor of a non-zero element $f \in H^0(E_1 + E_2 + mP)$. This divisor is linearly equivalent to 0 and is of the shape $E_{1,2} - E_1 - E_2 - mP$ where $E_{1,2}$ is effective. This gives rise to the representation $E_{1,2} - (m + m_1 + m_2)P$ for $E_1 + E_2 - (m_1 + m_2)P$. This can also be done directly using Magma’s built-in function `DivisorReduction`.

2.4.2 Over the complex numbers

From this point on, we assume C to have genus 3.

For a Zariski dense subset of divisor classes in $J(\mathbb{C})$, the representation described in the previous Subsection is of the shape $-3P + E$ with E effective of degree 3. Indeed, by Riemann–Roch $h^0(D + 3P) \geq 1$, so $m \leq 3$, while at the same time the space of divisor classes of the shape $-mP + E$ with $0 \leq m \leq 2$ and E effective of degree m has dimension at most 2 inside the 3-dimensional Jacobian.

Our (potential) torsion points, being special points on the Jacobian, quite regularly have a representation with $m < 3$. When we are doing numerical computations on $J(\mathbb{C})$, this often causes numerical instability for our algorithms. Luckily, there is an abundance of points on $J(\mathbb{C})$ and therefore we can use the following alternative presentation for elements of $J(\mathbb{C})$.

We represented them as $Q_1 + Q_2 + Q_3 - P_1 - P_2 - P_3$, where $P_1, P_2, P_3 \in C(\mathbb{C})$ are three arbitrary points that are chosen in advance and $Q_1, Q_2, Q_3 \in C(\mathbb{C})$ is a triple of points depending on the divisor class. We will now show why we have a practical guarantee that the set $\{Q_1, Q_2, Q_3\}$ will be unique for any divisor class that we encounter in our computation.

Proposition 2.6. *Let $D \in J(\mathbb{C})$ be any divisor class. Then for generic $P_1, P_2, P_3 \in C(\mathbb{C})$, the class D has a unique representation $Q_1 + Q_2 + Q_3 - P_1 - P_2 - P_3$, where two representations are called the same if the Q_i are the same up to reordering.*

Proof. If D has two such representations, this implies that $h^0(D + P_1 + P_2 + P_3) > 1$, as the associated Riemann–Roch space must have two functions with distinct zeros which are therefore linearly independent. The dimension $h^0(D + P_1 + P_2 + P_3)$ is upper semicontinuous as a function in P_1, P_2, P_3 by [11, Theorem III.12.8, page 288] applied to the map $f: C^4 \rightarrow C^3$ projecting away from the first coordinate and the sheaf $\mathcal{F} = \pi_1^* \mathcal{O}(D) \otimes \mathcal{O}(\Delta_1) \otimes \mathcal{O}(\Delta_2) \otimes \mathcal{O}(\Delta_3)$ where $\pi_1: C^4 \rightarrow C$ is the projection on the first coordinate, and $\Delta_i = \{(x_0, x_1, x_2, x_3) \in C^4 : x_0 = x_i\}$. By Riemann–Roch, we know that $h^0(D + P_1 + P_2 + P_3) \geq 1$. If $D + P_1 + P_2 + P_3 \sim 3Q$, where $Q \in C(\mathbb{C})$ is a non-Weierstraß point (i.e., a point P for which $h^0(mP) = 1$ for all $0 \leq m \leq g$, see for example [1, Chapter 1]), then this dimension is equal to 1. In particular, as $C(\mathbb{C})^3$ is irreducible, for all but a codimension 1 set of $(P_1, P_2, P_3) \in C(\mathbb{C})^3$, the dimension must equal 1 and the representation must be unique. $\square$

Remark 2.7. In fact, it follows that the map

$$\mathrm{Sym}^3 C \rightarrow J : (Q_1, Q_2, Q_3) \mapsto Q_1 + Q_2 + Q_3 - P_1 - P_2 - P_3$$

is smooth and injective outside of a subset of codimension 1. More details and alternative proofs can be found in [1, Chapter 1].

The representation also has the following useful property.

Proposition 2.8. *Let $D \in J(\mathbb{C})$ be a nonzero divisor class. Then for generic elements $P_1, P_2, P_3 \in C(\mathbb{C})$, the unique representation $Q_1 + Q_2 + Q_3 - P_1 - P_2 - P_3$ for D has the property that $\{Q_1, Q_2, Q_3\} \cap \{P_1, P_2, P_3\} = \emptyset$.*

Proof. It suffices to show that for generic $P_1, P_2 \in C(\mathbb{C})$, the class D is not equivalent to $Q_1 + Q_2 - P_1 - P_2$ for any $Q_1, Q_2 \in C(\mathbb{C})$. Equivalently, we need to show that $h^0(D + P_1 + P_2) = 0$ generically, as the existence of Q_1 and Q_2 with $Q_1 + Q_2 - P_1 - P_2 - D \sim 0$ is equivalent to the existence of a function f with $\mathrm{div}(f) + D + P_1 + P_2$ effective. Suppose that this is not the case, then this dimension must be at least 1 for any choice of P_1 and P_2 by the semicontinuity in [11, Theorem III.12.8, page 288]. In particular, for any distinct $P_1, P_2, P_3 \in C(\mathbb{C})$, we now have three ways of representing D :

$$Q_1 + Q_2 + P_3 - \sum_i P_i, \quad R_1 + P_2 + R_3 - \sum_i P_i, \quad \text{and} \quad P_1 + S_2 + S_3 - \sum_i P_i,$$

for certain $Q_1, Q_2, R_1, R_3, S_2, S_3 \in C(\mathbb{C})$. By the uniqueness of the representation, which we may assume for generic P_1, P_2, P_3 by Proposition 2.6, we now must have that $\{Q_1, Q_2, P_3\} = \{R_1, P_2, R_3\} = \{P_1, S_2, S_3\} = \{P_1, P_2, P_3\}$. In particular, $D = 0$, which is a contradiction. $\square$

Corollary 2.9. *Let $D_1, \dots, D_n \in J(\mathbb{C})$ be nonzero divisor classes. Then for all but at most a codimension 1 subset of $(P_1, P_2, P_3) \in C(\mathbb{C})^3$, the representations $Q_{i,1} + Q_{i,2} + Q_{i,3} - P_1 - P_2 - P_3$ of D_i satisfy the conditions in Propositions 2.6 and 2.8 for all $i = 1, \dots, n$.*

We will from now on assume that all divisor classes that we encounter in the computation of the torsion subgroup satisfy the conditions in Propositions 2.6 and 2.8.

To add two points, we use the following algorithm, which is a modified version of the algorithm in [9].

Algorithm 2.10.

Input: two triples of points Q_1, Q_2, Q_3 and R_1, R_2, R_3 representing points $Q = \sum_i Q_i - \sum_i P_i$ and $R = \sum_i R_i - \sum_i P_i$ on $J(\mathbb{C})$.

Output: a triple of points S_1, S_2, S_3 representing the point $Q + R = \sum_i S_i - \sum_i P_i$.

1. Pick (another) random point $B \in C(\mathbb{C})$.
2. Find the line ℓ through P_1 and P_2 , and compute the residual intersection A of this line with C , i.e., A is an effective divisor of degree 2 such that C intersects ℓ in $P_1 + P_2 + A$.
3. Find the cubic c through $Q_1, Q_2, Q_3, R_1, R_2, R_3, A$, and B , and compute the residual intersection E of this cubic with C , i.e., E is an effective divisor of degree 3 such that C intersects c in $\sum_i Q_i + \sum_i R_i + A + B + E$.

4. Find the conic n through B, P_3 , and E and compute the residual intersection S of this conic with C , i.e, S is an effective divisor of degree 3 such that C intersects n in $B + P_3 + E + S$.
5. Output the three points S_1, S_2 , and S_3 of which S consists.

Proposition 2.11. *The output of Algorithm 2.10 is correct.*

Proof. Consider the rational function $\frac{c}{\ell n}$. By construction, its associated principal divisor is

$$\begin{aligned} \left(\frac{c}{\ell n}\right) &= \sum_i Q_i + \sum_i R_i + A + B + E - P_1 - P_2 - A - B - P_3 - E - S \\ &= \sum_i Q_i + \sum_i R_i - \sum_i P_i - \sum_i S_i. \end{aligned}$$

In particular, we see that $\sum_i S_i - \sum_i P_i$ is equivalent to $\sum_i Q_i + \sum_i R_i - 2 \sum_i P_i$. $\square$

Remark 2.12. To find the residual intersection of a line/conic/cubic with f numerically, using the root finding algorithms described in Subsection 2.3, it is beneficial for these residual intersections to not have any double points, and no points in common with the rest of the divisor. In general, we expect the divisor $P_1 + P_2$ to not contain A and the divisor $\sum_i Q_i + \sum_i R_i + A + B$ to have no point in common with E , and E to not have any double points. Indeed, P_1 and P_2 were chosen generically, so ℓ should generically not be tangent to C . Similarly, Q_1, Q_2, Q_3 behave like generic points, because P_1, P_2, P_3 are generic. Note that Q could equal R . So c is either a cubic going through eight generic points on C , or a cubic going through five generic points on C , of which three with multiplicity 2. In both cases, the cubic c has no geometric reason to go through any of these generic points again, or through any other double point.

Because of this, the computation of A and E in Steps 2 and 3 is numerically stable and fast. In Step 4, there could be an issue when computing the residual divisor S . The divisor S could namely contain P_3 , but according to Proposition 2.8, this only happens in the case $Q + R = 0$. In all other cases, there is generally no double point and our algorithm to compute S will be numerically stable and fast.

Another way that we will use to represent points in $J(\mathbb{C})$ is by the means of an element in a complex torus $\mathbb{C}^3/\Lambda$. The computation of a period lattice Λ and an Abel–Jacobi map $\iota: J(\mathbb{C}) \rightarrow \mathbb{C}^3/\Lambda$ mapping $Q_1 + Q_2 + Q_3 - P_1 - P_2 - P_3$ to a corresponding point in the complex torus has been implemented in Magma by Neurohr, see also [23]. We will also write $\iota(Q_1, Q_2, Q_3)$ for $\iota(Q_1 + Q_2 + Q_3 - P_1 - P_2 - P_3)$.

In order to go back from a point in $\mathbb{C}^3/\Lambda$ to a divisor class, we use the following algorithm to invert the Abel–Jacobi map.

Algorithm 2.13.

Input: an element $x \in \mathbb{C}^3/\Lambda$.

Output: a triple of points $Q_1, Q_2, Q_3 \in C(\mathbb{C})$ such that $\iota(Q_1, Q_2, Q_3)$ is close to x .

1. Pick some integer n . We found that $n = 14$ worked well in practice for our examples.

2. After picking some local coordinate for C that is invertible around the points P_1, P_2 , and P_3 , the map ι , locally around (P_1, P_2, P_3) , can be considered as an analytic map $\mathbb{C}^3 \rightarrow \mathbb{C}^3$. Use Newton–Raphson (see Subsection 2.3) with starting point (P_1, P_2, P_3) to numerically approximate a solution to $\iota(Q_{1,n}, Q_{2,n}, Q_{3,n}) = \frac{1}{2^n} \cdot x$.
3. Add $Q_{1,n} + Q_{2,n} + Q_{3,n} - \sum_i P_i$ to itself using Algorithm 2.10. The output of this addition is an approximate solution to $\iota(Q_{1,n-1}, Q_{2,n-1}, Q_{3,n-1}) = \frac{1}{2^{n-1}} \cdot x$. We then use Newton–Raphson to increase the precision of this solution $(Q_{1,n-1}, Q_{2,n-1}, Q_{3,n-1})$. Decrease n by 1 and repeat this step until $n = 0$.
4. Use Newton–Raphson to refine $(Q_{1,0}, Q_{2,0}, Q_{3,0})$ to the desired precision and output the triple.

The reason for choosing an n and dividing by 2^n first, is to make sure that the starting point (P_1, P_2, P_3) in Step 2 is close enough to the solution for the Newton–Raphson method to actually converge. Because of the assumptions made and by Remark 2.12, the addition in Step 3 is expected to be numerically stable and to only introduce a moderate error. This means that the Newton–Raphson method in Steps 3 and 4 can reasonably be expected to converge again.

2.4.3 Changing base points

In Subsection 2.4.1, we worked with divisors represented as $E - 3P$ with E effective of degree 3, while in Subsection 2.4.2, we represented the same divisors as $Q_1 + Q_2 + Q_3 - P_1 - P_2 - P_3$. The following describes an algorithm, over $\mathbb{C}$ to convert the second representation back into the first representation. It is a modified version of Algorithm 2.10.

Algorithm 2.14.

Input: a triple of points Q_1, Q_2, Q_3 representing a point Q on $J(\mathbb{C})$ as $\sum_i Q_i - \sum_i P_i$.

Output: a triple of points S_1, S_2, S_3 representing the same point Q as $\sum_i S_i - 3P$.

1. Pick (another) random point $B \in C(\mathbb{C})$.
2. Find the line ℓ through P_1 and P_2 , and compute the residual intersection A of this line with C , i.e., A is an effective divisor of degree 2 such that C intersects ℓ in $P_1 + P_2 + A$.
3. Find the cubic c through Q_1, Q_2, Q_3, A , and B with multiplicity 1, and through P with multiplicity 3. Compute the residual intersection E of this cubic with C , i.e., E is an effective divisor of degree 3 such that C intersects c in $\sum_i Q_i + 3P + A + B + E$.
4. Find the conic n through B, P_3 , and E and compute the residual intersection S of this conic with C , i.e., S is an effective divisor of degree 3 such that C intersects n in $B + P_3 + E + S$.
5. Output the three points S_1, S_2 , and S_3 of which S consists.

Proposition 2.15. *The output of Algorithm 2.14 is correct.*

Proof. The proof is similar to that of Proposition 2.11. Consider the rational function $\frac{c}{\ell n}$. By construction, its associated principal divisor is

$$\begin{aligned} \left(\frac{c}{\ell n}\right) &= \sum_i Q_i + 3P + A + B + E - P_1 - P_2 - A - B - P_3 - E - S \\ &= \sum_i Q_i + 3P - \sum_i P_i - \sum_i S_i. \end{aligned}$$

In particular, we see that $\sum_i Q_i - \sum_i P_i$ is equivalent to $\sum_i S_i - 3P$. $\square$

Remark 2.16. What is said in Remark 2.12 still holds for Steps 1 and 2. The cubic in Step 3 does have a triple zero at P , but that is not a problem, because we already know what P is exactly, and we can use derivative functions to check if c actually goes through P four (or more) times. The same can be applied in Step 4 in case one or more of the S_i equals P . Note that the points in $B + P_3 + E$ behave like generic points and are generically not equal to one of the S_i . However, it could still happen that two of the S_i in Step 4 are equal to each other, but not equal to P . This is the only case in which I observed some significant precision loss occurring in the method, with the number of correct digits decreasing by a factor of 3 in the worst case.

3 Fake torsion points

Let $P \in J(\mathbb{Q})$ be a point and let ℓ be a prime number. We define

$$D_\ell(P) = \{Q \in J(\overline{\mathbb{Q}}) \mid \ell \cdot Q = P\} \quad \text{and} \quad D_{\ell,p}(P) = \{Q \in J_p(\overline{\mathbb{F}}_p) \mid \ell \cdot Q = \overline{P}\}$$

for every odd prime $p \neq \ell$ of good reduction. This is a torsor under the action of $J[\ell](\overline{\mathbb{Q}})$ or $J_p[\ell](\overline{\mathbb{F}}_p)$, respectively. We already saw in the introduction that it could happen that the set of $\mathbb{Q}$ -rational points in $D_\ell(P)$ is smaller than any of the sets of $\mathbb{F}_p$ -rational points in $D_{\ell,p}(P)$. In case this happens, we say that P has a *fake ℓ -divisor*.

In order to understand this phenomenon better, one considers the action of the absolute Galois group $\text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ on $D_\ell(P)$. Pick a not necessarily rational $Q_0 \in D_\ell(P)$ as base. Then for any $g \in G$ the action of g on D_ℓ can be described by

$$Q_0 + x \mapsto g(Q_0) + M \cdot x,$$

where $M \in \text{GL}(2g, \mathbb{F}_\ell)$. Because g has to respect the symplectic form on $J[\ell]$, the matrix M actually lies in $\text{Sp}(2g, \mathbb{F}_\ell)$, and we get a map $G \rightarrow \text{AGSp}(2g, \mathbb{F}_\ell)$. Let H be the image of this map. In the case $P = 0$, we actually have that H is a subgroup of the smaller group $\text{GSp}(2g, \mathbb{F}_\ell)$ since we can pick $Q_0 = 0$. (The same can be done if we can find a rational $Q_0 \in D_\ell(P)$, but the most interesting case for our method, is the case when such Q_0 does not exist.)

For each odd prime $p \neq \ell$ of good reduction, there is a conjugacy class Frob_p of H which describes how Frobenius acts on $D_{\ell,p}(P)$. Using these, we can exactly determine for which H the point P has a fake ℓ -divisor.

Proposition 3.1. *The point P has a fake ℓ -divisor if and only if for every element $h \in H$ we have*

$$D_\ell(P) \supset \text{Fix}(h) \supsetneq \text{Fix}(H) := \{x \in D_\ell(P) \mid \forall h \in H : h(x) = x\}.$$

Proof. The set $\text{Fix}(H)$ is exactly the set of $\mathbb{Q}$ -rational points in $D_\ell(P)$. For each odd prime $p \neq \ell$ of good reduction, the set of points in $D_\ell(P)$ reducing to an $\mathbb{F}_p$ -rational point in $D_{\ell,p}(P)$ is exactly $\text{Fix}(h)$ for some $h \in \text{Frob}_p$. Because of the Chebotarev density theorem, every conjugacy class will occur as Frob_p for some odd prime $p \neq \ell$, which concludes the proof of the proposition. $\square$

Looking at the group H , one can not only determine whether there is a fake torsion point, but also the degrees of the actual torsion points. By enumerating all the appropriate subgroups of $\text{AGSp}(2g, \mathbb{F}_\ell)$, we get the following result that shows that in certain cases the nonexistence of rational ℓ -divisors of P can be explained by points of degree at most 12.

Proposition 3.2. *Suppose that $g = 3$ and either $\ell = 2$, or both $\ell = 3$ and $P = 0$. Then there exist points $Q_1, \dots, Q_k \in D_\ell(P)$ such that $[\mathbb{Q}(Q_i) : \mathbb{Q}] \leq 12$ and a prime number p with the following properties. If $P \neq 0$, then $D_{\ell,p}(P) = \{Q_1 \bmod p, \dots, Q_k \bmod p\}$. If $P = 0$, then $D_{\ell,p}(P) = \langle Q_1 \bmod p, \dots, Q_k \bmod p \rangle$.*

Proof. This is a big group theoretic computation, enumerating all the appropriate subgroups of $\text{GSp}(6, \mathbb{F}_\ell)$ or $\text{AGSp}(6, \mathbb{F}_\ell)$, and figuring out the degrees of the fake torsion points needed. The code can be found at [3, extra/subgroups.m]. $\square$

4 Methods

In this section, we explain the main result of this paper: two methods to find torsion points over number fields. For the first method, we use the Chinese remainder theorem, taking torsion points modulo p_i for different primes p_i and trying to combine them into one torsion point over a number field. For the second method, we use a complex analytic approach, computing a complex approximation of torsion points up to high enough precision to reconstruct them algebraically. One could also imagine a third method, where one uses Hensel lifting to try to construct torsion points using methods from [17], but this approach has not been implemented as of now.

4.1 Algebraic reconstruction

Given a rational number $\alpha = \frac{r}{s}$ and its residue class modulo some suitable $N \gg \max(r^2, s^2)$, one could wonder if it is possible to construct α from this residue class. This question has been answered positively in [34, 35] with a fast algorithm using the Euclidean algorithm.

In this section, we will consider an algebraic number $\alpha \in \overline{\mathbb{Q}}$, its associated number field $K = \mathbb{Q}(\alpha)$ and prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_k$, such that $v_{\mathfrak{p}_i}(\alpha) \geq 0$ for $i = 1, \dots, k$. Then we can reduce α modulo each $\mathfrak{p}_i$ and we get finite field elements $\alpha_i \in \mathbb{F}_{\mathfrak{p}_i}$. The question

one can ask now is: can we reconstruct α from the α_i ? We will describe an algorithm that attempts to do this. Even though it is still practical for our purpose, the algorithm is definitely not as efficient as the rational reconstruction algorithm mentioned before.

For each i , let p_i be the residue class field characteristic of $\mathfrak{p}_i$, and let $f_i \in \mathbb{Z}[x]$ be a lift of the minimum polynomial of α_i over $\mathbb{F}_{p_i}$. Then we can consider the ideal $I_i = (f_i, p_i)$ of $\mathbb{Z}[x]$. The minimum polynomial f of α is an element of I_i for each i and hence of the intersection $I := \bigcap_i I_i$. The idea of our approach is to find a small element in I .

Algorithm 4.1.

Input: prime numbers p_i and polynomials f_i as described above.

Output: candidate minimum polynomial f for α .

1. Compute a Gröbner basis G for the ideal $I = \bigcap_i (f_i, p_i) \subset \mathbb{Z}[x]$.
2. Set $d := 1$, the degree of the candidate polynomial f that we are currently considering.
3. For each $g \in G$, compute $B_g^d := \{x^i \cdot g \mid i \in \mathbb{Z}_{\geq 0} \text{ such that } \deg(x^i \cdot g) \leq d\}$. Let $B^d := \bigcup_g B_g^d$ and $\Lambda^d \subset \mathbb{R}\{x^0, \dots, x^d\}$ be the lattice generated by B^d .
4. Find a short vector $f \in \Lambda^d$. Compute $|f|$, the maximum of the absolute values of the coefficients of f .
5. If $(2|f|)^{d+1}$ is significantly smaller than $\text{lcm}(\{p_i\})$ and $f \not\equiv 0 \pmod{p_i}$ for any i , then return f , otherwise set $d := d + 1$ and return to Step 3.

For Step 4 of the algorithm one could use any algorithm to find short vectors. In our implementation we used the LLL algorithm by Lenstra, Lenstra, and Lovász, see [16]. In Step 5, we do a heuristic check to see if the polynomial f that we are currently considering is small enough. For this purpose, we compare the number of polynomials of the same degree with coefficients of equal or smaller size with the product of the primes p over which we have information about $f \pmod{p}$. If the latter is much greater than the former, this suggests that the polynomial that we are currently considering might be the correct one.

Example 4.2. Suppose that $k = 2$, $p_1 = 1009$, $p_2 = 1019$, $f_1 = x - 55$ and $f_2 = x - 241$. Then we find $G = \{x + 635615, 1028171\}$. For $d = 1$, the shortest vector that we can find is $x - 392556$, which is too big to pass the test in Step 5. For $d = 2$, we find the short vector $x^2 + 2$, which we will output as f .

4.2 Finding torsion points: the CRT method

In this section, we will describe how to find torsion points using the Chinese remainder theorem. We assume that ℓ is prime and that we have some ℓ -power torsion point $Q \in J(\mathbb{Q})$. Our goal is to find points $R \neq 0$ such that $\ell R = Q$. In this section, all points will be represented using the representation described in Subsection 2.4.1. We will first give an outline of the method.

The idea is to look modulo different primes p_i for points R_{p_i} having the property that $\ell \cdot R_{p_i} = Q \bmod p_i$. These torsion points combined, give a point modulo the product $\prod p_i$, and using algebraic reconstruction (see Subsection 4.1), we attempt to find a torsion point R over a number field that reduces to the R_i . While every torsion point over a number field would theoretically be found eventually using this method, the sheer number of possible combination of torsion points to consider, might make this impractical. The algorithm, as described below, also uses several tricks in an attempt to keep this number of possible combinations as small as possible, so that we actually manage to get some non-trivial output in practice.

Algorithm 4.3.

Input: a prime number ℓ , a subgroup K of known torsion points inside $J[\ell](\mathbb{Q})$, and an ℓ -power torsion point $Q \in J(\mathbb{Q})$.

Output: a (possibly empty) list of nonzero points $R \in J(\overline{\mathbb{Q}})$ such that $\ell R = Q$.

1. Pick some medium size ($\approx 10^6$) auxiliary prime numbers $p_1, \dots, p_k$, such that C has good reduction at these primes.
2. For each p_i , compute the Weil polynomial P_{p_i} modulo p_i of the reduction J_{p_i} as described in Subsection 2.2. Using inequalities for the coefficients of P_{p_i} found in [10], construct a finite set B containing all the possible values of $N_i := \#J_{p_i}(\mathbb{F}_{p_i}) = P_{p_i}(1)$.
3. Take a random point $S \in J_{p_i}(\mathbb{F}_{p_i})$ and use a baby step giant step approach to identify all $b \in B$ such that $b \cdot S = 0$. Discard all other elements of B . Repeat this step until $\#B = 1$, which must mean that $B = \{N_i\}$.¹
4. For each p_i , decompose N_i as $\ell^{e_i} \cdot q_i$, where q_i has no factors ℓ . Then generate a bunch of random points S in $J_{p_i}(\mathbb{F}_{p_i})$ and compute $q_i \cdot S$, which is an element of $J_{p_i}[\ell^\infty]$. Keep finding new points, until there are enough points to generate the ℓ -power torsion of $J_{p_i}(\mathbb{F}_{p_i})$.
5. For each p_i , find the set D_{p_i} of points $R_i \in J_{p_i}(\mathbb{F}_{p_i})$ such that $\ell R_i = Q \bmod p_i$, and compute the image K_{p_i} of K inside $J_{p_i}(\mathbb{F}_{p_i})$. Discard some of the primes p_i for which the set D_{p_i} is relatively large.
6. For each finite set $I \subset \{1, \dots, k\}$ for which $D_I := \prod_{i \in I} D_{p_i} / K_{p_i}$ is not too large, enumerate representatives $(R_i)_{i \in I}$ for all elements of D_I and execute the next three steps for each such element. After finishing that, continue to Step 10.
7. For each $i \in I$ and $V \in K$ compute a representation

$$R_i + \overline{V} = m_{i,0} \overline{P} + \sum_{m=1}^{-m_{i,V}} R_{i,V,m}, \quad \text{where} \quad R_{i,V,m} \in C_p(\overline{\mathbb{F}}_p)$$

¹Even though it never happened in our dataset, theoretically, it could happen that $\#B$ never becomes equal to 1. This problem is studied in more detail in [28], in which there is a resolution to this problem called Step 3. For our purposes, it would always be possible to change the primes chosen in Step 1 to avoid this.

as in Subsection 2.4.1. If the multisets $\{m_{i,V} : V \in K\}$ are not all equal for the different $i \in I$, disregard this element of D_I . Otherwise, compute the polynomials $P_{x,i} = \prod_{m,R}(T - x(R_{i,V,m}))$ and $P_{y,i} = \prod_{m,R}(T - y(R_{i,V,m}))$ inside $\mathbb{F}_{p_i}[T]$.

8. Use algebraic reconstruction, as described in Subsection 4.1, to try to lift the matching coefficients of the $P_{x,i}$ and $P_{y,i}$ for the different $i \in I$ to elements of a number field. If the coefficients lift, and we get polynomials $P_x, P_y \in \overline{\mathbb{Q}}[T]$, apply the next step to them.
9. For all possible combinations of the roots of P_x and P_y see which ones give points on $C(\overline{\mathbb{Q}})$. Then try all combinations of m of these points to see if we can find an $R \in J(\overline{\mathbb{Q}})$ such that $\ell R = Q$. Use the Jacobian arithmetic described in Subsection 2.4.1 to verify this.
10. After finishing the loop described in Step 5, output all R with $\ell R = Q$ that we found in Step 9 during the computation.

The following example, which uses hyperelliptic curves instead of nonhyperelliptic curves and primes that are a bit smaller, demonstrates the method.

Example 4.4. Suppose C is the hyperelliptic curve $y^2 = x^7 + x^5 - 4x^3 - 4x$. Suppose that $K = \{0\}$, $\ell = 2$, $p_1 = 1009$ and $p_2 = 1019$. Over $\mathbb{F}_{p_1}$, we find 15 nonzero 2-torsion points, of which one is $(55, 0) - (0, 0)$. Over $\mathbb{F}_{p_2}$, we find 15 nonzero 2-torsion points, of which one is $(241, 0) - (0, 0)$. Using algebraic reconstruction on the x -coordinates of these divisors, as in Step 8, we quickly realise that these are the reductions of the torsion point $(\sqrt{-2}, 0) - (0, 0)$ on $\text{Jac}(C)(\mathbb{Q}(\sqrt{-2}))$, cf. Example 4.2.

Note that $\text{Jac}(C)$ actually has 8 rational 2-torsion points, so of the $15 \cdot 15 = 225$ combinations of 2-torsion points in Step 6, only $7 \cdot (1 \cdot 1) + 4 \cdot (2 \cdot 2) = 23$ of them are the reductions of a common point. If we had already found some of these 8 rational 2-torsion points, then the combinatorics in Step 6 would have been better.

Steps 1 through 4 are precomputation steps that only need to be done once for each curve. In most cases, the CRT method was the faster method to find torsion points over number fields. The biggest bottleneck of the method is the combinatorial explosion that can take place in Steps 6 through 9; the sets D_I can become very big in cases where there is a lot of fake torsion.

Remark 4.5. Let us give a very rough heuristic analysis of the method above through a fictive example in order to demonstrate this combinatorial explosion. Suppose that we found two points $P_1, P_2 \in J(\mathbb{Q})$ generating a subgroup isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ and we are looking for a fake 4-torsion point Q such that $2Q = P$. Suppose that Q is defined over a field K of degree 12 and that its coordinates have minimal polynomials with coefficients of size at most 100.

Note that there are about $200^{13} \approx 10^{30}$ polynomials of degree 12 with coefficients of size at most 100. In order to have a good chance to reconstruct Q from its reductions modulo different primes p_i in Step 8, we want the modulus $N = \prod_i p_i$ to be significantly greater than 10^{30} . That means that we should take at least 6 medium sized primes $p_1, \dots, p_6$.

We need Q to actually be defined over $\mathbb{F}_{p_i}$ rather than an extension field. The probability of a random prime p_i having this property is the same as the probability that a random element of the Galois group of K fixes one of the 12 roots of the minimal polynomial of the coordinates. In the worst case, this probability is $\frac{1}{12}$, meaning that we actually needed to pick about 72 medium sized primes and find the right combination of the 6 primes among the $\binom{72}{6} \approx 10^8$ possibilities.

For each prime p_i there are at least 4 candidate points in D_{p_i} , namely $\overline{Q}$, $\overline{Q + P_1}$, $\overline{Q + P_2}$, and $\overline{Q + P_1 + P_2}$. Of the at least $4^6 = 4096$ combinations of points we can pick, only 4 of them can be used to successfully reconstruct a point of degree 12. This part of the combinatorial explosion has been mitigated by the replacement of D_{p_i} by D_{p_i}/K_{p_i} in Step 6, but there is still the problem that some of the other fake torsion points will show up in some of the D_{p_i} .

4.3 Finding torsion points: the analytic method

The following analytic method to find torsion points has the advantage that there will be no combinatorial explosion of trying to combine torsion points modulo different primes into a torsion point over a number field. The downside is that we cannot utilise the fact that $J_{p_i}(\mathbb{F}_{p_i})[\ell^n]$ is typically much smaller than $J(\mathbb{C})[\ell^n]$. Recall that we assumed the existence of a point $P \in C(\mathbb{Q})$ and that we picked such a point at the start.

Algorithm 4.6.

Input: a prime number ℓ , a subgroup K of known torsion points inside $J[\ell](\mathbb{Q})$, and a point $Q \in J(\mathbb{Q})$ as described above.

Output: a (possible empty) list of nonzero points $R \in J(\overline{\mathbb{Q}})$ such that $\ell R = Q$.

1. Choose some $P_1, P_2, P_3 \in C(\mathbb{C})$ as in Subsection 2.4.2. Compute an Abel–Jacobi map $\iota: J(\mathbb{C}) \rightarrow \mathbb{C}^3/\Lambda$ with base points P_1, P_2, P_3 and compute the image $\iota(Q)$, using the methods in Magma implemented by Neurohr, see [23].
2. Pick an element t in each class in $(\frac{1}{\ell}\iota(Q) + \frac{1}{\ell}\Lambda)/\iota(K)$ and apply the following three steps for each element.
3. Use Algorithm 2.13 to find points $R'_1, R'_2, R'_3 \in C(\mathbb{C})$ such that $\iota(R'_1, R'_2, R'_3)$ is close to t . Use Algorithm 2.14 to write $\iota(R'_1, R'_2, R'_3)$ as $R_1 + R_2 + R_3 - 3P$ for some $R_1, R_2, R_3 \in C(\mathbb{C})$.
4. Compute Mumford-like coordinates for R , i.e., compute the product polynomial $P_x := \prod_i (T - x(R_i))$ in $\mathbb{C}[T]$ and a polynomial P_y of degree 2 such that $P_y(x(R_i)) = y(R_i)$.
5. Use a short lattice vector algorithm to try to find algebraic relations for the coefficients of P_x and P_y . If this succeeds, reconstruct the corresponding point in $J(\overline{\mathbb{Q}})$, which we call R_t .
6. After finishing the loop described in Step 2, output all R_t with $\ell R_t = Q$ that we found in Step 5 during the computation.

In practice, to recognise torsion points over number fields, we need several hundreds of digits of precision. This together with the sheer number of potential points we need to try (typically ℓ^6) makes the method slow in practice and only practical for $\ell = 2$ or $\ell = 3$.

5 Results

The algorithm has been implemented by the author in Magma and is publicly available at [3]. It has been run on a data set consisting of 82 240 plane quartic curves found by Andrew Sutherland, see [32]. As a result, the rational torsion subgroup has been computed successfully for 81 357 of the Jacobians of these curves. See Table 1. The total runtime for this computation, using version 2.25-7 of Magma, was approximately 8 core months and has been done in parallel, using at most 20 GB of memory per process, on a machine of the Simons Collaboration at Massachusetts Institute of Technology having a AMD EPYC 7713 CPU. For each computed torsion group a proof has been stored in the form of a list of primes, and a list torsion points over $\mathbb{Q}$ and over some number fields which together can be used to prove the completeness of the list of rational torsion points using Lemma 1.1. The vast majority of the time spent in Algorithms 4.3 and 4.6 was going through a huge number of potential torsion points over number fields. Hence, after these points have been found, the proofs can be verified significantly faster than it took to construct them. The proofs are available at [3, extra/proofs.tar.xz].

We also kept track of the number of cases in which we needed to find a fake torsion point in order to prove the upper bound for the torsion group. For 3 440 of the curves, we needed one or more fake torsion points. See Table 2.

For the majority of the 883 missing plane quartics, the reason that we could not compute their torsion subgroup was the failure to find a rational point on the curve. For 735 of these curves, we could verify the nonexistence of rational points by proving that there are no points over some local field. For the remaining 148 curves, which might give rise to counterexamples for the Hasse principle, we did not attempt to verify the nonexistence of rational points. This problem has recently been studied in more detail in [6].

To conclude this section we exhibit an example where we managed to find a torsion point over a degree 12 number field in order to certify the correctness of the computed rational torsion subgroup.

Example 5.1. Consider the smooth plane quartic $C: f = 0$ with

$$f = x^3y - xy^3 + y^4 + x^3z + 2x^2yz + 2xy^2z - y^3z + x^2z^2 + 2xyz^2 + y^2z^2 - 2xz^3 - yz^3 + z^4.$$

Its Jacobian J modulo 11 has 1772 points, and J modulo 67 has 274 944 points. As the primes 11 and 67 are both primes of good reduction, this implies that the torsion subgroup of J can have at most order $\gcd(1772, 274944) = 4$.

Besides 0, we find a second rational torsion point

$$T_2 = \left(\frac{1}{19}(-10\theta + 4) : -1 : 1 \right) + \left(\frac{1}{19}(-10\bar{\theta} + 4) : -1 : 1 \right) - 2 \cdot (1 : 0 : 0),$$

1	2	3	4	2,2	5	6	7	8	4,2	2,2,2	9
58 702	8 855	5 101	2 404	291	1 106	1 435	616	431	264	2	379
3,3	10	11	12	6,2	13	14	15	16	8,2	4,4	4,2,2
73	214	51	324	58	42	130	78	21	67	35	6
17	18	6,3	19	20	10,2	21	22	23	24	12,2	6,2,2
7	90	37	30	43	8	55	17	2	55	30	1
25	5,5	26	27	9,3	28	14,2	29	30	31	32	16,2
11	3	14	9	10	23	8	1	24	3	1	3
8,4	8,2,2	4,4,2	33	35	36	18,2	12,3	6,6	38	39	40
14	1	1	12	5	9	6	12	6	4	7	9
20,2	41	42	22,2	45	15,3	48	24,2	49	50	51	52
6	1	15	2	6	1	4	10	2	2	2	2
54	18,3	56	28,2	57	60	30,2	62	16,4	8,4,2	65	66
1	1	1	3	4	7	1	1	1	5	1	3
70	24,3	12,6	75	15,5	40,2	84	96	24,4	14,7	105	40,4
3	2	3	1	1	2	1	1	2	1	1	1

Table 1: The 96 different structures of the computed torsion groups of Jacobians of curves in the dataset [32]. The notation $n_1, n_2, \dots$ in the top row refers to the group $\mathbb{Z}/n_1\mathbb{Z} \times \mathbb{Z}/n_2\mathbb{Z} \times \dots$ while the number on the bottom row indicates how many curves yield this group structure.

none	deg. 2	deg. 3	deg. 4	deg. 6	deg. 8	deg. 9	deg. 10	deg. 12
77 917	1 386	191	478	52	70	1	10	1 217

Table 2: Number of curves for which our computation used fake torsion points, as a function of the maximum degree of those points.

where θ and $\bar{\theta}$ are zeros of $x^2 - \frac{27}{10}x - \frac{1713}{100}$. Looking at the reduction modulo 11, we easily find that there are no other points of order 2. After looking at a lot of different primes and seeing that T_2 has a 2-divisor modulo each of these primes, we suspect that T_2 might have a (fake) 2-divisor.

After about an hour of computation time, our program finds a torsion point T_4 over a degree 12 number field K defined by adjoining to $\mathbb{Q}$ a root of

$$x^{12} - 5x^{10} - 2x^9 - 20x^8 - 20x^7 + 7x^6 - 50x^5 + 26x^4 - 40x^3 - 58x^2 - 24x - 15.$$

This point satisfies $2T_4 = T_2$. As the prime 67 splits into four primes of residue degrees 67, 67, 67^2 , and 67^8 in the ring of integers of K , the point T_4 explains two of the 2-divisors of T_2 modulo 67. As there are only two 2-divisors of T_2 in $J \bmod 67$, we conclude that T_2 doesn't have a 2-divisor over $\mathbb{Q}$, and $\{0, T_2\}$ is the full torsion subgroup of J .

References

- [1] Enrico Arbarello, Maurizio Cornalba, Phillip A. Griffiths, and Joseph D. Harris. *Geometry of Algebraic Curves, Volume I*. Volume 267 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1985. ISBN: 9781441928252.
- [2] Bryan J. Birch and H. Peter F. Swinnerton-Dyer. Notes on elliptic curves II. *Journal für die reine und angewandte Mathematik*, 218:79–108, 1965. DOI: 10.1515/crll.1965.218.79.
- [3] Raymond van Bommel. genus3torsion. Magma package, 2023. <https://github.com/rbommel/genus3torsion>
- [4] Raymond van Bommel, David Holmes, and J. Steffen Müller. Explicit arithmetic intersection theory and computation of Néron–Tate heights. *Mathematics of Computation*, 89(321):395–410, 2020. DOI: 10.1090/mcom/3441.
- [5] Wieb Bosma, John Cannon, and Catherine Playoust. The Magma algebra system I: the user language. *Journal of Symbolic Computation*, 24(3–4):235–265, 1997. DOI: 10.1006/jSCO.1996.0125.
- [6] Nils Bruin and Brendan Creutz. Explicit Brauer–Manin obstructions on plane quartics. Preprint, 2026. DOI: 10.48550/arXiv.2601.16975.
- [7] David G. Cantor. Computing in the Jacobian of a hyperelliptic curve. *Mathematics of Computation*, 48(177):95–101, 1987. DOI: 10.1090/S0025-5718-1987-0866101-0.
- [8] Edgar Costa. *Effective computations of Hasse–Weil zeta functions*. PhD thesis, New York University, 2015. ISBN: 9781321953923.
- [9] Stéphane Flon, Roger Oyono, and Christophe Ritzenthaler. Fast addition on non-hyperelliptic genus 3 curves. In *Algebraic geometry and its applications*, volume 5 of *Number Theory and Its Applications*, pages 1–28. World Scientific, 2008. DOI: 10.1142/9789812793430_0001.

- [10] Safia Haloui. The characteristic polynomials of abelian varieties of dimensions 3 over finite fields. *Journal of Number Theory*, 130(12):2745–2752, 2010. DOI: 10.1016/j.jnt.2010.06.008.
- [11] Robin Hartshorne. *Algebraic geometry*. Volume 52 of *Graduate Texts in Mathematics*. Springer, 1977. ISBN: 9780387902449.
- [12] Nicholas M. Katz. Galois properties of torsion points on abelian varieties. *Inventiones mathematicae*, 62(3):481–502, 1981. DOI: 10.1007/BF01394256.
- [13] Kamal Khuri-Makdisi. Linear algebra algorithms for divisors on an algebraic curve. *Mathematics of Computation*, 73(245):333–357, 2004. DOI: 10.1090/S0025-5718-03-01567-9.
- [14] Kamal Khuri-Makdisi. Asymptotically fast group operations on Jacobians of general curves. *Mathematics of Computation*, 76(260):2213–2239, 2007. DOI: 10.1090/S0025-5718-07-01989-8.
- [15] Kamal Khuri-Makdisi. On Jacobian group arithmetic for typical divisors on curves. *Research in Number Theory*, 4(1):3, 2018. DOI: 10.1007/s40993-018-0101-6.
- [16] Arjen K. Lenstra, Hendrik W. Lenstra, jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982. DOI: 10.1007/BF01457454.
- [17] Nicolas Mascot. Hensel-lifting torsion points on Jacobians and Galois representations. *Mathematics of Computation*, 89(323):1417–1455, 2020. DOI: 10.1090/mcom/3484.
- [18] Barry Mazur. Modular curves and the Eisenstein ideal. *Publications Mathématiques de l’IHÉS*, 47:33–186, 1977. DOI: 10.1007/BF02684339.
- [19] Loïc Merel. Bornes pour la torsion des courbes elliptiques sur les corps de nombres. *Inventiones mathematicae*, 124(1–3):437–449, 1996. DOI: 10.1007/s002220050059.
- [20] Louis J. Mordell. On the rational solutions of the indeterminate equations of the third and fourth degrees. *Proceedings of the Cambridge Philosophical Society*, 21:179–192, 1922.
- [21] J. Steffen Müller and Berno Reitsma. Computing torsion subgroups of Jacobians of hyperelliptic curves of genus 3. *Research in Number Theory*, 9(2):23, 2023. DOI: 10.1007/s40993-023-00429-x.
- [22] David Mumford. *Abelian varieties*. Volume 5 of *Tata Institute of Fundamental Research Studies in Mathematics*. Oxford University Press, 1970.
- [23] Christian Neurohr. *Efficient integration on Riemann surfaces & applications*. PhD thesis, Carl von Ossietzky Universität Oldenburg, 2018. <https://oops.uni-oldenburg.de/3607/1/neueff18.pdf>.

- [24] Michel Raynaud. Schémas en groupes de type $(p, \dots, p)$. *Bulletin de la Société Mathématique de France*, 102:241–280, 1974. DOI: 10.24033/bsmf.1779.
- [25] Norbert Schappacher and René Schoof. Beppo Levi and the arithmetic of elliptic curves. *Mathematical Intelligencer*, 18(1):57–69, 1996.
- [26] René Schoof. Elliptic curves over finite fields and the computation of square roots mod p . *Mathematics of Computation*, 44(170):483–494, 1985. DOI: 10.1090/S0025-5718-1985-0777280-6.
- [27] René Schoof. Counting points on elliptic curves over finite fields. *Journal de Théorie des Nombres de Bordeaux*, 7(1):219–254, 1995. DOI: 10.5802/jtnb.142.
- [28] Jia Shi. Lifting L -polynomials of genus 3 curves. Preprint, 2026. DOI: 10.48550/arXiv.2602.00965.
- [29] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*. Volume 106 of *Graduate Texts in Mathematics*. Springer, 2009. ISBN: 9780387094939.
- [30] Michael Stoll. On the height constant of genus two II. *Acta Arithmetica*, 104(2):165–182, 2002. DOI: 10.4064/aa104-2-6.
- [31] Michael Stoll. An explicit theory of heights for hyperelliptic Jacobians of genus three. In *Algorithmic and Experimental Methods in Algebra, Geometry, and Number Theory*, pages 665–715. Springer, 2018. DOI: 10.1007/978-3-319-70566-8_29.
- [32] Andrew V. Sutherland. A database of nonhyperelliptic genus 3 curves over $\mathbb{Q}$. In *Proceedings of the Thirteenth Algorithmic Number Theory Symposium*, volume 2 of *The Open Book Series*, pages 443–459. Mathematical Sciences Publishers, 2019. DOI: 10.2140/obs.2019.2.443.
- [33] John Tate. Endomorphisms of abelian varieties over finite fields. *Inventiones mathematicae*, 2:134–144, 1966. DOI: 10.1007/BF01404549.
- [34] Paul S. Wang. A p -adic algorithm for univariate partial fractions. In *Proceedings of the fourth ACM symposium on Symbolic and algebraic computation*, pages 212–217. Association for Computing Machinery, 1981. DOI: 10.1145/800206.806398.
- [35] Paul S. Wang, Michael J. T. Guy, James H. Davenport. P -adic reconstruction of rational numbers. *ACM SIGSAM Bulletin*, 16(2):2–3, 1982. DOI: 10.1145/1089292.1089293.
- [36] André Weil. L’arithmétique sur les courbes algébriques. *Acta Mathematica*, 52:281–315, 1929. DOI: 10.1007/BF02592688.