

COMPUTING ZETA FUNCTIONS OF ARTIN–SCHREIER CURVES OVER FINITE FIELDS

ALAN G. B. LAUDER AND DAQING WAN

Abstract

The authors present a practical polynomial-time algorithm for computing the zeta function of certain Artin–Schreier curves over finite fields. This yields a method for computing the order of the Jacobian of an elliptic curve in characteristic 2, and more generally, any hyperelliptic curve in characteristic 2 whose affine equation is of a particular form. The algorithm is based upon an efficient reduction method for the Dwork cohomology of one-variable exponential sums.

1. Introduction

We present a low-degree polynomial-time algorithm for computing the zeta function of certain Artin–Schreier curves defined over finite fields. One consequence is a practical method for computing the order of the Jacobian of an elliptic curve in characteristic 2, and (more generally), any hyperelliptic curve whose affine equation is of a particular form. Hyperelliptic curves have been proposed for use in public key cryptosystems by Koblitz [2, 13]. Our algorithm provides the first method of finding ‘random’ hyperelliptic curves of arbitrary genus, defined over large finite fields of characteristic 2, whose Jacobians have orders suitable for cryptographic use. Our method can be extended to more general curves, and we plan to present one such generalisation in a sequel paper.

We now introduce some notation that will allow us to explain our results. Let p denote a prime number, and a a positive integer. Define $q = p^a$, and denote by \mathbb{F}_q the finite field with q elements. Fix an algebraic closure $\overline{\mathbb{F}}_q$ of \mathbb{F}_q , and let \mathbb{F}_{q^k} be the unique subfield of order q^k . We write $\overline{\mathbb{F}}_q^*$ for the set of non-zero elements in $\overline{\mathbb{F}}_q$. The Artin–Schreier curves over \mathbb{F}_q that we consider in this paper are defined by an equation of the form

$$Z^p - Z = f(X), \tag{1}$$

where $f \in \mathbb{F}_q[X, X^{-1}]$ is a Laurent polynomial. Specifically, we denote by C_f the curve embedded in $\overline{\mathbb{F}}_q^* \times \overline{\mathbb{F}}_q$ with equation (1), and we let \tilde{C}_f be the unique smooth projective curve that is birational to C_f . Let d denote the largest absolute value of any exponent that occurs in a non-zero term of f . For example, if $f \in \mathbb{F}_q[X]$, this is just the degree. Our main theorem is as follows.

Theorem 1. *The zeta function of the smooth projective curve \tilde{C}_f may be computed deterministically in $\tilde{O}(p^4 a^3 d^{5+\delta})$ bit operations. Here $\delta = 0$ for $p > 2$, and $\delta = 1$ for $p = 2$.*

Alan Lauder gratefully acknowledges the support of the EPSRC (Grant GR/N35366/01) and St John’s College, Oxford. Daqing Wan is partially supported by the NSF (DMS 9970417) and the NNSF of China (Grant 10128103). Received 21 November 2001, revised 26 March 2002; published 13 May 2002.
 2000 Mathematics Subject Classification 11T99, 11Y16, 14Q05.
 © 2002, Alan G. B. Lauder and Daqing Wan

Here we use the Soft-Oh notation $\tilde{\mathcal{O}}$, which ignores logarithmic factors, as in [14, Section 6.3]. More details of the complexity when using different methods of arithmetic, and also the space complexity, can be found in Section 8.2.

We now explain how our algorithm may be applied to certain hyperelliptic curves in characteristic 2. Let C denote the affine curve with equation

$$Y^2 + X^m Y = h(X),$$

where $h(X) \in \mathbb{F}_{2^a}[X]$ is of degree $2g + 1$ and m is a non-negative integer not greater than g . Let \tilde{C} be the unique smooth projective curve birational to C . Then \tilde{C} is birational to an Artin–Schreier curve, as explained in Note 5, and thus one may compute the zeta function of \tilde{C} in the complexity bounds of Theorem 1. From this, the next corollary follows.

Corollary 2. *The order of the Jacobian of the curve \tilde{C} may be computed deterministically in $\tilde{\mathcal{O}}(a^3 g^6)$ bit operations.*

This algorithm for hyperelliptic curves in characteristic 2 has been implemented by Vercauteren. With regard to the dependence on a , we note that our method, when restricted to elliptic curves, has comparable time complexity to [17]. Moreover, it is the first practical algorithm for hyperelliptic curves in characteristic 2 that has polynomial-time growth in both the field size and the genus. (The problem of polynomial-time computability for arbitrary varieties in small characteristic has already been solved in [14], but the general algorithm there is not very practical. Also, a practical algorithm for hyperelliptic curves in odd characteristic is presented in [12] using different, though related, methods.) We refer to the references in [2] for the large literature on point counting, including [7, 18, 19], and the more recent work in [8, 9, 10, 11, 12, 16, 17, 22, 23, 25].

Sections 2, 3, 4 and 5 lay the mathematical foundation of our algorithm: it is based mainly upon an extension of the work of Dwork [6], due to Adolphson and Sperber [1]. Section 6 contains a statement of the algorithm for what we call ‘Type 1 Artin–Schreier curves’, and Section 7 describes exactly how to perform the main steps. In particular, we present an efficient reduction method for the Dwork cohomology of one-variable exponential sums over the torus. This lies at the heart of our point-counting algorithm, and is the main original contribution of the paper. The complexity analysis is tied up in Section 8, and Section 9 discusses the remaining type of Artin–Schreier curve in a more condensed fashion. As in [14], we aim to give a largely self-contained presentation.

2. L-functions and Artin–Schreier curves

2.1. General theory

Let $\overline{\mathbb{Q}}$ denote an algebraic closure of the rationals \mathbb{Q} . Let $\Psi : \mathbb{F}_p \rightarrow \overline{\mathbb{Q}}$ be a non-trivial additive character, and let $\text{Tr}_k : \mathbb{F}_{q^k} \rightarrow \mathbb{F}_p$ be the absolute trace map. A specific Ψ will be constructed in Section 4.2, but for now it may be arbitrary. Define $\Psi_k : \mathbb{F}_{q^k} \rightarrow \overline{\mathbb{Q}}$ to be the non-trivial additive character $\Psi \circ \text{Tr}_k$.

For $f \in \mathbb{F}_q[X, X^{-1}]$, define

$$S_k^*(f, \Psi) := \sum_{x \in \mathbb{F}_{q^k}^*} \Psi_k(f(x)) \tag{2}$$

$$L^*(f, \Psi, T) := \exp \left(\sum_{k=1}^{\infty} \frac{S_k^*(f, \Psi)}{k} T^k \right). \tag{3}$$

For simplicity, we shall omit the Ψ in this notation. Let C_f be the curve embedded in $\overline{\mathbb{F}}_q^* \times \overline{\mathbb{F}}_q$ with equation

$$Z^p - Z = f(X).$$

Let \tilde{C}_f denote the unique smooth projective curve birational to C_f .

Lemma 3. *For each $x \in \mathbb{F}_{q^k}^*$, there are exactly n_x points of the form $(x, z) \in \mathbb{F}_{q^k}^* \times \mathbb{F}_{q^k}$ on C_f , where*

$$n_x = \begin{cases} p, & \text{if } \text{Tr}_k(f(x)) = 0, \\ 0, & \text{if } \text{Tr}_k(f(x)) \neq 0. \end{cases}$$

Proof. This follows from [15, Theorem 2.25]. □

Denote by $C_f(\mathbb{F}_{q^k})$ the set of \mathbb{F}_{q^k} -rational points on C_f . From Lemma 3, one deduces that

$$\#(C_f(\mathbb{F}_{q^k})) = \sum_{j=0}^{p-1} \sum_{x \in \mathbb{F}_{q^k}^*} \Psi_k(jf(x)) = \sum_{\theta \in G} \theta(S_k^*(f)) + (q^k - 1),$$

where G is the Galois group of $\mathbb{Q}(\zeta)$ over \mathbb{Q} , with $\zeta := \Psi(1)$ a primitive p th root of unity. Writing $Z(C_f, T)$ for the zeta function of C_f (see [14, Section 1]), it follows that

$$Z(C_f, T) = \frac{\{\prod_{\theta \in G} \theta(L^*(f, T))\}(1 - T)}{(1 - qT)}. \tag{4}$$

Here, G acts on the power series $\mathbb{Q}(\zeta)[[T]]$ coefficient-wise, fixing the monomials T^k .

To proceed further, it is necessary to split the possible Laurent polynomials f into three types.

2.2. Three types of Artin–Schreier curves

Type 1: Assume that $f \in \mathbb{F}_q[X]$ has degree d not divisible by p . Here we can also define

$$S_k(f) := \sum_{x \in \mathbb{F}_{q^k}} \Psi_k(f(x)) = S_k^*(f) + \Psi(k \text{Tr}_1(f(0))); \tag{5}$$

$$L(f, T) := \exp\left(\sum_{k=1}^{\infty} \frac{S_k(f)}{k} T^k\right) = L^*(f, T)(1 - \Psi(\text{Tr}_1(f(0)))T)^{-1}. \tag{6}$$

Hence, from (4),

$$Z(C_f, T) = \frac{\{\prod_{\theta \in G} \theta(L(f, T))\}(1 - T)}{(1 - qT)\{\prod_{\theta \in G} \theta(1 - \Psi(\text{Tr}_1(f(0)))T)\}}.$$

(The latter term on the denominator is $(1 - T)^{p-1}$ or $1 + T + \dots + T^{p-1}$, depending upon whether $\text{Tr}_1(f(0))$ is zero or not, although this does not concern us.) By Weil, the L-function $L(f, T)$ is a polynomial of degree $d - 1$, and all of its reciprocal roots have complex absolute value \sqrt{q} under all complex embeddings of \mathbb{Q} . Thus $\prod_{\theta \in G} \theta(L(f, T))$ is a polynomial of degree $(p - 1)(d - 1)$, pure of weight 1. (A complex number has weight i for $i = 0, 1, 2$ if it has absolute value $q^{i/2}$, and a rational function with algebraic integer coefficients is pure of weight i if its roots all have weight i under any complex embedding [24, Section 3].)

Let $Z(\tilde{C}_f, T)$ denote the zeta function of the smooth projective curve \tilde{C}_f . Let g be the genus of \tilde{C}_f . Again, by Weil, we know that

$$Z(\tilde{C}_f, T) = \frac{P(\tilde{C}_f, T)}{(1 - T)(1 - qT)},$$

where the numerator is a polynomial of degree $2g$, pure of weight 1. Since C_f and \tilde{C}_f are birational, they differ by a finite number of points, and hence their zeta functions differ by a factor of weight 0. Comparing the pure weight 1 parts in $Z(C_f, T)$ and $Z(\tilde{C}_f, T)$, we deduce that

$$P(\tilde{C}_f, T) = \prod_{\theta \in G} \theta(L(f, T)). \tag{7}$$

In particular, the genus g of the curve \tilde{C}_f is given by the formula

$$g = (p - 1)(d - 1)/2. \tag{8}$$

Type 2: Assume that $f \in \mathbb{F}_q[X^{-1}]$ has negative degree d^- , not divisible by p . That is, d^- is the lowest exponent occurring in f . Then C_f is birational to C_{f^*} , where $f^* := f(X^{-1})$, and we have reduced to Type 1.

Type 3: Assume that $f \in \mathbb{F}_q[X, X^{-1}]$, but that $f \notin \mathbb{F}_q[X] \cup \mathbb{F}_q[X^{-1}]$. Let the negative degree be d^- and the positive degree be d^+ . Assume that p does not divide d^-d^+ . In this case, by Weil, $L^*(f, T)$ is a polynomial of degree $d^+ - d^-$, pure of weight 1. By comparing the pure weight 1 parts in $Z(C_f, T)$ and $Z(\tilde{C}_f, T)$, we find that

$$P(\tilde{C}_f, T) = \prod_{\theta \in G} \theta(L^*(f, T)), \tag{9}$$

is the numerator of $Z(\tilde{C}_f, T)$. In particular, the genus g of the curve \tilde{C}_f is given by the formula

$$g = (p - 1)(d^+ - d^-)/2.$$

Thus in all cases the computation of the zeta function of the smooth projective curve \tilde{C}_f reduces to the evaluation of the L-function of certain one-variable exponential sums.

Note 4. The degree conditions on f are essential for cohomological arguments; however, given a Laurent polynomial f that does not satisfy them, one may replace it by a new polynomial \tilde{f} which does, such that C_f and $C_{\tilde{f}}$ are isomorphic. For each term $a_{jp}X^{jp}$, the isomorphism $Z \rightarrow Z + b_jX^j$ ($X \rightarrow X$) shows that we can replace the term $a_{jp}X^{jp}$ with b_jX^j , where $b_j := a_{jp}^{1/p} \in \mathbb{F}_q$. Repeat this procedure until no term has a non-zero exponent divisible by p . The resulting polynomial \tilde{f} has no terms with non-zero exponents divisible by p , and thus it certainly satisfies any necessary degree restrictions.

Note 5. Let C be the curve in Corollary 2. Then C is birational to the curve C_f with equation $Z^2 + Z = f(X)$, where $f := X^{-2m}h$. This can be seen by making the change of variable $Y = ZX^m$. Thus \tilde{C} , as in Corollary 2, is birational to \tilde{C}_f , and the zeta function of \tilde{C}_f can then be computed using Note 4 (to get f in the correct form) and the algorithm that we shall present later. By the special value formula for the zeta function at $T = 1$, the order of the Jacobian is the numerator of the zeta function evaluated at $T = 1$ [2, p. 175].

Similar comments apply to curves over \mathbb{F}_q of the form $Y^p - X^m Y = h(X)$, where $q = p^a$ and $p - 1$ divides m . Also, general hyperelliptic curves in characteristic 2 are birational to Artin–Schreier curves of the form $Z^2 + Z = f(X)$ for f a rational function. As such, they may be tackled using a generalisation of our approach.

2.3. Type 1 Artin–Schreier curves

In Sections 3, 4, 5, 6, 7 and 8, we shall denote by f a polynomial in $\mathbb{F}_q[X]$, of degree d , not divisible by p . We write $f = \sum_{j \in J} a_j X^j$, where $a_j \neq 0$. In these sections we shall explain how to compute the zeta function of the smooth projective curve \tilde{C}_f . That is, we shall cover classical ‘Type 1’ Artin–Schreier curves. In Section 9, we shall discuss the modifications required when $f \in \mathbb{F}_q[X, X^{-1}]$ is a Laurent polynomial with both positive and negative exponents. These are ‘Type 3’ Artin–Schreier curves. As mentioned before, ‘Type 2’ curves may be reduced to ‘Type 1’, so we shall not discuss them again.

3. p -adic theory

3.1. p -adic rings

Let \mathbb{Q}_p denote the p -adic numbers with ring of integers \mathbb{Z}_p . Fix Ω the completion of an algebraic closure of \mathbb{Q}_p . Denote by ϵ a primitive $(q - 1)$ th root of unity in Ω , and by $\pi \in \Omega$ an element that satisfies $\pi^{p-1} = -p$. Define

$$A := \mathbb{Z}_p[\epsilon, \pi]. \tag{10}$$

In particular, A has residue field \mathbb{F}_q . Elements in A may be represented via π -adic expansions whose coefficients are taken from some distinct set of representatives for the quotient $A/(\pi)$ of size q . By binomial expansion and Hensel’s lemma, one sees that the equation $(1 + \pi t)^p = 1$ has exactly p distinct solutions t in $\mathbb{Z}_p[\pi]$. Thus for any primitive p th root of unity ζ , we have $\mathbb{Z}_p[\zeta] = \mathbb{Z}_p[\pi]$, and so A contains the exponential sums defined in the previous section. Let G denote the group of automorphisms $\theta_j : \mathbb{Q}_p(\pi) \rightarrow \mathbb{Q}_p(\pi)$ for $1 \leq j \leq p - 1$, where each θ_j fixes \mathbb{Q}_p and

$$\theta_j(\pi) := \eta^j \pi. \tag{11}$$

Here $\eta \in \mathbb{Z}_p$ is a primitive $(p - 1)$ th root of unity. Then G is the Galois group of the extension $\mathbb{Q}_p(\pi)/\mathbb{Q}_p$. (Its action on ζ is $\theta_j : \zeta \mapsto \zeta^{\eta^j \bmod p}$, although we shall not need this explicitly.)

It will be convenient to work in complete rings that contain arbitrary roots π^r of π , for r a rational number. To this end, let $\Pi = \cup_{z \in \mathbb{N}} \{\pi^{1/z}\}$, and let \tilde{A} denote the completion of the ring $\mathbb{Z}_p[\epsilon, \Pi]$. Here \mathbb{N} denotes the positive integers.

Denote by τ the endomorphism on \tilde{A} defined as

$$\tau(\epsilon) := \epsilon^p, \quad \tau \text{ fixes } \mathbb{Z}_p[\Pi] \text{ and is continuous.} \tag{12}$$

Let ord and $|\cdot|_p$ denote the p -adic valuation and norm on \tilde{A} normalised so that $\text{ord}(p) = 1$ and $|p|_p = 1/p$.

Note that the ring A can be easily constructed and computed; see [14, Section 3] for more details on p -adic fields. (The larger ring \tilde{A} is introduced only for mathematical convenience; all of our computations are performed in A .)

3.2. A weight function

As in [14, Section 4] we define a weight function: for each non-negative integer $u \in \mathbb{Z}_{\geq 0}$, let

$$\text{wt}(u) := \lceil u/d \rceil.$$

Here $\lceil * \rceil$ is the least integer not less than $*$, and d is the degree of the polynomial $f \in \mathbb{F}_q[X]$. Define $\tilde{\text{wt}}(u) := u/d$, and so $\lceil \tilde{\text{wt}}(u) \rceil = \text{wt}(u)$. Notice that $\pi^{\text{wt}(r)} \in A$ and $\pi^{\tilde{\text{wt}}(r)} \in \tilde{A}$ for every $r \in \mathbb{Z}_{\geq 0}$.

3.3. Banach modules

Let B denote a complete subring of Ω . A Banach module over B is an ultrametrically normed complete module E over B , such that $\|re\| = |r|_p \|e\|$ for $r \in B$ and $e \in E$, where $\|\cdot\|$ is the module norm. An orthonormal basis for E is a set $\{e_i \mid i \in \mathbb{N}\}$ such that every element in E can be written uniquely in the form $\sum_i b_i e_i$ where $b_i \in B$ with $|b_i|_p \rightarrow 0$ as $i \rightarrow \infty$. (See [20, Section 1] and [4, Section A].) In the case that B is a field, we call it a *Banach space*.

Definition 6. For each rational number $\delta > 0$, let $\tilde{L}(\delta)$ be the Banach module over \tilde{A} whose orthonormal basis consists of all terms $\pi^{\tilde{\text{wt}}(r)\delta} X^r$ for $r \in \mathbb{Z}_{\geq 0}$. Let $A\{X\}$ and $\tilde{A}\{X\}$ denote the Banach module over A or \tilde{A} , respectively, whose orthonormal basis consists of all terms X^r for $r \in \mathbb{Z}_{\geq 0}$.

Note that $\tilde{L}(\delta') \subset \tilde{L}(\delta)$ for $\delta' > \delta$, and all the above spaces lie in $\tilde{A}\{X\}$. One may check that all the above spaces are closed under multiplication, and are in fact rings. Extend τ to act on each power series in the ring $\tilde{A}\{X\}$ by taking

$$\tau(X) := X, \quad \tau \text{ is a continuous endomorphism.} \tag{13}$$

4. Analytic representation of additive characters

4.1. Dwork’s splitting functions

We now present the analytic construction of an additive character due to Dwork (see [5, Section 1] and [6, pp. 55-57], referring to [14] for more details). Let $\theta(t)$ denote the splitting function [14, Section 4.1]

$$\theta(t) := \theta_1(t) = \exp(\pi(t - t^p)). \tag{14}$$

Write

$$\hat{f} := \sum_{j \in J} \hat{a}_j X^j \tag{15}$$

for the polynomial over A obtained by taking the Teichmüller lifting of each coefficient of f .

Lemma 7. For each term $\hat{a}_j X^j$ in \hat{f} , we have $\theta(\hat{a}_j X^j) \in \tilde{L}(\delta)$ for any

$$\delta < \left(\frac{p-1}{p} \right)^2.$$

Proof. Writing $\theta(t) =: \sum_{r=0}^{\infty} \lambda_r t^r$, we see from [14, Section 4] that

$$\text{ord}(\lambda_r) \geq (p-1)r/p^2.$$

Terms in $\theta(\hat{a}_j X^j)$ are of the form $\lambda_r \hat{a}_j^r X^{jr}$ for r a non-negative integer. Now $\text{ord}(\lambda_r \hat{a}_j^r) \geq (p-1)r/p^2$ since $\text{ord}(\hat{a}_j) = 0$. Also, $0 \leq j \leq d$, and so $\tilde{\text{wt}}(jr) = jr/d \leq r$. Thus

$$\text{ord}(\lambda_r \hat{a}_j^r) \geq \frac{p-1}{p^2} r \geq \frac{p-1}{p^2} \tilde{\text{wt}}(jr).$$

Also, for $\delta < ((p-1)/p)^2$ we have

$$\text{ord}(\pi^{\delta \tilde{\text{wt}}(jr)}) < \frac{p-1}{p^2} \tilde{\text{wt}}(jr),$$

and the result follows. □

Definition 8. Let F and $F^{(a)}$ be defined as follows:

$$F := \prod_{j \in J} \theta(\hat{a}_j X^j); \tag{16}$$

$$F^{(a)} := \prod_{i=0}^{a-1} \tau^i(F(X^{p^i})). \tag{17}$$

Recall here that $q = p^a$. Both F and $F^{(a)}$ are one-way infinite power series in $A\{X\}$. By Lemma 7 and the fact that each $\tilde{L}(\delta)$ is a ring, we have the next lemma.

Lemma 9. *The power series $F \in \tilde{L}(\delta)$ for any*

$$\delta < \left(\frac{p-1}{p}\right)^2.$$

4.2. Dwork’s additive character

For $k \geq 1$, define

$$\Phi_k(t) := \prod_{i=0}^{ak-1} \theta(t^{p^i}) \in \mathbb{Z}_p[\pi][[t]],$$

and let

$$\Phi(t) := \prod_{i=0}^{a-1} \theta(t^{p^i}) \in \mathbb{Z}_p[\pi][[t]].$$

Denote by ‘Teich’ the Teichmüller lifting map from $\overline{\mathbb{F}}_q$ to $\mathbb{Z}_p^{\text{unram}}$, where $\mathbb{Z}_p^{\text{unram}}$ is the unramified integral closure of \mathbb{Z}_p in Ω . Then

$$\Psi_k := \Phi_k \circ \text{Teich} \quad \text{and} \quad \Psi := \Phi \circ \text{Teich}$$

are non-trivial characters from \mathbb{F}_{q^k} and \mathbb{F}_p , respectively, to $\mathbb{Z}_p[\pi]$. We see that $\Psi_k(x) = \Psi(\text{Tr}_k(x))$, where Tr_k is the trace map from \mathbb{F}_{q^k} to \mathbb{F}_p . (See [14, Lemma 6].)

The following lemma is proved exactly as [14, Proposition 9]. It gives an analytic expression for the exponential sum $S_k^*(f) := S_k^*(f, \Psi)$.

Lemma 10. *Let $S_k^*(f)$ be the exponential sum defined in equation (2) using Dwork’s additive character Ψ , and let $F^{(a)}(X)$ be the power series from Definition 8. Then*

$$S_k^*(f) = \sum_{x^{q^k-1}=1} F^{(a)}(x) F^{(a)}(x^q) \dots F^{(a)}(x^{q^{k-1}}).$$

Here, the sum is over the Teichmüller liftings in $\mathbb{Z}_p^{\text{unram}}$ of the points on the torus $\mathbb{F}_{q^k}^*$.

4.3. Completely continuous maps

The next step is to introduce operators on $\tilde{A}\{X\}$ so that the right-hand side of the above expression can be interpreted as the ‘trace’ of a map on a certain Banach module.

Definition 11. Let ψ_p be the map on $\tilde{A}\{X\}$ that acts on monomials as

$$\psi_p(X^r) := \begin{cases} X^{r/p}, & \text{if } p \text{ divides } r, \\ 0, & \text{otherwise,} \end{cases}$$

and extends to all of $\tilde{A}\{X\}$ by τ^{-1} -linearity and continuity. Specifically, $\psi_p(\sum_r A_r X^r) = \sum_{r, p|r} \tau^{-1}(A_r) X^{r/p}$. Write $\psi_q := \psi_p^a$, a linear map since τ^{-a} is the identity on \tilde{A} . Let $\alpha := \psi_p \circ F$ and let $\alpha_a := \psi_q \circ F^{(a)}$. Precisely, α is multiplication by F followed by the map ψ_p , and likewise for α_a (see [14, Definitions 20, 21]).

Lemma 12. The maps satisfy $\alpha_a = \alpha^a$.

Proof. This is proved exactly as in [14, Lemma 22]. □

Lemma 13. The map α is stable on $\tilde{L}(\delta)$ for

$$\delta < \frac{(p-1)^2}{p}.$$

Proof. We first claim that $\psi_p(\tilde{L}(\delta)) \subseteq \tilde{L}(p\delta)$ for any rational $\delta > 0$. For this, it is enough to observe that

$$\psi_p(\pi^{\tilde{w}(r)\delta} X^r) = \pi^{\tilde{w}(r/p)p\delta} X^{r/p}$$

for any r divisible by p . Now let $G \in \tilde{L}(\delta)$, where δ satisfies the necessary inequality, depending on p . Then, by Lemma 9, $F \in \tilde{L}(\delta/p)$, and so (since also $G \in \tilde{L}(\delta/p)$ by closure under multiplication) we find that $FG \in \tilde{L}(\delta/p)$. Hence $\psi_p(FG) \in \tilde{L}(\delta)$; that is, $\alpha(G) \in \tilde{L}(\delta)$, as required. □

Definition 14. Let L be defined as

$$L := \begin{cases} \tilde{L}(1) \cap A\{X\}, & \text{if } p > 2, \\ (\tilde{L}(\gamma) \cap A\{X\}) \otimes \mathbb{Q}, & \text{if } p = 2. \end{cases}$$

Here γ may be taken to be any rational number in the range $1/(2 + (1/2d)) < \gamma < 1/2$. The key point is that α is stable on $\tilde{L}(\gamma)$, since $\gamma < 1/2$, and the space $\tilde{L}(\gamma)$ is small enough such that an ad hoc argument that we shall present later (Lemma 28) works. Thus for $p > 2$, we see that L is just the Banach module over A with orthonormal basis the terms $\pi^{\tilde{w}(u)} X^u$ for non-negative integers u . For $p = 2$ it is the Banach space over $A \otimes \mathbb{Q}$ with orthonormal basis the terms $\pi^{\lceil \gamma \tilde{w}(u) \rceil} X^u$ for non-negative u .

Lemma 15. The maps α and α_a are stable on L .

Proof. First, suppose that $p > 2$. Then α is stable on $\tilde{L}(1)$, by Lemma 13. Certainly, α is stable on the ring of convergent power series $A\{X\}$, since $F \in A\{X\}$. Thus α is stable on $\tilde{L}(1) \cap A\{X\} = L$. That α_a is stable on L now follows from Lemma 12. Second, consider the case $p = 2$. Putting $p = 2$ in Lemma 13, we find that α is stable on $\tilde{L}(\gamma) \cap A\{X\}$, since $\gamma < 1/2$. Given $G \in L$, we have $mG \in \tilde{L}(\gamma) \cap A\{X\}$ for some p -adic integer m . The result now follows easily. □

Note 16. The rings \tilde{A} and $\tilde{L}(\delta)$ and the function $\tilde{w}t$ were introduced to prove the above result in as simple a manner as possible; we shall have little further need for them, working from now on mainly with A and L .

For certain classes of linear maps on Banach modules, the trace and determinant are defined. This is done in the usual way, via matrices for the maps with respect to an orthonormal basis. We refer to [20, Section 5] and [4, Section A2] for definitions. The key result is the chain-level Dwork trace formula.

Theorem 17. With $S_k^*(f)$ and $L^*(f, T)$ the exponential sum and the L -function defined as in (2) and (3) using Dwork’s additive character Ψ (Section 4.2), and α_a the map on L given in Definition 11, we have

$$S_k^*(f) = (q^k - 1) \text{Tr}(\alpha_a^k|L).$$

Thus we have

$$L^*(f, T) = \frac{\det(1 - T\alpha_a|L)}{\det(1 - Tq\alpha_a|L)}.$$

Here the trace and the determinant are defined via matrices for the maps with respect to the orthonormal basis of L .

Proof. This is in essence a case of [14, Theorem 25] (with $n = 0$). Note that the matrix M_a in [14, Theorem 25] is that for the map α_a with respect to a ‘formal basis’ [14, Section 5.2] of the form $\{X^i \mid i \in \mathbb{Z}_{\geq 0}\}$. For the above formulae we require a matrix with respect to the orthonormal basis $\{\pi^{\text{wt}(i)} X^i \mid i \in \mathbb{Z}_{\geq 0}\}$ (when $p > 2$, with a slightly different basis for $p = 2$). One may verify that the traces of the powers of these two matrices are the same. □

These formulae may be used to compute the zeta function in a similar fashion to [14]. In other words, a finite matrix may be computed that represents the map α acting on some appropriate modular reduction of L . This matrix is then used to compute the characteristic polynomial of α_a itself, up to a necessary p -adic accuracy. (This algorithm has been implemented by Vercauteren). However, this ‘chain-level’ method results, for example, in a time complexity of $\tilde{O}(a^{4.38})$ with space $\tilde{O}(a^4)$ bits, using the fastest methods for matrix multiplication and ring arithmetic. Using some homological algebra, one can derive a better ‘cohomological’ formula, leading to an improved algorithm. That is what we do in this paper.

At this stage, since f is an ordinary polynomial rather than a Laurent polynomial with negative and positive terms, we can do a little more work to derive a better chain-level formula, as follows. Let $L_{>0}$ denote the Banach module comprising those power series in L with zero constant term. For $p > 2$, the module $L_{>0}$ is defined over A , and for $p = 2$, over $A \otimes \mathbb{Q}$. The next lemma follows easily from Lemma 15.

Lemma 18. The maps α and α_a are stable on $L_{>0}$.

Now one may check via a matrix for the map α_a with respect to the orthonormal basis $\{\pi^{\text{wt}(r)} X^r\}_{r \geq 0}$ for $p > 2$, and $\{\pi^{\lceil \gamma \tilde{w}t(r) \rceil} X^r\}_{r \geq 0}$ for $p = 2$, that

$$\det(1 - T\alpha_a|L) = (1 - F^{(a)}(0)T) \det(1 - T\alpha_a|L_{>0}).$$

Here

$$F^{(a)}(0) = \Psi(\text{Tr}_1(a_0))$$

is a root of unity, where Ψ is Dwork’s additive character from \mathbb{F}_p to $\mathbb{Z}_p[\pi]$.

Let $S_k(f)$ and $L(f, T)$ be as in (5) and (6). Then one has $S_k(f) = S_k^*(f) + \Psi(k \operatorname{Tr}_1(a_0))$, and so $L(f, T) = L^*(f, T)(1 - F^{(a)}(0)T)^{-1}$. Hence we have the following theorem.

Theorem 19. *Let $L(f, T)$ be the L -function for the exponential sum over the affine line from equation (6). Then*

$$L(f, T) = \frac{\det(1 - T\alpha_a|_{L_{>0}})}{\det(1 - Tq\alpha_a|_L)}.$$

5. Dwork cohomology

Let H be the polynomial in L defined as

$$H = \pi \hat{f}. \tag{18}$$

(In fact, $H \in \tilde{L}(1) \cap A\{X\}$ in all cases, and this latter ring equals L for $p > 2$, and lies strictly within L for $p = 2$.)

Let D be the operator on L defined as

$$D = X \frac{d}{dX} + H_X,$$

where

$$H_X := X \frac{dH}{dX}.$$

Here d/dX is the usual differential operator on polynomials extended to power series by continuity, with X and H_X just acting by multiplication. Now $X(d/dX)$ is stable on L , and L is a ring. From this it follows that D is stable on L , and in fact maps L to $L_{>0}$.

Note 20. We pause to explain the motivation behind the above definitions. Define $\hat{\theta}(t) := \prod_{i=0}^{\infty} \theta(t^{p^i})$. One may check that $\hat{\theta}(t) = \exp(\pi t)$. Now $F = \prod_{j \in J} \theta(\hat{a}_j X^j)$. Defining $\hat{F}(X) := \prod_{j \in J} \hat{\theta}(\hat{a}_j X^j)$, we find that this equals $\exp(\pi \hat{f})$, which is just $\exp(H)$, with H as in (18). Since $\hat{a}_j^q = \hat{a}_j$ and $\tau^i(\hat{a}_j) = \hat{a}_j^{p^i}$ for each $j \in J$, it follows from Definition 8 that $F^{(a)}(X) = \hat{F}(X)/\hat{F}(X^q)$. From this, we see that $\alpha_a = \psi_q \circ F^{(a)} = \exp(-H) \circ \psi_q \circ \exp(H)$. Define the operator $E := X(d/dX)$, and so $E \circ (q\psi_q) = \psi_q \circ E$. Then, with D as above, one may check that $D = E + E(H) = \exp(-H) \circ E \circ \exp(H)$. Thus D and α_a are obtained from E and ψ_q by some kind of twisting; also, it now follows that $D \circ (q\alpha_a) = \alpha_a \circ D$, which is the crucial relation. (See [6, pp. 55–60] and [21, pp. 267–270] for more details.)

Let \mathcal{L} be the complex

$$0 \longrightarrow L \xrightarrow{D} L_{>0} \longrightarrow 0.$$

This is a complex of A -modules when $p > 2$ and $A \otimes \mathbb{Q}$ -spaces for $p = 2$. Denote by H_1 and H_0 the kernel and co-kernel of the map D . In particular, $H_0 := L_{>0}/D(L)$.

Proposition 21. *The map D is injective, and so $H_1 = 0$. Moreover, for $p > 2$ and $p = 2$, H_0 is respectively a finite free A -module or an $A \otimes \mathbb{Q}$ -space, of rank $d - 1$. A basis for H_0 may be taken as the set of terms*

$$\{\pi^{\operatorname{wt}(i)} X^i \mid 0 < i < d\} = \{\pi X, \pi X^2, \dots, \pi X^{d-1}\}.$$

Proof. Over the formal power series ring $\Omega[[X]]$, the formal solutions of the first-order linear differential equation $D = 0$ is the one-dimensional subspace generated by $\exp(-H)$.

But the power series $\exp(-H) \notin L$ (the decay rate of coefficients is too slow). This shows that the restriction of the operator D to L is injective. The second part of the proposition follows from the normal form computations in Section 7. \square

By Note 20, we have

$$D \circ q\alpha_a = \alpha_a \circ D.$$

Thus the map α_a defines a chain map on \mathcal{L} :

$$\begin{array}{ccccccc} 0 & \longrightarrow & L & \xrightarrow{D} & L_{>0} & \longrightarrow & 0 \\ & & \downarrow q\alpha_a & & \downarrow \alpha_a & & \\ 0 & \longrightarrow & L & \xrightarrow{D} & L_{>0} & \longrightarrow & 0. \end{array}$$

Denote by $H_0(\alpha_a)$ and $H_1(q\alpha_a)$ the maps induced on the homologies H_0 and H_1 by this chain map, and by $\det(1 - H_0(\alpha_a)T)$ and $\det(1 - H_1(q\alpha_a)T)$ the corresponding determinants.

Theorem 22. *The L-function from Theorem 19 satisfies*

$$L(f, T) = \det(1 - H_0(\alpha_a)T).$$

Proof. We have

$$\frac{\det(1 - T\alpha_a|L_{>0})}{\det(1 - Tq\alpha_a|L)} = \frac{\det(1 - H_0(\alpha_a)T)}{\det(1 - H_1(q\alpha_a)T)}.$$

This identity is proved in the same way that [20, Proposition 9] is derived from [20, Lemma 2]. Now $H_1 = 0$, and so the denominator on the right-hand side is 1. The expression for the L-function now follows from Theorem 19. \square

Corollary 23. *The zeta function $Z(\tilde{C}_f, T)$ of the smooth projective curve \tilde{C}_f birational to the affine curve with equation $Z^p - Z = f(X)$ satisfies*

$$Z(\tilde{C}_f, T) = \frac{\prod_{j=1}^{p-1} \theta_j(\det(1 - H_0(\alpha_a)T))}{(1 - T)(1 - qT)}.$$

Here θ_j are the automorphisms (11) of $\mathbb{Z}_p[\pi]$ extended to act on polynomials by fixing T . The numerator is a polynomial of degree $(p - 1)(d - 1)$.

Proof. This follows from Theorem 22 and equation (7). \square

Thus the strategy of the algorithm is to compute the determinant of the map $H_0(\alpha_a)$ on the zeroth homology H_0 , up to a suitable modular precision. This may be done efficiently via the following lemma, which is an immediate consequence of Lemma 12.

Lemma 24. *Let $H_0(\alpha)$ denote the map induced on H_0 by α . Then $H_0(\alpha_a) = H_0(\alpha)^a$.*

It will be enough to compute the coefficients of the characteristic polynomial of Frobenius modulo p^N for

$$N = \lfloor (p - 1)(d - 1)(1 + a/2) + 1 \rfloor.$$

This follows since the L-function of the exponential sum $L(f, T)$ has reciprocal roots whose complex absolute values are \sqrt{q} . Thus the coefficient of T^k in the polynomial $\prod_{\theta_j \in G} \theta_j(L(f, T))$ are integers of absolute value at most $\binom{2g}{k} p^{ak/2} \leq 2^{2g} p^{ak/2}$. Since the polynomial $\prod_{\theta_j \in G} \theta_j(L(f, T))$ has degree $2g = (p - 1)(d - 1)$, it follows that determining

the coefficients modulo p^N for $N > (p - 1)(d - 1)(1 + a/2)$ is sufficient. (Due to a certain ‘loss of accuracy’ when one performs the homological reduction, it is initially necessary to compute the coefficients of F modulo $p^{\varepsilon(N+1)}$, for a small positive integer ε , whose precise value we shall determine.)

6. The algorithm

We now present our point-counting algorithm for Type 1 Artin–Schreier curves (see Section 2 for our classification of Artin–Schreier curves).

Algorithm 25 (Artin–Schreier Type 1).

Input: An equation $Z^p - Z = f(X)$ over \mathbb{F}_q , where $f \in \mathbb{F}_q[X]$ and $q = p^a$.

Output: The zeta function $Z(\tilde{C}_f, T)$ of the unique smooth projective curve birational to the affine curve defined by this equation.

Step 0: Replace f by a polynomial all of whose terms have exponents not divisible by p , in the manner explained in Note 4. Denote this new polynomial also by f . This will not change the zeta function. Set $N := \lfloor (p - 1)(d - 1)(1 + a/2) + 1 \rfloor$, where d is the degree of f . Let $\varepsilon := 4$ when $p > 2$, and $\varepsilon := (4d + 1)$ when $p = 2$. We shall compute the coefficients of the numerator of the zeta function modulo p^N .

Step 1: Compute the power series F given in Definition 8 with coefficients determined modulo $p^{\varepsilon(N+1)}$. Let α be the map on the ring L (Definition 6), defined as $\alpha = \psi_p \circ F$ (Definition 11). Let $H_0(\alpha)$ be the map induced on the zeroth homology H_0 of the complex \mathcal{L} by α .

Step 2: Let $\pi X, \pi X^2, \dots, \pi X^{d-1}$ be the basis for the zeroth homology H_0 . For each basis element e , compute the image $H_0(\alpha)(e) \in H_0$ with coefficients determined modulo p^N . Construct M , defined as the matrix representing the map $H_0(\alpha)$ with respect to the basis, with coefficients determined modulo p^N . Specifically, $M = (m_{ij})$, where i is the row index and j the column index, and $H_0(\alpha)(\pi X^j) = \sum_{i=1}^{d-1} m_{ij}(\pi X^i) \pmod{p^N}$ for $1 \leq j \leq d - 1$.

Step 3: Compute

$$M_a := M\tau^{-1}(M)\tau^{-2}(M) \dots \tau^{-(a-1)}(M)$$

modulo p^N , where the map τ is the lifting of Frobenius to A as given in (12). Thus M_a is a matrix for the map $H_0(\alpha_a)$.

Step 4: Output the rational function

$$Z(\tilde{C}_f, T) := \frac{\prod_{j=1}^{p-1} \theta_j(\det(I - M_a T))}{(1 - T)(1 - qT)},$$

where θ_j is the automorphism from (11) extended to act on $\mathbb{Z}_p[\pi][T]$ by fixing monomials.

The correctness of the algorithm follows from Corollary 23 and Lemma 24, along with the discussion of the choice of N at the end of Section 5, and the choice of ε from Lemmas 27 and 31. The matrix M is called the *absolute Frobenius matrix*, and $\det(I - M_a T)$ the *characteristic polynomial of Frobenius*. In Section 7 we shall describe exactly how this first matrix is computed, allowing us to give a complexity analysis of the algorithm in Section 8. This will complete the proof of Theorem 1 for Types 1 and 2 Artin–Schreier curves. We present the algorithm for Type 3 curves in Section 9.

Note 26. The above algorithm can be improved in practice by using the functional equation

$$T^{2g}q^g P\left(\tilde{C}_f, \frac{1}{qT}\right) = P(\tilde{C}_f, T).$$

This functional equation shows that it is enough to determine the coefficients modulo $p^{N'}$, where $N' := \lfloor (p - 1)(d - 1)(1 + a/4) + 1 \rfloor$, computing only the first half of the coefficients directly in $P(\tilde{C}_f, T)$, and then recovering the second half by the functional equation. Moreover, our choice of ε is perhaps rather large, especially in the case $p = 2$. It may be enough in practice, as observed by Vercauteren, to work initially to p -adic accuracy $N' + \delta$, where δ is some small variable that can be determined ‘experimentally’ (see also Note 32).

7. Performing the main steps

We shall work with elements in L with coefficients determined modulo $p^{\varepsilon(N+1)}$. (In the case $p = 2$ by this we mean that, given $G \in L$, we have $G = G' + p^{\varepsilon(N+1)}G''$, where G' is a known polynomial over $A \otimes \mathbb{Q}$ and G'' is a power series with coefficients in A .) If an element is given to this accuracy, we say that it lies in $L \bmod p^{\varepsilon(N+1)}$. Similarly, H_0 is the free module over A when $p > 2$, and $A \otimes \mathbb{Q}$ the free module when $p = 2$, spanned by the basis monomials

$$\{\pi X, \pi X^2, \dots, \pi X^{d-1}\}. \tag{19}$$

We write $s \in H_0 \bmod p^N$ if an element s is given in H_0 with coefficients modulo p^N . In the next two sections we shall explain how, given an element $G \in L_{>0} \bmod p^{\varepsilon(N+1)}$, we can compute $s \in H_0 \bmod p^N$ such that

$$G = D(r) + s \bmod p^N$$

for some $r \in L$. We call this process ‘finding a normal form’ in $L_{>0} \bmod p^{\varepsilon(N+1)}$ and say that s is ‘cohomologous’ to G . (Here, we identify H_0 with a subspace of L . Also, for an arbitrary $G \in L_{>0}$ the choice of ε would in fact depend upon the decay rate of the coefficients of G . As such, in what follows ε should be thought of as a variable; in Lemmas 27 and 31 we determine which values for ε suffice in the cases of interest.)

7.1. Normal forms in $L_{>0}/D(L)$: Case $p > 2$

In the case $p > 2$, because of the decay rate of power series in L , working in $L_{>0}/D(L)$ is particularly simple. From (15) and (18),

$$H = \pi \sum_{j=0}^d \hat{a}_j X^j, \quad \text{and thus} \quad H_X = \pi \sum_{j=1}^d \hat{a}_j j X^{j-1},$$

where $\text{ord}(\hat{a}_d d) = 0$ since $d \neq 0$ in \mathbb{F}_q . Consider the basis monomial $\pi^{\text{wt}(u)} X^u$ for the Banach module L , where $u \geq d$. We have the trivial identity

$$\begin{aligned} \pi^{\text{wt}(u)} X^u &= \left(\pi \left(\sum_{j=1}^d \hat{a}_j j X^j \right) + X \frac{d}{dX} \right) ((\hat{a}_d d)^{-1} \pi^{\text{wt}(u)-1} X^{u-d}) \\ &\quad - \left(\pi \left(\sum_{j=1}^{d-1} \hat{a}_j j X^j \right) + X \frac{d}{dX} \right) ((\hat{a}_d d)^{-1} \pi^{\text{wt}(u)-1} X^{u-d}) \\ &=: D(r) + r', \end{aligned} \tag{20}$$

where $r \in L$ and $r' \in L_{>0}$. Moreover, r' is a sum of monomials of degree less than u and greater than or equal to $u - d + 1$. Now let $G \in L_{>0} \bmod p^{\varepsilon(N+1)}$, and assume that $b\pi^{\text{wt}(u)}X^u$ is the highest term that occurs in G for some $b \in A$. (Note that by the decay rate on the coefficients of elements in L , we see that $u = \mathcal{O}(\varepsilon Npd)$.) We may suppose for our purposes that $u \geq d$. We can write this term as $D(br) + br'$, where br' is a sum of monomials of degree less than u but not less than $u - d + 1$. Write $G = G' + b\pi^{\text{wt}(u)}X^u$. Then in $L_{>0}/D(L)$ we find that G is ‘cohomologous’ to $G_1 := G' + br'$. To compute G_1 requires d multiplications in $A \bmod p^{\varepsilon(N+1)}$, and the same number of additions (plus a little precomputation, which can be ignored). Now continue in this way until the highest term in some G_m has degree less than d . Precisely, we need m at most $u - d + 1 = \mathcal{O}(\varepsilon Npd)$.

In this way we may find a ‘normal form’ for any element in $L_{>0}$. That is, given $G \in L_{>0} \bmod p^{\varepsilon(N+1)}$, we can write it as

$$G = D(r) + s,$$

where $r \in L$, and $s \in H_0 \bmod p^N$ is a linear combination with coefficients in $A \bmod p^N$ of the basis monomials

$$\{\pi X, \pi X^2, \dots, \pi X^{d-1}\}.$$

The process above has time complexity

$$\tilde{\mathcal{O}}((\varepsilon Npd)d(Npa)^c) \tag{21}$$

bit operations, where c is the exponent for multiplication as defined in Section 8.1. This complexity estimate lies at the heart of our proof of Theorem 1. (Strictly speaking, using the method that we describe, the final factor in the bracket should be $(\varepsilon Npa)^c$. However, a simple analysis based upon the proof of the next lemma shows that for each coefficient $c_u X^u$ in G where c_u is given modulo $p^{\varepsilon(N+1)}$, one must keep track of only the first N terms in the p -adic expansion of c_u , not including the leading zero terms.)

The next lemma justifies the choice of ε in the case $p > 2$.

Lemma 27. *Let $e \in \{\pi X, \dots, \pi X^{d-1}\}$ with $\alpha(e)$ cohomologous to s , a linear combination of the basis elements. To determine $s \bmod p^N$, it is enough to compute the coefficients of $\alpha(e)$ modulo $p^{4(N+1)}$.*

Proof. Let $e = \pi X^j$. By Lemma 9, terms in $F(e)$ are of the form $c_v X^{v+j}$, where $\text{ord}(c_v) \geq ((p-1)/p)^2(v/d)$. By the action of ψ_p , terms in $\alpha(e)$ are of the form $c'_u X^u$, where $\text{ord}(c'_u) \geq ((p-1)/p)^2((pu-j)/d)$. Equality (20) shows that for u such that

$$\left(\frac{p-1}{p}\right)^2 \left(\frac{up-j}{d}\right) - \frac{u}{d} \geq N,$$

the normal form of the term $c'_u X^u$ vanishes modulo p^N . Thus for

$$\frac{u}{d} \geq 3(N+1) \geq \frac{N+1}{((p-1)^2/p) - 1}$$

the term $c'_u X^u$ does not contribute to $s \bmod p^N$. It is now easy to check that computing the coefficients $c'_u X^u$ for u/d less than this bound with coefficients determined modulo p^b , where

$$b = 3(N+1) + (N+1) = 4(N+1),$$

is enough to determine the normal form of $\alpha(e)$ modulo p^N . □

7.2. Normal forms in $L_{>0}/D(L)$: Case $p = 2$

The reduction process for $p = 2$ is in essence the same; in other words, one uses a trivial identity to reduce the degree at each step. However, the justification that it works is somewhat more involved, since the power series in L decay rather more slowly.

We shall assume that $d\hat{a}_d = 1$; the more general situation just involves some notational complications, the essential point being that we always have $\text{ord}(d\hat{a}_d) = 0$. Let $J_1 := J - \{d\}$ be the support set J of f excluding the element d , and let $b_j := j\hat{a}_j$. We may assume that all non-constant terms in f have odd degree (Note 4), so each integer in J_1 is odd. (The argument below in fact works, provided only that d is odd, the key point being that in any case $\text{ord}(b_j) = \text{ord}(j)$ for all $j \in J_1$, and this is at least 1 for j even.) In a similar manner to that shown above, we have the identity

$$X^u = D\left(\frac{1}{\pi}X^{u-d}\right) - \left(\frac{(u-d)}{\pi}X^{u-d} + \sum_{j \in J_1} b_j X^{u-d+j}\right). \tag{22}$$

This is used to reduce a power series given in finite precision in $L_{>0}$ to its normal form, that is, a polynomial over $A \otimes \mathbb{Q}$ of degree less than d with no constant term. As before, the complexity is (21).

We must also address one theoretical problem. Let $G := \sum_u c_u X^u \in L$, and suppose that $c_u X^u = D(r_u) + s_u$ with s_u a polynomial of degree less than d over $A \otimes \mathbb{Q}$. Then $G = D(\sum_u r_u) + \sum_u s_u$, provided that $|r_u|, |s_u| \rightarrow 0$ as $u \rightarrow \infty$. To show that these sequences indeed converge, and to get a bound on their p -adic orders, requires a more careful analysis, which we now perform in a rather ad hoc fashion.

Lemma 28. *We may write any monomial X^u in the form $X^u = D(r) + s$, where r is a polynomial of degree at most $u - d$, and s is a linear combination of the monomials X, X^2, \dots, X^{d-1} , with coefficients in $A \otimes \mathbb{Q}$. Moreover, the coefficients of s have p -adic order at least $-mu - 1$, and for any v the coefficient of X^{u-v} in r has order at least $-mv - 2$. Here,*

$$m := \frac{1}{2d + (1/2)}.$$

Proof. Our approach will be to show that $X^u = D(r') + s'$, where r' satisfies the conditions in the statement of the lemma, and s' has the following property: it is a sum of terms of the form $c_v X^{u-v}$ where $v \geq 1$ and either $\text{ord}(c_v) \geq -mv$, or $\text{ord}(c_v) \geq -mv - 1$ with $u - v < d$. (In particular, if $v \geq 2d + 1$ with $\text{ord}(c_v) \geq -1$, or $v \geq 4d + 1$ with $\text{ord}(c_v) \geq -2$, then the term $c_v X^{u-v}$ is of the required form.) The result then follows by induction. For simplicity, we shall consider only the remainder term s' ; one may verify that the other term, r' , which we abbreviate as ‘*’, has the required properties in all cases.

If $u < d$ there is nothing to prove, so we assume that $u \geq d$. By (22), if $u - d$ is even, we have finished, since then $(u - d)/\pi$ is integral.

Assume then that $u - d$ is odd. Applying (22), we are reduced to considering the term $((u - d)/\pi)X^{u-d}$. If $u < 2d$, then $u - d < d$, and once again we have finished. Thus we assume that $u \geq 2d$. Applying (22) to this new term, we find that $((u - d)/\pi)X^{u-d}$ equals $D(*)$ plus

$$- \frac{(u-d)(u-2d)}{\pi^2} X^{u-2d} - \frac{u-d}{\pi} \sum_{j \in J_1} b_j X^{u-d-(d-j)}. \tag{23}$$

We first examine the terms in the final summation of (23). If $u - d - (d - j) < d$, we have finished. Otherwise, applying (22), we find that $((u - d)/\pi)X^{u-d-(d-j)}$ equals $D(*)$ plus

$$-\frac{(u - d)(u - 3d + j)}{\pi^2}X^{u-3d+j} - \frac{u - d}{\pi} \sum_{i \in J_1} b_i X^{u-d-(d-j)-(d-i)}. \tag{24}$$

For the first term, $(u - 3d + j)$ is divisible by 2, and also $u - 3d + j \leq u - (2d + 1)$. Thus this term is of the required form ‘ $c_v X^{u-v}$ ’, where $\text{ord}(c_v) \geq -1$ and $v \geq 2d + 1$. For the terms in the summation in (24), one repeatedly applies (22). The first time that one uses the middle term on the right-hand side of (22), one gets the term $c_v X^{u-v}$, say. Here, the order of the coefficient c_v is -1 , since $(u - 2d - (d - j) - (d - i) - \dots)$ is even. Moreover, we have $v = 2d + (d - j) + (d - i) + \dots \geq 2d + 1$, and once again this term is of the required form. If one never uses the middle term, then the resulting term $c_v X^{u-v}$, say, has coefficient c_v of order at least -1 and $u - v < d$, and once again we have finished.

It remains to consider the first term of (23). We may assume that $u \geq 3d$, (for otherwise it is already of the correct form). Applying (22) to this term, we get $D(*)$ plus

$$\frac{(u - d)(u - 2d)(u - 3d)}{\pi^3}X^{u-3d} + \frac{(u - d)(u - 2d)(u - 3d)}{\pi^2} \sum_{j \in J_1} b_j X^{u-2d-(d-j)}. \tag{25}$$

The terms in the last summation are of the required form, since $(u - 2d)$ is even and $2d + (d - j) \geq 2d + 1$. For the first term of (25), applying (22), we get $D(*)$ minus

$$\frac{(u - d) \dots (u - 4d)}{\pi^4}X^{u-4d} + \frac{(u - d)(u - 2d)(u - 3d)}{\pi^3} \sum_{j \in J_1} b_j X^{u-3d-(d-j)}. \tag{26}$$

The first term is of the required form, since $(u - 2d)(u - 4d)$ is divisible by 2^3 and also $4d \geq 2d + 1$. For the terms in the summation of (26), one repeatedly applies (22). The first time that one uses the middle term on the right-hand side of (22), one gets $c_v X^{u-v}$, say. Here the order of the coefficient c_v is -2 , since both $(u - 2d)$ and $(u - 4d - (d - j) - (d - i) - \dots)$ are divisible by 2. Also, $v = 4d + (d - j) + (d - i) + \dots \geq 4d + 1$, which shows that this term is of the required form. If one never uses the middle term, then the resulting term $c_v X^{u-v}$, say, has coefficient c_v of order at least -1 , since $(u - 2d)$ is even, and $u - v < d$, and once again we have finished. □

Corollary 29. For $p = 2$, the set $\{\pi X, \dots, \pi X^{d-1}\}$ is a basis for $L_{>0}/D(L)$.

Proof. Since $\gamma > 1/(2 + (1/2d))$, by Lemma 28 and a straightforward continuity argument we know that this set spans $L_{>0}/D(L)$. Any basis cannot have fewer than $d - 1$ elements, by consideration of the degree of the L-function $L(f, T)$ (using (8)), and so it must be a basis. □

The above result is primarily of theoretical interest. The next lemma shows that in the case in which we are interested, denominators do not in fact occur.

Lemma 30. Let $p = 2$ and $e \in \{\pi X, \pi X^2, \dots, \pi X^{d-1}\}$. Then $\alpha(e)$ is cohomologous to an element $\sum_{j=1}^{d-1} m_j (\pi X^j)$, where $m_j \in A$.

Proof. Dividing through by π , we show that $\alpha(\pi^{-1}e)$ is cohomologous to an element $\sum_{j=1}^{d-1} m_j X^j$ with $m_j \in A$. Let $e = X^j$, where $1 \leq j \leq d - 1$. Then the terms in $F(e)$ are of the form $c_v X^{v+j}$, where $\text{ord}(c_v) \geq v/4d$ (from Lemma 9). By the action of ψ_2

the terms in $\alpha(X^j)$ are of the form $c'_u X^u$, where $\text{ord}(c'_u) \geq \lceil (u/2d) - (j/4d) \rceil$. We claim that such terms are cohomologous to polynomials of degree at most $d - 1$ over A . If $u \leq 2d$, the result is true by one application of (22) and induction. For $u > 2d$, one proves the result using at most two applications of (22) and induction. \square

The next lemma justifies the choice of ε in the case $p = 2$.

Lemma 31. *To compute the normal form of $\alpha(e) \in L$ modulo p^N , it is sufficient to determine the coefficients of $\alpha(e)$ modulo $p^{(4d+1)(N+1)}$.*

Proof. Lemma 28 shows that for u such that

$$\text{ord}(c_u) - \frac{u}{2d + (1/2)} - 1 \geq N,$$

the normal form of the term $c_u X^u$ vanishes modulo p^N . Since $\alpha(e) \in L$, terms in this power series are of the form $c_u X^u$, say, where $\text{ord}(c_u) \geq u/2d$. For $u/2d := (4d + 1)(N + 1)$, we have $u/(2d + (1/2)) = 4d(N + 1)$ and $(u/2d) - (u/(2d + (1/2))) - 1 = N$, and for $u' > u$ we get a strict inequality in the latter. Thus, it is enough to work modulo p^b where

$$b = 4d(N + 1) + (N + 1) = (4d + 1)(N + 1). \quad \square$$

7.3. Computing the absolute Frobenius matrix

We now describe how to perform the main step of the algorithm—that is, constructing the matrix for the absolute Frobenius map with respect to the basis $\{\pi X, \dots, \pi X^{d-1}\}$. First, one may compute F with the coefficients determined modulo $p^{\varepsilon(N+1)}$ directly from the formula in Definition 8 and the expression for $\theta(t)$ in (14). Working with coefficients modulo $p^{\varepsilon(N+1)}$, for each basis element e the polynomial $\psi_p \circ F(e) \bmod p^{\varepsilon(N+1)}$ may be constructed. The reduction method of Sections 7.1 and 7.2 is then used to write this as a linear combination of the basis elements $\pi X, \dots, \pi X^{d-1}$. In this way the matrix M is found, with coefficients determined modulo p^N . (Note that the entries in M are p -adic integers for all p , by Lemma 30.)

7.4. Finding the characteristic polynomial of Frobenius

One may compute the matrix M_a via the formula

$$M_a = \prod_{i=0}^{a-1} \tau^{-i}(M). \tag{27}$$

This is proved from Lemma 24 in the same way as [14, Lemma 26]. (See also the sentence following that lemma for an alternative approach.) The characteristic polynomial may then be found deterministically by computing $\text{Tr}(M_a^k)$ for $1 \leq k \leq d$ and using the Newton identity

$$\det(I - M_a T) = \exp\left(-\sum_{k=1}^{\infty} \frac{\text{Tr}(M_a^k)}{k} T^k\right).$$

(Alternatively, one could use an interpolation method.) The numerator of the zeta function may now be found by computing the conjugates (11) and taking a product to get a polynomial in $\mathbb{Z}_p[T]$.

8. Complexity analysis

8.1. Exponents for ring multiplication

Let $A/(p^r)$ be some modular reduction of A , and let $A/(p^r)[X]$ be the ring of polynomials in one variable over $A/(p^r)$. Denote by c the (deterministic) exponent for multiplication in both rings. Precisely, polynomials of degree δ in $A/(p^r)[X]$ can be multiplied in $\tilde{\mathcal{O}}(\delta^c)$ operations in $A/(p^r)$, and elements in $A/(p^r)$ can be multiplied in $\tilde{\mathcal{O}}((rpa \log(p))^c)$ bit operations, where $rpa \log(p)$ is the logarithm of the size of the ring. Using classical methods, we take $c = 2$; Karatsuba’s algorithm gives $c = \log_2(3) < 1.59$, and $c = 1$ using Fast Fourier Transform (FFT) methods. (Note that we ignore logarithmic factors, and the space complexity is $\tilde{\mathcal{O}}(\delta)$, $\tilde{\mathcal{O}}(rpa \log(p))$ in all cases. For polynomial multiplication using FFT methods, we refer to [3]; we assume as in [12, Section 5] and [10, Section 4.3] that FFT methods may be applied to $A/(p^r)$, although we do not know a convenient reference for this.) Similarly, let ω denote the exponent for deterministic multiplication of matrices over $A/(p^r)$; thus two $\delta \times \delta$ matrices can be multiplied in $\tilde{\mathcal{O}}(\delta^\omega)$ operations in $A/(p^r)$. It will transpire that the choice of $\omega (\leq 3)$ does not affect the overall complexity.

8.2. Complexity of the point-counting algorithm

First, we must compute F with coefficients determined modulo $p^{\varepsilon(N+1)}$. This may be done by multiplying together the $\mathcal{O}(d)$ polynomials $\theta(\hat{a}_j X^j)$ in the ring $\tilde{L}(\delta) \bmod p^{\varepsilon(N+1)}$ for δ as in Lemma 9. (Note that each $\theta(t)$ can itself be constructed via multiplication of truncated power series of the form $\exp(*t^{p^j})$, as in [14, Lemma 29]. This part is dominated by the computation of F itself.) From the decay rate of the coefficients of the power series in $\tilde{L}(\delta)$, we see that polynomials $\theta(\hat{a}_j X^j) \bmod p^{\varepsilon(N+1)}$ are linear combinations of the monomials $\pi^{\lceil \tilde{w}t(i)\delta \rceil} X^i$ for $\tilde{w}t(i)$ bounded so that

$$\lceil \tilde{w}t(i)\delta \rceil \leq \varepsilon(N + 1)(p - 1).$$

There are $\mathcal{O}(\varepsilon Npd)$ such terms. Since the coefficients of the power series $\theta(\hat{a}_j X^j)$ lie in A , it follows that the construction of F may be done in

$$\tilde{\mathcal{O}}((\varepsilon Npd)^c (\varepsilon Npa \log p)^c) = \tilde{\mathcal{O}}((\varepsilon^2 N^2 p^2 ad)^c) \tag{28}$$

bit operations. (Here, $\mathcal{O}(\varepsilon Npa \log p)$ is the bit-size of elements in the ring $A \bmod p^{\varepsilon(N+1)}$.)

Second, finding $\psi_p \circ F(e)$ for all $d - 1$ basis monomials requires

$$\tilde{\mathcal{O}}(d(\varepsilon Npd)(\varepsilon Npa)) = \tilde{\mathcal{O}}(\varepsilon^2 N^2 p^2 ad^2) \tag{29}$$

bit operations. Here, we use the quasi-linear time method to compute the map τ^{-1} on $A \bmod p^{\varepsilon(N+1)}$ suggested in [12, Section 5] and [10, Step 2] (namely, precomputation of the map on the element ϵ by Newton iteration).

Third, we must compute a normal form for each such expression to find the coefficients in the matrix M . By the time estimate (21), we see that each column of M can be computed in

$$\tilde{\mathcal{O}}(\varepsilon N^{c+1} p^{c+1} d^2 a^c) \tag{30}$$

bit operations. We require $d - 1 = \mathcal{O}(d)$ such computations.

Fourth, computing the matrix for M_a may be done using equation (27) and the fast exponentiation method of [14, Lemma 31] in

$$\tilde{\mathcal{O}}(d^\omega (Npa)^c) \tag{31}$$

bit operations.

Finally, the computation of the characteristic polynomial takes

$$\tilde{\mathcal{O}}(d^{\omega+1}(Npa)^c) \tag{32}$$

bit operations. (Computation of the products of the conjugates using (11) is absorbed in the other estimates.) Adding (28), (29), $d \times$ (30), (31) and (32) together, and putting $N = \mathcal{O}(pad)$ with $\varepsilon = \mathcal{O}(1)$, $\mathcal{O}(d)$, we get

$$\begin{aligned} \tilde{\mathcal{O}}(p^{4c}a^{3c}d^{4+c}), & \quad \text{for } p > 2, \\ \tilde{\mathcal{O}}(p^{4c}a^{3c}d^{\max(5c,c+5)}), & \quad \text{for } p = 2, \end{aligned}$$

where c is the exponent for ring multiplication as discussed in Section 8.1. (Here we have assumed that $\omega \leq 3$.) Using FFT methods, we may take $c = 1$, giving the time complexity claimed in Theorem 1.

The space complexity is in all cases determined by the size of the ring $L \bmod p^{\varepsilon(N+1)}$ and also the polynomial $F \bmod p^{\varepsilon(N+1)}$. These are both

$$\tilde{\mathcal{O}}((\varepsilon Npd)(\varepsilon Npa)) = \begin{cases} \tilde{\mathcal{O}}(p^4a^3d^3), & \text{for } p > 2, \\ \tilde{\mathcal{O}}(a^3d^5), & \text{for } p = 2. \end{cases}$$

This completes the proof of Theorem 1 in the case of Type 1 Artin–Schreier curves, and also Type 2, since they are easily reduced to Type 1.

Note 32. It is possible to reduce the factors d^{5c} and d^5 in the time and space complexities, respectively, for $p = 2$ to d^{4c} and d^4 using a more careful analysis: we have not taken into account that only part (N terms) of the p -adic expansion of each coefficient in $F \bmod p^{\varepsilon(N+1)}$ needs to be computed. A much more detailed analysis of the action of D could perhaps reduce the value required for ε when $p = 2$, giving a uniform time estimate of $\tilde{\mathcal{O}}(p^4a^3d^5)$ with space $\tilde{\mathcal{O}}(p^4a^3d^3)$ for all p .

Note 33. We briefly describe one alternative approach for $p = 2$, which has also been implemented by Vercauteren. The idea is to use a more complicated splitting function to improve the decay rate of the coefficients in F . The result is that one may use $L := \tilde{L}(1) \cap A\{X\}$ for the case $p = 2$, and the proof that $L_{>0}/D(L)$ becomes easier. Specifically, setting $p = 2$, take $\theta := \theta_3$, where θ_3 is the splitting function defined in [6, p. 55]. One may compute the required element $\gamma_3 \in \mathbb{Z}_p$ in the following manner: let γ'_3 satisfy the relation

$$\gamma'_3 = 1 + 4\gamma_3'^3 - 8\gamma_3'^4 + 8\gamma_3'^5.$$

Then $\gamma_3 := 2\gamma'_3$. Let α be defined as before, but with the new θ . This time, it is stable on $L := \tilde{L}(1) \cap A\{X\}$. Define H to be

$$H := \sum_{j=0}^2 \gamma_{3,j} \tau^j(\hat{f}(X^{p^j})),$$

where

$$\gamma_{3,j} := \sum_{i=0}^j \frac{\gamma_3^{p^i}}{p^i}.$$

Then one must compute the action of α on the homology $L_{>0}/D(L)$, where $D := X(d/dX) + H_X$. To do this, for $j \geq 0$ define $\pi^j L_{>0}$ to be the Banach module over A with orthonormal basis $\pi^{\text{wt}(u)+j} X^u$ ($u > 0$). Then

$$H_X \equiv \pi X \frac{d\hat{f}}{dX} \bmod \pi L_{>0}.$$

Now one performs the reduction process from Section 7.1 coupled with a Hensel lifting argument to reveal r and s with $G = D(r) + s$, working modulo $\pi^j L_{>0}$ for increasing powers of j . Note that one may take any $N \geq (d - 1)(1 + (a/2)) + 1$ and $\varepsilon = 11/3$, and all computations are done with p -adic integers. The algorithm has complexity $\tilde{\mathcal{O}}(a^{\max(3c,4)} d^6)$, where c is the exponent for multiplication and space complexity $\tilde{\mathcal{O}}(a^3 d^3)$. This is slightly slower in terms of a due to a step in the Hensel lifting, for which we were unable to find a fast method.

9. Artin–Schreier covers of the torus

In this section we present in a much more condensed fashion the algorithm for computing the zeta function of the Type 3 Artin–Schreier curves.

Let $f \in \mathbb{F}_q[X, X^{-1}]$ have negative degree $d^- < 0$ and positive degree $d^+ > 0$. We shall explain how to compute the L-function $L^*(f, T)$. From this one may easily compute the zeta functions of Type 3 Artin–Schreier curves using (9).

Define a weight function wt on \mathbb{Z} by

$$\text{wt}(u) := \begin{cases} \lceil u/d^- \rceil, & \text{if } u < 0, \\ \lceil u/d^+ \rceil, & \text{if } u \geq 0, \end{cases}$$

where $\lceil \cdot \rceil$ is the usual (not p -adic) absolute value. For $p > 2$, let L be the Banach module over A with orthonormal basis $\pi^{\text{wt}(r)} X^r$ for $r \in \mathbb{Z}$. For $p = 2$, let L denote the Banach space over $A \otimes \mathbb{Q}$ with orthonormal basis $\pi^{\lceil \gamma \text{wt}(r) \rceil} X^r$ for $r \in \mathbb{Z}$. (Here γ is any rational number with $1/2 > \gamma > 1/(2 + 1/(2d'))$, where $d' := \max(|d^-|, d^+)$.) Note that L contains two-way infinite power series, and in fact still forms a ring because of the decay conditions on coefficients. Now define the power series F and $F^{(a)}$ in exactly the same manner as before. In this case they are two-way infinite power series. With α and α_a defined exactly as before, we find that both maps are stable on L . We have the formula

$$L^*(f, T) = \frac{\det(1 - T\alpha_a|L)}{\det(1 - Tq\alpha_a|L)}.$$

In this case one cannot remove any unit root factors. Let \mathcal{L} be the (slightly ‘larger’) complex of modules

$$0 \longrightarrow L \xrightarrow{D} L \longrightarrow 0.$$

These are A -modules for $p > 2$ and $A \otimes \mathbb{Q}$ -spaces for $p = 2$. Here D is defined in exactly the same manner as before (in Section 5). In this case, $H_0 = L/D(L)$ and $H_1 = 0$, and one recovers the cohomological trace formula

$$L^*(f, T) = \det(1 - H_0(\alpha_a)T).$$

Once again we have the crucial relation $H_0(\alpha_a) = H_0(\alpha)^a$. A basis for H_0 can be taken as

$$\{\pi X^{d^-}, \pi X^{d^-+1}, \dots, \pi X^{-1}, 1, \pi X, \dots, \pi X^{d^+-1}\},$$

with $d^+ - d^-$ the dimension of this space. Computation of normal forms in $L/D(L)$ is done in a similar manner to before. Precisely, one first uses the method in Sections 7.1 and 7.2 to find an element that is cohomologous to a given element whose leading term has degree less than d^+ . One then performs a similar process to increase the degree of the lowest term so that it is not less than d^- . Note that we need a final p -adic accuracy of any N greater than $(p - 1)(d^+ - d^-)(a/2 + 1)$. For $p > 2$, the factor ε can be taken to be 4, and for $p = 2$, it is $4d' + 1$. The complexity of the above algorithm may be checked to be identical to that in

the case when f is just an ordinary polynomial. More precisely, if we write $d := d^+ - d^-$, then all the bounds in Section 8 are still true. This completes the proof of Theorem 1, and Corollary 2 follows from Note 5.

Acknowledgements. The research was partly undertaken when the authors were visiting the Lorentz Center, University of Leiden, and the second author the Institute for Mathematical Sciences, National University of Singapore. We wish to thank Hendrik Lenstra and Harald Niederreiter, respectively, for their support during these visits, and also Pierrick Gaudry and Frederik Vercauteren for helpful comments on an earlier version of the paper.

References

1. A. ADOLPHSON and S. SPERBER, ‘Exponential sums and Newton polyhedra: cohomology and estimates’, *Ann. of Math.* 130 (1989) 367–406. 35
2. I. BLAKE, G. SEROUSSI and N. SMART, *Elliptic curves in cryptography*, LMS Lecture Note Ser. 265 (Cambridge Univ. Press, 1999). 34, 35, 37
3. D. CANTOR and E. KALTOFEN, ‘On fast multiplication of polynomials over arbitrary algebras’, *Acta Inform.* 28 (1991) 693–701. 51
4. R. F. COLEMAN, ‘ p -adic Banach spaces and families of modular forms’, *Invent. Math.* 127 (1997) 417–479. 39, 42
5. B. DWORK, ‘On the rationality of the zeta function of an algebraic variety’, *Amer. J. Math.* 82 (1960) 631–648. 39
6. B. DWORK, ‘On the zeta function of a hypersurface’, *Inst. Hautes Études Sci. Publ. Math.* 12 (1962). 35, 39, 43, 52
7. N. ELKIES, ‘Elliptic and modular curves over finite fields and related computational issues’, *Computational perspectives in number theory: Proceedings of a conference in honour of A. O. L. Atkin*, AMS/IP Stud. Adv. Math. 7 (ed. D. A. Buell and J. T. Teitelbaum, Amer. Math. Soc., Providence, RI, 1998) 21–76. 35
8. M. FOUQUET, P. GAUDRY and R. HARLEY, ‘An extension of Satoh’s algorithm and its implementation’, *J. Ramanujan Math. Soc.* 15 (2000) 281–318. 35
9. M. FOUQUET, P. GAUDRY and R. HARLEY, ‘Finding secure curves with the Satoh–FGH algorithm and an early abort strategy’, *Advances in Cryptology - EUROCRYPT 2001*, Lecture Notes in Comput. Sci. 2045 (ed. B. Pfitzmann, Springer, 2001) 14–29. 35
10. P. GAUDRY and N. GÜREL, ‘An extension of Kedlaya’s algorithm for counting points on superelliptic curves’, preprint, 2001. 35, 51, 51
11. P. GAUDRY and R. HARLEY, ‘Counting points on hyperelliptic curves over finite fields’, *Advances in Cryptology - EUROCRYPT 2000*, Lecture Notes in Comput. Sci. 1807 (ed. B. Preneel, Springer, 2000) 19–34. 35
12. K. S. KEDLAYA, ‘Counting points on hyperelliptic curves using Monsky–Washnitzer cohomology’, preprint, 2001. 35, 35, 51, 51
13. N. KOBLITZ, ‘Hyperelliptic cryptosystems’, *J. Cryptology* 1 (1989) 139–150. 34
14. A. G. B. LAUDER and D. WAN, ‘Counting points on varieties over finite fields of small characteristic’, preprint, 2001. 35, 35, 35, 36, 38, 39, 39, 39, 39, 40, 40, 41, 41, 42, 42, 42, 42, 50, 51, 51
15. R. LIDL and H. NIEDERREITER, *Introduction to finite fields and their applications* (Cambridge Univ. Press, 1986). 36

16. B. POONEN, ‘Computational aspects of curves of genus at least 2’, *Algorithmic number theory II*, Lecture Notes in Comput. Sci. 1122, (ed. H. Cohen, Springer, 1996) 283–306. 35
17. T. SATOH, ‘The canonical lift of an ordinary elliptic curve over a finite fields and its points counting’, *J. Ramanujan Math. Soc.* 15 (2000) 247–270. 35, 35
18. R. SCHOOF, ‘Elliptic curves over finite fields and the computation of square roots mod p ’, *Math. Comp.* 44 (1985) 483–494. 35
19. R. SCHOOF, ‘Counting points on elliptic curves over finite fields’, *J. Théor. Nombres Bordeaux* 7 (1998) 219–254. 35
20. J.-P. SERRE, ‘Endomorphismes complètement continus des espaces de Banach p -adique’, *Inst. Hautes Études Sci. Publ. Math.* 12 (1962) 69–85. 39, 42, 44, 44
21. S. SPERBER, ‘On the p -adic theory of exponential sums’, *Amer. J. Math.* 108 (1983) 255–296. 43
22. F. VERCAUTEREN, B. PRENEEL and J. VANDEWALLE, ‘A memory efficient version of Satoh’s algorithm’, *Advances in Cryptology - EUROCRYPT 2001*, Lecture Notes in Comput. Sci. 2045 (ed. B. Pfitzmann, Springer, 2001) 1–13. 35
23. D. WAN, ‘Computing zeta functions over finite fields’, *Contemp. Math.* 225 (1999) 131–141. 35
24. D. WAN, ‘Pure L-functions from algebraic geometry over finite fields’, *Finite fields and applications* (ed. D. Jungnickel and H. Niederreiter, Springer, 2000) 437–461. 36
25. D. WAN, ‘Algorithmic theory of zeta functions over finite fields’, preprint, 2001. 35

Alan G. B. Lauder alan.lauder@comlab.ox.ac.uk
<http://web.comlab.ox.ac.uk/oucl/work/alan.lauder/>

Computing Laboratory
Oxford University
Oxford OX1 3QD

Daqing Wan dwan@math.uci.edu
<http://www.math.uci.edu/~dwan/Overview.html>

Department of Mathematics
University of California
Irvine, CA 92697
USA