Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT®
Journal of Complexity 20 (2004) 331–349
Journal of COMPLEXITY
http://www.elsevier.com/locate/jco

Computing zeta functions of Artin-Schreier curves over finite fields II ☆

Alan G.B. Lauder[a,*] and Daqing Wan[b,c]

[a] *Mathematical Institute, Oxford University, 24-29 St. Giles, Oxford OX1 3LB, UK*
[b] *Department of Mathematics, University of California, Irvine, CA 92697, USA*
[c] *Institute of Mathematics, Chinese Academy of Sciences, Beijing, People's Republic of China*

Received 2 October 2002; accepted 12 August 2003

Abstract

We describe a method which may be used to compute the zeta function of an arbitrary Artin-Schreier cover of the projective line over a finite field. Specifically, for covers defined by equations of the form $Z^p - Z = f(X)$ we present, and give the complexity analysis of, an algorithm for the case in which $f(X)$ is a rational function whose poles all have order 1. However, we only prove the correctness of this algorithm when the field characteristic is at least 5. The algorithm is based upon a cohomological formula for the L-function of an additive character sum. One consequence is a practical method of finding the order of the group of rational points on the Jacobian of a hyperelliptic curve in characteristic 2.

Keywords: Artin-Schreier curve; Hyperelliptic curve; Finite field; Zeta function; Algorithm

1. Introduction

Starting with the work of Schoof [20], the computation of zeta functions for varieties over finite field has been a central strand in algorithmic number theory. An extensive reference of the papers in this area can be found in [3], and more recent work includes [8,9,19,22]. Various fast methods have been developed for elliptic curves, but until recently general curves, let alone higher dimensional varieties,

☆ Alan Lauder is a Royal Society University Research Fellow. Daqing Wan is partially supported by the NSF and the NNSF of China (10128103).
*Corresponding author.
E-mail addresses: lauder@maths.ox.ac.uk (A.G.B. Lauder), dwan@math.uci.edu (D. Wan).

0885-064X/$ - see front matter © 2003 Elsevier Inc. All rights reserved.
doi:10.1016/j.jco.2003.08.009

seemed beyond reach both theoretically and practically. In [13] the present authors proved polynomial-time computability of the zeta function of an arbitrary variety of fixed dimension over a finite field of small characteristic. This result was based upon Dwork's proof of the rationality of the zeta function of a variety [5], but unfortunately does not lead to practical algorithms. Dwork's proof does not use cohomology, but it was the starting point for the development of several partial p-adic cohomological theories, unified in rigid cohomology [2]. Independently of the present authors, Kedlaya proposed a fast method for counting points on a special class of curves based upon part of this theory, namely Monsky–Washnitzer cohomology for smooth affine varieties [10]. The authors have also developed fast methods for other classes of curves using the so-called Dwork cohomology, and more generally "Dwork-Reich cohomology" [12,14]. The two approaches may be explained in a common language: Kedlaya uses a formula from rigid cohomology for the L-function of the constant sheaf on the curve, whereas we use one for the L-function of non-constant sheaves (Artin-Schreier and Kummer) on open subsets of the projective line. Thus in our approach the complication lies in the sheaf, while with Kedlaya it lies in the base variety. In this paper we extend the work of [14] to arbitrary Artin-Schreier covers of the projective line. We only give details for a restricted class of curves, but these are sufficient to cover a large class of hyperelliptic curves in characteristic 2, which is the main motivation behind the work. (Such curves are of interest in cryptography [11], and the fast computation of their Jacobian orders was a long-standing open problem, first resolved for arbitrary genus in a special case in [13], and with no restrictions in the recent work [4,21].) Unfortunately, we are only able to prove the correctness of our algorithm for the case $p \geqslant 5$.

We now describe the main theorem. Denote by $\mathbb{F}_q$ the finite field with q elements, where $q = p^a$ and p is prime. Fix $\bar{\mathbb{F}}_q$ an algebraic closure of $\mathbb{F}_q$, and for each $k \geqslant 1$ denote by $\mathbb{F}_{q^k}$ the unique subfield of order q^k. Let $\bar{f} \in \mathbb{F}_q(X)$ be a rational function. Let $\tilde{C}_{\bar{f}}$ to be the unique smooth projective curve birational to the curve with equation $Z^p - Z = \bar{f}$ in $\{x \in \bar{\mathbb{F}}_q \,|\, \bar{f}(x) \neq \infty\} \times \bar{\mathbb{F}}_q$. We prove the following result.

Theorem 1. *Assume that $p \geqslant 5$, and that the poles of $\bar{f}$ include zero and infinity and all have order one. Let d_g be the number of finite poles of $\bar{f}$. The zeta function of the smooth projective curve $\tilde{C}_{\bar{f}}$ may be computed deterministically in $\tilde{\mathcal{O}}(p^4 a^3 d_g^6)$ bit operations.* (*See the paragraph above for an explanation of the notation used.*)

By "zero" and "infinity" we mean the zero and pole, respectively, of the function X on the projective line. We use soft-Oh notation which ignores logarithmic factors, as in [13, Section 6.3]. We believe the algorithm also works when $p < 5$, although we do not give a proof of this. In particular, the analysis in the case $p = 2$ requires further work which we have not undertaken. Also, it should be possible to improve the complexity dependence on d_g to fourth power, although our analysis at present is not

good enough to show this. Regarding the pole order restrictions, the method extends easily to arbitrary rational functions whose finite poles all have a common order. One should be able to tackle the mixed-pole order case using partial fraction decompositions, but this seems rather complicated and we have not worked out the details.

The case $p = 2$ is of particular interest. Let C be the unique smooth projective curve birational to that with affine equation

$$Y^2 + r(X)\,Y = s(X) \tag{1}$$

where $r, s \in \mathbb{F}_{2^a}[X]$ with $\deg(s) = 2\gamma + 1$, $\deg(r) = \gamma$, and $r(X)$ squarefree with $r(0) = 0$. This curve can be transformed easily into the form $Z^2 + Z = \bar{f}$ for some rational function $\bar{f}$ with poles of order one. Assuming that the algorithm we present also works when $p = 2$, and that one can improve the complexity dependence on d_g as suggested above, this would yield an algorithm for determining the order of the group of rational points on the Jacobian of C with running time $\tilde{\mathcal{O}}(a^3\gamma^4)$ bit operations. It would be of great interest to give a rigorous proof of this, and implement the algorithm in this case. Note that in [13, Corollary 2], the complexity given when $r(X)$ has a single root is sixth power in the genus (although one might be able to improve this to fifth power). In this case our approach is a little slower than when $r(X)$ has distinct roots.

The article is organised in the following manner. In Sections 2 and 3 we derive a p-adic cohomological formula for the zeta function of an arbitrary Artin-Schreier curve. In these sections essentially no restrictions are placed upon the rational function $\bar{f}$, and the formula itself is not original. In Section 4 we state the simplifying assumptions we shall make. Note that at this stage the assumption $p \geqslant 5$ is still not used, only that on the pole orders. Section 5 contains the algorithm, and Section 6 our method for reduction in the cohomology space. This method is original. Up until this stage we have not at any point used the assumption $p \geqslant 5$; however, it is required in Section 7 where we justify the choice of p-adic accuracy in the algorithm. Section 8 contains comments on some non-trivial subroutines that much be performed in the algorithm, and the complexity analysis is given in Section 9. Thus the only aspect of the algorithm which we do not justify for $p < 5$ is that the choice of p-adic accuracy is sufficient to recover the zeta function.

2. Zeta functions and L-functions

We begin by recalling a classical relationship between the zeta function of an Artin-Schreier curve, and the L-function of an associated additive character (see for example [15, Section 4.6.2]).

Let $\bar{f} \in \mathbb{F}_q(X)$ have a pole at infinity. Let $C_{\bar{f}}$ denote the curve embedded in $\{x \in \bar{\mathbb{F}}_q \mid \bar{f}(x) \neq \infty\} \times \bar{\mathbb{F}}_q$ with equation

$$Z^p - Z = \bar{f}(X).$$

Denote by $\tilde{C}_{\bar{f}}$ the unique smooth projective curve birational to $C_{\bar{f}}$. By Weil the zeta function of $\tilde{C}_{\bar{f}}$ has the form

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

where $P(T)$ is a polynomial of degree twice the genus whose roots have complex absolute value $q^{1/2}$.

Let ζ be a primitive pth root of unity in $\bar{\mathbb{Q}}$. For $x \in \mathbb{F}_p$ define $\Psi(x) = \zeta^x$, where the exponent is thought of as an integer. Let $\Psi_k : \mathbb{F}_{q^k} \to \bar{\mathbb{Q}}$ be defined as $\Psi \circ \mathrm{Tr}_k$ where $\mathrm{Tr}_k : \mathbb{F}_{q^k} \to \mathbb{F}_p$ is the absolute trace map.

Define the exponential sum

$$S_k(\bar{f}, \Psi) := \sum_{x \in \mathbb{F}_{q^k},\ \bar{f}(x) \neq \infty} \Psi_k(\bar{f}(x)).$$

Associate to this the L-function

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

Let $L^*(\bar{f}, T)$ denote the L-function for the sums $S_k^*(\bar{f}, \Psi)$ over $\{x \in \mathbb{F}_{q^k}^* \mid \bar{f}(x) \neq \infty\}$. Using the same approach as [14, Section 2] one may show that

$$P(T) = \prod_{i=1}^{p-1} \theta_i(L(\bar{f}, T)),$$

where $\theta_i : \zeta \mapsto \zeta^i$ are the automorphisms of the cyclotomic field $\mathbb{Q}(\zeta)$ (cf. with [15, Eq. (4.9)]). Thus the computation of the zeta function of the smooth projective curve $\tilde{C}_{\bar{f}}$ reduces to that of the L-function $L(\bar{f}, T)$.

3. A p-adic cohomological formula

In this section we present a p-adic cohomological formula for the L-function $L^*(\bar{f}, T)$ which is the basis of our algorithm. The formula in the incarnation we give is due to Robba [18], but it owes greatly to the work of Dwork [6] and Reich [16,17].

Let $\bar{f} \in \mathbb{F}_q(X)$ be a rational function. The only assumption we shall make in this section is that infinity is a pole of $\bar{f}$, and that all the poles of $\bar{f}$ have order not divisible by the characteristic p. Note that by applying a suitable isomorphism over $\mathbb{F}_q$ one may always find an equation for the curve $Z^p - Z = \bar{f}$ such that $\bar{f}$ takes this form. (Namely, extend the method of [14, Note 4] by considering the finite poles as well as the pole at infinity, using a partial fraction decomposition.)

Let $\mathbb{Q}_p$ be the field of p-adic numbers with ring of integers $\mathbb{Z}_p$. Fix Ω the completion of an algebraic closure of $\mathbb{Q}_p$. Let K_0 be the unique unramified extension of $\mathbb{Q}_p$ of degree a in Ω (recall that $q = p^a$). Denote by R_0 the ring of integers of K_0.

Fix a choice of $\pi \in \Omega$ with $\pi^{p-1} = -p$ and define $K_1 = \mathbb{Q}_p(\pi)$, which has ring of integers $R_1 = \mathbb{Z}_p(\pi)$. Let K denote the compositum field of K_0 and K_1, and R the compositum ring of R_0 and R_1. Let τ be the lifting of the Frobenius pth power map from $\mathbb{F}_q$ to K with $\tau(\pi) = \pi$. For a power series or polynomial $H(X)$, the notation $H^\tau(X)$ will mean τ acting on coefficients and fixing the indeterminate X.

Let $f \in K_0(X)$ be any lifting of $\bar{f}$. In Section 4 we take a specific lifting which makes computations and analysis easier, but for now it may be arbitrary. Let the pole set of $\bar{f}$ be $\{\bar{r}_0 = \infty, \bar{r}_1, \ldots, \bar{r}_{d_g}\}$, and denote by $\{r_i\}$ the corresponding pole set of f. Define

$$A := \{x \in \Omega \mid \mathrm{ord}(x) \geqslant 0,\ \mathrm{ord}(x - r_i) \leqslant 1 \text{ for } 1 \leqslant i \leqslant d_g\}.$$

This is just the projective line over Ω with unit disks around the poles of f removed. Denote by $\mathscr{H}^\dagger(A)$ the "ring of overconvergent analytic elements on A" [18, Section 5.1]. This space may be identified with the ring of formal power series

$$a_0 + \sum_{i=0}^{d_g} \sum_{j=1}^{\infty} a_{ij} T_i^j. \tag{2}$$

Here for $1 \leqslant i \leqslant d_g$ we take $T_i = 1/(X - r_i)$, and $T_0 = X$. Also, the coefficients a_0, a_{ij} lie in Ω and for each such power series (2) there exists $\epsilon > 0$ such that $\mathrm{ord}(a_{ij}) - \epsilon j \geqslant 0$ for all i and sufficiently large j.

Define the map $\psi_p : \mathscr{H}^\dagger(A) \to \mathscr{H}^\dagger(A)$ as follows. For $\xi \in \mathscr{H}^\dagger(A)$, the function $\psi_p(\xi)$ is defined for $x \in A$ as

$$\psi_p(\xi)(x) := \frac{1}{p} \tau^{-1} \left(\sum_{z^p = x} \xi(z) \right).$$

This is a τ^{-1}-linear map under which $\mathscr{H}^\dagger(A)$ is stable. Define $\psi_q := \psi_p^a$, a τ^{-a}-linear map on $\mathscr{H}^\dagger(A)$. The following properties of ψ_p will be used on several occasions. They may be proved directly from the definition.

Lemma 2. *For $\alpha(X)$ and $\beta(X)$ power series of the form* (2) *in $\mathscr{H}^\dagger(A)$ we have*

$$\psi_p(\alpha^\tau(X^p)\beta(X)) = \alpha(X)\psi_p(\beta(X)).$$

The action of ψ_p on monomials is given by

$$\psi_p(X^u) = \begin{cases} X^{u/p} & \text{if } p|u, \\ 0 & \text{otherwise.} \end{cases}$$

The first property may be used to compute the action of ψ_p on rational functions (see Section 8.1 and also [1, Lemma 1]).

Let $\theta(X)$ denote the formal power series with coefficients in R_1 obtained by expansion from the relation

$$\theta(X) := \exp(\pi(X - X^p)).$$

From, for example, [13, Section 4] we have that

$$\theta(X) = \sum_{r=0}^{\infty} \lambda_r X^r, \quad \mathrm{ord}(\lambda_r) \geqslant \frac{(p-1)}{p^2} r. \tag{3}$$

Let F and $F^{(a)}$ be the functions in $\mathscr{H}^\dagger(A)$ defined for $x \in A$ via composition of functions as

$$F(x) := \exp(\pi(f(x)^p - f^\tau(x^p)))\theta(f(x)),$$

$$F^{(a)}(x) := \exp(\pi(f(x)^q - f(x^q)))\theta(f(x))\theta(f(x)^p)\cdots\theta(f(x)^{p^{a-1}}).$$

We have the easily checked relation

$$F^{(a)}(x) = \prod_{i=0}^{a-1} F^{\tau^i}(x^{p^i}) \tag{4}$$

for any $x \in A$. Viewing elements in $\mathscr{H}^\dagger(A)$ as formal power series, the functions F and $F^{(a)}$ can be written formally as

$$F(X) = \exp(\pi(f(X) - f^\tau(X^p))),$$

$$F^{(a)}(X) = \exp(\pi(f(X) - f(X^q))). \tag{5}$$

However, the functions cannot be evaluated at a point $x \in A$ using these equations, since $\exp(\pi(X))$ does not have a large enough radius of convergence.

Let $\alpha := \psi_p \circ F$ and $\alpha_a := \psi_q \circ F^{(a)}$ be τ^{-1}-linear and τ^{-a}-linear maps, respectively, on the space $\mathscr{H}^\dagger(A)$. For example, for $H \in \mathscr{H}^\dagger(A)$, the function $\alpha(H)$ maps $x \in A$ to $\psi_p(FH)(x)$. From Eq. (4) and the second part of Lemma 2, one sees that $\alpha_a = \alpha^a$. We have the "chain level" formula [18, Eq. (6.3.11)]

$$L^*(\bar{f}, T) = \frac{\det(1 - T\alpha_a \mid \mathscr{H}^\dagger(A))}{\det(1 - Tq\alpha_a \mid \mathscr{H}^\dagger(A))}.$$

Here the determinant is in the sense of Monsky, viewing α_a as a "nuclear" operator on $\mathscr{H}^\dagger(A)$. (More simply one may just use a matrix for the map with respect to the "basis" used in (2).) This formula may be used to compute the zeta function efficiently, but it is not very fast. By introducing a "differential operator" one can derive a much more useful cohomological formula.Writing $\hat{F} := \exp(\pi(f(X)))$ we see that formally (viewing $\mathscr{H}^\dagger(A)$ as a power series ring)

$$\alpha_a = \hat{F}^{-1} \circ \psi_q \circ \hat{F}.$$

Observing that the operators $X\frac{d}{dX}$ and ψ_q commute up to a factor of q, this suggests defining

$$D := \hat{F}^{-1} \circ X\frac{d}{dX} \circ \hat{F} = X\frac{d}{dX} + \pi X\frac{df}{dX}. \tag{6}$$

We find that $\alpha_a \circ D = qD \circ \alpha_a$ and thus the map α_a defines a chain map on the complex:

$$\begin{array}{ccccccc} 0 & \to & \mathscr{H}^\dagger(A) & \stackrel{D}{\to} & \mathscr{H}^\dagger(A) & \to & 0 \\ & & \downarrow q\alpha_a & & \downarrow \alpha_a & & \\ 0 & \to & \mathscr{H}^\dagger(A) & \stackrel{D}{\to} & \mathscr{H}^\dagger(A) & \to & 0. \end{array}$$

The kernel of D on $\mathscr{H}^\dagger(A)$ is trivial, since $\exp(-\pi(f))$ does not define a function in $\mathscr{H}^\dagger(A)$. Thus moving to homology one deduces (see [18, p. 235])

$$L^*(\bar{f}, T) = \det(1 - T\alpha_a \mid \mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))). \tag{7}$$

Under the further assumption that zero is also a pole of $\bar{f}$, say $\bar{r}_1 = 0$, the space $\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))$ has basis (see [18, Eq. (8.1.1)])

$$\{X^{-n_1}, X^{-n_1+1}, \ldots, X^{-1}, 1, X, \ldots, X^{n_0-1}\} \cup \{1/(X - r_i)^j \mid 2 \leqslant i \leqslant d_g,\ 1 \leqslant j \leqslant n_i + 1\}.$$

Here as before r_i are the finite poles of f, and for each $0 \leqslant i \leqslant d_g$ we denote by n_i the order of the pole of f at r_i. (Recall that $r_0 = \infty$ and we assume p does not divide n_i.) Formula (7) along with the relation $\alpha_a = \alpha^a$ may be used to compute the L-function $L(\bar{f}, T)$ for any rational function $\bar{f} \in \mathbb{F}_q(T)$.

4. Simplifying assumptions

For the remainder of the paper we shall make the following assumptions. First, the poles of $\bar{f}$ include zero and infinity. Second, all the poles of $\bar{f}$ have order 1. Third, we assume $p \geqslant 5$. All of the pole order assumptions may be dropped, and the only one that significantly simplifies the algorithm is the uniformity of the finite pole orders. As mentioned before, we believe the algorithm still works when $p < 5$, although the arguments we give are not refined enough to prove this.

Recall that d_g denotes the number of finite poles. We may write

$$\bar{f} = \bar{h} + \frac{\bar{k}}{\bar{g}},$$

where $\bar{h}, \bar{k}, \bar{g} \in \mathbb{F}_q[X]$ with $\deg(\bar{h}) = 1$, $\gcd(\bar{k}, \bar{g}) = 1$ and $0 \leqslant \deg(\bar{k}) < d_g$, and $\deg(\bar{g}) = d_g$ with $\bar{g}(0) = 0$ and $\bar{g}$ squarefree. Let h, k, g be the following liftings of $\bar{h}, \bar{k}, \bar{g}$: let h and k be obtained by taking Teichmüller liftings of coefficients, and g by taking Teichmüller liftings of roots (thus $\bar{g} = \prod_{i=1}^{d_g} (X - \bar{r}_i)$ and $g = \prod_{i=1}^{d_g} (X - r_i)$ where $r_i = \mathrm{Teich}(\bar{r}_i)$). The assumptions on h and k are used to analyse the decay rate of coefficients of f, and that on g seems helpful when computing ψ_p.

The basis for $\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))$ can now be taken as

$$\left\{\frac{X}{g^2}, \frac{X}{g^2}, \ldots, \frac{X^{2d_g}}{g^2}\right\}.$$

Since zero is a pole of f, we have that $L(\bar{f}, T) = L^*(\bar{f}, T)$ and thus (7) gives an expression for the whole L-function.

Note that all computations take place with power series of the form

$$\sum_{j=0}^{\infty} H_j X^j + \sum_{j=-1}^{-\infty} H_j(X) g^j,$$

where $H_j \in K$ for $j \geqslant 0$, and for $j < 0$ each $H_j(X) \in K[X]$ has degree less than d_g. This is because each basis element has this form, as does F, the ring of all such power series is stable under the map ψ_p, and our reduction formulae "keep" series in this form.

5. The algorithm

We now present our point counting algorithm. The remainder of the paper will be taken up explaining how to perform the required subroutines, proving the correctness, and estimating the complexity.

Algorithm 3. *Input*: An equation $Z^p - Z = \bar{f}(X)$ over $\mathbb{F}_q$, where $\bar{f} \in \mathbb{F}_q(X)$ and $q = p^a$. We assume that $\bar{f}$ satisfies the conditions of Section 4 and $p \geqslant 5$.

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

Step 0: Let d_g be the number of finite poles of $\bar{f}$. Let $N_0 := \lfloor 2(p-1)d_g(1 + a/2) + 1 \rfloor$, $N_1 = N_0 + 6ad_g^2$ and $N_2 = 2N_1 + d_g$ (these choices are not optimal). We shall compute the coefficients of the numerator of the zeta function modulo p^{N_0}.

Step 1: Compute the power series F given in (5) with coefficients determined modulo p^{N_2} (see Section 8.2). Let α be the map on the ring $\mathscr{H}^\dagger(A)$ defined as $\alpha = \psi_p \circ F$.

Step 2: Let

$$\mathscr{E} = \left\{\frac{X}{g^2}, \frac{X^2}{g^2}, \ldots, \frac{X^{2d_g}}{g^2}\right\}$$

be the basis for the homology space $\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))$. For each basis element e, compute $\alpha(e)$ in $\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))$ as a K-linear combination of elements in $\mathscr{E}$ with coefficients determined modulo p^{N_1} (see Section 6). Construct M, defined as the matrix representing the map α acting on $\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))$ with respect to the basis, with coefficients determined modulo p^{N_1}.

Step 3: Compute

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

modulo p^{N_1}, where the map τ is the lifting of Frobenius to R and acts on the matrix coefficients. Thus M_a is a matrix for the map α_a on $\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))$. Compute $\det(1 - M_a T)$ with coefficients determined modulo p^{N_0}.

Step 4: Output the rational function

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

where θ_j is the cyclotomic field automorphism from [14, Eq. (11)] extended to act on $\mathbb{Z}_p[\pi][T]$ by fixing monomials.

Note that in Steps 2 and 3 the p-adic numbers might have small negative p-adic order. Determining their coefficients modulo p^{N_*} should just be understood in the obvious, though slightly non-mathematical, sense.

6. Cohomological reduction

In this section we present explicit reduction formulae which may be used to perform computations in the factor space $\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))$. The method is similar to Hermite reduction for the integration of rational functions [7, Section 22.2]. We shall also estimate the complexity of this process in terms of field operations, although of course all actual computations take place with "truncated" p-adic numbers.

At this stage we shall use the assumption that zero is a pole of $\bar{f}$, and so $g(0) = 0$, to simplify the reduction method. Under this assumption, the element X is invertible in $\mathscr{H}^\dagger(A)$ and so we have that

$$\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A)) \cong \mathscr{H}^\dagger(A)/X^{-1}D(\mathscr{H}^\dagger(A)).$$

For any $H \in \mathscr{H}^\dagger(A)$ we see that

$$H \bmod D = X(X^{-1}H \bmod X^{-1}D). \tag{8}$$

Thus it suffices to find a reduction method for the space $\mathscr{H}^\dagger(A)/X^{-1}D(\mathscr{H}^\dagger(A))$ where $X^{-1}D = (d/dX) - \pi(df/dX)$. (When $g(0) \neq 0$ one can derive slightly different reduction formulae working directly with D.)

Recall that

$$f = h + \frac{k}{g},$$

where $h, k \in R_0[X]$ with $\deg(h) = 1$, $\deg(g) = d_g \geqslant 1$, and $\deg(k) < d_g$ with $\gcd(k, g) = 1$ and g squarefree. Thus

$$X^{-1}D = \frac{d}{dX} + \pi\left(h' + \frac{k'}{g} - \frac{kg'}{g^2}\right).$$

We consider the action of $X^{-1}D$ on a rational function $X^n g^j$ for some $0 \leqslant n < d_g$ and $j \in \mathbb{Z}$. We see zero is equivalent modulo $X^{-1}D(\mathscr{H}^\dagger(A))$ to

$$X^{-1}D(X^n g^j) = nX^{n-1}g^j + jg^{j-1}g'X^n + \pi h'X^n g^j + \pi k'X^n g^{j-1} - \pi k g' g^{j-2}X^n. \tag{9}$$

Recall that $\gcd(k, g) = \gcd(g, g') = 1$. Given any polynomial $T(X) \in K[X]$ with $\deg(T) < d_g$ we may write $T = Ukg' + Vg$ for $U, V \in K[X]$ with $\deg(U) < \deg(g) + \deg(T) < 2d_g$ and $\deg(V) < \deg(kg') + \deg(T) < 3d_g - 2$. Replacing j by $-j$ and X^n by $U(X)$ in (9) we see that

$$X^{-1}D(U(X)/g^j) = \frac{U'}{g^j} - \frac{jUg'}{g^{j+1}} + \pi\frac{h'U}{g^j} + \pi\frac{k'U}{g^{j+1}} - \pi\frac{kg'U}{g^{j+2}}.$$

Hence

$$\begin{aligned}\frac{T}{g^{j+2}} &= \frac{Ukg'}{g^{j+2}} + \frac{Vg}{g^{j+2}} \equiv \frac{U'}{\pi g^j} + \frac{h'U}{g^j} - \frac{jUg'}{\pi g^{j+1}} + \frac{k'U + V}{g^{j+1}} \\ &= \frac{U' + \pi h'U}{\pi g^j} + \frac{-jUg' + \pi k'U + \pi V}{\pi g^{j+1}}.\end{aligned} \tag{10}$$

This equation may be used to reduce any rational function T/g^{j+2} for $j > 0$ where $T(X) \in R[X]$ with $\deg(T) < \deg(d_g)$ to the form S/g^2 where $S \in \pi^{-j}R[X]$ with $\deg(S) < 3d_g$. Each reduction step requires the computation of one gcd plus a fixed number of additions and multiplications of polynomials of degree $\mathcal{O}(d_g)$. Thus it can be done in $\tilde{\mathcal{O}}(d_g)$ operations in K. The total time to reduce T/g^{j+2} to this form is therefore $\tilde{\mathcal{O}}(d_g j)$ operations in K.

The term with highest degree, $jd_g + n$, in (9) is $\pi h'X^n g^j$. Thus all polynomials of degree $m \geqslant \deg(h') = 0$ can be reduced using the relation

$$h'X^n g^j \equiv -\pi^{-1}nX^{n-1}g^j - \pi^{-1}jg^{j-1}g'X^n - k'X^n g^{j-1} + kg'g^{j-2}X^n \tag{11}$$

where $0 \leqslant n < d_g$ and $j \geqslant 0$ are chosen so that $m = jd_g + n$. Specifically, the left-hand side in (11) has degree m. Given the g-adic expansion of a polynomial $T(X)$ of degree m, Eq. (11) allows one to compute the g-adic expansion of a homologous rational function $W(X)/g^2$ such that $\deg(W) - 2d_g < m$. This takes $\tilde{\mathcal{O}}(d_g)$ operations in the field K. Notice that if $T(X) \in R[X]$ then the coefficients of $W(X)/g^2$ lie in $\pi^{-1}R$. Repeating this we can reduce any polynomial $T(X) \in R[X]$ to a rational function $W(X)/g^2$ where $\deg(W) < 2d_g$ has coefficients in $\pi^{-\deg(T)}R$. The complexity of this reduction process is $\tilde{\mathcal{O}}(d_g \deg(T))$ operations in K.

Thus one may reduce all rational functions in $\mathscr{H}^\dagger(A)$ modulo $X^{-1}D$ to a linear combination of the set

$$\left\{\frac{1}{g^2}, \frac{X}{g^2}, \ldots, \frac{X^{2d_g-1}}{g^2}\right\}.$$

Hence by (8) one may reduce all rational functions in $\mathscr{H}^\dagger(A)$ modulo D to a linear combination of the basis set

$$\mathscr{E} := \left\{\frac{X}{g^2}, \frac{X^2}{g^2}, \ldots, \frac{X^{2d_g}}{g^2}\right\}.$$

To prove that $\mathscr{H}^\dagger(A)/D(\mathscr{H}^\dagger(A))$ is also isomorphic to the space spanned by $\mathscr{E}$ one must show that the denominators introduced in the above process have p-adic order bounded by a function with "sub-linear" growth (presumably the logarithm function). This is apparently shown in [18], as the result is stated on [18, Eq. (8.1.1)]. However, for the purposes of our algorithm we do not need to prove this. Rather we just give a linear bound on the growth of denominators which is enough to determine the p-adic accuracy "lost" during the reduction process. Note that the bound we give is actually enough to establish (for $p \geqslant 5$) an alternative version of the cohomological formula (7), in which the space $\mathscr{H}^\dagger(A)$ is replaced by a p-adic Banach space of functions which converge on the projective line over Ω with disks around the poles of f of radius Δ removed, for any $\Delta < 1/(p-1)$. Such a formula is more in the spirit of the original work of Dwork, involving p-adic Banach spaces rather than dagger spaces. (Indeed, in [14] we derive such a formula in the special case in which f has only a pole at infinity.)

7. Analysis of the reduction method

This section is devoted solely to a justification of the various accuracies to which one must work at different stages of the computation. The reader mainly interested in the practical details of the algorithm may move directly to Section 8. Note that we shall also use some of the estimates in this section in the determination of the complexity of the algorithm in Section 9.

7.1. Choice of final p-adic accuracy

The justification for the choice of N_0 is similar to that given in the paragraph following Lemma 24 in [14], and we do not comment further on this (see also [14, Note 26]). The necessity to initially work to different higher p-adic accuracies is more subtle, and is a consequence of a loss of accuracy and introduction of small denominators during the reduction process. We address the choice of N_1 and N_2 in Section 7.5.

7.2. The decay rate of coefficients

We first determine the decay rate of the coefficients of the power series F.

Definition 4. Let Δ, ε be positive real numbers or infinity, and δ a real number. Define the R-module $L(\Delta; \varepsilon, \delta)$ to be

$$\left\{ \sum_{r=0}^{\infty} H_r X^r + \sum_{j=-1}^{-\infty} H_j(X) g^j \mid H_r \in K, H_j(X) \in K[X], \deg(H_j(X)) < d_g, \right.$$
$$\left. \operatorname{ord}(H_r) \geqslant \Delta r, \operatorname{ord}(H_j(X)) \geqslant \varepsilon |j| + \delta \right\}.$$

Here the order of a polynomial is the minimum order of its coefficients.

This is a subspace of $\mathscr{H}^{\dagger}(A)$.

Lemma 5. *Write*

$$F = \sum_{r=0}^{\infty} F_r X^r + \sum_{j=-1}^{-\infty} F_j(X) g^j,$$

where $F_r \in R$ *and* $F_j(X) \in R[X]$ *with* $\deg(F_j) < \deg(g)$. *Then*

$$\operatorname{ord}(F_r) \geqslant \frac{(p-1)}{p^2} r, \quad \operatorname{ord}(F_j) \geqslant \frac{|j|}{2p}.$$

That is, $F \in L((p-1)/p^2; 1/2p, 0)$.

Proof. The set of all power series in the R-module $R[[X, Y]]/(gY - 1)$ which satisfy the decay estimates claimed for F is a complete ring. Thus it suffices to find a suitable factorization for F and to show that each factor satisfies the estimates.

Write $k(X) = \sum_{i=0}^{d_g - 1} k_i X^i$. We have

$$\begin{aligned} F &= \exp(\pi(f(X) - f^{\tau}(X^p))) \\ &= \exp(\pi(h - h^{\tau}(X^p))) \prod_{i=0}^{d_g - 1} \exp\{\pi((k_i X^i / g) - (k_i^{\tau} X^{pi} / g^{\tau}(X^p)))\}. \end{aligned} \tag{12}$$

Since the coefficients of h are Teichmüller points, the first factor in (12) is

$$\prod_{j=0}^{1} \theta(h_j X^j), \ \theta(T) := \exp(\pi(T - T^p)) \in R_1[[T]],$$

where $h = \sum_{j=0}^{1} h_j X^j$. By the decay rate of the coefficients of θ from (3) we have that the first factor is an expansion $\sum_{i=0}^{\infty} u_i X^i$ where $\operatorname{ord}(u_i) \geqslant (p-1)i/p^2$. We shall deal with each of the factors in the product in (12) separately.

Write the ith factor in the product in (12) as

$$\exp(\pi((k_i X^i / g)^p - (k_i^{\tau} X^{pi} / g^{\tau}(X^p))) \theta(k_i X^i / g). \tag{13}$$

Here the second factor is of the form $\sum_{j=1}^{\infty} v_j(X)/g^j$ where $v_j \in R[X]$ with $\deg(v_j) < d_g$ and $\operatorname{ord}(v_j) \geqslant (p-1)j/p^2 \geqslant j/2p$. To handle the first factor in (13)

observe that

$$\frac{1}{g^\tau(X^p)} = \frac{1}{g(X)^p}\left(1 - p\frac{r(X)}{g(X)^p}\right)^{-1}$$

where $g^\tau(X^p) = g(X)^p - pr(X)$ with $r(X) \in R_0[X]$ and $\deg(r) \leqslant pd_g$. Since the coefficients of k are Teichmüller points, we have $k_i^\tau = k_i^p$. The first factor in (13) is $\exp(\pi *)$ where $*$ is

$$\frac{k_i^p X^{pi}}{g^p} - \frac{k_i^p X^{pi}}{g^p}\left(1 - p\frac{r}{g^p}\right)^{-1} = -\sum_{j=1}^{\infty} p^j \frac{r^j k_i^p X^{pi}}{g^{p(j+1)}}. \tag{14}$$

This sum has the form $\sum_{j=1}^{\infty} t_j/g^j$ where $t_j \in R_0[X]$ with $\deg(t_j) < d_g$ and

$$\operatorname{ord}(t_j) \geqslant \max(\lceil j/p \rceil - 1, 1) \geqslant j/2p.$$

The first factor in (13) therefore has the form

$$\prod_{j=1}^{\infty} \exp(\pi(t_j/g^j)). \tag{15}$$

Since the power series $\exp(\pi(T)) \in R_1[T]$ the result now follows. □

(An easy modification of the above proof gives decay estimates in the case that all the finite poles of $\bar{f}$ have common order d, say, and the pole at infinity takes any order d_h, say. Specifically, one gets the lower bounds $(p-1)/p^2 d_h$ and $|j|/2pd$.)

7.3. The action of $\psi_p \circ F$ on the basis set

We now examine the decay rate of the coefficients of the power series obtained when $\psi_p \circ F$ acts on an element in the basis $\mathscr{E}$.

Lemma 6. *Let $L(\Delta; \varepsilon, \delta)$ be the R-module in Definition 4. The space $L(\Delta; \varepsilon, \delta)$ can also be viewed as module over $L(\Delta; \varepsilon, 0)$. In particular*

$$L(\Delta; \varepsilon, 0)L(\Delta; \varepsilon, \delta) \subseteq L(\Delta; \varepsilon, \delta).$$

Moreover, if $p\varepsilon < 1$ then

$$\psi_p(L(\Delta; \varepsilon, \delta)) \subseteq L(p\Delta; p\varepsilon, \delta - \varepsilon(p-1)^2).$$

Proof. The first part is straightforward. For the second, the factor p improvement of Δ follows from the explicit formulae for the action of ψ_p on monomials in Lemma 2. To understand the change in ε and δ, we consider the action of ψ_p on a element X^i/g^j for $j \geqslant 1$. Write $j = pk + \ell$ where $0 \leqslant \ell < p$, and $g^p = g^\tau(X^p) + pr(X)$ for some $r(X) \in A[X]$ with $\deg(r) \leqslant pd_g$. Using the property of ψ_p from the first part of Lemma

2 we find

$$\psi_p\left(\frac{X^i}{g^{pk+\ell}}\right) = \frac{1}{g^k}\sum_{m=0}^{\infty}\binom{-k}{m}p^m\frac{1}{g^m}\psi_p\left(\frac{X^i r^m}{g^\ell}\right).$$

(Note that $\binom{-k}{m} \in \mathbb{Z}_p$.) We have that

$$\psi_p\left(\frac{X^i r^m}{g^\ell}\right) = \frac{1}{g^\ell}\psi_p\left(X^i r^m\left\{\frac{g^\tau(X^p)}{g(X)}\right\}^\ell\right).$$

By the choice of the lifting of g, the operand of ψ_p is a polynomial of degree $\leqslant i + mpd_g + (p-1)d_g\ell$. Using again the explicit action of ψ_p on polynomials we find that

$$\psi_p\left(\frac{X^i}{g^{pk+\ell}}\right) = \frac{1}{g^{k+\ell}}\sum_{m=0}^{\infty}\binom{-k}{m}p^m\frac{R_m}{g^m} \tag{16}$$

where $R_m \in R[X]$ has degree

$$\leqslant d_g(m+\ell) + \frac{i - d_g\ell}{p}.$$

Now assume that $c \in K$ is such that $cX^i/g^{pk+\ell} \in L(\infty;\varepsilon,\delta)$; so $\mathrm{ord}(c) \geqslant \varepsilon(pk+\ell)+\delta$. We show its image lies in the appropriate space, and the result then follows by continuity. From (16) we see

$$\psi_p\left(\frac{cX^i}{g^{pk+\ell}}\right) = \frac{c^{\tau^{-1}}}{g^{k+\ell}}(G(X)+H(X))$$

where $G(X) \in R[X]$ has degree $\leqslant d_g\ell(1-(1/p)) + (i/p)$, and $H(X) \in L(\infty;1,0)$. Since $p\varepsilon < 1$ the space $L(p\Delta; p\varepsilon, \delta - \varepsilon(p-1)^2)$ is closed under multiplication by elements of $L(\infty;1,0)$. Now $c^{\tau^{-1}}/g^{k+\ell} \in L(p\Delta; p\varepsilon, \delta-\varepsilon(p-1)^2)$ since $\mathrm{ord}(c^{\tau^{-1}}) \geqslant p\varepsilon(k+\ell) + (\delta - \varepsilon\ell(p-1))$ and $\ell \leqslant (p-1)$. By the previous comment, multiplication by $H(X)$ does not remove it from this space, and certainly neither does that by $G(X)$. The proof is complete. □

Lemma 7. *For each $e \in \mathscr{E}$ we have $\psi_p(Fe) \in L(1/2;1/2,-p)$.*

Proof. By Lemma 5, $F \in L((p-1)/p^2; 1/2p, 0)$ and one may check that $e \in L(\infty; 1/2p, -1/p)$. Thus by Lemma 6, $Fe \in L((p-1)/p^2; 1/2p, -1/p)$ and $\psi_p(Fe) \in L((p-1)/p; 1/2, (-1/p) - (p-1)^2/2p) \subset L(1/2;1/2,-p)$. □

7.4. Loss of accuracy during reduction

During the reduction process powers of π are introduced in the denominators of rational functions. The next lemma quantifies this loss of accuracy.

Lemma 8. *Let $r \geqslant 0$ and $0 \leqslant i < d_g = \deg(g)$. Denote by $W_r(X)$ and $W_{r,i}(X)$ the unique expressions for X^r and X^i/g^r, respectively, as a K-linear combination of elements in the basis set $\mathscr{E}$. Then $W_r(X)$ becomes integral on multiplication by π^{r-1}, and $W_{r,i}(X)$ becomes integral on multiplication by π^{r+d_g-2}.*

Proof. The first result follows from the penultimate sentence of the paragraph after Eqs. (11) and (8). For the second, dividing by X we are reduced to considering either X^{i-1}/g^r or $(g/X)/g^{r+1}$. From the comment following Eq. (10) either of these may be reduced modulo $X^{-1}D$ to a rational function S/g^2 with $\deg(S) < 3d_g$ and S integral upon multiplication by π^{r-1}. Writing $S/g^2 = S_1 + (S_2/g^2)$ where $\deg(S_2) < 2d_g$, we need only now reduce the polynomial $S_1 \in \pi^{-r+1}R[X]$. The reduction of S_1 modulo $X^{-1}D$ will be integral upon multiplication by $\pi^{r-1+\deg(S_1)}$, which completes the proof since $\deg(S_1) < d_g$. □

7.5. Choice of intermediate accuracies

We now put together the estimates in Sections 7.3 and 7.4.

Lemma 9. *Assume that $p \geqslant 5$. For each basis element $e \in \mathscr{E}$, to determine the coefficients of $(\psi_p \circ F(e) \bmod D)$ modulo p^{N_1} it is enough to compute F modulo p^{N_2} where*

$$N_2 = 2N_1 + d_g\left(> \frac{(p-1)N_1 + 2p + d_g - 2}{(p-3)} \right).$$

Proof. The value for N_2 is obtained by finding a neat integral solution of

$$N_2 - \left(\frac{2(N_2 + p) + d_g - 2}{p-1} \right) \geqslant N_1.$$

Precisely, this choice of N_2 ensures that, given the decay rate of the coefficients of $\psi_p(Fe)$ from Lemma 7 and the denominators which are introduced in the reduction process in Lemma 8, computing $\psi_p(Fe) \bmod p^{N_2}$ is enough to recover the (possibly non-integral) coefficients in $\psi_p(Fe) \bmod D$ to the accuracy p^{N_1}. □

From Lemmas 7 and 8 one sees that each coefficient of $\psi_p(Fe) \bmod D$ has p-adic order at least

$$-\frac{(d_g - 2 + 2p)}{p-1}.$$

We know that the coefficients of $\det(I - M_a T)$ are integral. Each coefficient of this polynomial is obtained by summing certain products of exactly $2ad_g$ elements in the matrix M (and its conjugates under τ). Thus to compute the polynomial modulo p^{N_0} it is sufficient to find the entries in M modulo p^{N_1}

for any integral N_1 with

$$N_1 - \frac{2ad_g(d_g - 2 + 2p)}{p-1} > N_0.$$

For example, $N_1 = N_0 + 6ad_g^2$. This motivates the choice of N_1. Note that the authors do not believe that the increase from N_0 to N_1 in accuracy should be necessary, but it precisely this which alters the complexity dependence on d_g from fourth to sixth power.

8. Auxiliary routines

In this section we explain to how perform within the desired complexity estimates two essential but non-trivial tasks.

8.1. Computation of the map ψ_p

A method of computing the map ψ_p on rational functions is given in [12, Section 7.3]. Precisely, one should replace the notation "f" and "d" in this section by our "g" and "d_g". The method there computes the linear map $\tau(\psi_p)$ in soft-Oh linear time in the input size for any rational function with denominator a power of g. (See the final paragraph of Section 9 for why it is better to compute $\tau(\psi_p)$ rather than ψ_p.)

8.2. Expanding F

By the decay estimates in Lemma 5, the power series F reduces to a rational function $F \bmod p^{N_2}$ of "degree in g" at most $2pN_2 = \mathcal{O}(pN_2)$, with polynomial part of degree at most $p^2N_2/(p-1) = \mathcal{O}(pN_2)$. By "degree in g" we mean the power of g which occurs in the denominator when it is written as a quotient of two coprime polynomials. In this section we explain one method of computing the power series quickly. More precisely, in time $\tilde{\mathcal{O}}(d_g^2 pN_2)$ operations in the ring $R/(p^{N_2})$.

First note that $\theta(t) \bmod p^{N_2}$ is a polynomial of degree $\mathcal{O}(pN_2)$ which can be easily computed (see [13, Lemma 29]). Thus both the first factor in (12), and the second factor in (13) (there are d_g such factors), may be computed by first explicitly finding $\theta(t) \bmod p^{N_2}$ and making an appropriate substitution in each case. This requires $\tilde{\mathcal{O}}(d_g^2 pN_2)$ operations in $R/(p^{N_2})$. The difficultly lies in the first factor in (13), which must also be computed d_g times. We compute it via the product in (15). By the estimate $\mathrm{ord}(t_j/g^j) \geqslant j/2p$ we see that we need only compute t_j/g^j for $1 \leqslant j < 2pN_2$, and these may be found directly from (14) in $\tilde{\mathcal{O}}(d_g pN_2)$ operations in R_0/p^{N_2}. Since $\exp(\pi T) \in R_1[T]$ it follows that $\exp(\pi(t_j/g^j)) \bmod p^{N_2}$ is a rational function of "degree in g" $\lceil 2pN_2/j \rceil - 1$. It may be computed by direct substitution of $T = t_j/g^j$ in the first $\mathcal{O}(pN_2/j)$ terms of $\exp(\pi T)$. This takes $\tilde{\mathcal{O}}(d_g pN_2/j)$ operations in R/p^{N_2}. Thus the time taken to compute all factors $\exp(\pi(t_j/g^j))$

for $1 \leqslant j < 2pN_2$ is

$$\tilde{\mathcal{O}}\left(\sum_{j=1}^{2pN_2-1} 2d_g pN_2/j\right) = \tilde{\mathcal{O}}(d_g pN_2).$$

Multiplying these truncated power series together starting with the one of smallest degree, we can compute the product in time

$$\tilde{\mathcal{O}}\left(\sum_{j=1}^{2pN_2-1} (2d_g pN_2/j)\log(2d_g pN_2/j)\right) = \tilde{\mathcal{O}}(d_g pN_2).$$

Here we use the fact that for $i<j$ the product $\exp(\pi(t_j/g^j))\exp(\pi(t_i/g^i)) \bmod p^{N_2}$ has "degree in g" of $\mathcal{O}(pN_2/i)$. Thus the truncated power series F modulo p^{N_2} can be computed in $\tilde{\mathcal{O}}(d_g^2 pN_2)$ operations in R/p^{N_2}. Note that throughout this argument we have used soft-Oh linear algorithms for polynomial multiplication. Also, one needs to use a soft-Oh linear method for converting between different basis representations of polynomials to express all rational functions as linear combinations of the set $\{X^i/g^j\}_{0 \leqslant i < d_g,\, j \geqslant 0}$ (see [7, Section 9.2]).

9. Proof of theorem

The correctness of the algorithm follows from our justification of the choice of N_0, N_1 and N_2 in Section 7, and the results of Robba quoted in Section 3. It remains to estimate the complexity of the algorithm. From Section 8.2 the rational function $F \bmod p^{N_2}$ may be computed in $\tilde{\mathcal{O}}(d_g^2 pN_2)$ operations in the ring $R/(p^{N_2})$. By Section 8.1 computation of $\psi_p \circ F(e)$ for any basis element $e \in \mathscr{E}$ may be achieved in $\tilde{\mathcal{O}}(d_g pN_2)$ operations in $R/(p^{N_2})$. By Lemma 7, $\psi_p(Fe) \bmod p^{N_2}$ is a rational function with polynomial part of degree $\mathcal{O}(N_2)$ and "degree in g" $\mathcal{O}(N_2)$. By the complexity estimates in Section 6, it can be reduced in $\tilde{\mathcal{O}}(d_g N_2)$ operations in the field K "modulo p^{N_2}", with elements of p-adic order at least $-(d_g - 2 + 2p)/(p-1)$. That is, equivalently $\tilde{\mathcal{O}}(d_g N_2)$ operations in $R/(p^{N_2+*})$, where here $* = \lceil (d_g - 2 + 2p)/(p-1) \rceil$. Repeating this $2d_g$ times allows one to construct the matrix M, with coefficients determined modulo p^{N_1}. This matrix can therefore be found in $\tilde{\mathcal{O}}(N_2 d_g^2)$ operations in $R/(p^{N_2+*})$, with $*$ as immediately above. To compute M_a from M we can use the "fast exponentiation" method in [13, Lemma 31] in $\tilde{\mathcal{O}}(\log(a) d_g^3)$ operations in $R/(p^{N_1+*})$, where now $* = \lceil a(d_g - 2 + 2p)/(p-1) \rceil$. The extra exponent accounts for any growth in denominators. The algorithm from, for example, [4, Section 4] can be used to compute the characteristic polynomial of M_a in $\tilde{\mathcal{O}}(d_g^3)$ operations in $R/(p^{N_1+*})$, where this time $* = \lceil 2ad_g(d_g - 2 + 2p)/(p-1) \rceil$. Finally, the "norm" of this p-adic integral polynomial can be found in $\tilde{\mathcal{O}}(pd_g)$ further operations in $R/(p^{N_0})$ (see [14, Eq. (11)]).

By "operations in $R/(p^N)$" for some N in the above paragraph we mean arithmetic, which may be done in Soft-Oh linear time i.e. in $\tilde{\mathcal{O}}(paN)$ bit operations. Putting $N_0 = \mathcal{O}(pad_g)$, $N_1, N_2 = \mathcal{O}(pad_g^2)$ in the above paragraph gives a total complexity of

$$\tilde{\mathcal{O}}((d_g^2 pN_2)(paN_2)) = \tilde{\mathcal{O}}(p^4 a^3 d_g^6)$$

bit operations.

We have above ignored the contributions from the computation of the map τ^{-1}. By the method in [10, Section 5], this may be done on $R/(p^N)$ in $\tilde{\mathcal{O}}(a(paN))$ bit operations. In the computation of M_a from M, we require $\mathcal{O}(\log(a))$ applications of τ^{-1} to a matrix of side $\mathcal{O}(d_g)$ with entries in K "modulo" p^{N_1} and p-adic order at least $-2ad_g(d_g - 2 + 2p)/(p-1)$. One may check that this is absorbed in the final bit estimate above. One also needs to compute τ^{-1} when constructing $\psi_p \circ F(e) \bmod p^{N_2}$ for each basis element (in the application of the map ψ_p to Fe). Doing this directly would increase the complexity estimate above by a factor a. As such, we make the following minor modification to the algorithm: instead of reducing $\psi_p \circ F(e)$ modulo D, reduce $\tau(\psi_p \circ F)(e)$ modulo $\tau(D)$, and apply τ^{-1} to the answer. Reduction modulo $\tau(D)$ just involves replacing f by f^τ in the formulae in Section 5, and the linear map $\tau \circ \psi_p$ can be computed in soft-Oh linear time. (An alternative approach is to compute $F^{\tau^{-1}}$ rather than F by replacing f by $f^{\tau^{-1}}$ in the equations in the proof of Lemma 5 and using the method of Section 8.2. This requires soft-Oh linear time. Then $\psi_p \circ F(e) = \tau \circ \psi_p(F^{\tau^{-1}} e^{\tau^{-1}})$ and here the map $\tau \circ \psi_p$ can be computed in soft-Oh linear time.) Using either approach, the total bit complexity from the computation of the map τ^{-1} is absorbed in the estimate above. The proof of Theorem 1 is complete. (Note that in [14, Section 8.2] we made the incorrect claim that the method of [10, Section 5] gives a soft-Oh linear time algorithm for computing τ^{-1}. This minor error in the complexity estimate in [14, Section 8.2] can be rectified by applying either of the tricks above, i.e., reducing $\tau(\psi_p \circ F)(e)$ modulo $\tau(D)$ and applying τ^{-1} to the reduced form, or initial computing $F^{\tau^{-1}}$.)

References

[1] A. Adolphson, S. Sperber, Exponential sums on the complement of a hypersurface, Amer. J. Math. 102 (3) (1980) 461–487.

[2] P. Berthelot, Géométrie rigid et cohomologie des variétés algébrique de caractéristique p, Mém, Soc. Math. France (N.S.) 23 (1986) 7–32.

[3] I. Blake, G. Seroussi, N. Smart, Elliptic Curves in Cryptography, in: LMS Lecture Note Series, Vol. 265, Cambridge University Press, Cambrdige, 1999.

[4] J. Denef, F. Vercauteren, An extension of Kedlaya's algorithm to Artin-Schreier curves in characteristic 2, in: C. Fieker, D.R. Kohel (Eds.), Algorithmic Number Theory, Fifth International Symposium (ANTS-V), Lecture Notes in Computer Science, Vol. 2369, Springer, Berlin, 2002, pp. 308–323.

[5] B. Dwork, On the rationality of the zeta function of an algebraic variety, Amer. J. Math. 82 (1960) 631–648.
[6] B. Dwork, On the zeta function of a hypersurface, Pub. Math. IHES 12 (1962) 5–68.
[7] J. von zur Gathen, J. Gerhard, Modern Computer Algebra, Cambridge University Press, Cambridge, 1999.
[8] P. Gaudry, N. Gürel, An extension of Kedlaya's algorithm for counting points on superelliptic curves, in: C. Boyd (Ed.), Advances in Cryptology—ASIACRYPT 2001, Lecture Notes in Computer Science, Vol. 2248, Springer, Berlin, 2001, pp. 480–494.
[9] P. Gaudry, R. Harley, Counting points on hyperelliptic curves over finite fields, in: B. Preneel (Ed.), Advances in Cryptology—EUROCRYPT 2000, Lecture Notes in Computer Science, Vol. 1807, Springer, Berlin, 2000, pp. 19–34.
[10] K.S. Kedlaya, Counting points on hyperelliptic curves using Monsky–Washnitzer cohomology, J. Ramanujan Math. Soc. 16 (2001) 323–338.
[11] N. Koblitz, Hyperelliptic cryptosystems, J. Cryptol. 1 (1989) 139–150.
[12] A.G.B. Lauder, Computing zeta functions of Kummer curves via multiplicative characters, Found. Comput. Math. 3 (2003) 273–295.
[13] A.G.B. Lauder, D. Wan, Counting points on varieties over finite fields of small characteristic, in: J.P. Buhler, P. Stevenhagen (Eds.), Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography (Mathematical Sciences Research Institute Publications), Cambridge University Press, Cambridge, to appear, available at: http://web.comlab.ox.ac.uk/oucl/work/alan.lauder/
[14] A.G.B. Lauder, D. Wan, Computing zeta functions of Artin-Schreier curves over finite fields, London Math. Soc. JCM 5 (2002) 34–55.
[15] C.J. Moreno, Algebraic curves over finite fields, in: Cambridge Tracts in Mathematics, Vol. 97, Cambridge University Press, Cambridge, 1993.
[16] D. Reich, p-Adic function spaces and the theory of the zeta function, Ph.D. Thesis, Princeton University, 1966.
[17] D. Reich, A p-adic fixed point formula, Amer. J. Math. 91 (1969) 835–850.
[18] P. Robba, Index of p-adic differential operators III. Applications to twisted exponential sums, Astérique 119–120 (1984) 191–266.
[19] 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.
[20] R. Schoof, Elliptic curves over finite fields and the computation of square roots mod p, Math. Comput. 44 (170) (1985) 483–494.
[21] F. Vercauteren, Computing zeta functions of hyperelliptic curves over finite fields of characteristic 2, in: M. Yung (Ed.), Advances in Cryptology—CRYPTO 2002, Lecture Notes in Computer Science, Vol. 2442, Springer, Berlin, 2002, pp. 369–384.
[22] D. Wan, Algorithmic theory of zeta functions, in: J.P. Buhler, P. Stevenhagen (Eds.), Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography (Mathematical Sciences Research Institute Publications), Cambridge University Press, Cambridge, to appear, available at: http://www.math.uci.edu/~dwan/preprint.html