

Computing zeta functions in families of $C_{a,b}$ curves using deformation

Wouter Castryck², Hendrik Hubrechts^{1*}, and Frederik Vercauteren^{2*}

¹ Department of Mathematics, University of Leuven,
Celestijnenlaan 200B, B-3001 Leuven-Heverlee, Belgium
`hendrik.hubrechts@wis.kuleuven.be`

² Department of Electrical Engineering, University of Leuven
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium
`firstname.lastname@esat.kuleuven.be`

Abstract. We apply deformation theory to compute zeta functions in a family of $C_{a,b}$ curves over a finite field of small characteristic. The method combines Denef and Vercauteren's extension of Kedlaya's algorithm to $C_{a,b}$ curves with Hubrechts' recent work on point counting on hyperelliptic curves using deformation. As a result, it is now possible to generate $C_{a,b}$ curves suitable for use in cryptography in a matter of minutes.

1 Introduction

The development of algorithms that compute the Hasse-Weil zeta function of a curve over a finite field has witnessed several revolutions in the past 20 years, partly motivated by applications in cryptography. The first was the Schoof-Elkies-Atkin algorithm [16] to compute the number of points on an elliptic curve over a finite field. Although this algorithm readily generalises to higher genus, it is not really practical except in the genus 2 case for moderately sized finite fields [6]. The second revolution was the canonical lift approach introduced by Satoh [15] and reinterpreted by Mestre [14] using the AGM. Extensions and improvements of this algorithm (an overview is given in [2]) resulted in very efficient point counting methods for ordinary elliptic and hyperelliptic curves over finite fields of small characteristic. The third revolution was the p -adic cohomological approach introduced by Kedlaya [9] and Lauder and Wan [11]. Although the resulting algorithms are polynomial time for fixed characteristic, they are only practical for hyperelliptic curves. Finally, the fourth revolution consists of two components, deformation and fibration, and was introduced by Lauder [12, 13] to compute the zeta function of higher dimensional hypersurfaces.

Despite the efforts of many researchers, the ultimate goal of having a set of algorithms that can handle any given curve of genus g over any finite field \mathbb{F}_q where q^g is limited to having several hundred bits, is still far off. In fact, up to the time of writing, only the case of elliptic curves (both in large and small

* Postdoctoral Fellow of the Research Foundation - Flanders (FWO)

characteristic) and the case of hyperelliptic and superelliptic [5] curves in small characteristic have a satisfactory solution.

Although tiny steps towards tackling the large characteristic case have been made [4], handling all curves over finite fields of small characteristic looks much more feasible. In the latter case, there has been partial progress to include $C_{a,b}$ [3] and non-degenerate curves [1], but these algorithms are not sufficiently practical. Although the approach is similar to Kedlaya's algorithm for hyperelliptic curves, these algorithms use a different Frobenius lifting technique, which makes them slow.

The goal of this paper is to remedy this situation by taking a totally different approach based on deformation theory. Although this theory was primarily introduced for high dimensional hypersurfaces, Hubrechts [8, 7] showed it to be efficient in the hyperelliptic case.

The advantage of using deformation for the broader classes of $C_{a,b}$, non-degenerate or even more general curves is twofold: firstly, it avoids the explicit computation of the Frobenius lift that makes the algorithms in [3] and [1] slow and secondly, the core of the algorithms, i.e. solving a p -adic differential equation, is always the same. Only the computation of the so-called connection matrix differs for each class of curves, but is in itself a much easier problem than developing an efficient differential reduction method as needed in Kedlaya's approach.

In this paper we present a detailed version of this method for $C_{a,b}$ curves, which should readily extend to non-degenerate curves. Our algorithm is used in two applications: firstly, given a random $C_{a,b}$ curve over a finite field \mathbb{F}_q , compute its zeta function and secondly, given a finite field \mathbb{F}_q , generate $C_{a,b}$ curves whose Jacobian has nearly prime order for use in cryptography. The speed-up over known techniques for the second application is remarkable: after a precomputation, computing the zeta function of each member of a family with a Jacobian of 160-bit order only takes a few seconds. As a result, generating cryptographically useful $C_{a,b}$ curves now is feasible in a matter of minutes.

The remainder of this paper is organised as follows: Section 2 reviews p -adic cohomology and deformation for general curves and Section 3 covers the necessary background on $C_{a,b}$ curves. Section 4 studies relative Monsky-Washnitzer cohomology for a family of $C_{a,b}$ curves, resulting in a practical algorithm described and analysed in Section 5. Finally, Section 6 reports on a preliminary Magma implementation of this algorithm.

2 p -Adic cohomology and deformation

Throughout this section, the survey paper on Monsky-Washnitzer cohomology by van der Put [17] and the AWS 2007 lecture notes by Kedlaya [10, Chapter 3] are implicit references.

2.1 Zeta functions and cohomology.

Let \mathbb{F}_q be a finite field of characteristic p with q elements. The zeta function of a polynomial $\overline{C}(x, y) \in \mathbb{F}_q[x, y]$ defining a non-singular affine curve is determined

by the action of the Frobenius endomorphism $\overline{\mathcal{F}}_q : \overline{A} \rightarrow \overline{A} : a \mapsto a^q$ on a certain cohomology space $H_{MW}^1(\overline{A}/\mathbb{Q}_q)$. Here \overline{A} denotes $\mathbb{F}_q[x, y]/(\overline{C}(x, y))$ and $H_{MW}^1(\overline{A}/\mathbb{Q}_q)$ is constructed as follows. Let \mathbb{Q}_q be an unramified extension of the field of p -adic numbers \mathbb{Q}_p , with valuation ring \mathbb{Z}_q and residue field \mathbb{F}_q . Let $C(x, y) \in \mathbb{Z}_q[x, y]$ be such that it reduces to $\overline{C}(x, y) \pmod{p}$ and consider the \mathbb{Z}_q -algebra

$$A^\dagger = \frac{\mathbb{Z}_q\langle x, y \rangle^\dagger}{(C(x, y))}$$

where $\mathbb{Z}_q\langle x, y \rangle^\dagger$ is the *weak completion* of $\mathbb{Z}_q[x, y]$. It consists of power series $\sum a_{i,j} x^i y^j \in \mathbb{Z}_q[[x, y]]$ for which there is a $\rho \in]0, 1[$ such that $|a_{i,j}|_p / \rho^{i+j} \rightarrow 0$ as $i+j \rightarrow \infty$. The idea behind this convergence condition is that $\mathbb{Z}_q\langle x, y \rangle^\dagger$ should be closed under integration. Let $D^1(A^\dagger)$ be the universal module of differentials on A^\dagger over \mathbb{Z}_q and let $d : A^\dagger \rightarrow D^1(A^\dagger)$ be the usual exterior derivation. Then $H_{MW}^1(\overline{A}/\mathbb{Q}_q)$ is defined as

$$\frac{D^1(A^\dagger)}{d(A^\dagger)} \otimes_{\mathbb{Z}_q} \mathbb{Q}_q,$$

which turns out to be the right object for the following theorem to hold.

Theorem 1 (Monsky, Washnitzer). *There exists a \mathbb{Z}_q -algebra endomorphism $\mathcal{F}_q : A^\dagger \rightarrow A^\dagger$ that lifts $\overline{\mathcal{F}}_q$ in the sense that $\overline{\mathcal{F}}_q \circ \pi = \pi \circ \mathcal{F}_q$, where $\pi : A^\dagger \rightarrow \overline{A}$ is reduction mod p . For any such lift, the induced map $\mathcal{F}_q^* : D^1(A^\dagger) \rightarrow D^1(A^\dagger)$ is well-defined modulo $d(A^\dagger)$ and acts on $H_{MW}^1(\overline{A}/\mathbb{Q}_q)$ as an invertible \mathbb{Q}_q -vector space morphism, which does not depend on the choice of \mathcal{F}_q . Moreover, the zeta function of \overline{C} is given by*


$$Z_{\overline{C}}(T) = \frac{\det(\mathbb{I} - q\mathcal{F}_q^{*-1} | H_{MW}^1(\overline{A}/\mathbb{Q}_q))}{1 - qT}.$$

2.2 Relative cohomology

Let $\overline{C}(x, y, t) \in \mathbb{F}_q[t][x, y]$ define a family of smooth curves over an open dense subset $\text{Spec } \overline{S}$ of the affine t -line. Thus $\overline{S} = \mathbb{F}_q[t, \overline{r}(t)^{-1}]$ for some nonzero $\overline{r}(t) \in \mathbb{F}_q[t]$. Write $\overline{A} = \overline{S}[x, y]/(\overline{C}(x, y, t))$ and, for every $\overline{t}_0 \in \mathbb{F}_q$ where $\overline{r}(t)$ does not vanish, write $\overline{A}_{\overline{t}_0}$ for $\overline{A}/(t - \overline{t}_0)$, the coordinate ring of the fiber at \overline{t}_0 . Then the aim of relative cohomology is to describe how the action of Frobenius on $H_{MW}^1(\overline{A}_{\overline{t}_0}/\mathbb{Q}_q)$ alters as \overline{t}_0 varies. Let $C(x, y, t) \in \mathbb{Z}_q[t][x, y]$ and $r(t) \in \mathbb{Z}_q[t]$ be such that they reduce mod p to $\overline{C}(x, y, t)$ and $\overline{r}(t)$ respectively. Define $S^\dagger = \mathbb{Z}_q\langle t, r(t)^{-1} \rangle^\dagger = \mathbb{Z}_q\langle t, z \rangle^\dagger / (zr(t) - 1)$ along with the S^\dagger -module

$$A^\dagger = \frac{\mathbb{Z}_q\langle t, r(t)^{-1}, x, y \rangle^\dagger}{(C(x, y, t))}$$

(the weak completion being realised as in the bivariate case). Note that there is a well-defined p -adic valuation on S^\dagger and A^\dagger . Let $D_t^1(A^\dagger)$ be the universal module of differentials on A^\dagger over S^\dagger and let $d_t : A^\dagger \rightarrow D_t^1(A^\dagger)$ be the corresponding



exterior derivation. Thus in all this, t is left constant. Write $S_{\mathbb{Q}_q}^\dagger = S^\dagger \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$. Then our object of interest is the $S_{\mathbb{Q}_q}^\dagger$ -module

$$H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger) = \frac{D_t^1(A^\dagger)}{d_t(A^\dagger)} \otimes_{\mathbb{Z}_q} \mathbb{Q}_q.$$

As above, one can show that there exists a \mathbb{Z}_q -algebra endomorphism \mathcal{F}_q on A^\dagger that lifts the Frobenius action $\bar{\mathcal{F}}_q$ on \bar{A} . Moreover, one can realise that $\mathcal{F}_q(t) = t^q$ (we will illustrate this in Section 4 in our specific families of $C_{a,b}$ curves). The induced map \mathcal{F}_q^* on $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)$ is well-defined, though in general it is not an $S_{\mathbb{Q}_q}^\dagger$ -module endomorphism.

Let $\bar{t}_0 \in \mathbb{F}_q$ be a non-zero of $\bar{r}(t)$ and let $\hat{t}_0 \in \mathbb{Z}_q$ be its Teichmüller lift, i.e. the unique root of $X^q - X \in \mathbb{Z}_q[X]$ that reduces to $\bar{t}_0 \bmod p$. Then one sees that $H_{MW}^1(\bar{A}_{\bar{t}_0}/\mathbb{Q}_q)$ can be identified with $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)/(t - \hat{t}_0)$, and that \mathcal{F}_q^* induces a well-defined map on $H_{MW}^1(\bar{A}_{\bar{t}_0}/\mathbb{Q}_q)$ which exactly matches with the Frobenius action described in Theorem 1.

In summary, the action of Frobenius on a single fiber can be obtained from the relative Frobenius action by substituting for t a suitable Teichmüller representative. So one could think of the relative Frobenius action as an interpolation of the Frobenius actions on all fibers in the family.

2.3 The Gauss-Manin connection

In addition to the notation from above, we introduce $D^1(A^\dagger)$, denoting the module of differentials on A^\dagger over \mathbb{Z}_q (so t is no longer left constant), and $D^2(A^\dagger) = \wedge^2 D^1(A^\dagger)$, denoting the corresponding module of 2-forms. Let d be the usual exterior derivation, both on A^\dagger and $D^1(A^\dagger)$, giving rise to the Monsky-Washnitzer complex

$$0 \rightarrow A^\dagger \xrightarrow{d} D^1(A^\dagger) \xrightarrow{d} D^2(A^\dagger) \rightarrow 0$$

of the surface $\text{Spec } \bar{A}$ over \mathbb{F}_q . Note that we have a natural surjective morphism $D^1(A^\dagger) \rightarrow D_t^1(A^\dagger) : dt \mapsto 0$, thus we can identify $D_t^1(A^\dagger)$ with $D^1(A^\dagger)/(dt)$.

Definition 1. *The Gauss-Manin connection*

$$\nabla : H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger) \rightarrow H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger) : \omega \mapsto \nabla(\omega)$$

is constructed as follows. For large enough $e \in \mathbb{N}$, let $p^e \omega$ be represented by a 1-form $\tilde{\omega} \in D^1(A^\dagger)$ and take its exterior derivative $d(\tilde{\omega}) \in D^2(A^\dagger)$, which one can always write as $\tilde{\varphi} \wedge dt$ for some $\tilde{\varphi} \in D^1(A^\dagger)$. Reduce $\tilde{\varphi}$ modulo (dt) to end up in $D_t^1(A^\dagger)$. Then reduce modulo $d_t(A^\dagger)$ and tensor with p^{-e} , so that one ends up in $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)$: this is $\nabla(\omega)$.

We leave it to the reader to show that the above is well-defined, i.e. $\nabla(\omega)$ does not depend on the choice of e , $\tilde{\omega}$ and $\tilde{\varphi}$. Remark that the above construction does not result in a geometric connection in the usual sense of the word, in which case ∇ should take values in $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger) \otimes D^1(S_{\mathbb{Q}_q}^\dagger)$. But for our purposes, we prefer to think of the Gauss-Manin connection as mapping $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)$ into itself. Then the following observation is the key towards deformation theory.

Theorem 2. *One has $\nabla \circ \mathcal{F}_q^* = qt^{q-1} \circ \mathcal{F}_q^* \circ \nabla$, where qt^{q-1} denotes the corresponding multiplication map on $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)$.*

Proof. (sketch only) This follows from the commutativity of the diagram of \mathbb{Z}_q -module morphisms

$$\begin{array}{ccc} D^1(A^\dagger) & \xrightarrow{d} & D^2(A^\dagger) \\ \downarrow \mathcal{F}_q^* & & \downarrow \mathcal{F}_q^* \\ D^1(A^\dagger) & \xrightarrow{d} & D^2(A^\dagger). \end{array}$$

2.4 Deformation.

Suppose that $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)$ is finitely generated and free over $S_{\mathbb{Q}_q}^\dagger$, having a basis that for any $\bar{t}_0 \in \mathbb{F}_q$ for which $\bar{r}(\bar{t}_0) \neq 0$, reduces mod $(t - \hat{t}_0)$ to a basis of $H_{MW}^1(\bar{A}_{\bar{t}_0}/\mathbb{Q}_q)$. Here, \hat{t}_0 is the Teichmüller lift of \bar{t}_0 . In Section 4 we will prove this assumption for our concrete families of $C_{a,b}$ curves.

Let s_1, \dots, s_d be an $S_{\mathbb{Q}_q}^\dagger$ -basis of $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)$ and let $F = (F_{i,j}), G = (G_{i,j})$ be $(d \times d)$ -matrices with entries in $S_{\mathbb{Q}_q}^\dagger$ such that

$$\mathcal{F}_q^*(s_j) = \sum_{i=1}^d F_{i,j} s_i, \quad \nabla(s_j) = \sum_{i=1}^d G_{i,j} s_i$$

for $j = 1, \dots, d$. Then the quasi-commutativity of the Gauss-Manin connection with the Frobenius action gives rise to a first-order differential equation

$$G \cdot F - \frac{d}{dt} F = qt^{q-1} \cdot F \cdot G(t^q).$$

This allows one to compute F from an initial value. Typically, this is the matrix of Frobenius acting on $H_{MW}^1(\bar{A}_{\bar{t}_0}/\mathbb{Q}_q) = H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)/(t - \hat{t}_0)$ for some ‘easy’

fiber $\text{Spec } \overline{A}_{\hat{t}_0}$. When expressed with respect to the basis $s_1(\hat{t}_0), \dots, s_d(\hat{t}_0)$ of $H_{MW}^1(\overline{A}_{\hat{t}_0}/\mathbb{Q}_q)$, this exactly matches with $F(\hat{t}_0)$. If $n := \log_p q$ is large, a substantial speed-up in the algorithms can be achieved by working with the matrix F_p of the p^{th} power Frobenius $\overline{\mathcal{F}}_p : \overline{A} \rightarrow \overline{A}$ suitably acting on $H_{MW}^1(\overline{A}/S_{\mathbb{Q}_q}^\dagger)$, and then reconstructing F as $F_p^{\sigma^{n-1}} \cdot F_p^{\sigma^{n-2}} \cdots F_p^\sigma \cdot F_p$. Here $\sigma : S_{\mathbb{Q}_q}^\dagger \rightarrow S_{\mathbb{Q}_q}^\dagger$ maps t to t^p , acts on \mathbb{Q}_q by Frobenius substitution and extends by linearity and continuity. Furthermore, F_p can be computed from an initial $F_p(\hat{t}_0)$ as the solution to the differential equation

$$G \cdot F_p - \frac{d}{dt} F_p = pt^{p-1} \cdot F_p \cdot G^\sigma(t^p). \quad (1)$$

3 Generalities on $C_{a,b}$ curves

3.1 Definition and first properties.

Let a and b be coprime integers ≥ 2 and let k be any field. An algebraic curve C/k is said to be $C_{a,b}$ if it admits a non-singular affine ‘Weierstrass model’

$$C(x, y) = y^a + c_{b,0}x^b + \sum_{ai+bj < ab} c_{i,j}x^i y^j \in k[x, y] \quad (c_{b,0} \neq 0). \quad (2)$$

Such a model has a unique, generally singular point at infinity. One can prove that this point is dominated by a single place P on the non-singular model, and the pole divisors of x and y are aP and bP respectively. Since a and b are coprime, this allows us to determine the pole divisor of any function $f(x, y)$ in the affine coordinate ring $A = k[x, y]/(C)$. Indeed, using $C(x, y) = 0$ one can write

$$f(x, y) = \sum_{j=0}^{a-1} \sum_{i=0}^{\deg_x f} f_{i,j} x^i y^j,$$

in which no two monomials have the same pole order at P . Hence $-\text{ord}_P(f) = \max\{ai + bj \mid i = 0, \dots, \deg_x f; j = 0, \dots, a-1, f_{i,j} \neq 0\}$, and the Weierstrass semigroup

$$\{-\text{ord}_P(f) \mid f \in k(C) \setminus \{0\}\} \subset \mathbb{N}$$

of P equals $a\mathbb{N} + b\mathbb{N}$. From the Riemann-Roch theorem it follows that the geometric genus of C equals $g = (a-1)(b-1)/2$. Hyperelliptic curves of genus g having a rational Weierstrass point are $C_{2,2g+1}$, and are therefore special instances of $C_{a,b}$ curves.

Let $\Delta \subset \mathbb{R}^2$ be the convex hull of $(0, 0)$, $(b, 0)$ and $(0, a)$. It contains (and generically equals) the Newton polytope of $C(x, y)$. Then the following property is a key feature of $C_{a,b}$ curves. One can copy the proof of [3, Lemma 1], replacing \mathbb{F}_q and \mathbb{Z}_q with k and R respectively.

Lemma 1 (Effective Nullstellensatz). *Let R be a discrete valuation ring or a field with maximal ideal \mathfrak{m} . Let $\overline{C}(x, y) \in \frac{R}{\mathfrak{m}}[x, y]$ define a $C_{a,b}$ curve. Let*

$C(x, y) \in R[x, y]$ be such that it reduces to $\overline{C}(x, y) \pmod{\mathfrak{m}}$ and such that it is again supported on Δ . Then there exist polynomials $\alpha, \beta, \gamma \in R[x, y]$ that are supported on 2Δ , such that $1 = \alpha C + \beta C_x + \gamma C_y$. In particular, if $C(x, y)$ was chosen to be monic in y , it defines a $C_{a,b}$ curve over the fraction field of R .

Here C_x (resp. C_y) denotes $\partial C / \partial x$ (resp. $\partial C / \partial y$).

3.2 Cohomology.

Write $A = k[x, y]/(C)$ and suppose first that $\text{char}(k) = 0$. Then in [3], it is shown that

$$\{x^r y^s dx \mid r = 0, \dots, b-2; s = 1, \dots, a-1\} \quad (3)$$

is a basis for the k -vector space $H_{DR}^1(A/k) = D^1(A)/d(A)$. The proof moreover gives an explicit procedure to express a differential form $\omega \in D^1(A)$ in terms of this basis: using $C(x, y) = 0$ and the exactness of forms of the type $d(x^r y^s)$, one immediately sees that $H_{DR}^1(A/k)$ is generated by $x^r y^s dx$ for $0 < s < a$. These generators are totally ordered by $-\text{ord}_P$ and as long as $r \geq b-1$, each of them can be rewritten in terms of forms $x^r y^s dx$ having strictly smaller pole order. This is because

$$\omega_{r,s} = x^{r-(b-1)} y^s dC - d \left(x^{r-(b-1)} \left(\frac{a}{a+s} y^{a+s} + \sum_{ai+bj < ab} \frac{j c_{i,j}}{s+j} x^i y^{s+j} \right) \right)$$

is exact, and after expanding and reducing mod $C(x, y)$ one can check that its pole order is determined by the term $\lambda x^r y^s dx$, where

$$\lambda = \left(b + (r - b + 1) \frac{a}{a+s} \right) c_{b,0} \neq 0.$$

Therefore, subtracting $\omega_{r,s}/\lambda$ from $x^r y^s dx$ reduces the pole order. Continuing in this way will reduce everything onto the basis.

Next, suppose that $k = \mathbb{F}_q$. Let $\tilde{C}(x, y) \in \mathbb{Z}_q[x, y]$ be monic in y and supported on Δ , such that it reduces to $C(x, y) \pmod{p}$. Let \tilde{A}^\dagger be the weak completion of $\tilde{A} = \mathbb{Z}_q[x, y]/(\tilde{C})$. Then from [3, Lemma 4], it follows that the canonical map

$$H_{DR}^1(\tilde{A}/\mathbb{Q}_q) = \frac{D^1(\tilde{A})}{d(\tilde{A})} \otimes \mathbb{Q}_q \longrightarrow H_{MW}^1(A/\mathbb{Q}_q) = \frac{D^1(\tilde{A}^\dagger)}{d(\tilde{A}^\dagger)} \otimes \mathbb{Q}_q$$

is an isomorphism, more precisely the above reduction process converges. Since \tilde{C} is $C_{a,b}$ by Lemma 1, the set given in (3) is a \mathbb{Q}_q -basis for $H_{MW}^1(A/\mathbb{Q}_q)$.

3.3 Families of $C_{a,b}$ curves.

Let k be any field. Let $C(x, y, t) \in k[t][x, y]$ be supported on Δ and suppose that the coefficient of y^a is 1. Let $c_{b,0}(t) \in k[t]$ be the coefficient of x^b . Denote

the monic polynomial generating the $k[t]$ -ideal $(C, C_x, C_y) \cap k[t]$ by $f(t)$ and let $\mathfrak{r}(t) = c_{b,0}(t)f(t)$. One can then check that for any $t_0 \in \bar{k}$, $C(x, y, t_0) \in \bar{k}[x, y]$ defines a $C_{a,b}$ curve (in Weierstrass form) if and only if $\mathfrak{r}(t_0) \neq 0$. We will say that $C(x, y, t)$ defines a (one-dimensional) family of $C_{a,b}$ curves if $\mathfrak{r}(t) \neq 0$; the polynomial $\mathfrak{r}(t)$ will be referred to as the *resultant* of the family. Any $C(x, y, t)$ defining a family of $C_{a,b}$ curves gives rise to a flat family of smooth curves

$$\text{Spec } \frac{k[t][x, y]}{(C)} \rightarrow \text{Spec } k[t, \mathfrak{r}(t)^{-1}].$$

A condition equivalent to $\mathfrak{r}(t) \neq 0$ is: $C(x, y, t)$ defines a $C_{a,b}$ curve over the function field $k(t)$. Indeed, consider the system of equations $C = C_x = C_y = z\mathfrak{r}(t) - 1$, where z is a new variable. It has no solutions over \bar{k} , and therefore there are polynomials $\alpha, \beta, \gamma, \delta \in k[x, y, z, t]$ for which

$$1 = \alpha C + \beta C_x + \gamma C_y + \delta(z\mathfrak{r}(t) - 1).$$

If $\mathfrak{r}(t) \neq 0$, we can replace z by $1/\mathfrak{r}(t)$, to get an expansion

$$1 = \alpha' C + \beta' C_x + \gamma' C_y, \quad \alpha', \beta', \gamma' \in k(t)[x, y]. \quad (4)$$

Together with $c_{b,0}(t) \neq 0$ this implies that $C(x, y, t)$ indeed defines a $C_{a,b}$ curve over $k(t)$. Conversely, for any expansion (4), $f(t)$ must divide the least common multiple of the denominators appearing in α', β' and γ' and can therefore not be zero. Together with $c_{b,0}(t) \neq 0$ this gives $\mathfrak{r}(t) \neq 0$.

The above observation allows us to bound the degree of the resultant.

Lemma 2. *Let $C(x, y, t)$ define a family of $C_{a,b}$ curves and let $\mathfrak{r}(t) \in k[t]$ be its resultant. Then $\deg \mathfrak{r}(t) \leq (9g + 6(a + b) - 1) \deg_t C$.*

Proof. Let $\alpha, \beta, \gamma \in k(t)[x, y]$ be as in the effective Nullstellensatz (Lemma 1) applied over $k(t)$. The coefficients of α, β and γ can be obtained by solving a system of $\#(3\Delta \cap \mathbb{Z}^2)$ equations in $3\#(2\Delta \cap \mathbb{Z}^2)$ unknowns. Both numbers are bounded by $9g + 6(a + b) - 2$. Now by Cramer's theorem, the denominators of α, β and γ can be chosen to be the determinant $d(t)$ of some fixed minor matrix of our system; since $d(t)$ contains $f(t)$ as a factor, $\deg f(t)$ is clearly bounded by $(9g + 6(a + b) - 2) \deg_t C$. Together with $\deg c_{b,0}(t) \leq \deg_t C$, this gives the desired result.

Lemma 3. *Let R be a discrete valuation ring with maximal ideal \mathfrak{m} and suppose that $\bar{C}(x, y, t) \in (R/\mathfrak{m})[t][x, y]$ defines a family of $C_{a,b}$ curves. Let $C(x, y, t) \in R[t][x, y]$ be supported on Δ , such that it reduces to $\bar{C}(x, y, t)$ mod \mathfrak{m} and such that the coefficient of y^a is 1. Then $C(x, y, t)$ defines a family of $C_{a,b}$ curves (over the fraction field K of R).*

Proof. This follows from Lemma 1, when applied over the discrete valuation ring $R[t]_{\mathfrak{m}R[t]}$, i.e. the subring of $K(t)$ consisting of rational functions that can be written as a quotient of two integral polynomials whose denominator does not reduce to zero modulo \mathfrak{m} .

4 Relative MW cohomology of a family of $C_{a,b}$ curves

Let $\overline{C}(x, y, t) \in \mathbb{F}_q[t][x, y]$ define a family of $C_{a,b}$ curves. Let $C(x, y, t) \in \mathbb{Z}_q[t][x, y]$ lift $\overline{C}(x, y, t)$ such that it is monic in y and again supported on Δ . By Lemma 3, $C(x, y, t)$ defines a family of $C_{a,b}$ curves over \mathbb{Q}_q .

Instead of the resultant $\mathfrak{r}(t)$ of $C(x, y, t)$, we will work with a possibly larger polynomial $r(t) = c_{b,0}(t)d(t)$, where $d(t)$ is obtained as in the proof of Lemma 2 by linear algebra over the discrete valuation ring $\mathbb{Z}_q[t]_{p\mathbb{Z}_q[t]}$ (see also the proof of Lemma 3). In particular, $d(t)$ has p -adic valuation 0 and there exists a completely integral Nullstellensatz expansion

$$r(t) = \alpha C + \beta C_x + \gamma C_y \quad (5)$$

where α, β and γ are supported (in x and y) on 2Δ and where $\deg r(t)$, $\deg_t \alpha$, $\deg_t \beta$ and $\deg_t \gamma$ are bounded by $(9g + 6(a + b) - 1)\tau$, with $\tau := \deg_t C(x, y, t)$.

Let $\overline{r}(t)$ be the reduction modulo p . Since (5) is integral, it follows that $\overline{C}(x, y, t)$ defines a family of smooth curves over $\text{Spec } \mathbb{F}_q[t, \overline{r}(t)^{-1}]$, so the theory explained in Section 2 applies. We inherit the notation introduced there, where for simplicity we drop the lower indices from d_t and D_t . Below, we give a basis for $H_{MW}^1(\overline{A}/S_{\mathbb{Q}_q}^\dagger)$ and discuss the action of Frobenius on it. We will intensively make use of [1] and [3], so the proof-verifying reader should take these references at hand. The following lemma is easily proved.

Lemma 4. *Let $f(t, z) \in S_{\mathbb{Q}_q}^\dagger$ have p -adic valuation ν . There is only a finite number of Teichmüller elements \hat{t}_0 in $\overline{\mathbb{Z}}_q$ for which both $r(\hat{t}_0) \neq 0$ and the p -adic valuation of $f(\hat{t}_0, r(\hat{t}_0)^{-1})$ is $> \nu$.*

Lemma 5. *Let $r, s \in \mathbb{N}$ with $0 \leq s < a$. Then in $D^1(A^\dagger)$, $x^r y^s dx$ can be rewritten as*

$$\sum_{j=1}^{a-1} \sum_{i=0}^{b-2} \alpha_{i,j}(t, z) x^i y^j dx + d \left(\sum_{j=0}^{a-1} \sum_{i=0}^{r+b+1} \beta_{i,j}(t, z) x^i y^j \right),$$

where

1. $\alpha_{i,j}$ and $\beta_{i,j}$ are polynomial expressions of degree $\leq (a+b)(9g+7a+6b-1)\tau$ in t , and of degree $\leq ar + b$ in z ;
2. $p^m \alpha_{i,j}$ and $p^m \beta_{i,j}$ are integral, with $m = \lfloor \log_p((r+1)a + sb) \rfloor + 4(a-1)b \lfloor \log_p(2a-1) \rfloor$.

Proof. One can follow the procedure described in Section 3.2. The factor $1/c_{b,0}(t)$ that is introduced in each reduction step can be rewritten as $\frac{r(t)}{c_{b,0}(t)}z$, which is a polynomial expression of degree at most $(9g+6(a+b)-2)\tau$ in t and of degree 1 in z . The $\alpha_{i,j}(t, z)$ are obtained by subsequently (i) expanding $x^r y^s dx - \omega_{r,s}/\lambda(t)$ and (ii) consecutively substituting $y^a - C(x, y, t)$ for y^a until only monomial

forms of the type $x^i y^j dx$ with $j < a$ remain, so that one can start over again. The corresponding operations to compute the $\beta_{i,j}(t, z)$ are (i) computing

$$x^{r-(b-1)} \left(\frac{a}{a+s} y^{a+s} + \sum_{ai+bj < ab} \frac{j c_{i,j}(t)}{s+j} x^i y^{s+j} \right)$$

and (ii) substituting $y^a - C(x, y, t)$ for y^a until only monomial forms of the type $x^i y^j$ with $j < a$ remain. Since there are at most $ar+bs-a(b-2)-b(a-1) < ar+b$ reduction steps, the degree bounds follow.

The $r+b+1$ bound on the degree in x in the $d(\dots)$ -part follows from the fact that all terms that are introduced have pole order $\leq (r-(b-1))a + (2a-1)b$.

The bound on the p -adic valuations follows from the above lemma, together with [3, Lemma 4].

Corollary 1. $\{x^r y^s dx \mid r = 0, \dots, b-2; s = 1, \dots, a-1\}$ is an $S_{\mathbb{Q}_q}^\dagger$ -module basis of $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)$.

Proof. Linear independence follows from the corresponding statement in the absolute case. To see that it is a generating set, note that the above lemma implies the convergence of the reduction process described in Section 3.2.

Next, we determine the action of the p^{th} power Frobenius on this basis. To this end, we construct a \mathbb{Z}_p -algebra endomorphism $\mathcal{F}_p : A^\dagger \rightarrow A^\dagger$ (along with explicit bounds on its rate of convergence) that lifts $\bar{A} \rightarrow \bar{A} : \bar{a} \mapsto \bar{a}^p$. The concrete aim is to find polynomials $\delta_x, \delta_y \in \mathbb{Z}_q[x, y, t, z]$ and overconvergent series $W, Z \in p\mathbb{Z}_q\langle x, y, t, z \rangle^\dagger$ such that

$$\mathcal{F}_p : \begin{cases} x \mapsto x^p(1 + \delta_x W) \\ y \mapsto y^p(1 + \delta_y W) \\ t \mapsto t^p \\ z \mapsto z^p + Z \end{cases}$$

(acting on \mathbb{Z}_q by Frobenius substitution σ) extends by linearity and continuity to a well-defined map $A^\dagger \rightarrow A^\dagger$, i.e. modulo the relations $C(x, y, t) = 0$ and $r(t)z - 1 = 0$. Using Newton iteration, the latter relation allows one to determine Z , which should satisfy $r^\sigma(t^p)(z^p + Z) - 1 = 0$. As we are not interested in its rate of convergence, we move on to the determination of W .

The former relation implies that W should satisfy

$$H(W) := C^\sigma(x^p(1 + \delta_x W), y^p(1 + \delta_y W), t^p) = 0$$

over A^\dagger . We try to find δ_x and δ_y such that this equation can be solved using Newton iteration, starting from the approximate solution $W = 0$. From (5) it follows that

$$1 = z\alpha C + z\beta C_x + z\gamma C_y - (r(t)z - 1),$$

so we can take $\delta_x = z^p \beta^p$ and $\delta_y = z^p \gamma^p$: indeed, then $H(W) = 0$ satisfies the initial conditions for Newton iteration over A^\dagger . To find a unique representative however, we will instead solve

$$\tilde{H}(W) := H(W) - C^p + (\tau^p z^p - 1 - z^p \alpha^p C^p)W = 0,$$

for which these conditions are satisfied over the base ring $\mathbb{Z}_q \langle x, y, t, z \rangle^\dagger$.

If we expand $\tilde{H}(W) = \sum h_k W^k$, one verifies that the polynomials $h_k \in \mathbb{Z}_q[x, y][t][z]$ are supported on

$$(2k+1)p\Delta \times (\chi k + \tau)p[0, 1] \times kp[0, 1] \subset \mathbb{R}^4$$

where $\chi = \max\{\deg_t \alpha, \deg_t \beta, \deg_t \gamma\} \leq (9g + 6(a+b) - 1)\tau$. This is contained in $(k+1)p\Delta_{t,z}$, where $\Delta_{t,z} = 2\Delta \times [0, \chi] \times [0, 1]$. Proceeding as in [1], we finally find that $1 + \delta_x W \bmod p^N, 1 + \delta_y W \bmod p^N$ are supported on $5p(N+1)\Delta_{t,z}$, for any $N \in \mathbb{N}$.

Lemma 6. *Let $F_p(t, z)$ be a matrix of the induced action of \mathcal{F}_p on $H_{MW}^1(\bar{A}/S_{\mathbb{Q}_q}^\dagger)$, with respect to the basis $\{x^r y^s dx \mid r = 0, \dots, b-2; s = 1, \dots, a-1\}$. Then for any $N \in \mathbb{N}$ we can represent any entry of $F_p(t, z)$ modulo p^N as a polynomial of degree $\leq 7p(a+b)(ab+1)(N+\theta+1)\kappa\tau$ in t and of degree $\leq 7p(a+b)(ab+1)(N+\theta+1)$ in z . Here $\kappa = 9g + 7a + 6b - 1$ and θ is the smallest positive integer satisfying $\theta \geq \log_p(8pab(a+b)(N+\theta+1)) + 4ab \log_p(2a)$.*

Proof. Write $\mu = p(1 + 5(a+b-2)(N+\theta+1))$. Then one can check that the differential form $x^i y^j dx$ is mapped to an expression $x^{p-1} f dx$, where f is supported modulo $p^{N+\theta}$ on $\mu\Delta_{t,z}$ (use that $(i, j, 0, 0) \in \Delta_{t,z}$ and that $i+j+1 \leq a+b-2$). Rewrite the polynomial $x^{p-1} f \bmod p^{N+\theta}$ as $\sum_{j=0}^{a-1} \sum_{i=0}^r f_{i,j}(t, z) x^i y^j$ by subsequently substituting $y^a - C(x, y, t)$ for y^a . Since there are less than $a\mu$ substitution steps, this adds at most $a\tau\mu$ to the degree in t . Therefore $\deg_t f_{i,j} \leq \chi\mu + a\tau\mu = \kappa\tau\mu$ and $\deg_z f_{i,j} \leq \mu$. By pole order arguments, one finds $r \leq p-1 + b\mu$. Following Lemma 5, this reduces further to

$$x^{p-1} f dx \equiv \sum_{j=1}^{a-1} \sum_{i=0}^{b-2} f'_{i,j}(t, z) x^i y^j dx,$$

where the congruence is valid modulo p^N since the valuations of the denominators introduced during reduction are bounded by

$$[\log_p((p+b\mu)a + (a-1)b)] + 4(a-1)b[\log_p(2a-1)] \leq \theta.$$

Moreover, $\deg_z f'_{i,j} \leq a(p-1) + (ab+1)\mu + b$ and $\deg_t f'_{i,j} \leq (a(p-1) + (ab+1)\mu + b)\kappa\tau$. One can verify that $a(p-1) + (ab+1)\mu + b \leq 7p(a+b)(ab+1)(N+\theta+1)$.

5 Our deformation algorithm

We follow the strategy explained in Section 2.4, applied to families of the type considered in Section 4. In this, we aim for two applications: (i) computing the

zeta function of a given $C_{a,b}$ curve over a given finite field \mathbb{F}_q , and (ii) generating $C_{a,b}$ curves having an (almost) prime order Jacobian over a given finite field \mathbb{F}_q for given a and b .

As for (i), let $\overline{C}_1(x, y) \in \mathbb{F}_q[x, y]$ be the $C_{a,b}$ curve of interest and let $\overline{C}_0(x, y)$ be a $C_{a,b}$ curve defined over the prime subfield \mathbb{F}_p . E.g. one can take $\overline{C}_0(x, y) = y^a - x^b + \varphi(x, y)$ where $\varphi(x, y) = 1$ if $p \nmid a, b$, and $\varphi(x, y) = y$ resp. $\varphi(x, y) = x$ if $p \mid a$ resp. $p \mid b$. Then our family of interest is $\overline{C}(x, y, t) = t\overline{C}_1(x, y) + (1 - t)\overline{C}_0(x, y)$ and the goal is to compute $F_p(1)$ from $F_p(0)$ by solving equation (1).

In (ii), we take a ‘random’ family $\overline{C}(x, y, t) \in \mathbb{F}_p[t][x, y]$ and compute $F_p(t)$ from $F_p(0)$ by solving equation (1). Afterwards, we substitute various Teichmüller elements $\hat{t}_0 \in \mathbb{Z}_q$ until we find a curve with an (almost) prime order Jacobian. We remark that some special families are unsuited for this application, such as the supersingular family $y^2 = x^3 + tx$ with $t \in \mathbb{F}_q$ and $q \equiv 3 \pmod{4}$.

First we compute the polynomial $r(t)$ as explained in the beginning of Section 4. Note that $\bar{r}(t)$ contains the actual resultant as an in general non-trivial factor, so it may accidentally happen that e.g. $\bar{r}(0) = 0$ or $\bar{r}(1) = 0$. We will assume that this is *not* the case, i.e. all fibers of interest correspond to non-roots of $\bar{r}(t)$.

Before describing the main steps of the algorithm we define several constants. As before, $\tau = \deg_t C(x, y, t)$, and we define $\rho := \deg r(t)$, so that $\rho = \mathcal{O}(g\tau)$. We will (see [3]) have to compute both $F_p(t)$ and $F_p(1)$ modulo p^m with

$$m := \left\lceil \log_p \left(2 \binom{2g}{g} q^{g/2} \right) \right\rceil + (g+1)ng \log_p a.$$

Let $\alpha := (2g - 1)g \log_p a + g$ and $\gamma := 2g^2 \log_p a + g$, and choose θ and κ as in Lemma 6, where the accuracy N is now equal to m . Now we define $M := 7p(a+b)(ab+1)(m+\theta+1)$ and $\ell := \kappa\tau M + \rho M + 1$. The matrices $F_p(0)$ and G will be computed with p -adic accuracy $\varepsilon := m + (5\gamma + 1)\lceil \log_p \ell \rceil + 12\alpha$ and all computations are modulo t^ℓ .

5.1 Step I: Computing $F_p(0)$

In all instances, we can reduce to the case where the 0-fiber $\overline{C}_0(x, y)$ is defined over the prime subfield \mathbb{F}_p . Computing the Frobenius matrix of such a curve is of course easier, but note that we need the Frobenius matrix $F_p(0)$ up to a much higher precision than required for computing the zeta function of $\overline{C}_0(x, y)$.

Currently, we use two very basic methods. The first method consists of computing $F_p(0)$ using the $C_{a,b}$ -algorithm described in [3]. For the basic forms of the 0-fiber suggested above, the action of Frobenius has much nicer properties than in the general case, thereby circumventing the problem we originally set out to solve. The second method is more efficient and relies on the extension of Kedlaya’s algorithm to superelliptic curves [5]. Note that all basic forms suggested in application (i) above fall in the category of superelliptic curves, so the algorithm of [5] applies. However, the basis used in [5] is different from ours, so we need to apply a basis transformation obtained by reducing our basis onto the basis of [5] using the reduction procedure given there.

5.2 Step II: Computing G

To compute the Gauss-Manin connection ∇ , we simply apply Definition 1. For each basis differential $x^i y^j dx$ with $i = 0, \dots, b-2$ and $j = 1, \dots, a-1$, we rewrite $d(x^i y^j dx)$ as $\varphi_{i,j} \wedge dt$ to obtain $\nabla(x^i y^j dx) = \varphi_{i,j}$.

Define $\beta' = \beta/r(t)$ and $\gamma' = \gamma/r(t)$ with β, γ as in Equation (5), i.e. $1 \equiv \beta' C_x + \gamma' C_y \pmod{C(t)}$, then a short computation shows that

$$d(x^i y^j dx) = x^i j y^{j-1} (\beta' C_x + \gamma' C_y) dy \wedge dx = x^i j y^{j-1} (\gamma' dx - \beta' dy) C_t \wedge dt.$$

So all that remains to do is to apply the reduction formulae given in Section 3.2 to $x^i j y^{j-1} (\gamma' dx - \beta' dy) C_t$, the result of which gives a column of G .

Note that $x^i j y^{j-1} (\gamma' dx - \beta' dy) C_t$ can be rewritten as $h_{i,j} dx$ where $h_{i,j}$ is supported (in x and y) on 4Δ . So the pole order is at most $4ab$ and we can write $h_{i,j}$ in terms of $x^k y^\ell$ with $0 \leq \ell < a$ and $0 \leq k \leq 4b$. From Lemma 5 it follows that the entries of G are of degree $\leq (4ab + b)(9g + 7a + 6b - 1)\tau$ in t and of degree $\leq 4ab + b$ in z . The p -adic valuations of the denominators are $\tilde{\mathcal{O}}(g)$.

5.3 Step III: Solving the differential equation

We first reformulate the differential equation in a way that ensures that the coefficients as well as the solution modulo p^m of the equation are all polynomials, rather than just rational functions or power series in t . From Lemma 6 above, it follows that $K(t) := r(t)^M \cdot F_p(t) \pmod{p^m}$ has polynomial entries of degree less than ℓ . Let $d_G(t) \in \mathbb{Z}_q[t]$ be a factor of some power of $r(t)$ such that $d_G(t)G(t)$ consists of polynomials. As follows from the end of Section 5.2, we can take $\deg d_G(t) = \mathcal{O}(g^2\tau)$. Rewriting equation (1) using $K(t)$ gives

$$r(t) \frac{dK(t)}{dt} - \left(M \frac{dr(t)}{dt} + r(t)G(t) \right) \cdot K(t) + K(t) \cdot (pt^{p-1}r(t)G^\sigma(t^p)) = 0. \quad (6)$$

After multiplying this equation with $d_G(t)d_G^\sigma(t^p)$ we find an equation of the form $A \frac{dK}{dt} B + AKX + YKB$, where $A(t) := r(t)d_G(t)$, $B(t) := d_G^\sigma(t^p)$,

$$X(t) := pt^{p-1}d_G^\sigma(t^p)G^\sigma(t^p) \quad \text{and} \quad Y(t) := -Md_G(t) \frac{dr(t)}{dt} - r(t)d_G(t)G(t),$$

all consisting of polynomials of degree bounded by $\mathcal{O}(g^2\tau)$. In [7, Theorem 2], it is explained how to solve this equation for $K(t)$ with precision (p^m, t^ℓ) respectively $K(1) \pmod{p^m}$, given that the initial precision p^ε is large enough. From [3] it follows that $\text{ord}_p(K(t)) \geq -g \log_p a$, and as shown in [8, Lemma 18] this implies that $\text{ord}_p(K^{-1}(t)) \geq -\alpha$. Let $C(t) = \sum_i C_i t^i$ and $D(t) = \sum_i D_i t^i$ be matrices in $\mathbb{Q}_q[[t]]^{2g \times 2g}$ that satisfy $A \frac{dC}{dt} + YC = 0$, $C(0) = \mathbb{I}$, and $\frac{dD}{dt} B + DX = 0$, $D(0) = \mathbb{I}$ respectively. Denote with C'_i and D'_i the respective coefficients of t^i in $C(t)^{-1}$ and $D(t)^{-1}$. As shown in [8, Proposition 20] for $C(t)$ and $C(t)^{-1}$ and in [7, Section 3.2] for $D(t)$ and $D(t)^{-1}$ we then have that

$$\text{ord}_p(C_i), \text{ord}_p(C'_i), \text{ord}_p(D_i), \text{ord}_p(D'_i) \geq -\gamma \cdot \lceil \log_p(i+1) \rceil - 2\alpha.$$

These properties, together with straightforward estimates on the valuation of A , A^{-1} , B , B^{-1} , X and Y , guarantee that working modulo p^ε suffices for finding the correct result modulo p^m as proved in [7, Theorem 2].

For application (ii) we now have to compute the Teichmüller lift \hat{t}_0 and compute $F_p(\hat{t}_0) = K(\hat{t}_0)r(\hat{t}_0)^{-M}$. Application (i) requires us only to compute $F_p(1) = K(1)r(1)^{-M}$. As final steps the calculation of the q^{th} power Frobenius and the characteristic polynomial of Frobenius are needed, but for this we can refer to Steps 9 and 10 of the algorithm in [1]. The loss of precision in these steps is easily seen to be at most $(g+1)ng \log_p a$, where $n = \log_p q$, so that working modulo p^m guarantees correctness of the zeta function modulo p to the power $\left\lceil \log_p \left(2 \binom{2g}{g} q^{g/2} \right) \right\rceil$. The latter precision allows us to determine the zeta function correctly, as follows from the Weil conjectures, see [3, Section 4].

5.4 Complexity analysis

We will throughout suppose that asymptotically fast arithmetic is used [18]. We see that $m = \mathcal{O}(g^2 n \log a)$ and $\varepsilon = \tilde{\mathcal{O}}(g^2 n)$. From the analysis in [3] it is clear that Step I requires both time and space $\tilde{\mathcal{O}}(g^6 n^2)$. For Step II and application (ii) we need to reduce $2g$ basis elements (each one requiring $\mathcal{O}(g)$ steps) and the objects have size $\mathcal{O}(g\tau\varepsilon)$, whence working over the prime field requires time $\mathcal{O}(g^5 n\tau)$. In the situation of application (i) this step needs time $\mathcal{O}(g^5 n^2\tau)$.

Next we need an estimate on

$$\zeta := \max\{\deg A + \deg B, \deg A + \deg X + 1, \deg Y + \deg B + 1\}.$$

From the estimates in Step II we see that $\zeta = \mathcal{O}(g^2\tau)$. Now Theorem 2 from [7] shows that the computation of $F_p(t)$ requires time $\tilde{\mathcal{O}}(\ell\zeta g^\omega \varepsilon) = \tilde{\mathcal{O}}(g^{9+\omega} n^2 \tau^2)$ (with ω as an exponent for matrix multiplication, e.g. $\omega = 2.376$ [18]) and space $\mathcal{O}(\ell g^2 \varepsilon) = \tilde{\mathcal{O}}(g^9 n^2 \tau)$. Note that in the estimates in [7] we have to take $n = 1$ as we are working over the field \mathbb{Q}_p .

For application (i) the time requirements are $\tilde{\mathcal{O}}(\ell\zeta g^\omega n\varepsilon) = \tilde{\mathcal{O}}(g^{9+\omega} n^3 \tau^2)$ and we need $\mathcal{O}(\zeta g^2 n\varepsilon) = \tilde{\mathcal{O}}(g^6 n^2 \tau)$ space. Finally for the computation of the matrix of the q^{th} power Frobenius and the zeta function we can follow [1], needing $\tilde{\mathcal{O}}((n+g)n^2 g^3)$ time and $\mathcal{O}(n^2 g^3)$ space. Taking the maximum over all these steps gives the following result for the respective applications:

- (i) time $\tilde{\mathcal{O}}(g^{9+\omega} n^3 \tau^2)$ and space $\tilde{\mathcal{O}}(g^9 n^2 \tau)$,
- (ii) time $\tilde{\mathcal{O}}(g^{9+\omega} n^3 \tau^2)$ and space $\tilde{\mathcal{O}}(g^6 n^2 \tau)$.

The complexity in g seems bad but in all concrete examples, multiplication with $p^{\mathcal{O}(\log g)}$ suffices to make the matrices of the p^{th} as well as the q^{th} power Frobenius integral. If we take this into account, we can remove at least a factor g^2 . Moreover, the implementation results below show that the algorithm performs quite well for relatively high genera.

We note that for the second application, where we compute zeta functions within families defined over the prime field, it is possible to achieve a time complexity of $\tilde{\mathcal{O}}(n^{2.667})$ (where g is fixed) by computing a suitable defining polynomial for \mathbb{Q}_g . For more details we refer to Section 6.3 of [8].

6 Preliminary implementation results

In this section, we briefly report on some experiments with application (ii), i.e. with families defined over prime fields, using the computer algebra system Magma V2.13-14 running on a Pentium IV 2.4 GHz. From a cryptographic viewpoint, the goal is a curve whose Jacobian order has a prime factor $> 2^{160}$. We can achieve this by trying many curves over a suitable field and verifying whether this condition holds. A consequence is that if we fix a family and vary the parameter in a field \mathbb{F}_q , we can consider Steps I, II and the computation of $K(t)$ in Step III as precomputation. The results of our experiments are given in Table 1. For Step I, we used the algorithm described in [5] in the column ‘G.-G.’, and the algorithm presented in [3] in the column ‘D.-V.’. The column ‘Precomp.’ accounts for the precomputations other than $F_p(0)$, and ‘t/c’ gives the time required for each curve after these precomputations.

Table 1. Running times (in seconds) and memory usage to compute the zeta function of a fiber in a family over a prime field

Equation $C(X, Y, t)$	\mathbb{F}_{p^n}	g	G.-G.	D.-V.	Precomp.	t/c	Memory
$Y^3 + X^4 + X + X^3 + t(XY^2 + 1)$	2^{59}	3	6.83	46.38	379	12.31	59 MB
$Y^3 + X^5 + X^2 + t + 1$	2^{43}	4	11.81	261.37	27	6.94	43 MB
$Y^4 + X^5 + Y + t(XY + 1)$	2^{27}	6	10.96	10.45	1080	57.52	126 MB
$Y^3 + X^9 + 1 + tX^4Y$	2^{20}	8	2.29	37.56	25	7.6	60 MB
$Y^3 - X^4 + Y + tXY$	3^{37}	3	2.86	4.77	10	10.63	46 MB
$Y^3 + (t+2)X^5 + (t+1)Y + t$	3^{29}	4	4.30	6.22	11	2.42	28 MB
$Y^4 - X^5 + tXY + tY - 1$	3^{21}	6	1.64	21.83	876	77.30	102 MB
$Y^3 - X^4 + tX^2 + t - 1$	5^{23}	3	0.75	7.82	4.5	1.27	69 MB
$Y^4 - X^5 - X - t(X + Y)$	5^{12}	6	9.76	7.14	4260	77.21	290 MB

These results have to be compared with [3], where, for curves comparable to the first line in this table, each curve required 5000 to 7000 seconds of computing time (albeit on a somewhat slower AMD XP 1700+) and 130 to 147 MB of memory.

Acknowledgements

The authors would like to thank an anonymous referee for his/her detailed verification of the article and useful suggestions.

References

1. W. Castryck, J. Denef, and F. Vercauteren. Computing zeta functions of nondegenerate curves. *IMRP Int. Math. Res. Pap.*, pages Art. ID 72017, 57, 2006.

2. H. Cohen, G. Frey, R. Avanzi, C. Doche, T. Lange, K. Nguyen, and F. Vercauteren, editors. *Handbook of elliptic and hyperelliptic curve cryptography*. Discrete Mathematics and its Applications (Boca Raton). Chapman & Hall/CRC, Boca Raton, FL, 2006.
3. J. Denef and F. Vercauteren. Counting points on C_{ab} curves using Monsky-Washnitzer cohomology. *Finite Fields Appl.*, 12(1):78–102, 2006.
4. B. Edixhoven, J.-M. Couveignes, R. de Jong, F. Merkl, and J. Bosman. On the computation of coefficients of a modular form. Available at <http://arxiv.org/abs/math/0605244>, 2006.
5. P. Gaudry and N. Gürel. An extension of Kedlaya’s point-counting algorithm to superelliptic curves. In C. Boyd, editor, *Advances in Cryptology, Proceedings Asiacrypt 2001*, volume 2248 of *Lecture Notes in Comput. Sci.*, pages 480–494. Springer, Berlin, 2001.
6. P. Gaudry and É. Schost. Construction of secure random curves of genus 2 over prime fields. In *Advances in cryptology—EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Comput. Sci.*, pages 239–256. Springer, Berlin, 2004.
7. H. Hubrechts. Memory efficient hyperelliptic curve point counting. Preprint available at <http://arxiv.org/abs/math/0609032>, 2006.
8. H. Hubrechts. Point counting in families of hyperelliptic curves. To appear in *Foundations of Computational Mathematics*, ‘Online first’.
9. K. S. Kedlaya. Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology. *J. Ramanujan Math. Soc.*, 16(4):323–338, 2001.
10. K. S. Kedlaya. p -adic cohomology: from theory to practice. Arizona Winter School 2007 lecture notes, 2007.
11. A. G. B. Lauder and D. Wan. Counting points on varieties over finite fields of small characteristic. In J.P. Buhler and P. Stevenhagen, editors, *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography*, Mathematical Sciences Research Institute Publications 44, 2007.
12. A. G. B. Lauder. Deformation theory and the computation of zeta functions. *Proc. London Math. Soc. (3)*, 88(3):565–602, 2004.
13. A. G. B. Lauder. A recursive method for computing zeta functions of varieties. *LMS J. Comput. Math.*, 9:222–269 (electronic), 2006.
14. J.-F. Mestre. Lettre adressée à Gaudry et Harley, December 2000. Available at <http://www.math.jussieu.fr/~mestre/>.
15. T. Satoh. The canonical lift of an ordinary elliptic curve over a finite field and its point counting. *J. Ramanujan Math. Soc.*, 15(4):247–270, 2000.
16. R. Schoof. Counting points on elliptic curves over finite fields. *J. Théor. Nombres Bordeaux*, 7(1):219–254, 1995. Les Dix-huitièmes Journées Arithmétiques (Bordeaux, 1993).
17. M. van der Put. The cohomology of Monsky and Washnitzer. *Mém. Soc. Math. France (N.S.)*, 23:4, 33–59, 1986. Introductions aux cohomologies p -adiques (Luminy, 1984).
18. J. von zur Gathen, and J. Gerhard. *Modern computer algebra*. Cambridge University Press, New York, 1999.