

The Mordell–Weil sieve: proving non-existence of rational points on curves

Nils Bruin and Michael Stoll

ABSTRACT

We discuss the Mordell–Weil sieve as a general technique for proving results concerning rational points on a given curve. In the special case of curves of genus 2, we describe quite explicitly how the relevant local information can be obtained if one does not want to restrict to mod p information at primes of good reduction. We describe our implementation of the Mordell–Weil sieve algorithm and discuss its efficiency.

Supplementary materials are available with this article.

1. Introduction

The Mordell–Weil sieve uses knowledge about the Mordell–Weil group of the Jacobian variety of a curve, together with local information (obtained by reduction mod p , say, for many primes p), to obtain strong results on the rational points on the curve.

The most obvious application, which also provided the original motivation for this work, is the possibility of verifying that a given curve does not have any rational points. This is done by deriving a contradiction from the various bits of local information, using the global constraint that a rational point on the curve must map into the Mordell–Weil group. This idea is simple enough (see Section 2), but its implementation in the form of an algorithm that will run in reasonable time on a computer is not completely straightforward. The relevant algorithms are discussed in Section 3, and our concrete implementation is described in Section 7. Section 8 contains a discussion of the efficiency of the implementation and gives some timings.

The idea of using this kind of ‘Mordell–Weil sieve’ computation to prove that a given curve does not have rational points appeared for the first time in the thesis [23] of Scharaschkin, who used it in a few examples involving twists of the Fermat quartic. It was then taken up by Flynn [14] in a more systematic study of genus 2 curves; his selection of examples was somewhat biased, however (in favor of curves he was able to compute with). In our ‘small curves’ project [5], we applied the procedure systematically and successfully to all genus 2 curves

$$y^2 = f_6x^6 + \cdots + f_1x + f_0$$

with $f_i \in \{-3, -2, -1, 0, 1, 2, 3\}$ that do not possess rational points.

In this situation, it is not strictly necessary to know a full generating set of the Mordell–Weil group. It is sufficient to know generators of a finite-index subgroup whose index is coprime to a certain set of primes. This can be checked again by using only local information. In fact, the necessary information is usually part of the input for the sieve procedure. This remark is relevant since one needs to be able to compute canonical heights and to enumerate points on the Jacobian up to a given bound for the canonical height if one wants to obtain generators for the full Mordell–Weil group. The necessary algorithms are currently available only for curves of genus 2; see [25, 27]. Nevertheless, we can use the Mordell–Weil sieve to show that there

Received 15 June 2009; revised 30 November 2009.

2000 Mathematics Subject Classification 11D41, 11G30, 11Y50 (primary), 14G05, 14G25, 14H25, 14H45, 14Q05 (secondary).

Research by the first author was supported by NSERC.

are no rational points on a given curve, even when the genus is 3 or more. Of course, we still need to know the Mordell–Weil rank and the right number of independent points. See [22] for an example where this technique is applied with a curve of genus 3 to show that there are no rational points satisfying certain congruence conditions.

The approach can be modified so that it can be used to verify that there are no rational points satisfying a given set of congruence conditions or mapping into a certain coset in the Mordell–Weil group. This is what was used in [22]. If we can show in some way that, moreover, in each of the cosets or residue classes under consideration there can be at most one rational point, then this provides a way of determining the set of rational points on the curve. Specifically, if a given coset or residue class contains a rational point, we will eventually find it, and we then also know that there are no other rational points in this coset or class; on the other hand, if there is no rational point in this coset or residue class, then we can hope to verify this by an application of the Mordell–Weil sieve. In this situation, the above remark that it is sufficient to know a finite-index subgroup still applies.

There is one case in which we can actually prove that for a suitable choice of prime p , no residue class mod p on the curve can contain more than one rational point. This is the ‘Chabauty situation’, where the Mordell–Weil rank is less than the genus. We can (hope to) find a suitable p , and then we can (hope to) determine the rational points on our curve as outlined above. This yields a procedure whose termination is not (yet) guaranteed, since it relies on some conjectures. However, the procedure itself is correct: if it terminates, and it has done so in all the examples we tried, then it gives the exact set of rational points on the curve. In the Chabauty context, the sieving idea has already been used in [4] to rule out the presence of rational points in certain cosets. See also [22] for some more examples and [3] for an example that uses ‘deep’ information.

Even when the rank is too large to apply the idea we just mentioned, the sieve can still be used to show that any rational point on the curve which we have not already found must be astronomically huge. This provides at least some kind of moral certainty that there are no other points. In conjunction with (equally huge) explicit bounds for the size of *integral* points, this allows us to show that we know at least all the integral points on our curve; see [8]. For this application, however, we really need to know the full Mordell–Weil group; so, with currently available technology, this is restricted to curves of genus 2.

We discuss these various applications in some detail in Section 4.

In Sections 5 and 6, we discuss how to extract local information that can be used for the sieve, when we do not want to restrict ourselves to just information mod p for primes p of good reduction. In these sections, we assume that the curve is of genus 2 and that we are working over \mathbb{Q} .

As to the theoretical background, we remark here that under a mild finiteness assumption on the Shafarevich–Tate group of the curve’s Jacobian variety, the information that can be obtained via the Mordell–Weil sieve is equivalent to the Brauer–Manin obstruction; see [23] or [29].

2. The idea

Let C/\mathbb{Q} be a smooth projective curve of genus $g \geq 2$ with Jacobian variety J . (In [30, 31] we consider, more generally, a subvariety of an abelian variety; the idea is the same, however.)

Our goal is to show that a given curve C/\mathbb{Q} does not have rational points. For this, we consider the following commuting diagram, where v runs through the (finite and infinite) places of \mathbb{Q} .

$$\begin{array}{ccc}
 C(\mathbb{Q}) & \xrightarrow{\iota} & J(\mathbb{Q}) \\
 \downarrow & & \downarrow \alpha \\
 \prod_v C(\mathbb{Q}_v) & \xrightarrow{\iota} & \prod_v J(\mathbb{Q}_v)
 \end{array}$$

We assume that we know an embedding $\iota : C \rightarrow J$ defined over \mathbb{Q} (that is, we know a \mathbb{Q} -rational divisor class of degree one on C) and that we know generators of the Mordell–Weil group $J(\mathbb{Q})$. If $C(\mathbb{Q})$ is empty, then the images of α and the lower ι are disjoint, and conversely.

However, since the sets and groups involved are infinite, we are not able to compute this intersection. Therefore, we replace the groups by finite approximations. Let S be a finite set of places of \mathbb{Q} and let $N \geq 1$ be an integer. Then we consider the following diagram.

$$\begin{array}{ccc} C(\mathbb{Q}) & \xrightarrow{\iota} & J(\mathbb{Q})/NJ(\mathbb{Q}) \\ \downarrow & & \downarrow \alpha \\ \prod_{v \in S} C(\mathbb{Q}_v) & \xrightarrow{\beta} & \prod_{v \in S} J(\mathbb{Q}_v)/NJ(\mathbb{Q}_v) \end{array}$$

Under the assumptions made, now we can compute the images of α and β and check whether they are disjoint. If $C(\mathbb{Q}) = \emptyset$, then, according to [29, Main Conjecture] and the heuristic given in [21], the two images should be disjoint when S and N are large enough. Note that (as shown in [29]) the two images will be disjoint for some choice of S and N if and only if $\prod_v \iota(C(\mathbb{Q}_v))$ does not meet the topological closure of $J(\mathbb{Q})$ in $\prod_p J(\mathbb{Q}_p) \times J(\mathbb{R})/J(\mathbb{R})^0$, where $J(\mathbb{R})^0$ denotes the connected component of the origin. This is a stronger condition than the requirement that $\prod_v \iota(C(\mathbb{Q}_v))$ miss the image of $J(\mathbb{Q})$. The conjecture claims that the two statements are, in fact, equivalent.

As a further simplification, we can just use a set S of primes of good reduction and replace the above diagram by the following simpler one.

$$\begin{array}{ccc} C(\mathbb{Q}) & \xrightarrow{\iota} & J(\mathbb{Q})/NJ(\mathbb{Q}) \\ \downarrow & & \downarrow \alpha \\ \prod_{p \in S} C(\mathbb{F}_p) & \xrightarrow{\beta} & \prod_{p \in S} J(\mathbb{F}_p)/NJ(\mathbb{F}_p) \end{array}$$

Poonen originally formulated his heuristic for this case. However, in practice it appears to be worthwhile to use also ‘bad’ information (coming from primes of bad reduction) and ‘deep’ information (involving parts of the kernel of reduction) to keep the running time of the actual sieve computation within reasonable limits. In Sections 5 and 6 below, we show how to obtain ‘bad’ and ‘deep’ information for curves of genus 2 over \mathbb{Q} .

3. Algorithms

In what follows, we will work with the simpler version of diagram involving only reduction mod p , as described at the end of Section 2.

Let r denote the rank of the Mordell–Weil group $J(\mathbb{Q})$. For a given set S and a parameter N , denote by $A(S, N) \subset J(\mathbb{Q})/NJ(\mathbb{Q})$ the subset of elements that map into the image of $C(\mathbb{F}_p)$ in $J(\mathbb{F}_p)/NJ(\mathbb{F}_p)$ for all $p \in S$; in symbols,

$$A(S, N) = \{a \in J(\mathbb{Q})/NJ(\mathbb{Q}) : \alpha(a) \in \text{im}(\beta_{N,p}) \text{ for all } p \in S\}.$$

Here, $\beta_{N,p} : C(\mathbb{F}_p) \rightarrow J(\mathbb{F}_p)/NJ(\mathbb{F}_p)$ denotes the composition of $\iota : C(\mathbb{F}_p) \rightarrow J(\mathbb{F}_p)$ and the canonical epimorphism $J(\mathbb{F}_p) \rightarrow J(\mathbb{F}_p)/NJ(\mathbb{F}_p)$.

The procedure splits into three parts.

(1) *Choice of S .*

In the first step, we have to choose a set S of primes such that we can be reasonably certain that the combined information obtained from reduction mod p for all $p \in S$ is

sufficient to give a contradiction (or, more generally, to have $A(S, N)$ equal to the image of $C(\mathbb{Q})$, for suitable N). In § 3.1, we explain a criterion that tells us whether S is likely to be good for our purposes. The actual computation of the relevant local information is also part of this step. For each prime $p \in S$, we find the abstract finite abelian group G'_p representing $J(\mathbb{F}_p)$ (or some other finite quotient of $J(\mathbb{Q}_p)$) and the image $X'_p \subset G'_p$ of $\iota : C(\mathbb{F}_p) \rightarrow J(\mathbb{F}_p)$. We also compute the homomorphism $\phi_p : J(\mathbb{Q}) \rightarrow G'_p$. We let G_p denote the image of ϕ_p and write $X_p = X'_p \cap G_p$.

In what follows, we will use $\#X_p/\#G_p$ as a measure of how much information about rational points on C can be obtained at p . Note that it is possible to have $G_p \subset X'_p \subsetneq G'_p$. In that case, $\#X'_p/\#G'_p < 1$, but no element of the Mordell–Weil group can be precluded from coming from $C(\mathbb{Q})$, based on the information at p . If we were to use this quantity, we would obtain erroneous estimates in the second step. This can then lead to huge sets $A(S, N)$ in the third step and even to a failure of the computation.

(2) *Choice of N .*

In the second step, we fix a target value of N and determine a way to compute $A(S, N)$ efficiently. We do so by finding an ordered factorization $N = q_1 q_2 \cdots q_m$ such that none of the intermediate sets $A(S, q_1 \cdots q_k)$ becomes too large. This is explained in § 3.2.

(3) *Computation of $A(S, N)$.*

Finally, we have to actually compute $A(S, N)$ in a reasonably efficient way. We explain in § 3.3 how this can be done.

The last two steps can be considered independently of the Mordell–Weil sieve context. Basically, we need a procedure that, given a finite family of surjective group homomorphisms $\phi_i : \Gamma \rightarrow G_i$ and subsets $X_i \subset G_i$, $i \in I$, attempts to prove that for every $a \in \Gamma$ there is some $i \in I$ such that $\phi_i(a) \notin X_i$. Here Γ is a finitely generated abelian group and the G_i are finite abelian groups. In our application, Γ is the Mordell–Weil group, the index set is S , G_p is the image of $J(\mathbb{Q})$ in $J(\mathbb{F}_p)$, and $X_p = \iota(C(\mathbb{F}_p)) \cap G_p$.

We give some more details of our actual implementation in Section 7.

3.1. Choice of S

The first task of the algorithm is to come up with a suitable set S of places. We will restrict to finite places (that is, primes), but in principle one could also include information at infinity, which would mean considering the connected components of $J(\mathbb{R})$ that meet the image of $C(\mathbb{R})$ under the embedding ι .

It is clear that the only possibility of getting some interaction between the information at various primes p (and eventually a contradiction) is when the various group orders $\#J(\mathbb{F}_p)$ have common factors. This is certainly more likely when these common factors are relatively small. We therefore look for primes p (of good reduction) such that the group order $\#J(\mathbb{F}_p)$ is B -smooth (that is, with all prime divisors no greater than B), for some fixed value of B ; in practice, values such as $B = 100$ or $B = 200$ lead to good results.

For each such prime, we compute the group structure of $J(\mathbb{F}_p)$, that is, an abstract finite abelian group G'_p together with an explicit isomorphism $J(\mathbb{F}_p) \cong G'_p$. We also compute the images of the generators of $J(\mathbb{Q})$ in G'_p and the image of $C(\mathbb{F}_p)$ in G'_p . In order to do so, we need to solve roughly p discrete logarithm problems in G'_p . Since G'_p has smooth order, we can use Pohlig–Hellman reduction [20] to reduce to a number of small discrete log problems. Therefore, in practice this part of the computation is essentially linear in p . We do need to compute reasonably efficiently in $J(\mathbb{F}_p)$, though. If C is a curve of genus 2, Cantor reduction [9] gives us a way of doing that. To fix notation, let W denote an effective canonical divisor on C . Cantor reduction takes as input a degree-zero divisor in the form $D - dW$, where D is

an effective divisor of degree $2d$, and computes a unique divisor D_0 of degree two such that

$$[D - dW] = [D_0 - W],$$

with the convention that if $D - dW$ is principal, then $D_0 = W$. Adding two divisor classes $[D_1 - W]$ and $[D_2 - W]$ can be accomplished by feeding the divisor $(D_1 + D_2) - 2W$ into the reduction algorithm.

Cantor reduction also allows us to map elements from $C(\mathbb{F}_p)$ into $J(\mathbb{F}_p)$. If ι is given by a rational base point $P_0 \in C(\mathbb{Q})$ (that is, $\iota(P) = [P - P_0]$) and \bar{P}_0 is the reduction of P_0 modulo p , then for each $\bar{P} \in C(\mathbb{F}_p)$ we have $\iota(\bar{P}) = [\bar{P} + \bar{P}'_0 - \bar{W}]$, where \bar{P}'_0 is the hyperelliptic involute of \bar{P}_0 . In this case, we already get $\iota(\bar{P})$ as a reduced divisor class. Otherwise, ι is given by $\iota(P) = [P - D_3 + W]$, where D_3 is a rational effective divisor of degree three; then we can compute a reduced representative of $\iota(\bar{P})$ by performing Cantor reduction on $(\bar{P} + \bar{D}'_3) - 2W$.

As mentioned above, we finally replace G'_p by $G_p = \phi_p(J(\mathbb{Q}))$ and let C_p be the intersection of the image of $C(\mathbb{F}_p)$ in G'_p with G_p . We then use ϕ_p to denote the surjective homomorphism $\phi_p : J(\mathbb{Q}) \rightarrow G_p$.

In order to determine whether we have collected enough primes, we compute the expected size of the set $A(S, N)$, where S is the set of all primes p collected so far and N is a suitable value which we will specify below. We follow Poonen [21] and assume that the images of the $C(\mathbb{F}_p)$ in $J(\mathbb{F}_p)$ are random and independent for the various p . This leads to the expected value

$$n(S, N) = \#(J(\mathbb{Q})/NJ(\mathbb{Q})) \prod_{p \in S} \frac{\#C_{N,p}}{\#(G_p/NG_p)}$$

where $C_{N,p}$ is the image of C_p in G_p/NG_p .

In principle, we would like to find the value of N that minimizes $n(S, N)$ for the given set S . However, this would lead to much too involved a computation. We therefore propose to proceed as follows. Write

$$\prod_{p \in S} J(\mathbb{F}_p) \cong \mathbb{Z}/N_1\mathbb{Z} \times \mathbb{Z}/N_2\mathbb{Z} \times \cdots \times \mathbb{Z}/N_l\mathbb{Z}$$

where N_j divides N_{j+1} for each $j = 1, 2, \dots, l - 1$. Then we take $N = N_{l-r-1-j}$, $j = 0, 1, 2, 3$, as values that are likely to produce a small $n(S, N)$. The reason for this choice is the following. Usually the target groups will be essentially cyclic, and the kernel of the homomorphism $J(\mathbb{Q}) \rightarrow J(\mathbb{F}_p)$ will be a random subgroup of index $\#J(\mathbb{F}_p)$ and more or less cyclic quotient. If we take a prime number q for N and the Mordell–Weil rank is r , then we obtain a random codimension-one subspace of \mathbb{F}_q^r . Unless q is very small, it is rather unlikely for these subspaces to intersect in a non-trivial way, unless there are more than r of them. So, for every prime power dividing our N , we want to have more than r factors in the product above that have order divisible by the prime power. Thus we should restrict to divisors of N_{l-r-1} . Taking $N_{l-r-1-j}$ with $j > 0$ ensures that we will get even more independent factors.

By the same token, any subgroup $L \subset J(\mathbb{Q})$ such that we can expect to get sufficient information on the image of $C(\mathbb{Q})$ in $J(\mathbb{Q})/L$ will be very close to $NJ(\mathbb{Q})$ for some N : as soon as the various bits of information interact, we will have exhausted all ‘directions’ in the dual of $J(\mathbb{Q})/NJ(\mathbb{Q})$, and the intersection of the kernels of the relevant maps will be close to $NJ(\mathbb{Q})$. This also explains why our approach to the computation of $A(S, N)$, which will be described below in § 3.3, works quite well.

Note that by taking S (and perhaps also B) large, we will get large values for the number l of factors. Once $l \gg r$, the image of the Mordell–Weil group $J(\mathbb{Q})$ in this product will be rather small, so that we can expect it to eventually miss the image of the curve. Poonen’s heuristic [21] makes this argument precise.

We continue collecting primes into S until we find a sufficiently small $n(S, N)$. In practice, it appears that $n(S, N) < \varepsilon = 10^{-2}$ is sufficient. Note that if the final sieve computation is unsuccessful (and does not lead to the discovery of a rational point on C), then we can enlarge S until $n(S, N)$ gets small enough and repeat the sieve computation.

3.2. *Choice of N*

Once S is chosen and the relevant information is computed, we can forget about the original context and consider the following more abstract situation.

We are given a finitely generated abstract abelian group Γ of rank r , together with a finite family $(G_i, \phi_i, X_i)_{i \in I}$ of triples, where each G_i is a finite abstract abelian group, $\phi_i : \Gamma \rightarrow G_i$ is a surjective homomorphism and $X_i \subset G_i$ is a subset. In practice, each of Γ and G_i is given as a product of cyclic groups, ϕ_i is given by the images of the generators of Γ , and X_i is given by enumerating its elements. The following definition generalizes $A(S, N)$.

DEFINITION 3.1. Let $L \subset \Gamma$ be a subgroup of finite index. We set $G_{L,i} = G_i/\phi_i(L)$, write $X_{L,i}$ for the image of X_i in $G_{L,i}$, and denote by $\phi_{L,i}$ the induced homomorphism $\Gamma/L \rightarrow G_{L,i}$. We let

$$A(L) = \{\gamma \in \Gamma/L : \phi_{L,i}(\gamma) \in X_{L,i} \text{ for all } i \in I\}$$

and define its expected size to be

$$n(L) = \#(\Gamma/L) \prod_{i \in I} \frac{\#X_{L,i}}{\#G_{L,i}}.$$

Now the task is as follows.

PROBLEM 3.2.

- (i) Find a number N such that $A(N\Gamma)$ has a good chance of being empty and such that $A(N\Gamma)$ can be computed efficiently.
- (ii) Compute $A(N\Gamma)$.

In our application, $\Gamma = J(\mathbb{Q})$, $I = S$, G_p and ϕ_p for $p \in S$ are as before, and $X_p = C_p$.

Since we may have to take N fairly large ($N \approx 10^6$ is not uncommon, and values on the order of 10^{12} or even 10^{100} do occur in practice in our applications), it would not be a good idea to enumerate the roughly N^r elements of $\Gamma/N\Gamma$ and check whether each satisfies the conditions. Instead, we build up N multiplicatively in stages: we compute $A(N_j\Gamma)$ successively for a sequence of values

$$N_0 = 1, \quad N_1 = q_1, \quad N_2 = N_1q_2, \quad N_3 = N_2q_3, \quad \dots, \quad N_m = N_{m-1}q_m = N,$$

where the q_k are the prime divisors of N . We want to choose the sequence (q_k) , and therefore N , in such a way that the intermediate sets $A(N_k\Gamma)$ are likely to be small. For this, we use again the expected size $n(N_k\Gamma)$ of $A(N_k\Gamma)$. By a best-first search, we find the sequence $(q_k)_{k=1, \dots, m}$ such that:

- (i) $n((\prod_{k=1}^m q_k)\Gamma)$ is less than a target value $\varepsilon_1 < 1$ (for example, 0.1); and
- (ii) $\max\{n(N_k\Gamma) : 0 \leq k \leq m\}$ is minimal, where $N_k = \prod_{j=1}^k q_j$.

From the first step, which provides the input, we can deduce a number M such that all reasonable choices for N should divide M (usually we take $M = N_{l-1-r-j}$ for some small value of j , in the notation used above). The following procedure returns a suitable sequence (q_1, \dots, q_m) .

FindQSequence:

```

c := {(), 1, 1.0} // () is an empty sequence of q_k, 1 is N, 1.0 is n(NΓ)
while c ≠ ∅:
  (s, N, n) := triple in c with minimal n
  remove this triple from c
  if n < ε: // success?
    return s
  end if
  // compute the possible extensions of s and add them to the list
  c := c ∪ {(append(s, q), Nq, n(NqΓ)) : q prime, Nq|M}
end while
// if we leave the while loop here, the target was not reached
return 'failure'

```

When we extend c , we can restrict to triples (s', N', n') such that N' does not occur as the second component of a triple already in c (since in this case we have already found a ‘better’ sequence leading to this N').

If the information given by $(G_i, \phi_i, X_i)_{i \in I}$ is sufficient (as determined in the first step), then this procedure usually does not take much time (compared to the computation of ‘local information’ such as the image of $C(\mathbb{F}_p)$ in $J(\mathbb{F}_p)$). In any case, if we made sure in the first step that there is some M such that $n(M\Gamma) < \varepsilon_1$, then FindQSequence will not fail.

In this step and also in the first step, it is a good idea to keep the orders of the cyclic factors of the groups G_i and the numbers N in factored form, and convert only the greatest common divisors of N with the relevant group orders into actual integers.

3.3. Computation of $A(N\Gamma)$

Now we have fixed the sequence $(q_k)_{j=1, \dots, m}$ of primes whose product is N . In the final stage of the algorithm, we have to compute the set $A(N\Gamma)$ (and hope to find that it is empty or sufficiently small, depending on the intended application).

This is done iteratively, by successively computing $A(N_k\Gamma)$ where $N_k = \prod_{j=1}^k q_j$. We start at $k = 0$ and initialize $A(N_0\Gamma) = A(\Gamma) = \{0\} \subset \Gamma/\Gamma$. Then, assuming we know $A(N_{k-1}\Gamma)$, we compute $A(N_k\Gamma)$ as follows.

We first find the triples (G_i, ϕ_i, X_i) that can possibly provide new information. The relevant condition is $v_{q_k}(e_i) \geq v_{q_k}(N_k)$, where e_i is the exponent of the group G_i . For these i , we compute the group $G_{N_k\Gamma, i}$, the image $X_{N_k\Gamma, i}$ of X_i in this group and the homomorphism $\phi_{N_k\Gamma, i} : \Gamma/N_k\Gamma \rightarrow G_{N_k\Gamma, i}$.

The most obvious approach now would be to take each $\gamma \in A(N_{k-1}\Gamma)$, run through its various lifts to $\Gamma/N_k\Gamma$ and check whether each lift is mapped into $X_{N_k\Gamma, i}$ under $\phi_{N_k\Gamma, i}$. The complexity of this procedure is $\#A(N_{k-1}\Gamma) \cdot q_k^r$ times the average number of tests we have to make (we disregard possible torsion in Γ , which will not play a role once N_{k-1} is large enough). Unless r is very small, the procedure will be rather slow when the intermediate sets $A(N_k\Gamma)$ get large.

In order to improve on this, we split the inclusion $N_k\Gamma \subset N_{k-1}\Gamma$ into several stages:

$$N_{k-1}\Gamma = L_0 \supset L_1 \supset \dots \supset L_t = N_k\Gamma.$$

Note that the quotient $N_{k-1}\Gamma/N_k\Gamma$ is isomorphic to $(\mathbb{Z}/q_k\mathbb{Z})^r$ (again disregarding torsion in Γ), so we can hope to get up to r intermediate steps. We now proceed as follows.

```

PrepareLift( $k$ ):
 $j := 0$ ;  $L_0 := N_{k-1}\Gamma$  // initialize
 $I' := \{i \in I : v_{q_k}(e_i) \geq v_{q_k}(N_k)\}$  // the relevant subset of  $I$ 
while  $I' \neq \emptyset$  do
   $j := j + 1$ 
  // list the possible subgroups for the next step
   $\Lambda := \{L_{j-1} \cap \ker(\phi_i) : i \in I'\}$ 
  for  $L \in \Lambda$  do
    // compute a measure of how 'good' each subgroup is
     $n(L_{j-1}, L) := (L_{j-1} : L) \prod_{i \in I'} \frac{\#X_{L,i}}{\#X_{L_{j-1},i}} \frac{1}{(\phi_i(L_{j-1}) : \phi_i(L))}$ 
  end for
   $L_j :=$  the  $L \in \Lambda$  that has the smallest  $n(L_{j-1}, L)$ 
   $I_j := \{i \in I' : \phi_i(L_j) \neq \phi_i(L_{j-1})\}$  // record the  $i \in I'$  that contribute to this step
   $I' := \{i \in I' : \phi_i(L_j) \not\subset N_k G_i\}$  // update  $I'$ 
end while
if  $L_j \neq N_k\Gamma$ :
  // fill the remaining gap to  $N_k\Gamma$ 
   $t := j + 1$ ;  $L_t := N_k\Gamma$ ;  $I_t := \emptyset$ 
else
   $t := j$ 
end if
    
```

The quantity $n(L_{j-1}, L)$ computed in the preceding algorithm is the expected number of ‘offspring’ that an element of $A(L_{j-1})$ generates in $A(L)$.

We then successively compute $A(L_1), \dots, A(L_t) = A(N_k\Gamma)$ in the same way as described above for the one-step procedure:

```

Lift( $k$ ):
// note that  $A(L_0) = A(N_{k-1}\Gamma)$ 
for  $j = 1, \dots, t$  do
   $A(L_j) := \emptyset \subset \Gamma/L_j$ 
  for  $a \in A(L_{j-1})$  do
     $a' :=$  a representative of  $a$  in  $\Gamma/L_j$ 
    for  $l \in L_{j-1}/L_j$  do
      if  $\forall i \in I_j : \phi_{L_j,i}(a' + l) \in X_{L_j,i}$ :
         $A(L_j) := A(L_j) \cup \{a' + l\}$ 
      end if
    end for
  end for
end for
// now  $A(N_k\Gamma) = A(L_t)$ 
return
    
```

In practice, `PrepareLift` and `Lift` together form one subroutine, whose input is $(N, q, A) = (N_{k-1}, q_k, A(N_{k-1}\Gamma))$ (along with the global data Γ and $(G_i, \phi_i, X_i)_{i \in I}$) and whose output is $A(Nq\Gamma)$ (with $Nq = N_k$).

The complexity of the lifting step is now

$$\sum_{j=1}^t \#A(L_{j-1})(L_{j-1} : L_j) \approx \#A(N_{k-1}\Gamma) \sum_{j=1}^t (L_{j-1} : L_j) \prod_{i=1}^{j-1} n(L_{i-1}, L_i).$$

In the worst case, we have $n(L_{j-1}, L_j) = (L_{j-1} : L_j)$; then the second factor cannot be larger than $q_k + q_k^2 + \dots + q_k^r < (q_k/(q_k - 1))q_k^r$; this is not much worse than the factor q_k^r we had before. Usually, however, and in particular when N_{k-1} is already fairly large, the numbers $n(L_{j-1}, L_j)$ will be much smaller than $(L_{j-1} : L_j)$; also, we should have $t = r$ and $(L_{j-1} : L_j) = q_k$, so that the complexity is essentially $\#A(N_{k-1}\Gamma)q_k$. As an additional benefit, we distribute the tests we have to make over the intermediate steps, so that the average number of tests in the innermost loop will be smaller than when going directly from $N_{k-1}\Gamma$ to $N_k\Gamma$.

In this way, it is possible to compute these sets even when r is not very small. For example, in order to find the integral solutions of $\binom{y}{2} = \binom{x}{5}$ (see [8]), it was necessary to perform this kind of computation for a group of rank six, and this was only made possible by our improvement of the lifting step. As another example, one of the two rank-four curves that had to be dealt with by the Mordell–Weil sieve in our experiment [5] took the better part of a day with the implementation we had at the time, which was based on the ‘obvious approach’ mentioned above. Now, with the new method, this computation takes less than 15 minutes.

If we find that $A(N_k\Gamma) = \emptyset$ for some $k \leq m$, then we stop. In the context of our application, this means that we have proved $C(\mathbb{Q}) = \emptyset$ as well. Otherwise, we can check whether the remaining elements in $A(N\Gamma)$ actually come from rational points, by computing the element of $J(\mathbb{Q})$ of smallest height that is in the corresponding coset. It is usually a good idea to first do some more mod p checks so that one can be certain that the point in $J(\mathbb{Q})$ really gives rise to a point in $C(\mathbb{Q})$. If we do not find a rational point on C in this way, then we can increase S and decrease ε and ε_1 and repeat the computation.

Let us also remark here that the lifting step can easily be parallelized, since we can compute the ‘offspring’ of the various $a \in A(N_{k-1}\Gamma)$ independently. After the preparatory computation in `PrepareLift` has been done, we can split $A(N_{k-1}\Gamma)$ into a number of subsets and give each of them to a separate thread to compute the resulting part of $A(N_k\Gamma)$. Then we collect the results and check to see if the new set is empty; if it is not, we repeat this procedure with the next lifting step.

4. Applications

4.1. Non-existence of rational points

The main application we had in mind (and, in fact, the motivation for developing the algorithm described in this paper) is in the context of our project on deciding the existence of rational points on all ‘small’ genus 2 curves; see the report [5].

Out of about 200 000 isomorphism classes of curves initially, 1492 turned out to be undecided after a search for rational points, a check for local points and a 2-descent [6]. We applied our algorithm to these curves and were able to prove for all of them that they do not have rational points. For some curves, we needed to assume the Birch and Swinnerton-Dyer conjecture for the correctness of the rank of the Mordell–Weil group.

For the curves whose Jacobians have rank at most two, we originally used only ‘good’ and ‘flat’ information, that is, groups $J(\mathbb{F}_p)$ for primes p of good reduction. For ranks three and four (no higher ranks occur), we also used ‘bad’ and ‘deep’ information, as described in Sections 5 and 6 below. The running time of the MAGMA implementation of the Mordell–Weil sieve algorithm we had at the time was about one day for all 1492 curves (on a 1.7 GHz machine with 512 MB of RAM). Two-thirds of that time was taken by one of the two rank-four curves, and most of the remaining time was used for the 152 rank-three curves.

With the current implementation discussed in Section 7 below, the overall running time (now on a 2.0 GHz machine with 4 GB of RAM) is about two and a half hours. For a detailed discussion of the timings, see Section 8.

4.2. Finding points

Besides proving that no rational points on C exist, we can also use the Mordell–Weil sieve idea to find rational points on C up to very large height. When the rank is less than the genus, we can even combine the Mordell–Weil sieve with Chabauty’s method in order to compute the set of rational points on C exactly; see § 4.4 below.

We want to find the rational points on C up to a certain (large) logarithmic height bound H . We assume that the height-pairing matrix for the generators of $J(\mathbb{Q})$ and a bound for the difference between naive and canonical height on $J(\mathbb{Q})$ are known. See [25, 27] for algorithms that provide these data in the case of genus 2 curves. From this information and the embedding $C \rightarrow J$, we can then compute constants δ and d such that $\hat{h}(\iota(P)) \leq dh(P) + \delta$ for all points $P \in C(\mathbb{Q})$. Here \hat{h} denotes the canonical height on $J(\mathbb{Q})$ and h denotes a suitable height function on the curve. The upshot of this is that $h(P) \leq H$ implies $\hat{h}(\iota(P)) \leq H' = dH + \delta$.

Note that in many cases where we want to find all rational points up to height H , we already know a rational point P_0 on C . Then we can just use $P \mapsto [P - P_0]$ for the embedding ι .

We now proceed as before: we find a suitable set S of primes together with a number N and compute $A(S, N) \subset J(\mathbb{Q})/NJ(\mathbb{Q})$. For the purposes of this application, we require N to be divisible by the exponent of the torsion group $J(\mathbb{Q})_{\text{tors}}$ and to be such that $N^2 > 4H'/m$ where m is the minimal canonical height of a non-torsion point in $J(\mathbb{Q})$. These conditions imply that if $Q, Q' \in J(\mathbb{Q})$ are such that $Q - Q' \in NJ(\mathbb{Q})$ and $\hat{h}(Q), \hat{h}(Q') \leq H'$, then $Q = Q'$. In other words, each coset of $NJ(\mathbb{Q})$ in $J(\mathbb{Q})$ contains at most one point of canonical height at most H' .

We do not necessarily expect $A(S, N)$ to be empty now. However, by the preceding discussion, each element of $A(S, N)$ corresponds to at most one point in $C(\mathbb{Q})$ of height at most H . Therefore we consider the elements of $A(S, N)$ in turn (we expect them to be few in number), and for each of them we do the following. First, we check whether there is an element Q in the corresponding coset of $NJ(\mathbb{Q})$ such that $\hat{h}(Q) \leq H'$. If this is not the case, we discard the element; otherwise, there is only one such Q , and we check some more primes $p \notin S$ for whether the image of Q in $J(\mathbb{F}_p)$ is in the image of $C(\mathbb{F}_p)$. Note that we can perform these tests quickly only based on the representation of Q as a linear combination of the generators of $J(\mathbb{Q})$: we reduce the generators mod p and compute the reduction of Q as a linear combination of the reduced generators. Depending on H' , we can determine such a set of primes beforehand, with the property that a point $Q \in J(\mathbb{Q})$ with $\hat{h}(Q) \leq H'$ that ‘survives’ all these tests must be in $\iota(C(\mathbb{Q}))$; see the lemma below. So if Q fails one of the tests, we discard it; otherwise, we compute Q as an explicit point and find its preimage in $C(\mathbb{Q})$ under ι .

LEMMA 4.1. Let $P_0 \in C(\mathbb{Q})$ and write $x(P_0) = (a : b)$ with coprime integers a and b . Let p_1, p_2, \dots, p_m be primes of good reduction with

$$p_1 p_2 \cdots p_m > e^{H'+\gamma} \max\{|a|, |b|\}^2$$

such that P_0 and its hyperelliptic conjugate \bar{P}_0 are distinct mod some p_{j_0} if they are distinct in $C(\mathbb{Q})$. Here γ is a bound for the difference $h - \hat{h}$ between naive and canonical height on $J(\mathbb{Q})$. We take $\iota : P \mapsto [P - P_0]$.

If $Q \in J(\mathbb{Q})$ satisfies $\hat{h}(Q) \leq H'$ and is such that the reduction of Q mod p_j is in $\iota(C(\mathbb{F}_{p_j}))$ for all $1 \leq j \leq m$, then $Q \in \iota(C(\mathbb{Q}))$.

Proof. Let $(k_1 : k_2 : k_3 : k_4)$, with coprime integers k_j , be the image of Q on the Kummer surface of J . If Q mod p_j is on the image of the curve, then p_j divides $k_1 b^2 - k_2 ab + k_3 a^2$. This integer has absolute value at most $e^{H'+\gamma} \max\{|a|, |b|\}^2$, so if it is divisible by p_1, \dots, p_m , it must be zero. This implies that $Q = [P - P_0]$ or $Q = [P - \bar{P}_0]$ for some $P \in C(\mathbb{Q})$. If $P_0 \neq \bar{P}_0$, these two cases can be distinguished mod p_{j_0} . □

Testing whether a given coset of $NJ(\mathbb{Q})$ contains a point of canonical height at most H' comes down to a ‘closest vector’ computation with respect to the lattice $(NJ(\mathbb{Q}), \hat{h})$. Depending on the efficiency of this operation, we can start eliminating elements from $A(S, N_k)$ already at some earlier stage of the computation of $A(S, N)$, thus reducing the effort needed for subsequent stages of the procedure.

If we want to reach a very large height bound, then we should at some point switch over to the variant of the sieving procedure described in § 4.3 below.

Of course, there is a simpler alternative, which is to enumerate all lattice points in $(J(\mathbb{Q})/J(\mathbb{Q})_{\text{tors}}, \hat{h})$ of norm H' or less and then checking all corresponding points in $J(\mathbb{Q})$ for whether they are in the image of ι . (For this test, one conveniently uses reduction mod p again, for a suitable set of primes p .) Which of the two methods will be more efficient depends on the curve in question and on the height bound H . If the curve is fixed, then we expect our Mordell–Weil sieve method to be more efficient than the short-vectors enumeration when H gets large. The reason for this is that once S and N are sufficiently large, the set $A(S, N)$ is expected to be uniformly small (most of its elements should come from rational points on C), and so the computation of $A(S, N)$ for large N will not take much additional time. On the other hand, the number of vectors of norm H' or less will grow like a power of H' , and the enumeration will eventually become infeasible.

4.3. Integral points on hyperelliptic curves

What the preceding application really gives us is a lower bound H for the logarithmic height of any rational point that we do not know (and therefore believe does not exist). If we can produce such a bound in the order of $H = 10^k$ with k in the range of several hundred, then we can combine this information with upper bounds for integral points that can be deduced by using linear forms in logarithms and thus determine the set of integral points on a hyperelliptic curve: if $C : y^2 = f(x)$ is a hyperelliptic curve over \mathbb{Q} , then it is possible to compute an upper bound $\log |x| \leq H$ that holds for integral points $(x, y) \in C$, where H is usually of a size comparable to that mentioned above. See [8, §§ 3–9].

With the procedure we have described here, it is feasible to reach values of N in the range of 10^{100} , corresponding to $H \approx 10^{200}$. However, this is usually not enough: the upper bounds provided by the methods described in [8] are more like 10^{600} . The part of the computation that dominates the running time is the computation of the image of $C(\mathbb{F}_p)$ in the abstract finite abelian group representing $J(\mathbb{F}_p)$. To close the gap, we therefore switch to a different sieving strategy that avoids having to compute all these roughly p discrete logarithms in $J(\mathbb{F}_p)$. We assume that we know a subgroup $L \subset J(\mathbb{Q})$ (initially this is $NJ(\mathbb{Q})$) such that the image of $C(\mathbb{Q})$ in $J(\mathbb{Q})/L$ is given by rational points we already know on C . We then try to find a smaller subgroup L' with the same property. Let q be a prime of good reduction, and recall the notation $\phi_q : J(\mathbb{Q}) \rightarrow J(\mathbb{F}_q)$ for the reduction homomorphism. Let $W \subset J(\mathbb{Q})$ be the image of the known rational points on C , let $L' = L \cap \ker \phi_q$, and take $R \subset J(\mathbb{Q})$ to be a complete set of representatives of the *non-trivial* cosets of L' in L . We can now check for each $w \in W$ and $r \in R$ whether $\phi_q(w + r) \notin \iota(C(\mathbb{F}_q))$. If this is the case, then W will also represent the image of $C(\mathbb{Q})$ in $J(\mathbb{Q})/L'$. Note that this test does not require the computation of a discrete logarithm. We still need to find the discrete logarithms of the images under ϕ_q of our generators of the Mordell–Weil group in order to find the kernel of ϕ_q , but this is a small fixed number of discrete log computations for each q .

The Weil conjectures tell us that $\#C(\mathbb{F}_q)/\#J(\mathbb{F}_q) \approx 1/q$ when C has genus 2, so our chance of success at replacing L with L' is in the range of $(1 - \frac{1}{q})^{((L:L')-1) \cdot \#W}$. This will be very small when $(L : L') \cdot \#W$ is much larger than q . Therefore we try to pick q such that $L/(L \cap \ker \phi_q)$ is non-trivial but comparable with q in size. A necessary condition for this to hold is that the part of the group order of the image of ϕ_q which is coprime to the index of L in $J(\mathbb{Q})$

should be at most some constant times q . Since it is much faster to compute $\#J(\mathbb{F}_q)$ than it is to compute ϕ_q and its image and kernel, we simply check $\#J(\mathbb{F}_q)$ instead. When q passes this test, we do the more involved computation of the group structure of $J(\mathbb{F}_q)$ and the images of the generators of $J(\mathbb{Q})$ in the corresponding abstract group, so that we can find the kernel of ϕ_q and check the condition on $(L : L')$. If q also passes this test, we check whether we can replace L by L' . Of course, we can abort this computation (and declare failure) as soon as we find some $w + r$ as above such that $w + r$ maps into $\iota(C(\mathbb{F}_q))$; see [8, § 11]. The idea for this second sieving stage is due to Samir Siksek.

If N is sufficiently large, then we will have a good chance of finding enough primes q that allow us to go to a subgroup of larger index. Also, once we have been successful with a number of primes, more primes might become available for future steps, since the index of $L \cap \ker \phi_q$ in L may have become smaller.

In the two examples treated in [8], this second stage of the sieving procedure was successful in reaching a subgroup of sufficiently large index (up to 10^{1800}) to be able to conclude that any putative unknown integral point must be so large as to violate the upper bounds obtained earlier.

4.4. Combination with Chabauty’s method

Chabauty originally came up with his method in [11] in order to prove a special case of Mordell’s conjecture. More recently, the method has been developed into a powerful tool that allows us to determine, in many cases, the set of rational points on a given curve; see, for example, [12, 13, 18, 28]. We can combine it with the Mordell–Weil sieve idea to obtain a very efficient procedure for determining $C(\mathbb{Q})$. Examples of Chabauty computations supported by sieving can be found in [3, 4, 22]. In these examples, it is the Chabauty part that is the focus of the computation, and sieving plays a helping role. This is in contrast to what we describe here, where sieving is at the core of the computation and the Chabauty approach is just used to supply us with a ‘separating’ number N such that $C(\mathbb{Q})$ injects into $J(\mathbb{Q})/NJ(\mathbb{Q})$.

Chabauty’s method is applicable when the rank of $J(\mathbb{Q})$ is less than the genus g of C . In this case, for every prime p there is a regular non-zero differential $\omega_p \in \Omega(C_{\mathbb{Q}_p})$ that annihilates the Mordell–Weil group under the natural pairing $J(\mathbb{Q}_p) \times \Omega(C_{\mathbb{Q}_p}) \rightarrow \mathbb{Q}_p$. If p is a prime of good reduction for C , then a suitable multiple of ω_p reduces mod p to a non-zero regular differential $\bar{\omega}_p \in \Omega(C_{\mathbb{F}_p})$. If $P \in C(\mathbb{F}_p)$ is a point such that $\bar{\omega}_p$ does not vanish at P (and $p \geq 3$), then there is at most one rational point on C that reduces mod p to P . See, for example, [28, § 6].

On the other hand, if N is divisible by the exponent of $J(\mathbb{F}_p)$, then the rational points on C mapping via ι into a given coset of $NJ(\mathbb{Q})$ in $J(\mathbb{Q})$ will all reduce mod p to the same point in $C(\mathbb{F}_p)$. So if $\bar{\omega}_p$ does not vanish at any point in $C(\mathbb{F}_p)$, then we know that each coset of $NJ(\mathbb{Q})$ can contain the image under ι of at most one point in $C(\mathbb{Q})$. If there is no such point and we assume [29, Main Conjecture], then we will be able to show, using the Mordell–Weil sieve, that no point of $C(\mathbb{Q})$ maps to this coset. If there is a point, we will eventually find it.

This leads to the following outline of the procedure.

- (1) Find a prime $p \geq 3$ of good reduction for C such that there is $\omega_p \in \Omega(C_{\mathbb{Q}_p})$ annihilating $J(\mathbb{Q})$ and such that $\bar{\omega}_p$ does not vanish on $C(\mathbb{F}_p)$.
- (2) Find a suitable set S of primes and a number N as described in §§ 3.1 and 3.2 above, satisfying the additional condition that the exponent of $J(\mathbb{F}_p)$ divides N .
- (3) Compute $A(S, N)$ as described in § 3.3 above.
- (4) For each element $a \in A(S, N)$, verify that it comes from a rational point on C . To do this, take the point of smallest canonical height in the coset of $NJ(\mathbb{Q})$ given by a , and check whether it comes from a rational point on C . If it does, record the point.
- (5) If the previous step is unsuccessful, enlarge S and/or increase N and compute a new $A(S, N)$ based on the unresolved members of the old $A(S, N)$. Then continue with step (4).

We have implemented this procedure in MAGMA and used it on a large number of genus 2 curves with Jacobian of Mordell–Weil rank one. It proved to be quite efficient: the computation usually takes less than two seconds and almost always less than five seconds. For this implementation, we assume that one rational point is already known and use it as a base-point for the embedding ι . In practice, this is no essential restriction, as there seems to be a strong tendency for small points (which can be found easily) to exist on C if there are rational points at all. Of course, we also need to know a generator of the free part of $J(\mathbb{Q})$, or at least a point of infinite order in $J(\mathbb{Q})$. If we only have a point P of infinite order, we also have to check that the index of $\mathbb{Z} \cdot P + J(\mathbb{Q})_{\text{tors}}$ in $J(\mathbb{Q})$ is prime to N . If P is not a generator, then in step (4) we could have the problem that the point we are looking for is not in the subgroup generated by $P \pmod{\text{torsion}}$. In this case, the smallest representative of a is likely to look large, and we should first try to see if some multiple of a is small, so that it can be recognized. A version of this procedure is used by the Chabauty function provided in recent releases of MAGMA.

As mentioned in the discussion above, steps (4) and (5) will eventually be successful if [29, Main Conjecture] holds for C . There is, however, an additional assumption we have to make, which is that step (1) will always be successful. We state this as a conjecture.

CONJECTURE 4.2. Let C/\mathbb{Q} be a curve of genus $g \geq 2$ such that its Jacobian is simple over \mathbb{Q} and the Mordell–Weil rank r is less than g . Then there are infinitely many primes p for which there exists a regular differential $\omega_p \in \Omega(C_{\mathbb{Q}_p})$ annihilating $J(\mathbb{Q})$ such that the reduction mod p of (a suitable multiple of) ω_p does not vanish on $C(\mathbb{F}_p)$.

Of course, this can easily be generalized to number fields in place of \mathbb{Q} .

We need to assume that the Jacobian is simple, because otherwise there could be a differential killing the Mordell–Weil group that comes from one of the simple factors. Such a differential could possibly vanish at a rational point on the curve, and then its reductions mod p would vanish at an \mathbb{F}_p -point for all p . For example, when C is a curve of genus 2 that covers two elliptic curves, one of rank zero and one of rank one, then the (essentially unique) differential killing the Mordell–Weil group will be the pull-back of the regular differential on one of the elliptic curves, and hence a global object. Of course, in such a case, we can instead work with one of the simple factors which still satisfies the ‘Chabauty condition’ that its Mordell–Weil rank is less than its dimension.

We give a heuristic argument which indicates that Conjecture 4.2 is plausible. We first prove a lemma.

LEMMA 4.3. Let C be a smooth projective curve of genus $g \geq 2$ over \mathbb{F}_p . The probability that a random non-zero regular differential $\bar{\omega}$ on C does not vanish on any point in $C(\mathbb{F}_p)$ is at least $\frac{1}{3} + O(gp^{-1/2})$.

Proof. First, assume that C is not hyperelliptic; then we can consider the canonical embedding $C \rightarrow \mathbb{P}^{g-1}$. We have to estimate the number n of hyperplane sections that do not meet the image of $C(\mathbb{F}_p)$. If $g = 3$, then $C \subset \mathbb{P}^2$ is a smooth plane quartic curve, and the non-zero regular differentials correspond to \mathbb{F}_p -defined lines in \mathbb{P}^2 (up to scaling). For $k = 0, 1, 2, 4$, let ℓ_k be the number of such lines that contain exactly k points of $C(\mathbb{F}_p)$ (counting multiplicity). We want to estimate ℓ_0 . In the following, we disregard lines that are tangent to C at an \mathbb{F}_p -rational point; their number is $O(p)$ and so the result is unaffected by them.

Fix a point $P \in C(\mathbb{F}_p)$ and consider the $p + 1$ lines through P . Projection away from P gives a covering $C \rightarrow \mathbb{P}^1$ of degree three, which can be Galois only for at most four choices of P (since a necessary condition is that five tangents at inflection points of C meet at P , and there are at most 24 such tangents). These potential exceptions do not affect our estimate. For the other points, the covering has Galois group S_3 , and by results in [19] we have the following

expressions for $\ell_{k,P}$, the number of lines through P meeting $C(\mathbb{F}_p)$ in exactly k points:

$$\ell_{1,P} = \frac{p}{3} + O(\sqrt{p}), \quad \ell_{2,P} = \frac{p}{2} + O(\sqrt{p}) \quad \text{and} \quad \ell_{4,P} = \frac{p}{6} + O(\sqrt{p}).$$

We obtain

$$\begin{aligned} \ell_1 &= \sum_P \ell_{1,P} + O(p) = \frac{p^2}{3} + O(p^{3/2}), \\ \ell_2 &= \frac{1}{2} \sum_P \ell_{2,P} + O(p) = \frac{p^2}{4} + O(p^{3/2}), \\ \ell_4 &= \frac{1}{4} \sum_P \ell_{4,P} + O(p) = \frac{p^2}{24} + O(p^{3/2}), \\ \ell_0 &= p^2 + p + 1 - (\ell_1 + \ell_2 + \ell_4) = \frac{3}{8}p^2 + O(p^{3/2}), \end{aligned}$$

which shows that the probability here is $\frac{3}{8} + O(p^{-1/2})$.

Now let $g \geq 4$ (still assuming that C is not hyperelliptic). Let t_3 denote the number of triples of distinct points in $C(\mathbb{F}_p)$ that are collinear in the canonical embedding. By the inclusion–exclusion principle we have, for the number n of hyperplane sections missing $C(\mathbb{F}_p)$,

$$\begin{aligned} n \geq \#\mathbb{P}^{g-1}(\mathbb{F}_p) - \#C(\mathbb{F}_p)\#\mathbb{P}^{g-2}(\mathbb{F}_p) + \binom{\#C(\mathbb{F}_p)}{2}\#\mathbb{P}^{g-3}(\mathbb{F}_p) \\ - \left(\binom{\#C(\mathbb{F}_p)}{3} - t_3 \right)\#\mathbb{P}^{g-4}(\mathbb{F}_p) - t_3\#\mathbb{P}^{g-3}(\mathbb{F}_p). \end{aligned}$$

A collinear triple is part of a one-dimensional linear system of degree three on C . It is known that there are at most two such linear systems when $g = 4$ (see, for instance, [17, Example IV.5.5.2]) and at most one when $g \geq 5$ (see, for instance, [24, Example I.3.4.3]). This implies $t_3 \leq 2(p + 1)$ and therefore that t_3 has no effect on the estimate below. Since $\#C(\mathbb{F}_p) = p + O(gp^{1/2})$, we find that

$$\frac{n}{\#\mathbb{P}^{g-1}(\mathbb{F}_p)} \geq 1 - 1 + \frac{1}{2} - \frac{1}{6} + O(gp^{-1/2}) = \frac{1}{3} + O(gp^{-1/2}).$$

If C is hyperelliptic, the problem is equivalent to the question of how likely it is for a random homogeneous polynomial of degree $g - 1$ in two variables not to vanish on the image X of $C(\mathbb{F}_p)$ in $\mathbb{P}^1(\mathbb{F}_p)$ under the hyperelliptic quotient map $C \rightarrow \mathbb{P}^1$. The number n in this case can be estimated by

$$n \geq \#\mathbb{P}^{g-1}(\mathbb{F}_p) - \#X\#\mathbb{P}^{g-2}(\mathbb{F}_p).$$

Since the size of X is $p/2 + O(gp^{1/2})$, here we even obtain

$$\frac{n}{\#\mathbb{P}^{g-1}(\mathbb{F}_p)} \geq 1 - \frac{1}{2} + O(gp^{-1/2}) = \frac{1}{2} + O(gp^{-1/2}). \quad \square$$

We expect that arguments similar to those used in the non-hyperelliptic genus 3 case can be used to show that the probability in question is

$$\alpha_g + O_g(p^{-1/2}) \quad \text{with} \quad \alpha_g = \sum_{k=0}^{2g-2} \frac{(-1)^k}{k!} \approx e^{-1}$$

in the non-hyperelliptic case. In the hyperelliptic case, the corresponding probability

$$\beta_g + O_g(p^{-1/2}) \quad \text{with} \quad \beta_g = \sum_{k=0}^{g-1} \frac{(-1)^k}{2^k k!} \approx e^{-1/2}$$

is obtained by an obvious extension of the argument used in the proof above.

We now consider a curve C/\mathbb{Q} as in Conjecture 4.2, with $r = g - 1$. It seems reasonable to assume that the reduction $\bar{\omega}_p$ of the unique (up to scaling) differential ω_p annihilating $J(\mathbb{Q})$ behaves like a random element of $\Omega^1(C/\mathbb{F}_p)$ as p varies. By Lemma 4.3, we would then even expect a set of primes p of positive density at least $1/3$ such that $\bar{\omega}_p$ does not vanish on $C(\mathbb{F}_p)$.

When $r \leq g - 2$, the situation should be much better. We have at least a pencil of differentials, giving rise to a linear system of degree $2g - 2$ and positive dimension on the curve over \mathbb{F}_p . Unless this linear system has a base-point in $C(\mathbb{F}_p)$, effective versions of the Chebotarev density theorem as in [19] show that there is a divisor in the system whose support does not contain rational points, at least when p is sufficiently large. However, we still have to exclude the possibility that the relevant linear system has a base-point in $C(\mathbb{F}_p)$ for (almost) every p .

If we mimic the set-up of Lemma 4.3 in the situation where $g - r = d \geq 2$, then we have to look at the Grassmannian of $(r - 1)$ -dimensional linear subspaces in \mathbb{P}^{g-1} : there is a d -dimensional linear space of differentials killing $J(\mathbb{Q})$, and the intersection of the corresponding hyperplanes in \mathbb{P}^{g-1} is an $(r - 1)$ -dimensional linear subspace. The set of such subspaces through a given point corresponds, via projection away from this point, to $\text{Gr}(\mathbb{P}^{r-2} \subset \mathbb{P}^{g-2})$; so, by the simplest case of the inclusion–exclusion inequality, we have that the number n of base-point-free subspaces satisfies

$$n \geq \# \text{Gr}(\mathbb{P}^{r-1} \subset \mathbb{P}^{g-1}) - \#C(\mathbb{F}_p) \# \text{Gr}(\mathbb{P}^{r-2} \subset \mathbb{P}^{g-2}),$$

and therefore the ‘density’ of such subspaces is

$$\frac{n}{\# \text{Gr}(\mathbb{P}^{r-1} \subset \mathbb{P}^{g-1})} \geq 1 - \#C(\mathbb{F}_p) \frac{\# \text{Gr}(\mathbb{P}^{r-2} \subset \mathbb{P}^{g-2})}{\# \text{Gr}(\mathbb{P}^{r-1} \subset \mathbb{P}^{g-1})} = 1 - O(p^{-(d-1)}).$$

When $d = 2$, one is thus led to expect an infinite but very sparse set of primes for which there is a base-point (since $\sum p^{-1}$ diverges), whereas for $d > 2$, one would expect only finitely many such primes.

If we modify the algorithm in such a way that it considers (arbitrarily) ‘deep’ information at p , then the requirement can be weakened to the following.

CONJECTURE 4.4. Let C/\mathbb{Q} be a curve of genus $g \geq 2$ such that its Jacobian is simple and has Mordell–Weil rank $r < g$. Then there is a prime $p \geq 3$ for which there exists a regular non-zero differential $\omega_p \in \Omega(C_{\mathbb{Q}_p})$ annihilating $J(\mathbb{Q})$ such that ω_p does not vanish on $C(\mathbb{Q})$.

Heuristically, the probability that ω_p does vanish at a rational point should be zero (except when there is a good reason for that to happen, as seen above), which lets us hope that the weaker conjecture may be amenable to proof. In fact, Tzanko Matev (a PhD student of the second author) has recently established a p -adic version of the ‘analytic subgroup theorem’ for abelian varieties (see [1] for the background); it states that when J is absolutely simple, then the p -adic logarithm of an algebraic point on J cannot be contained in a proper subspace of the tangent space $T_0J(\mathbb{Q}_p)$ that is generated by algebraic vectors. This implies that the statement of Conjecture 4.4 is true for every p when the Mordell–Weil rank is one.

5. Information at bad primes

This section and the next discuss how to extract the information that the Mordell–Weil sieve needs as input in the specific case where C is a curve of genus 2 over \mathbb{Q} (or a more general number field) and we are interested in $C(\mathbb{F}_p)$ and $J(\mathbb{F}_p)$ not merely when p is a prime of good reduction.

In particular, when the rank is large, which in practice means $r \geq 3$, it becomes important to use a sufficient amount of ‘local’ information to keep the sizes of the sets $A(S, N_j)$ reasonably small. A valuable source of such information is given by primes of bad reduction, as the group

orders of suitable quotients of $J(\mathbb{Q}_p)$ tend to be rather smooth. More precisely, we would like to make use of the top layers of the filtration given by the well-known exact sequences

$$0 \longrightarrow J^0(\mathbb{Q}_p) \longrightarrow J(\mathbb{Q}_p) \longrightarrow \Phi_p(\mathbb{F}_p) \longrightarrow 0$$

and

$$0 \longrightarrow J^1(\mathbb{Q}_p) \longrightarrow J^0(\mathbb{Q}_p) \longrightarrow \tilde{J}(\mathbb{F}_p) \longrightarrow 0.$$

Here Φ_p is the component group of the special fiber of the Néron model of J over \mathbb{Z}_p , \tilde{J} is the connected component of the special fiber, and $J^1(\mathbb{Q}_p)$ is the kernel of reduction.

In this section, we describe how this information can be obtained when C is a genus 2 curve, p is odd, and the given model of C is regular at p . Here and in what follows, we will use $J^1(\mathbb{Q}_p)$ and, later, $J^n(\mathbb{Q}_p)$ to denote the kernel of reduction and the ‘higher’ kernels of reduction *with respect to the given model of the curve*. If the model is not minimal in a suitable sense, then our kernel of reduction will be strictly contained in the kernel of reduction with respect to a Néron model. To be precise, for us, $J^1(\mathbb{Q}_p)$ denotes the subgroup of points in $J(\mathbb{Q}_p)$ whose reduction mod p on the projective model in \mathbb{P}^{15} as described in [10, Chapter 2] is the origin; see below. Of course, this then changes the meaning of the quotients in the sequences above.

But first we will establish some general facts. Let k be a field with $\text{char}(k) \neq 2$, and let

$$F(X, Z) = f_6X^6 + f_5X^5Z + f_4X^4Z^2 + f_3X^3Z^3 + f_2X^2Z^4 + f_1XZ^5 + f_0Z^6$$

be a homogeneous polynomial of degree six with coefficients in k . We do not assume that F is squarefree or even that $F \neq 0$.

DEFINITION 5.1.

- (i) Let C_F be the curve given by the equation

$$Y^2 = F(X, Z)$$

in the weighted projective plane with weights 1, 3 and 1 for the coordinates X, Y and Z , respectively.

- (ii) Denote by J_F the scheme in \mathbb{P}_k^{15} that is defined by the 72 quadrics described in [10, Chapter 2] (see [15, jacobian.variety/defining.equations] for explicit equations).

- (iii) Let K_F be the surface in \mathbb{P}_k^3 that is defined by the Kummer surface equation as given in [10, Chapter 3], and denote by $\delta_F = \delta = (\delta_1, \dots, \delta_4)$ the polynomials giving the duplication map on the Kummer surface; see [15, kummer/duplication].

If $\delta(P) \neq 0$, then we write $\mathbb{P}\delta(P) \in \mathbb{P}^3$ for the point with projective coordinates given by $(\delta_1(P) : \dots : \delta_4(P))$.

- (iv) Let $\tilde{D}_F \subset \mathbb{A}_k^3 \times \mathbb{A}_k^4 \times \mathbb{A}_k^5$ be the scheme of triples (A, B, C) such that $A \neq 0$ and

$$F(X, Z) = A(X, Z)C(X, Z) + B(X, Z)^2,$$

where for $A = (a_0, a_1, a_2)$, $B = (b_0, \dots, b_3)$ and $C = (c_0, \dots, c_4)$ we set

$$\begin{aligned} A(X, Z) &= a_2X^2 + a_1XZ + a_0Z^2, \\ B(X, Z) &= b_3X^3 + b_2X^2Z + b_1XZ^2 + b_0Z^3, \\ C(X, Z) &= c_4X^4 + c_3X^3Z + c_2X^2Z^2 + c_1XZ^3 + c_0Z^4. \end{aligned}$$

Let $D_F \subset \mathbb{P}_k^2 \times \mathbb{A}_k^4$ be the image of \tilde{D}_F under the projection to the first two factors followed by the canonical map $\mathbb{A}^3 \setminus \{0\} \rightarrow \mathbb{P}^2$ on the first factor.

When F is squarefree, then C_F is a smooth curve of genus 2, J_F is its Jacobian, and K_F is the associated Kummer surface. The scheme D_F then gives the possible Mumford representations of effective divisors of degree two on C_F ; it therefore maps onto $J_F \setminus \{O\}$. We will extend these relations to our more general setting.

The ‘origin’ $O = (1 : 0 : \dots : 0)$ is always a (smooth) point on J_F . The sixteen coordinates on J_F split into ten ‘even’ and six ‘odd’ ones; the even coordinates are given (up to a simple invertible linear transformation) by the monomials of degree two in the coordinates on K_F .

Let us first look at the relation between J_F and K_F .

LEMMA 5.2. *Projection to the ten even coordinates gives rise to a morphism $\kappa : J_F \rightarrow K_F$, which is a double cover.*

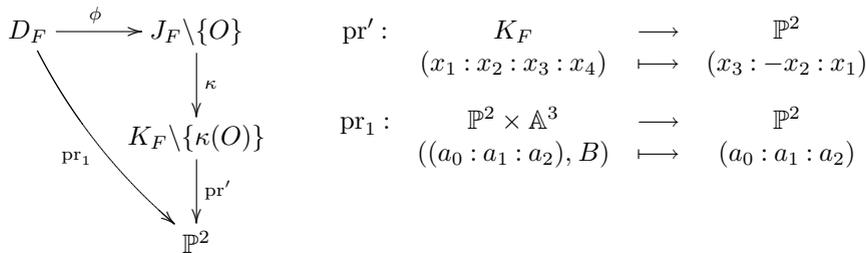
Proof. The monomials of degree two in the odd coordinates can be expressed as quadratic forms in the even coordinates. So if all the even coordinates vanish, the odd coordinates have to vanish, too. Therefore, projection to the \mathbb{P}^9 spanned by the even coordinates is a morphism. The relations between the even coordinates are exactly those arising from the fact that the even coordinates come from the monomials of degree two in the coordinates of the \mathbb{P}^3 containing K_F , together with the quadratic relation coming from the quartic equation defining K_F . Therefore the image of J_F in \mathbb{P}^9 is the image of K_F under the 2-uple embedding of \mathbb{P}^3 into \mathbb{P}^9 and hence is isomorphic to K_F . This gives the morphism κ . The fact stated in the first sentence of this proof then implies that κ is a (ramified) double cover. \square

Now let us consider the relation between \tilde{D}_F , D_F and J_F .

LEMMA 5.3. *There is a morphism*

$$\phi : D_F \rightarrow J_F \setminus \{O\}$$

that specializes to the representation of points on J_F mentioned above when F is squarefree. The morphism ϕ is surjective on k -points and makes the following diagram commute.



Furthermore, $\phi(A, B) = \phi(A', B')$ if and only if $A = A'$ and $B(X, Z) \equiv B'(X, Z) \pmod{A(X, Z)}$.

Proof. Let $(A, B) \in D_F$. Then ϕ can be given as

$$\begin{aligned}
 \phi(A, B) &= (* : * : * : * : * : * : * : * : * : * : * : * : * : * : * : *) \\
 &: -b_2a_0^2 + b_1a_0a_1 - b_0(a_1^2 - a_0a_2) : b_3a_0^2 - b_1a_0a_2 + b_0a_1a_2 \\
 &: -b_3a_0a_1 + b_2a_0a_2 - b_0a_2^2 : b_3(a_1^2 - a_0a_2) - b_2a_1a_2 + b_1a_2^2 \\
 &: a_0^2 : -a_0a_1 : a_0a_2 : -a_1a_2 : a_2^2 : a_1^2 - 4a_0a_2.
 \end{aligned}$$

One can check, using the defining equations of J_F given at [15], that the first six coordinates are uniquely determined by the last ten when the last six are not all zero. It is also possible to write down expressions for the first six coordinates in terms of A, B and C , where (A, B, C) is the point on \tilde{D}_F mapping to (A, B) . The image of this point under κ has the form $(a_2 : -a_1 : a_0 : *)$, which shows that $\text{pr}' \circ \kappa \circ \phi = \text{pr}_1$. It remains to show that ϕ is surjective on k -points. Let $P \in J_F(k) \setminus \{O\}$; then $A = \text{pr}'(\kappa(P)) \in \mathbb{P}^2(k)$ is defined. Consider the middle four coordinates on J (the seventh to the tenth). The above expression for $\phi(A, B)$ gives rise to a system of linear equations for B . The last six of the equations defining J_F ensure that the system has

a solution $B \in \mathbb{A}^4(k)$. Then $\phi(A, B)$ agrees with P in the last ten coordinates, and hence we must have $\phi(A, B) = P$.

To prove the final assertion, note first that $\phi(A, B) = \phi(A', B')$ implies $A = A'$ (apply $\text{pr}' \circ \kappa$). The kernel of the matrix giving the linear equations that determine B is spanned by the coefficient tuples of $ZA(X, Z)$ and $XA(X, Z)$. This shows that $\phi(A, B) = \phi(A, B')$ is equivalent to $A(X, Z) \mid B(X, Z) - B'(X, Z)$. □

By the above, the fibers of the map $\phi : D_F \rightarrow J_F \setminus \{O\}$ are isomorphic to \mathbb{A}^2 . We can remove this ambiguity at the cost of restricting to a subscheme.

LEMMA 5.4. *Let*

$$\begin{aligned} U_0 &= \{(A, B) \in D_F : a_0 = 1, b_0 = b_1 = 0\}, \\ U_1 &= \{(A, B) \in D_F : a_1 = 1, a_0 a_2 \neq 1, b_1 = b_2 = 0\}, \\ U_2 &= \{(A, B) \in D_F : a_2 = 1, b_2 = b_3 = 0\}. \end{aligned}$$

Then $\phi|_{U_j}$ is an isomorphism onto its image for each $j \in \{0, 1, 2\}$, and

$$\phi(U_0) \cup \phi(U_1) \cup \phi(U_2) = J_F \setminus \{O\}.$$

Proof. In each case, the linear system giving b_0, \dots, b_3 in terms of the middle four coordinates on J_F , together with the conditions $b_j = b_{j+1} = 0$, has a unique solution, yielding the inverse morphism $\phi(U_j) \rightarrow U_j$. The last statement then follows, since the images of the U_j in \mathbb{P}^2 cover \mathbb{P}^2 . □

Now we can describe the smooth locus of J_F .

PROPOSITION 5.5. *The origin O is always a smooth point on J_F . If $P \in J_F \setminus \{O\}$, write $P = \phi(A, B)$ with $(A, B) \in D_F$. Then P is a singular point on J_F if and only if:*

- (i) $A(X, Z)$ has a simple root (in \mathbb{P}^1) at a multiple root of F ; or
- (ii) $A(X, Z) = cL(X, Z)^2$ has a double root at a multiple root of F and $L(X, Z)^3$ divides $F(X, Z) - B(X, Z)^2$.

Note that the last condition means that the curve $Y = B(X, Z)$ is tangential to a branch of C_F at the singular point $L(X, Z) = Y = 0$.

Proof. The assertion that $O \in J_F$ is smooth is easily checked by using the explicit equations. The general statement is geometric, so we can assume k to be algebraically closed. Then there is a transformation $\sigma \in \text{GL}_2(k)$ such that $A^\sigma(X, Z) = XZ$ or X^2 . In the first case, we can take $Q \in U_1$ and easily check that Q is singular on U_1 if and only if $f_0 = f_1 = 0$ or $f_6 = f_5 = 0$, which means that F has a multiple root at one of the two simple roots of $A(X, Z)$, namely 0 or ∞ . In the second case, we can take $Q \in U_2$, and we find that Q is singular on U_2 if and only if $f_0 = f_1 = f_2 - b_1^2 = 0$, which means that F has a multiple root at the double root 0 of $A(X, Z)$ and that $X^3 = L(X, Z)^3$ divides $F(X, Z) - B(X, Z)^2$. Since ϕ is an isomorphism on U_j , P is singular on J_F if and only if Q is singular on U_j . □

DEFINITION 5.6. We denote by D'_F the locus of points $Q \in D_F$ such that $\phi(Q)$ is a smooth point on J_F , and we write J'_F for the subscheme of smooth points on J_F .

According to Proposition 5.5 above, the complement of D'_F in D_F consists of the points (A, B) satisfying one of the conditions in the proposition.

LEMMA 5.7. Assume that k is algebraically closed. Then J_F is reduced and irreducible except in the following two cases.

(i) $F = 0$. Then J_F has two irreducible components. One is $\phi(\mathbb{P}^2 \times \{0\})$ and is not reduced; the other contains O , and its remaining points are of the form $\phi(A, B)$ such that there is a linear form L with $A(X, Z) = cL(X, Z)^2$ and $L(X, Z)|B(X, Z)$.

(ii) $F = H(X, Z)^2$ is a non-zero square. Then J_F has three irreducible components, all of which are reduced. Two of them are given by $\phi(\mathbb{P}^2 \times \{\pm H\})$, while the third contains the origin O .

Proof. It is easy to check the claim for the two special cases. In all other cases, C_F is reduced and irreducible. Consider the symmetric square $C_F^{(2)}$. Let $S \subset C_F$ be the (finite) set of singular points (given by the multiple roots of F). Identify S with its image in $C_F^{(2)}$ under the diagonal map. There is a morphism

$$\psi : C_F^{(2)} \setminus S \rightarrow J_F$$

that can be defined using the expressions for the coordinates on the Jacobian given in [10, Chapter 2]. Its image is

$$J_F \setminus \{ \phi(A, B) : A(X, Z) = cL(X, Z)^2, L(P) = 0 \text{ for some } P \in S \},$$

which is dense in J_F . Since $C_F^{(2)}$ is irreducible, this implies that J_F is irreducible as well. The component containing the origin is always reduced, since the origin is a smooth point. \square

REMARK 5.8. If k is not algebraically closed, then there is an additional case to consider: $F = cH(X, Z)^2$ with $H \neq 0$ and a non-square $c \in k$. According to Lemma 5.7, J_F has three geometric components. One is defined over k and contains the origin, while the other two are conjugate over $k(\sqrt{c})$ and do not have any smooth k -points.

If we apply the argument used in the proof above to the case where $F = H^2 \neq 0$, then we see that C_F has two components and therefore $C_F^{(2)}$ has three; we thus deduce again that J_F has three (reduced) irreducible components.

From the description given in the proof, we see that ψ extends to a morphism

$$\tilde{\psi} : \text{Bl}'_S C_F^{(2)} \rightarrow J_F.$$

Here $\text{Bl}'_S C_F^{(2)}$ is obtained from $C_F^{(2)}$ by replacing each point in S by a \mathbb{P}^1 in such a way that locally near a point in S , $\text{Bl}'_S C_F^{(2)}$ is the closure of the graph of the rational map giving the slope (in a suitable affine chart) of the line that connects the two points in the divisor corresponding to a point in $C_F^{(2)}$. Let $\pi : C_F \rightarrow \mathbb{P}^1$ be the canonical map, and denote by π^* the induced map $\mathbb{P}^1 \rightarrow C_F^{(2)}$. Then $\tilde{\psi}$ is an isomorphism away from $\pi^*(\mathbb{P}^1)$ and contracts $\pi^*(\mathbb{P}^1)$ to the origin $O \in J_F$. We therefore have an isomorphism

$$\text{Bl}'_S C_F^{(2)} \cong \text{Bl}_O J_F.$$

This generalizes the standard fact that $C_F^{(2)} \cong \text{Bl}_O J_F$ if C_F is smooth.

DEFINITION 5.9. We denote by J_F^0 the component of the smooth part J'_F of J_F that contains the origin O . We write K_F^0 for the open subscheme of K_F on which $\delta \neq 0$. Let B_F denote the matrix of biquadratic forms as defined in [10, Chapter 3]; see [15, kummer/biquadratic.forms] for explicit expressions.

PROPOSITION 5.10. We have $\kappa(J_F^0) = K_F^0$. Equivalently, a point $P \in J_F$ is smooth and on the component of the origin if and only if $\delta(\kappa(P)) \neq 0$.

Proof. We can again assume that k is algebraically closed and that $\text{pr}'(\kappa(P))$ is one of $(0 : 1 : 0)$ or $(1 : 0 : 0)$. (Note that O is always smooth, and $\delta_4(\kappa(O)) \neq 0$.) We represent P as $\phi(Q)$ with $Q = ((0 : 1 : 0), (b_0, 0, 0, b_3))$ or $Q = ((1 : 0 : 0), (b_0, b_1, 0, 0))$, respectively. Then we can use the description of singular points given in Proposition 5.5 and the description of the components of J_F given in Lemma 5.7. Writing down the polynomials δ_j evaluated at $\kappa(\phi(Q))$, we conclude after some fairly straightforward manipulations that in the first case, $\delta = 0$ if and only if $f_0 = f_1 = 0$ or $f_6 = f_5 = 0$, or there are b_1 and b_2 such that $F = (b_3X^3 + b_2X^2Z + b_1XZ^2 + b_0Z^3)^2$. The first two conditions mean, as before, that there is a singularity at 0 or ∞ , and the third condition says that P is not on the right component. In the second case, we find in a similar way that $\delta = 0$ if and only if $X^2|F$ and $X^3|(F - (b_1XZ^2 + b_0Z^3)^2)$, or F is a square and does not vanish at 0. The first condition means that P is not smooth; the second says again that P is not on the right component. \square

This result is due (with a different proof) to Jan Steffen Müller, a PhD student of the second author.

We can now state and prove the main result of this section.

THEOREM 5.11. *The scheme J_F^0 is a commutative algebraic group in a natural way. If we represent its non-zero elements by pairs $(A, B) \in D^0(F)$, then composition in the group can be performed by Cantor composition and reduction [9], except when both polynomials $A(X, Z)$ vanish at the same singular point of C_F . Without loss of generality, take this point to be at $X = 0$; then we have*

$$\phi(X^2, \lambda XZ^2) + \phi(X^2, \mu XZ^2) = \phi\left(X^2, \frac{f_2 + \lambda\mu}{\lambda + \mu} XZ^2\right)$$

where $F(X, Z) = f_2X^2Z^4 + f_3X^3Z^3 + \dots + f_6Z^6$. If $\lambda + \mu = 0$, the result is the zero element in J_F^0 .

Proof. Let \mathcal{O} be a complete discrete valuation ring with uniformizer π , residue field k and field of fractions L . We can then find a homogeneous polynomial $\tilde{F} \in \mathcal{O}[X, Z]$ of degree six that is squarefree and whose reduction mod π is F . We denote reduction mod π by a bar. Let $G = J_{\tilde{F}}(L)$, $G^0 = \{P \in G : \bar{P} \in J_F^0(k)\}$ and $G^1 = \{P \in G : \bar{P} = O \in J_F(k)\}$. Then, for $P \in G^0$ and $Q \in G^1$, we have $\overline{P + Q} = \bar{P}$. To see this, note that the images of $P \pm Q$ under κ are given by $B_{\tilde{F}}(P, Q)$. Since $B_{\tilde{F}}(P, Q) = B_F(\bar{P}, \bar{Q}) \sim \bar{P}^\top \bar{P}$ (where we have abused notation by letting \bar{P} denote a vector of projective coordinates for \bar{P}), we must have $\kappa(\overline{P \pm Q}) = \kappa(\bar{P})$. This implies that $\overline{P + Q} = \bar{P}$ or $-\bar{P}$. The function $Q \mapsto \overline{P + Q}$ cannot take exactly two distinct values on the residue class of O , so we must have $\overline{P + Q} = \bar{P}$.

This implies that G^1 is a subgroup of G , that G^1 acts on G^0 , and that (at least as sets) $G^0/G^1 \cong J_F^0(k)$. By a similar argument, we see that G^0 is also a subgroup of G . (If $P, Q \in G^0$, then by [27, Proposition 3.1] we have $B_F(\bar{P}, \bar{Q}) \neq 0$, which in turn implies by [27, Lemma 3.2] that $\overline{P \pm Q} \in J_F^0(k)$.) This already shows that $J_F^0(k)$ has a group structure (and the same is true of $J_F^0(\ell)$ for every field extension ℓ of k).

To see that the group law on J_F^0 is given by Cantor’s algorithm, we can lift two given elements to G^0 in such a way that we stay in the same case in the algorithm, and then apply the algorithm over L (in fact, over \mathcal{O}) and reduce mod π . This works unless we are in the special case mentioned in the statement of the proposition. The formula in this case can be obtained by a suitable limit argument. This then also shows that J_F^0 is an algebraic group. \square

The upshot of this result is that we can do computations in the group $J_F^0(k)$, much in the same way as we compute in the Jacobian of C_F when C_F is smooth.

factorization	c	order if $\text{sq}(c)$	order otherwise
0	–	q^2	
$\ell_1^2 h_4$	$\text{Res}(\ell_1, h_4)$	$(q - 1) \# E(\mathbb{F}_q)$	$(q + 1) \# E(\mathbb{F}_q)$
$\ell_1^3 h_3$	–	$q \# E(\mathbb{F}_q)$	
$\ell_1^2 m_1^2 h_2$	$c = \text{Res}(\ell_1, h_2)$ $c' = \text{Res}(m_1, h_2)$	$\begin{cases} (q - 1)^2 & \text{if } \text{sq}(c') \\ q^2 - 1 & \text{else} \end{cases}$	$\begin{cases} q^2 - 1 & \text{if } \text{sq}(c') \\ (q + 1)^2 & \text{else} \end{cases}$
$g_2^2 h_2$	$\text{Res}(g_2, h_2)$	$q^2 - 1$	$q^2 + 1$
$cg_1^2 h_1^2 \ell_1^2$	leading coeff.	$(q - 1)^2$	$(q + 1)^2$
$cg_1^2 h_2^2$	leading coeff.	$q^2 - 1$	$q^2 - 1$
cg_3^2	leading coeff.	$q^2 + q + 1$	$q^2 - q + 1$
$g_1^3 \ell_1^2 h_1$	$\text{Res}(\ell_1, g_1 h_1)$	$q(q - 1)$	$q(q + 1)$
$\ell_1^4 h_2$	$\text{Res}(\ell_1, h_2)$	$q(q - 1)$	$q(q + 1)$
$cg_1^3 h_1^3$	–	q^2	
cg_2^3	–	q^2	
$cg_1^2 h_1^4$	leading coeff.	$q(q - 1)$	$q(q + 1)$
$g_1^5 h_1$	–	q^2	
cg_1^6	leading coeff.	q^2	q^2

FIGURE 1. Group orders $\#J_F^0(\mathbb{F}_q)$.

REMARK 5.12. If $k = \mathbb{F}_q$ and q is odd, one can work out the order of the group $J_F^0(k)$, depending on the factorization of F . This leads to the table shown in Figure 1. The subscripts give the degrees of the factors, which are assumed to be irreducible if they occur with multiplicity greater than one and to be pairwise coprime. In the table, E denotes the genus 1 curve $y^2 = h_4(x, 1)$ or $y^2 = h_3(x, 1)$, and ‘sq(c)’ means that c is a square in \mathbb{F}_q^\times .

If $F = H^2$ is a non-zero square, then, by Lemma 5.7, J_F splits into three components, with the two components not containing O being given by $\phi(\mathbb{P}^2 \times \{\pm H\})$. We denote their intersection with J_F' by J_F^\pm . In a similar way to how we argued above for the group structure of J_F^0 , we obtain well-defined maps

$$J_F^0 \times J_F^\pm \rightarrow J_F^\pm, \quad J_F^0 \times J_F^\mp \rightarrow J_F^\mp \quad \text{and} \quad J_F^\pm \times J_F^\mp \rightarrow J_F^0$$

which are compatible with the group structure of J_F^0 and show that J_F^\pm and J_F^\mp are principal homogeneous spaces under J_F^0 . Therefore the number of smooth points in $J_F(k)$ is three times the cardinality of $J_F^0(k)$. On the other hand, our addition is not defined on $J_F^\pm \times J_F^\pm$ or $J_F^\mp \times J_F^\mp$. (In this case, the B polynomial one obtains from Cantor’s algorithm vanishes along one of the components of C_F , and we get an undefined A .)

As in the proof of Theorem 5.11, we now consider the situation where \mathcal{O} is a complete discrete valuation ring with uniformizer π , residue field k such that $\text{char}(k) \neq 2$ and field of fractions L . We denote by $v : L^\times \rightarrow \mathbb{Z}$ the normalized valuation. Let $F \in \mathcal{O}[X, Z]$ be homogeneous of degree six and squarefree. The 72 quadrics defining J_F have coefficients in \mathcal{O} ; we obtain a flat scheme over $\text{Spec}(\mathcal{O})$. We abuse notation slightly and set

$$J_F^0(L) = \{P \in J_F(L) : \bar{P} \in J_F^0(k)\} \quad \text{and} \quad J_F^1(L) = \{P \in J_F(L) : \bar{P} = \bar{O}\}.$$

We will call $J_F^1(L)$ the *kernel of reduction*. The reader should be warned that this notion depends on the given model of the curve and need not coincide with the kernel of reduction defined in terms of a Néron model of the Jacobian.

LEMMA 5.13. Consider a pair $(A, B) \in D_F(L)$ with $A(X, Z) = X^2 + a_1 XZ + a_0 Z^2$ and $B(X, Z) = b_1 XZ^2 + b_0 Z^3$.

(i) If $a_0, a_1 \in \mathcal{O}$ but b_0 and b_1 are not both integral, then $P = \phi(A, B)$ is in the kernel of reduction.

(ii) Now assume that a_0, a_1, b_0 and b_1 are integral. If π divides f_0, f_1 and a_0 but π^2 does not divide f_0 , then π also divides a_1 and b_0 but does not divide $f_2 - b_1^2$.

Proof.

(i) We work in the affine chart given by $(X : Z) = (x : 1)$. Upon reducing $F(x, 1)$ modulo $A(x, 1) = x^2 + a_1x + a_0$, we obtain a relation $y^2 = \alpha_1x + \alpha_0$ that holds for the points in the divisor described by the pair of polynomials (A, B) . Since the coefficients of F as well as a_0 and a_1 are integral, the same holds for α_0 and α_1 . If we square the relation $y = B(x, 1)$ and reduce it mod $A(x, 1)$, we obtain

$$b_1(2b_0 - b_1a_1) = \alpha_1, \quad b_0^2 - b_1^2a_0 = \alpha_0.$$

The second relation shows that $v(b_0) < v(b_1)$ is impossible, so we must have $v(b_1) < 0$. Eliminating b_0 from the two equations above gives $(a_1^2 - 4a_0)b_1^2 \in \mathcal{O}$, so the discriminant of $A(x, 1)$ must be divisible by π^2 . Therefore the two points in the divisor reduce mod π to points with the same x -coordinate. If these points were not opposite, then $y = B(x, 1)$ would reduce to the equation of the (non-vertical) tangent line at the point on $C_F(k)$ that both points reduce to; but then b_0 and b_1 would be integral, which contradicts the assumptions. So the divisor must reduce to the sum of two opposite points, and hence P reduces mod π to the origin.

(ii) We know that $x^2 + a_1x + a_0$ divides $F(x, 1) - (b_1x + b_0)^2$. Write $f_0 = \pi f'_0, f_1 = \pi f'_1$ and $a_0 = \pi a'_0$. From

$$\begin{aligned} f_6x^6 + \dots + f_2x^2 + \pi f'_1x + \pi f'_0 - (b_1x + b_0)^2 \\ = (x^2 + a_1x + \pi a'_0)(c_4x^4 + c_3x^3 + c_2x^2 + c_1x + c_0) \end{aligned}$$

we get

$$\begin{aligned} \pi f'_0 - b_0^2 &= \pi a'_0c_0, \\ \pi f'_1 - 2b_0b_1 &= a_1c_0 + \pi a'_0c_1, \\ f_2 - b_1^2 &= c_0 + a_1c_1 + \pi a'_0c_2. \end{aligned}$$

The first of these equations implies that $b_0 = \pi b'_0$ for some $b'_0 \in \mathcal{O}$. Since f'_0 is not divisible by π (by assumption), we then also see that $\pi \nmid a'_0c_0$. The second equation then shows that π divides a_1 , and the third equation tells us that $f_2 - b_1^2 \equiv c_0 \not\equiv 0 \pmod{\pi}$. □

The above result allows us to get a description of the reductions of points not in the kernel of reduction when the curve is regular.

COROLLARY 5.14. *Assume that C_F/\mathcal{O} as above is regular. Consider $P \in J_F(L) \setminus J_F^1(L)$. If $P = \phi(A, B)$ with $A(X, Z) \in \mathcal{O}[X, Z]$ primitive, then after adding a suitable multiple of $A(X, Z)$, $B(X, Z)$ has coefficients in \mathcal{O} , and the reduction (\bar{A}, \bar{B}) of $(A, B) \pmod{\pi}$ is in $D'_{\bar{F}}(k)$, hence \bar{P} is a smooth point on $J_{\bar{F}}$. In particular, if \bar{F} is not a square, then $J_F(L) = J_F^0(L)$.*

Proof. First, assume that the coefficient of X^2 in $A(X, Z)$ is a unit. Then we can take $A(x, 1)$ to be monic and $B(x, 1)$ to be of degree at most one. The integrality of B is given by Lemma 5.13(i) If \bar{A} vanishes at a singularity of $C_{\bar{F}}$, then by a suitable shift we can assume that the singularity is at $x = 0$ (we may have to extend the field for that; note that the shift would be by an integral element). We then have that π divides f_0, f_1 and a_0 , which, by Lemma 5.13(ii), implies that π also divides a_1 and b_0 (note that $\pi^2 \nmid f_0$ because of the regularity assumption).

This shows that \bar{A} has a double root at the singularity (and hence that no field extension was necessary) and $\bar{B} = \lambda XZ^2$. We also know from the lemma that $\lambda^2 \neq \bar{f}_2$, which means exactly that the slope of the line described by \bar{b} does not coincide with the slope of a branch of the curve at the singularity. Hence $(\bar{A}, \bar{B}) \in D'_{\bar{F}}(k)$. This implies that $\bar{P} \in J'_{\bar{F}}$. If \bar{F} is not a square, then $J'_{\bar{F}}(k) = J^0_{\bar{F}}(k)$ and the last claim follows.

The case where the coefficient of X^2 in $A(X, Z)$ is not a unit can be reduced to the general case discussed above by a suitable change of coordinates. □

COROLLARY 5.15. *Assume that C_F/\mathcal{O} is regular and that \bar{F} is not a square. Then the following sequence is exact:*

$$0 \longrightarrow J^1_{\bar{F}}(L) \longrightarrow J_F(L) \xrightarrow{P \mapsto \bar{P}} J^0_{\bar{F}}(k) \longrightarrow 0.$$

Proof. By Corollary 5.14, we know that $J^0_F(L) = J_F(L)$, and by the proof of Theorem 5.11, we know that reduction mod π gives a group homomorphism $J^0_F(L) \rightarrow J^0_{\bar{F}}(k)$ with kernel $J^1_F(L)$. This homomorphism is surjective because of Hensel’s lemma (recall that the points in $J^0_{\bar{F}}(k)$ are smooth). □

REMARK 5.16. When C_F/\mathcal{O} is regular, then the scheme obtained from J_F/\mathcal{O} by removing the singular points from the special fiber $J_{\bar{F}}/k$ is the Néron model of J_F/L , and $J^0_{\bar{F}}/k$ is the connected component of the identity on the special fiber.

If C_F/\mathcal{O} is not regular, then the smooth part of J_F/\mathcal{O} still maps to the Néron model (by the universal property of the latter), but the image of $J^0_{\bar{F}}$ in the special fiber of the Néron model can be trivial or a one-dimensional subgroup.

We now consider a genus 2 curve $C = C_F$ over \mathbb{Q}_p given by a Weierstrass equation $Y^2 = F(X, Z)$ over \mathbb{Z}_p . We will drop the subscript F in the following. By the above, we have $J^0(\mathbb{Q}_p)/J^1(\mathbb{Q}_p) \cong J^0_{\bar{F}}(\mathbb{F}_p)$, and the map is given by reducing the standard representation modulo p (on elements that are not in the kernel of reduction).

This gives us a handle on the quotient $J(\mathbb{Q}_p)/J^1(\mathbb{Q}_p)$ when p is odd, the model is regular and the special fiber of C has just one component; cf. Corollary 5.15.

Since we have now established that we can use Cantor reduction on $J^0_{\bar{F}}(\mathbb{F}_p)$ in the same way as in the good reduction case, we can proceed and find the image $\iota(C(\mathbb{F}_p)) \subset J^0_{\bar{F}}(\mathbb{F}_p)$ in the same way as described in § 3.1.

Otherwise (that is, when $p = 2$, the model is not regular, or the special fiber has several components), we first need to find $J^0(\mathbb{Q}_p)$ or, rather (for our purposes), $J(\mathbb{Q}) \cap J^0(\mathbb{Q}_p)$. We can do this by an enumerative process.

In the following, A is a finitely generated free abelian group and t is a test that determines whether a given element of A is in the subgroup. In our application, $A = J(\mathbb{Q})$ and t tests whether or not a point P is in $J^0(\mathbb{Q}_p)$. According to Proposition 5.10, we can use

$$t(P) \iff v_p(\delta(\kappa(P))) = 4v_p(\kappa(P))$$

(with the same choice of projective coordinates for $\kappa(P)$ on both sides) or, in the notation of [27], $t(P) \iff \epsilon_p(P) = 0$.

```

GetSubgroup(A, t):
  g := ∅ // g will contain the generators of the subgroup
  A' := {0} ⊂ A // known part of quotient group
  for b ∈ Generators(A) do
    // find smallest multiple of b such that A' + b meets the subgroup
    j := 1; b' := b; while ¬∃a ∈ A' : t(b' + a) do j := j + 1; b' := b' + b end while
    g := g ∪ {b' + a}, where a ∈ A' satisfies t(b' + a) // note new subgroup generator
    // extend A' to get a set of representatives of the image of the group
    // generated by the first few generators of A in the quotient
    A' := {a + i · b : i ∈ {0, . . . , j - 1}, a ∈ A'}
  end for
  return ⟨g⟩ // a subgroup of A
    
```

This algorithm allows us to find $J(\mathbb{Q}) \cap J^0(\mathbb{Q}_p)$ and hence the image of $J(\mathbb{Q})$ in $J(\mathbb{Q}_p)/J^0(\mathbb{Q}_p)$. It remains to determine the image of $C(\mathbb{Q}_p)$ in this group. It is, however, better to find the image of $C(\mathbb{Q}_p)$ in $J(\mathbb{Q}_p)/J^1(\mathbb{Q}_p)$ directly or, rather, to find the subset of $J(\mathbb{Q})/(J(\mathbb{Q}) \cap J^1(\mathbb{Q}_p))$ which is in the image of $C(\mathbb{Q}_p)$. For this, we use the map to the dual Kummer surface described below in Section 6: for a representative $P \in J(\mathbb{Q})$ of each element of $J(\mathbb{Q})/(J(\mathbb{Q}) \cap J^1(\mathbb{Q}_p))$, we check whether its image on the dual Kummer surface satisfies $p^2 | \eta_4$ and $p | \eta_1 \eta_3 - \eta_2^2$.

The reason for working mod $J^1(\mathbb{Q}_p)$ and not mod $J^0(\mathbb{Q}_p)$ (which might appear to be more efficient) is that there does not seem to be a simple criterion that tells us whether or not we are in $\iota(C(\mathbb{Q}_p)) + J^0(\mathbb{Q}_p)$.

6. ‘Deep’ information

In this section, we work with genus 2 curves over \mathbb{Q} for simplicity. Everything can easily be generalized to genus 2 curves over arbitrary number fields.

Especially for small primes p , we can hope to gain valuable information by looking not just at $J(\mathbb{F}_p)$ or, more generally, $J(\mathbb{Q}_p)/J^1(\mathbb{Q}_p)$ but also into the kernel of reduction to some depth. If $J^n(\mathbb{Q}_p)$ (with $n \geq 1$) denotes the ‘ n th kernel of reduction’, that is, the subgroup of elements consisting of the $p^n \mathbb{Z}_p$ -points of the formal group, then we would like to determine (the image of $J(\mathbb{Q})$ in) $J(\mathbb{Q}_p)/J^n(\mathbb{Q}_p)$ and the image of $C(\mathbb{Q}_p)$ in this group.

The first step is to find $J(\mathbb{Q}) \cap J^n(\mathbb{Q}_p)$. This can be done with the help of the p -adic logarithm on the Jacobian. The power series of the formal logarithm up to terms of degree seven can be found at Victor Flynn’s website [15, local/log]. If higher precision is needed, we perform a p -adic numerical integration, as follows. We can express a given point in the kernel of reduction in the form $[P_1 - P_2]$, where P_1 and P_2 are points on the curve that reduce mod p to the same point. Assuming, for simplicity, that the points have p -adically integral coordinates and do not reduce to a Weierstrass point, we write $P_1 = (\xi + \delta, \eta_1)$ and $P_2 = (\xi - \delta, \eta_2)$. We then write the differentials $\omega_0 = dx/2y$ and $\omega_1 = x dx/2y$ as a power series in terms of the uniformizer $t = x - \xi$, times dt , and integrate this numerically from $t = -\delta$ up to $t = \delta$, to the desired precision (note that δ has positive valuation). Alternatively, we can use the fact that on the Kummer surface,

$$p^{n-1} \cdot P = (\lambda_1^2 p^{2n} + O(p^{3n}) : 2\lambda_1 \lambda_2 p^{2n} + O(p^{3n}) : \lambda_2^2 p^{2n} + O(p^{3n}) : 1)$$

where (λ_1, λ_2) is the logarithm of P . So, to compute the logarithm up to $O(p^n)$, we multiply the point by p^{n-1} on the Kummer surface to find the logarithm up to a sign. (If $p = 2$, we need a few more bits of precision here.) We then fix the sign by comparing this with the first-order

approximation obtained from the functions λ and μ on the Jacobian, in the notation of [10, Section 2].

Given that we are able to compute the logarithm

$$\log : J^1(\mathbb{Q}_p) \longrightarrow (p\mathbb{Z}_p)^2$$

to any desired accuracy, we compute the finite-index subgroup $K_n = J(\mathbb{Q}) \cap J^n(\mathbb{Q}_p)$ of $J(\mathbb{Q})$ as follows. Assume that K_1 is already given. We can therefore set up the group homomorphism

$$K_1 \xrightarrow{\log} (p\mathbb{Z}_p)^2 \longrightarrow \left(\frac{p\mathbb{Z}_p}{p^n\mathbb{Z}_p}\right)^2 \cong \left(\frac{\mathbb{Z}}{p^{n-1}\mathbb{Z}}\right)^2,$$

and then K_n is just its kernel.

The second and more time-consuming step is to find the image of $C(\mathbb{Q}_p)$ in $J(\mathbb{Q}_p)/J^n(\mathbb{Q}_p)$. We assume again that the ‘flat’ information (that is, the image of $C(\mathbb{Q}_p)$ in $J(\mathbb{Q}_p)/J^1(\mathbb{Q}_p)$) is already known. For each point in the intersection of the images of $C(\mathbb{Q}_p)$ and $J(\mathbb{Q})$ inside $J(\mathbb{Q}_p)/J^1(\mathbb{Q}_p)$, we then have to find all its ‘liftings’ to elements in the intersection of the images of $C(\mathbb{Q}_p)$ and $J(\mathbb{Q})$ in $J(\mathbb{Q}_p)/J^n(\mathbb{Q}_p)$.

One approach would be to take some lifting P_0 in $C(\mathbb{Q}_p)$, add representatives of $J(\mathbb{Q}_p)/J^n(\mathbb{Q}_p)$ to it and see which lie sufficiently close to C . One practical problem lies in the word ‘add’. By [10, Chapters 2 and 3], the Jacobian can be embedded into \mathbb{P}^{15} , and the sum $P + Q$ can be expressed in terms of biquadratic forms in the coordinates of P and Q . For a given curve C , these forms can be determined using interpolation, but they can have several thousand terms, so any subsequent computations based on them will be rather slow.

The usual method of adding points on J , following [9], essentially uses some affine part of the Jacobian. Problems with denominators make it not well-suited for p -adic fixed-precision calculations.

We instead propose to use the Kummer surface and its dual (see [10, Chapter 4]). The hyperelliptic involution on C induces an involution on the principal homogeneous space Pic_C^1 of J , and the quotient of Pic_C^1 by this involution is again a quartic surface in \mathbb{P}^3 . An explicit equation is given by

$$\psi(\eta_1, \eta_2, \eta_3, \eta_4) := \begin{vmatrix} 2f_0\eta_4 & f_1\eta_4 & \eta_1 & \eta_2 \\ f_1\eta_4 & 2f_2\eta_4 - 2\eta_1 & f_3\eta_4 - \eta_2 & \eta_3 \\ \eta_1 & f_3\eta_4 - \eta_2 & 2f_4\eta_4 - 2\eta_3 & f_5\eta_4 \\ \eta_2 & \eta_3 & f_5\eta_4 & 2f_6\eta_4 \end{vmatrix} = 0;$$

see [10, p. 33]. This model has the property that the natural image of C in Pic_C^1 is given by $\eta_4 = 0$. Furthermore, if $P \in \text{Pic}_C^1$ maps to $(\eta_1 : \eta_2 : \eta_3 : \eta_4)$ on the dual Kummer surface and $Q \in J$ maps to $(\xi_1 : \xi_2 : \xi_3 : \xi_4)$ on the Kummer surface, then $P \in C \pm Q$ if and only if $\xi_1\eta_1 + \xi_2\eta_2 + \xi_3\eta_3 + \xi_4\eta_4 = 0$. We will denote the Kummer surface by \mathcal{K} and the dual Kummer surface by \mathcal{K}^* .

The group law on J leaves its traces on \mathcal{K} . Consider two points $Q, R \in J$. We write $\mathbf{y} = (\xi_1(Q), \dots, \xi_4(Q))$ and $\mathbf{z} = (\xi_1(R), \dots, \xi_4(R))$ for the projective coordinates of their images on \mathcal{K} . Following [10, Chapter 3], there is a matrix of biquadratic forms $B(\mathbf{y}, \mathbf{z}) = (B_{ij})$ such that

$$2B_{ij} = \xi_i(Q + R)\xi_j(Q - R) + \xi_i(Q - R)\xi_j(Q + R).$$

The action of J on Pic_C^1 can be similarly described on \mathcal{K}^* . Suppose that $Q \in J$ and $P \in \text{Pic}_C^1$ and that $\mathbf{x} = \eta(P)$ and $\mathbf{y} = \xi(Q)$ are projective coordinates for their images on \mathcal{K}^* and on \mathcal{K} , respectively. There is now a symmetric matrix of biquadratic forms $A(\mathbf{x}, \mathbf{y}) = (A_{ij})$ such that

$$2A_{ij} = \eta_i(P + Q)\eta_j(P - Q) + \eta_i(P - Q)\eta_j(P + Q).$$

The following result lets us compute A from B rather easily. We assume that B has been scaled so that $B_{44}(0, 0, 0, 1; 0, 0, 0, 1) = 1$ and A has been scaled so that $A_{11}(1, 0, 0, 0; 0, 0, 0, 1) = 1$.

LEMMA 6.1. *Let \mathbf{x} be coordinates of the image of $P \in \text{Pic}_C^1$ on \mathcal{K}^* , and let \mathbf{y} and \mathbf{z} be coordinates of the images of $Q, R \in J$ on \mathcal{K} . Then, considering \mathbf{x}, \mathbf{y} and \mathbf{z} as row vectors, we have*

$$\mathbf{x} B(\mathbf{y}, \mathbf{z}) \mathbf{x}^\top = \mathbf{z} A(\mathbf{x}, \mathbf{y}) \mathbf{z}^\top.$$

Proof. Both sides are triquadratic forms in \mathbf{x}, \mathbf{y} and \mathbf{z} . Using the duality property mentioned above, it can be checked that each side vanishes if and only if

$$P \in C \pm Q \pm R \quad \text{for some choice of signs.}$$

This implies that the two sides are proportional to each other, and since they take the same value 1 when $\mathbf{x} = (1, 0, 0, 0)$ and $\mathbf{y} = \mathbf{z} = (0, 0, 0, 1)$, they must be equal (as there are no quadrics vanishing on either of the two surfaces). □

So, in order to find A , we construct the polynomial on the left-hand side and interpret it as a quadratic form in \mathbf{z} .

On the Kummer surface, we can use B to find the image of $P + Q$ if the images of P, Q and $P - Q$ are known. This is known as ‘pseudo-addition’ (see [16]) and can be extended to the computation of images of linear combinations $a_1P_1 + \dots + a_mP_m$ provided that the images of the 2^m points $e_1P_1 + \dots + e_mP_m$, where $e_j \in \{0, 1\}$, are known. It should be noted that the complexity of this procedure in terms of pseudo-additions is 2^m times the bit-length of the coefficients, so we should not use it to compute linear combinations of many points. One important feature of the procedure is that it works with projective coordinates and is therefore well-suited for p -adic arithmetic with fixed precision.

In a similar way, we can compute the image of $P + a_1P_1 + \dots + a_mP_m$ on the dual Kummer surface, if $P \in \text{Pic}_C^1$ and $P_1, \dots, P_m \in J$. We need to know the images of $P + e_1P_1 + \dots + e_mP_m$ (with $e_j \in \{0, 1\}$) in addition to $e_1P_1 + \dots + e_mP_m$, and in the pseudo-addition step B is replaced with A . The remark on complexity applies here as well. Below, we will take generators of the successive quotients K_{l-1}/K_l as the P_j ; in most cases, each such quotient is isomorphic to a subgroup of $(\mathbb{Z}/p\mathbb{Z})^2$, so that $m \leq 2$.

The following lemma tells us how to find the subset of $J(\mathbb{Q}_p)/J^n(\mathbb{Q}_p)$ consisting of elements such that the corresponding cosets of $J^n(\mathbb{Q}_p)$ meet the image of the curve.

LEMMA 6.2. *Let $P_0 \in C(\mathbb{Q}_p)$ and $Q \in J^n(\mathbb{Q}_p)$. Let $(\eta_1 : \eta_2 : \eta_3 : \eta_4)$ be coordinates of the image of $P_0 + Q$ on \mathcal{K}^* , normalized so that the minimal p -adic valuation is zero. Then we have*

$$v_p(\eta_1\eta_3 - \eta_2^2) \geq n \quad \text{and} \quad v_p(\eta_4) \geq 2n.$$

Proof. Let P be the image on \mathcal{K}^* of $P_0 \in C(\mathbb{Q}_p)$. If we make an invertible coordinate change over \mathbb{Z}_p on the \mathbb{P}^1 that C maps to, then this induces an invertible coordinate change over \mathbb{Z}_p on the ambient projective spaces of \mathcal{K} and \mathcal{K}^* , which leaves the valuations of $\eta_1\eta_3 - \eta_2^2$ and η_4 invariant. We can therefore assume without loss of generality that the point on the curve is at infinity. Then $P = (1 : 0 : 0 : 0)$.

Since $Q \in J^n(\mathbb{Q}_p)$, its image on \mathcal{K} has coordinates of the form

$$(\alpha p^{2n} : \beta p^{2n} : \gamma p^{2n} : 1) \quad \text{with } \alpha, \beta, \gamma \in \mathbb{Z}_p.$$

Denote the coordinates of the images of $P_0 \pm Q$ on \mathcal{K}^* by $(\eta_1 : \eta_2 : \eta_3 : \eta_4)$ and $(\eta'_1 : \eta'_2 : \eta'_3 : \eta'_4)$. If we evaluate the entries of the matrix A at the coordinates of P and Q , then by the definition of A we have (with suitable scaling)

$$2A(P, Q) = (\eta_1, \eta_2, \eta_3, \eta_4)^\top (\eta'_1, \eta'_2, \eta'_3, \eta'_4) + (\eta'_1, \eta'_2, \eta'_3, \eta'_4)^\top (\eta_1, \eta_2, \eta_3, \eta_4).$$

We obtain

$$\eta_1 \eta'_1 = A(P, Q)_{11} \equiv 1 \pmod{p^{2n}},$$

so that we can scale the coordinates to have

$$\eta_1 \equiv \eta'_1 \equiv 1 \pmod{p^{2n}}.$$

We then find that

$$\eta_4 + \eta'_4 \equiv 2A(P, Q)_{14} \equiv 0 \pmod{p^{2n}} \quad \text{and} \quad \eta_4 \eta'_4 = A(P, Q)_{44} \equiv 0 \pmod{p^{4n}},$$

and this implies

$$\eta_4 \equiv \eta'_4 \equiv 0 \pmod{p^{2n}}.$$

All entries in $A(P, Q)$, except A_{11} , have valuation at least $2n$. In a similar way it follows that

$$\eta_2, \eta'_2, \eta_3, \eta'_3 \equiv 0 \pmod{p^n}$$

and therefore

$$\eta_1 \eta_3 - \eta_2^2 \equiv \eta'_1 \eta'_3 - \eta_2'^2 \equiv 0 \pmod{p^n}$$

as claimed. □

Recall that we have fixed an embedding $\iota : C \rightarrow J$, given by some rational divisor (class) of degree one on C . This induces an isomorphism $\iota : \text{Pic}_C^1 \xrightarrow{\sim} J$. So, in order to test whether an element of $J(\mathbb{Q})/K_n$ is in the image of $C(\mathbb{Q}_p)$, we map a representative in $J(\mathbb{Q})$ to Pic_C^1 via ι^{-1} and then to the dual Kummer surface, and check whether the normalized coordinates of the image satisfy

$$v_p(\eta_4) \geq 2n \quad \text{and} \quad v_p(\eta_1 \eta_3 - \eta_2^2) \geq n.$$

Note that we can compute the image on the dual Kummer surface if we know the images of $e_1 P_1 + \dots + e_m P_m$ on \mathcal{K} and on \mathcal{K}^* , where the P_j are representatives of generators of $J(\mathbb{Q})/K_n$ (with $e_j \in \{0, 1\}$).

If we proceed as just described, then we need to enumerate $J(\mathbb{Q})/K_n$ (of size approximately p^{2n}) in order to find the image of C , which is of size approximately p^n . We can make several improvements in order to reduce the complexity to something closer to the lower bound of $O(p^n)$. One improvement is to compute the images successively for $n = 2, 3, \dots$. When we go from $n = m$ to $n = m + 1$, we only have to consider group elements that map into the image of the curve on the previous level; there will usually be p^2 of these for each of the roughly p^m elements in the previous image. This gives a complexity of p^{m+2} for this step, and a total complexity of $(p^2/(p - 1)) p^n$. This is still worse by a factor of $p^2/(p - 1)$ than what we would get if we could directly compute the images of points from $C(\mathbb{Q}_p)$ in $J(\mathbb{Q}_p)/J^n(\mathbb{Q}_p)$, but it is reasonably good for applications.

We can further improve on this in many cases. Let $P \in J(\mathbb{Q})$ be such that its image on \mathcal{K}^* satisfies $v_p(\eta_1 \eta_3 - \eta_2^2) \geq m$ as above. We work in an affine patch of \mathcal{K}^* such that the image of P has p -adically integral coordinates, and we write $h(P)$ for the function $\eta_1 \eta_3 - \eta_2^2$ evaluated at P in terms of these affine coordinates. The theory of formal groups implies that the map

$$J^m(\mathbb{Q}_p) \longrightarrow \mathbb{F}_p, \quad Q \longmapsto p^{-m}(h(P + Q) - h(P)) \pmod{p}$$

is linear, with kernel containing $J^{m+1}(\mathbb{Q}_p)$. We obtain a linear form $\ell_m : K_m/K_{m+1} \rightarrow \mathbb{F}_p$. If ℓ_m is non-zero, then we only need to evaluate it on a generating set of K_m/K_{m+1} in order to find the points $Q \in K_m$ such that $v_p(h(P + Q)) \geq m + 1$. Since K_m/K_{m+1} usually has two generators, this gives a complexity of order $(2 + p)p^m$: for each of the roughly p^m points P , we have to evaluate ℓ_m on the two generators and then compute the (usually) p lifts to the next level. Note that the linear form is non-zero on $J^m(\mathbb{Q}_p)/J^{m+1}(\mathbb{Q}_p)$ if and only if the reduction

mod p of the image of P on \mathcal{K}^* is non-singular. This will be the case unless $p = 2$ or the corresponding point in $C(\mathbb{F}_p)$ has vanishing y -coordinate. So if p is an odd prime such that the polynomial defining C is not divisible by p , there will be at most six ‘problematic’ classes mod p , contributing at most $6p^2$ to the complexity at each step. The overall complexity is therefore $O(p^n)$ for such primes, which is of the order of the obvious lower bound.

7. Implementation

In this section, we describe a concrete implementation of the Mordell–Weil sieve on genus 2 curves that can be used to prove that a given curve does not have a rational point. For this implementation, the MAGMA computer algebra system [2] was used. Our implementation is available as supplementary material enclosed with this paper [7].

Assume that we are given the following as input:

- (1) the polynomial $f(x)$ on the right-hand side of the equation $y^2 = f(x)$ of the curve C ;
- (2) generators of the Mordell–Weil group $J(\mathbb{Q})$, where J is the Jacobian variety of the curve;
- (3) a rational divisor D of degree three on the curve.

The latter is used to provide the embedding $\iota : C \rightarrow J$, which is given by sending a point $P \in C$ to the class of $P + W - D$ where W is a canonical divisor.

Elements of $J(\mathbb{Q})$ can be represented by divisors of degree two, and divisors can be represented by pairs (a, b) of polynomials as described in Section 5. We let r denote the rank of $J(\mathbb{Q})$.

In the first step, we have to provide the necessary input for the actual sieving procedure, which means that we have to determine the group structure of $J(\mathbb{F}_p)$, the reduction homomorphism $\phi_p : J(\mathbb{Q}) \rightarrow J(\mathbb{F}_p)$ and the image of $C(\mathbb{F}_p)$ in $J(\mathbb{F}_p)$ in terms of this group structure. This involves the computation of $r + t + \#C(\mathbb{F}_p)$ discrete logarithms in the group $J(\mathbb{F}_p)$, where r is the Mordell–Weil rank and t is the number of generators of the torsion subgroup of $J(\mathbb{Q})$. The first $r + t$ of these are needed to find ϕ_p , and the others are needed to find the image of $C(\mathbb{F}_p)$ in $J(\mathbb{F}_p)$, represented by the abstract group G_p . If we restrict to primes p such that $\#J(\mathbb{F}_p)$ is B -smooth, then we can use the Pohlig–Hellman reduction [20] for the computation of the discrete logarithms, so that the complexity of this step is about $r + t + \#C(\mathbb{F}_p)$ (assuming B to be fixed). Therefore, the total effort required for the computation in the first step is approximately

$$(r + t)\#S + \sum_{p \in S} \#C(\mathbb{F}_p) \approx (r + t)\#S + \sum_{p \in S} p \approx \left(r + t + \frac{1}{2} \max S\right) \#S.$$

In the last estimate we have made the simplifying assumption that the primes in S are distributed fairly regularly, so the factor $\frac{1}{2}$ will not be completely accurate. The point is that this is essentially quadratic in $\#S$ or $\max S$. So the relevant question is how far we have to go with $\max S$ in order to collect enough information to make success likely.

A reliable theoretical analysis of this question appears to be rather difficult, although one could try to get some information out of an approach along the lines of Poonen’s heuristic [21]. Therefore we take the following approach. We compute the relevant information for each prime p (such that $\#J(\mathbb{F}_p)$ is B -smooth) in turn. Then we compute the numbers $n(S, N_{l-1-r-j})$ for $j = 0, 1, 2, 3$ in the notation of § 3.1, where S is the set of primes used so far. This can be done incrementally, caching the values of $\#C_{N,p}/\#(G_p/NG_p)$ for later use (these depend only on the gcd of N and the exponent of G_p), and does not cost much time. We stop this part of the computation when

$$\min_j n(S, N_{l-1-r-j}) < \varepsilon$$

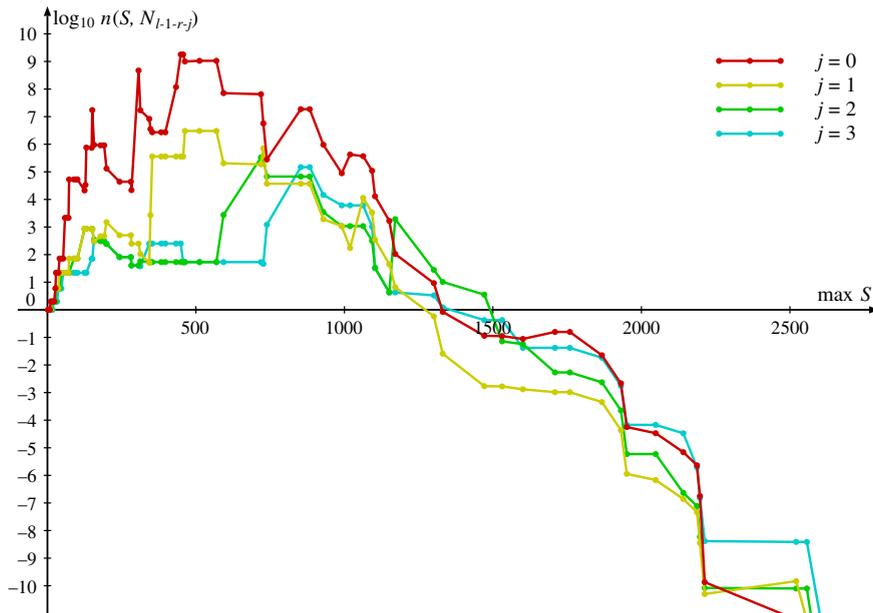


FIGURE 2. Expected sizes $n(S, N_{l-1-r-j})$ versus $\max S$.

for a given parameter $\varepsilon \ll 1$. Tests performed with the ‘small curves’ from [5] indicate that $\varepsilon = 0.01$ is a reasonable choice and that $B = 200$ leads to good results. Figure 2 shows the dependence of $n(S, N_{l-1-r-j})$ from $\max S$ in a fairly typical example (of rank three).

We include the computation of ‘bad’ and ‘deep’ information (as described in Sections 5 and 6 above) as we go along. We take $n = 2, 3, 4$ etc., and when $n = p$ is a prime we compute information mod p if $p < 10$, or if $p \leq B$ and the given model of C is regular at p such that C/\mathbb{F}_p has only one component, or if p is a prime of good reduction and $\#J(\mathbb{F}_p)$ is B -smooth. If $n = p^v$ is a prime power p^v , then we compute information mod p^m with $m = (v + 1)/2$ if v is odd. This scheme proved to give the best performance with our implementation. It hits a good balance between the effort required to compute the information (which is much greater than for ‘flat’ and ‘good’ information at primes $q \approx p^m$) and the gain in speed resulting from the additional information. The information mod p^m is therefore computed in the following order:

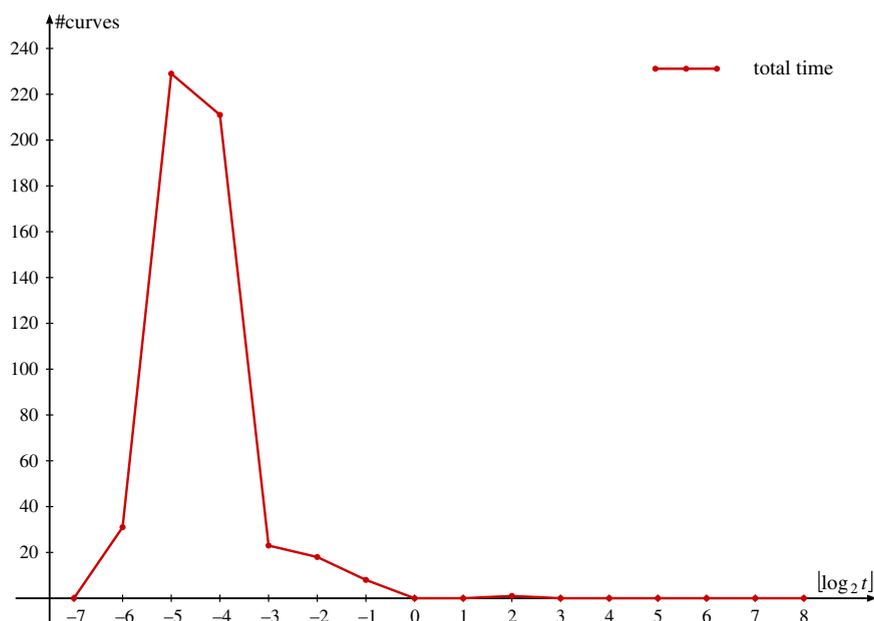
$$p^m = 2, 3, 5, 7, 2^2, 11, 13, 17, 19, 23, 3^2, 29, 31, 2^3, 37, \dots$$

After the information has been collected, we compute a ‘ q sequence’ as described in § 3.2, using a target value of ε_1 with $\varepsilon < \varepsilon_1 < 1$. We take $\varepsilon_1 = 0.1$ as the standard value of this parameter. Since $\varepsilon_1 > \varepsilon$, we know from the first part of the computation that a suitable sequence exists. If we take ε_1 not too close to ε , this second part of the computation is usually rather fast.

Finally, we use the collection $\{(G_p, \phi_p, C_p) : p \in S\}$ and the q sequence as input for the actual sieve computation. This computation is done as described in § 3.3. If it does not result in the desired contradiction, we divide the ε and ε_1 parameters by 10 and start over (keeping the local information we have already computed).

8. Efficiency

How long do our computations take? Let us look at the various steps that have to be performed, in the context of the first application discussed in Section 4 above: verifying that a given curve C of genus 2 over \mathbb{Q} does not have rational points. We assume that a Mordell–Weil basis is known.

FIGURE 3. Running times for $r = 1$.

Note that in practice, the part of the computation that determines this Mordell–Weil basis can be rather time-consuming; but this is a different problem, and we will not consider it here. See [25–27] for the relevant algorithms. We also assume that we know a rational divisor of degree three on C . Again, it may not be easy to find such a divisor in practice.

We consider the 1447 curves for which we had to perform a Mordell–Weil sieve computation in [5] so as to rule out the existence of rational points. The number of curves differs from the 1492 mentioned earlier because some curves had rank zero, while some others could be ruled out immediately by the information coming from the Birch and Swinnerton-Dyer conjecture. The timings mentioned below were obtained on a machine with 4 GB of RAM and a 2.0 GHz dual core processor. As before, r denotes the Mordell–Weil rank.

Among the 521 curves with $r = 1$, a contradiction was obtained for 514 already while collecting the information. This occurs when we find a prime p or prime power p^n such that the images of $J(\mathbb{Q})$ and $C(\mathbb{Z}/p^n\mathbb{Z})$ in $J(\mathbb{Z}/p^n\mathbb{Z})$ are disjoint. It is perhaps worth noting that without looking at ‘bad’ and ‘deep’ information, we obtain this kind of immediate contradiction only for 406 curves. The average computing time for a single curve was about 0.1 seconds, and the longest time was about 6.3 seconds. The distribution of running times is shown in Figure 3 (on a logarithmic scale).

The anonymous referee asked whether there is a heuristic explanation for the observation that information at one prime is almost always enough to rule out rational points. Here is an attempt at such an explanation. We use the following probabilistic model. Assume that $J(\mathbb{F}_p)$ is cyclic of order uniformly distributed in an interval around p^2 of length of the order of $p^{3/2}$, that the generator P_0 of $J(\mathbb{Q})$ (which we assume to be torsion-free of rank one) is mapped to a random element of $J(\mathbb{F}_p)$ and that the points in $C(\mathbb{F}_p)$ form a random subset of $J(\mathbb{F}_p)$. We are interested in the probability that $C(\mathbb{F}_p)$ and the image of $J(\mathbb{Q})$ in $J(\mathbb{F}_p)$ are disjoint. Note that the case where $J(\mathbb{F}_p)$ is cyclic is the worst case; if $J(\mathbb{F}_p)$ is not cyclic, then the cyclic image of $J(\mathbb{Q})$ will be more likely to be small.

LEMMA 8.1. *In the model described above, the probability that $C(\mathbb{F}_p)$ does not meet the image of $J(\mathbb{Q})$ is bounded below by a constant times $1/p$.*

Proof. Let $n = p^2 + O(p^{3/2})$ be the order of $J(\mathbb{F}_p)$, denote the index of the image of $J(\mathbb{Q})$ in $J(\mathbb{F}_p)$ by d , and let $m = p + O(p^{1/2})$ denote $\#C(\mathbb{F}_p)$. Then the conditional probability, given that the index is $d \geq 2$, is

$$\begin{aligned} q_d &= \frac{\binom{n-n/d}{m}}{\binom{n}{m}} = \prod_{k=0}^{m-1} \left(1 - \frac{1}{d(1-k/n)}\right) \\ &= \exp\left(\sum_{k=0}^{m-1} \log\left(1 - \frac{1}{d(1-k/n)}\right)\right) = \exp\left(-\sum_{k=0}^{m-1} \frac{1}{d(1-k/n)} + O(pd^{-2})\right) \\ &= \exp\left(-\frac{1}{d} \int_0^m \frac{dt}{1-t/n} + O(d^{-1}) + O(pd^{-2})\right) \\ &= \exp\left(\frac{n}{d} \log\left(1 - \frac{m}{n}\right) + O(d^{-1}) + O(pd^{-2})\right) \\ &= \exp\left(-\frac{m}{d} + O(d^{-1}) + O(pd^{-2})\right). \end{aligned}$$

Here $O(d^{-1})$ denotes a quantity that is bounded by a constant times d^{-1} , and $O(pd^{-2})$ denotes a quantity that is bounded by a constant times pd^{-2} for large p .

We restrict to the range $\alpha p \leq d \leq \beta p$ with fixed $0 < \alpha < \beta$. Then

$$q_d = e^{-m/d}(1 + O(p^{-1})) = e^{-p/d}(1 + O(p^{-1/2})).$$

We now have to estimate the probability that d has a given value d_0 in the range under consideration. Fix a generator Q of $J(\mathbb{F}_p)$ and write $\bar{P}_0 = k \cdot Q$, where \bar{P}_0 is the image of P_0 in $J(\mathbb{F}_p)$. Then the probability is

$$\begin{aligned} \Pr(d = d_0) &= \frac{\#\{(n, k) : 0 \leq k < n = p^2 + O(p^{3/2}), \gcd(n, k) = d_0\}}{\#\{(n, k) : 0 \leq k < n = p^2 + O(p^{3/2})\}} \\ &= \frac{6}{\pi^2 d_0^2} (1 + O(p^{-1/2+\epsilon})). \end{aligned}$$

So the total probability can be bounded below by

$$\begin{aligned} \sum_{\alpha p \leq d_0 \leq \beta p} \Pr(d = d_0)q_{d_0} &= \frac{6}{\pi^2} \sum_{\alpha p \leq d_0 \leq \beta p} \frac{1}{d_0^2} e^{-p/d_0} (1 + O(p^{-1/2+\epsilon})) \\ &= \frac{6}{\pi^2} \left(\int_{\alpha p}^{\beta p} e^{-p/t} \frac{dt}{t^2}\right) (1 + O(p^{-1/2+\epsilon})) \\ &= \frac{6}{\pi^2 p} \left(\int_{1/\beta}^{1/\alpha} e^{-u} du\right) (1 + O(p^{-1/2+\epsilon})) \\ &= \frac{6}{\pi^2 p} (e^{-1/\beta} - e^{-1/\alpha}) + O(p^{-3/2+\epsilon}). \end{aligned}$$

Letting $\alpha \rightarrow 0$ and $\beta \rightarrow \infty$, we obtain

$$\liminf_{p \rightarrow \infty} p \cdot \Pr(C(\mathbb{F}_p) \cap \langle \bar{P}_0 \rangle = \emptyset) \geq \frac{6}{\pi^2}. \quad \square$$

Since the cases with $d_0 \ll p$ and $d_0 \gg p$ are not likely to contribute anything in the limit, we would expect that in the model considered, we actually have

$$\Pr(C(\mathbb{F}_p) \cap \langle \bar{P}_0 \rangle = \emptyset) \sim \frac{6}{\pi^2} \cdot \frac{1}{p} \quad \text{as } p \rightarrow \infty.$$

Since $\sum_p p^{-1}$ diverges, we expect an infinite (but rather sparse) set of primes p for which information mod p proves that there are no rational points. This is consistent with the

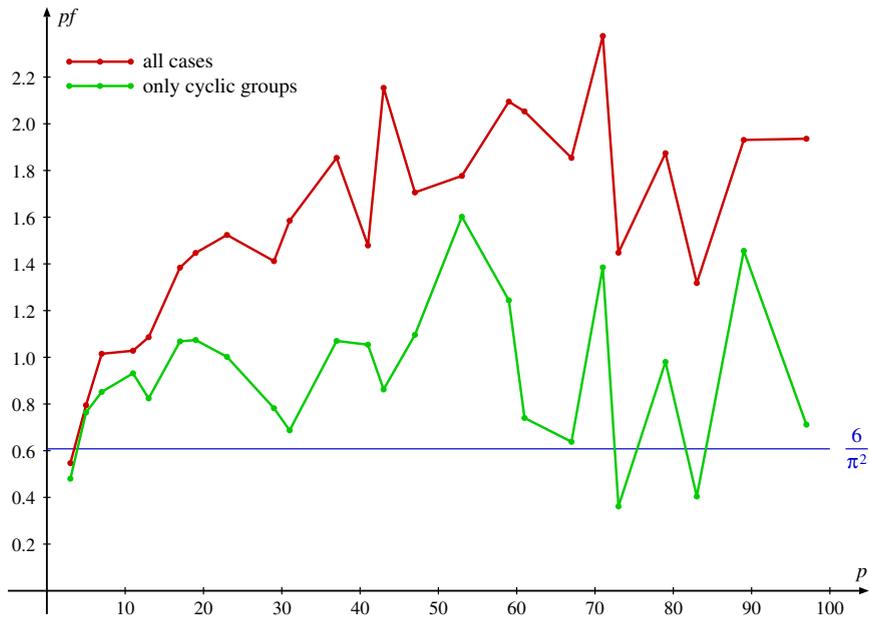


FIGURE 4. The prime p times the success frequency at p plotted against p .

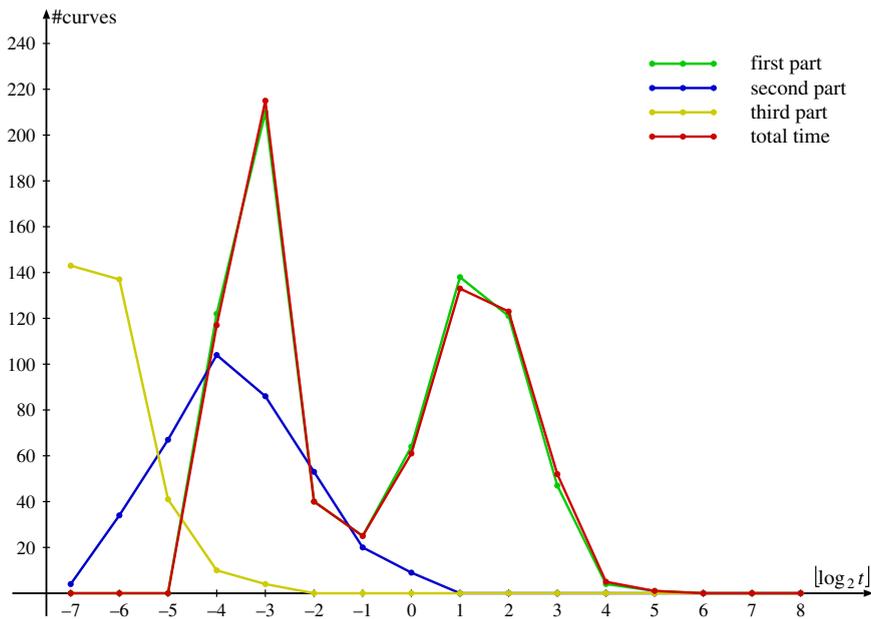


FIGURE 5. Running times for $r = 2$.

observations above. Figure 4 shows, in red, as a function of $2 < p < 100$, p times the fraction of curves in our data set where reduction mod p proves the absence of rational points among all curves with $r = 1$ and trivial torsion that have good reduction at p . We see that (except for $p = 3$) this value is considerably larger than $6/\pi^2$. The most likely explanation, that this is an effect of the occurrence of non-cyclic groups among the $J(\mathbb{F}_p)$, is confirmed by the data obtained from looking only at cases where $J(\mathbb{F}_p)$ is cyclic (plotted in green in Figure 4).

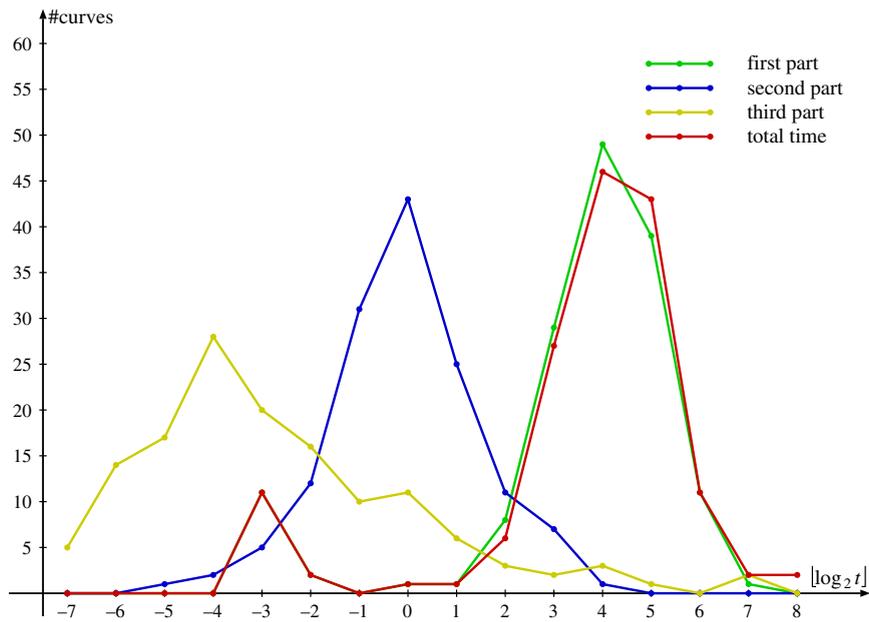


FIGURE 6. Running times for $r = 3$.

In general, a similar heuristic approach should give a success probability on the order of p^{-r} when the rank is r . This indicates that there is a positive probability for success at some single prime, but that this probability is less than 1 and decreases to zero as r increases. This is consistent with the observations described below.

There are 772 curves with $r = 2$. For 394 of them, we obtain a contradiction from one prime or prime power alone. The average computing time for these curves was 0.24s, with a maximum of 6.4s. For the remaining curves, the average total computing time was 4.9s, with a maximum of 51.8s. The distribution of the running times (overall and for the various parts of the computation) is shown in Figure 5. The two peaks essentially correspond to the two groups of curves. The largest size of a set $A(L)$ that occurred in the computation was 236, and the average of this maximum size in each computation was 6.1. Note that the inclusion of ‘bad’ and ‘deep’ information results in a speed-up by roughly a factor of two.

There are 152 curves with $r = 3$. For 14 curves we still find a contradiction from the local information at one prime alone. The average total time was 34.3 s, and the maximum was about 5.6 minutes. The first step took 28.1 s on average. For the curves where the second and third steps were performed, the second step took 2.3 s and the third step 4.6 s, on average. The distribution of the running times (overall and for the various steps) is shown in Figure 6. The largest size of a set $A(L)$ was 251 148 (occurring for the curve with the largest running time), and the average was 5049. For these curves, the computation is infeasible without using ‘bad’ and ‘deep’ information, since otherwise the sets $A(L)$ occurring in the last part of the computation get much too large.

There are only two curves with $r = 4$. One of them is ‘hard’ and the other is ‘easy’. For the ‘hard’ curve, the computation takes about 26 minutes with the standard settings (about 2 minutes for the first step, 10s for the second and the remaining 24 minutes for the sieving step). This is mostly due to the large size of the sets $A(L)$ (up to more than two million) occurring in this computation. If we change the parameters so that deep information mod p^n is used for all $p^n < 520$, then the computation takes less than 12 minutes (3 min + 10 s + 8.5 min), and the largest set $A(L)$ has a size of about 750 000 only. The ‘easy’ curve is dealt with in 47 s (44.5 s + 2 s + 0.5 s) using the standard settings.

From these data, we conclude that our current implementation works well for curves with Jacobians of Mordell–Weil rank $r \leq 3$. For larger rank, there is so far only sparse evidence from examples, suggesting that individual curves with r as large as 6 are still within the range of feasibility. In any case, it is clear that average running times increase quickly with r .

Our timings also show that the first part of the computation (gathering the local information) usually takes the lion’s share of the total time. Improvements in this part (and faster discrete logarithm computations in particular) would result in a noticeable speed-up of the procedure as a whole.

Acknowledgements. We would like to thank Victor Flynn and Bjorn Poonen for useful discussions related to our project. Further thanks go to the anonymous referee for some helpful remarks. For the computations, the MAGMA system [2] was used.

References

1. A. BAKER and G. WÜSTHOLZ, *Logarithmic forms and Diophantine geometry*, New Mathematical Monographs 9 (Cambridge University Press, Cambridge, 2007).
2. W. BOSMA, J. CANNON and C. PLAYOUST, ‘The Magma algebra system I: The user language’, *J. Symb. Comp.* 24 (1997) 235–265. Also see the MAGMA home page at <http://magma.maths.usyd.edu.au/magma/>.
3. N. BRUIN, ‘The arithmetic of Prym varieties in genus 3’, *Compositio Math.* 144 (2008) 317–338.
4. N. BRUIN and N. D. ELKIES, ‘Trinomials $ax^7 + bx + c$ and $ax^8 + bx + c$ with Galois groups of order 168 and $8 * 168$ ’, *Algorithmic number theory: 5th international symposium, ANTS-V (Sydney, Australia, July 2002) proceedings*, Lecture Notes in Computer Science 2369 (eds Claus Fieker and David R. Kohel; Springer, Berlin, 2002) 172–188.
5. N. BRUIN and M. STOLL, ‘Deciding existence of rational points on curves: an experiment’, *Experiment. Math.* 17 (2008) 181–189.
6. N. BRUIN and M. STOLL, ‘2-cover descent on hyperelliptic curves’, *Math. Comp.* 78 (2009) 2347–2370.
7. N. BRUIN and M. STOLL, ‘MWSieve-new.m’, MAGMA code for Mordell–Weil sieve computation, 2009, electronic appendix to ‘The Mordell–Weil sieve: proving non-existence of rational points on curves’, *LMS J. Comput. Math.* 13 (2010) 272–306, doi: 10.1112/S1461157009000187.
8. Y. BUGEAUD, M. MIGNOTTE, S. SIKSEK, M. STOLL and SZ. TENGELY, ‘Integral points on hyperelliptic curves’, *Algebra Number Theory* 2 (2008) 859–885.
9. D. G. CANTOR, ‘Computing in the Jacobian of a hyperelliptic curve’, *Math. Comp.* 48 (1987) 95–101.
10. J. W. S. CASSELS and E. V. FLYNN, *Prolegomena to a middlebrow arithmetic of curves of genus 2* (Cambridge University Press, Cambridge, 1996).
11. C. CHABAUTY, ‘Sur les points rationnels des courbes algébriques de genre supérieur à l’unité’, *C. R. Acad. Sci. Paris* 212 (1941) 882–885 (in French).
12. R. F. COLEMAN, ‘Effective Chabauty’, *Duke Math. J.* 52 (1985) 765–770.
13. E. V. FLYNN, ‘A flexible method for applying Chabauty’s theorem’, *Compositio Math.* 105 (1997) 79–94.
14. E. V. FLYNN, ‘The Hasse principle and the Brauer–Manin obstruction for curves’, *Manuscripta Math.* 115 (2004) 437–466.
15. E. V. FLYNN, FTP site with formulas relating to genus 2 curves, <http://people.maths.ox.ac.uk/~flynn/genus2/>.
16. E. V. FLYNN and N. P. SMART, ‘Canonical heights on the Jacobians of curves of genus 2 and the infinite descent’, *Acta Arith.* 79 (1997) 333–352.
17. R. HARTSHORNE, *Algebraic geometry*, Graduate Texts in Mathematics 52 (Springer, New York, 1977).
18. W. MCCALLUM and B. POONEN, The method of Chabauty and Coleman, Preprint, 2007, <http://www-math.mit.edu/~poonen/papers/chabauty.pdf>. Proceedings of the 2004 IHP Trimestre on Explicit Methods in Number Theory, in the ‘Panorama & Syntheses’ series of the SMF, to appear.
19. V. K. MURTY and J. SCHERK, ‘Effective versions of the Chebotarev density theorem for function fields’, *C. R. Acad. Sci. Paris Sér. I Math.* 319 (1994) 523–528.
20. G. C. POHLIG and M. E. HELLMAN, ‘An improved algorithm for computing logarithms over $GF(p)$ and its cryptographic significance’, *IEEE Trans. Inform. Theory* IT-24 (1978) 106–110.
21. B. POONEN, ‘Heuristics for the Brauer–Manin obstruction for curves’, *Experiment. Math.* 15 (2006) 415–420.
22. B. POONEN, E. F. SCHAEFER and M. STOLL, ‘Twists of $X(7)$ and primitive solutions to $x^2 + y^3 = z^7$ ’, *Duke Math. J.* 137 (2007) 103–158.
23. V. SCHARASCHKIN, ‘Local-global problems and the Brauer–Manin obstruction’, PhD Thesis, University of Michigan, 1999.
24. I. R. SHAFAREVICH (ed.), *Algebraic geometry I*, Encyclopaedia of Mathematical Sciences 23 (Springer, Berlin, 1994).

25. M. STOLL, ‘On the height constant for curves of genus two’, *Acta Arith.* 90 (1999) 183–201.
26. M. STOLL, ‘Implementing 2-descent on Jacobians of hyperelliptic curves’, *Acta Arith.* 98 (2001) 245–277.
27. M. STOLL, ‘On the height constant for curves of genus two, II’, *Acta Arith.* 104 (2002) 165–182.
28. M. STOLL, ‘Independence of rational points on twists of a given curve’, *Compositio Math.* 142 (2006) 1201–1214.
29. M. STOLL, ‘Finite descent obstructions and rational points on curves’, *Algebra Number Theory* 1 (2007) 349–391.
30. M. STOLL, ‘Applications of the Mordell–Weil sieve’, *Oberwolfach Rep.* 4 (2007) 1967–1970.
31. M. STOLL, ‘How to obtain global information from computations over finite fields’, *Higher-dimensional geometry over finite fields*, NATO Science for Peace and Security Series: Information and Communication Security 16 (eds D. Kaledin and Y. Tschinkel; IOS Press, Amsterdam, 2008) 189–196.

Nils Bruin
Department of Mathematics
Simon Fraser University
Burnaby, BC, V5A 1S6
Canada

nbruin@cecm.sfu.ca

Michael Stoll
Mathematisches Institut
Universität Bayreuth
95440 Bayreuth
Germany

Michael.Stoll@uni-bayreuth.de