THE OPEN BOOK SERIES 2

ANTS XIII
Proceedings of the Thirteenth Algorithmic Number Theory Symposium

Computing zeta functions of cyclic covers in large characteristic

Vishal Arul, Alex J. Best, Edgar Costa, Richard Magner, and Nicholas Triantafillou

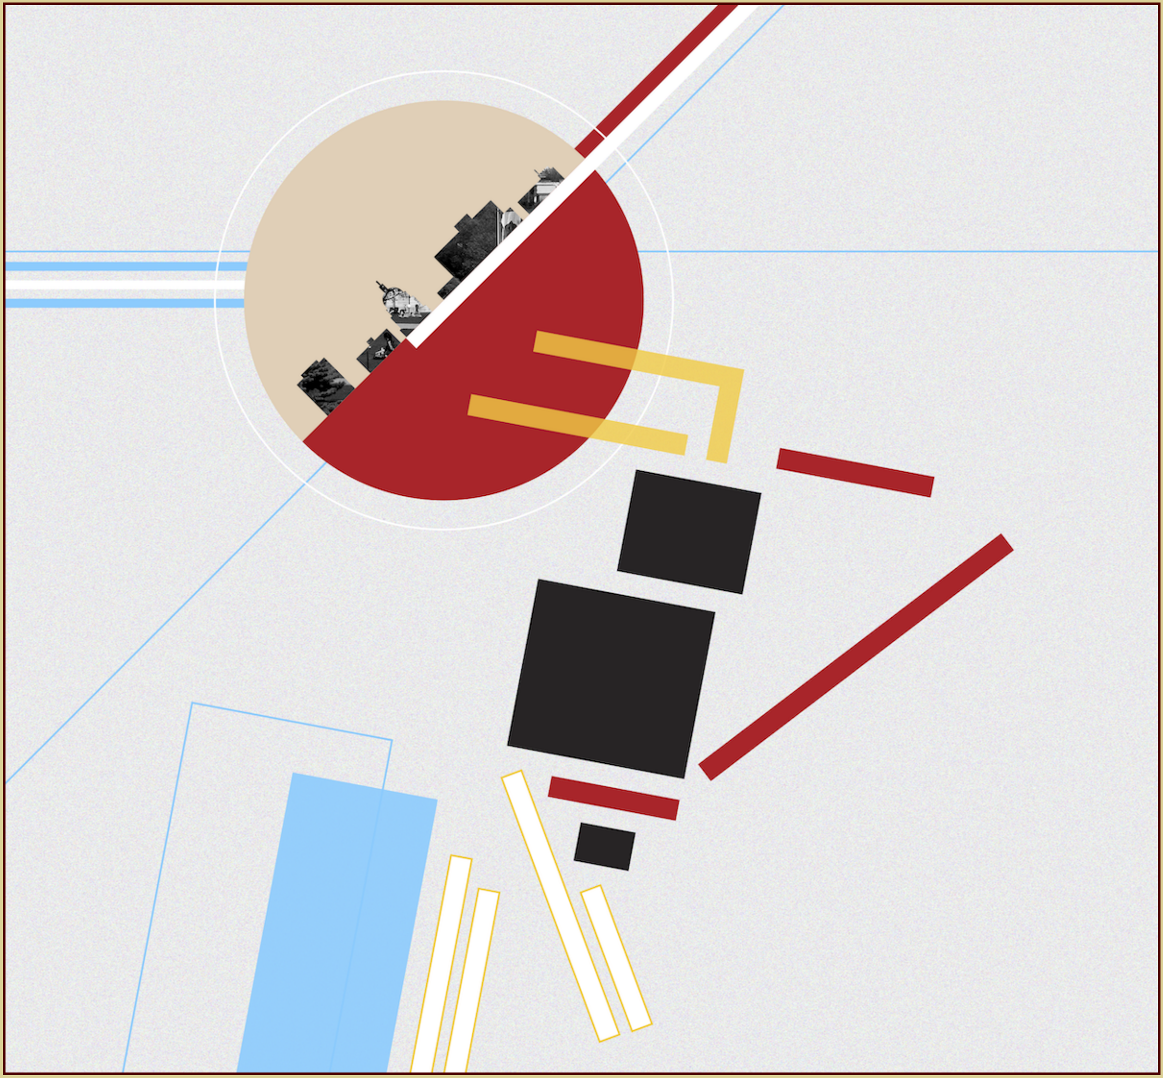

msp

THE OPEN BOOK SERIES **2** (2019)
Thirteenth Algorithmic Number Theory Symposium
dx.doi.org/10.2140/obs.2019.2.37

Computing zeta functions of cyclic covers in large characteristic

Vishal Arul, Alex J. Best, Edgar Costa, Richard Magner, and Nicholas Triantafillou

We describe an algorithm to compute the zeta function of a cyclic cover of the projective line over a finite field of characteristic p that runs in time $p^{1/2+o(1)}$. We confirm its practicality and effectiveness by reporting on the performance of our SageMath implementation on a range of examples. The algorithm relies on Gonçalves's generalization of Kedlaya's algorithm for cyclic covers, and Harvey's work on Kedlaya's algorithm for large characteristic.

1. Introduction

For $\mathcal{C}$ an algebraic curve of genus g over a finite field $\mathbb{F}_q$ of characteristic p and cardinality $q = p^n$, the zeta function of $\mathcal{C}$ is defined by

$$Z(\mathcal{C}, t) := \exp\left(\sum_{i=1}^{\infty} \#\mathcal{C}(\mathbb{F}_{q^i})\frac{t^i}{i}\right) = \frac{L(\mathcal{C}, t)}{(1-t)(1-qt)},$$

where $L(\mathcal{C}, t) \in 1 + t\mathbb{Z}[t]$ is a polynomial of degree $2g$, with reciprocal roots of complex absolute value $q^{1/2}$, and satisfies the functional equation $L(\mathcal{C}, t) = q^g t^{2g} L(\mathcal{C}, 1/(tq))$. In this paper, we address how to effectively compute $Z(\mathcal{C}, t)$ for a cyclic cover of $\mathbb{P}^1$ defined by $y^r = \bar{F}(x)$, where $\bar{F}(x)$ is squarefree and p is large in comparison to g, without any restrictions on r and $\deg \bar{F}$ sharing a common factor.

For curves of small genus, Schoof's method and its variants [Sch85; Pil90; GS04; GKS11; GS12] can compute $Z(\mathcal{C}, t)$ in time and space polynomial in $\log q$ and exponential in the genus. However, the practicality of these methods has only been shown for genus at most 2. These are known as ℓ-adic methods, as their efficiency derives from the realization of the ℓ-adic cohomology of the variety via torsion points.

The authors are grateful to the organizers of Sage Days 87, where this project began. We would also like to thank the reviewers for their many helpful comments. The second author was supported by the Simons Collaboration Grant #550023. The third author was partially supported by the Simons Collaboration Grant #550029. The fifth author was supported by the National Science Foundation Graduate Research Fellowship under Grant #1122374.

MSC2010: primary 11G20; secondary 11M38, 11Y16, 14G10.

Keywords: computational number theory, superelliptic, hyperelliptic, arithmetic geometry, Monsky–Washnitzer cohomology, zeta functions, p-adic.

Alternatively, Kedlaya [Ked01] showed that $Z(\mathcal{C}, t)$ can be determined in quasilinear time in p for an odd hyperelliptic curve, i.e., $r = 2$ and $\deg \bar{F} = 2g + 1$, by computing an approximation of the Frobenius matrix acting on p-adic cohomology (Monsky–Washnitzer cohomology). Kedlaya's algorithm and its variants are known as p-adic methods. In [Har07], Harvey improved the time dependence in p to $p^{1/2+o(1)}$. In [Har14], this improvement plays a major role in Harvey's algorithm for computing the p-local zeta functions of an odd hyperelliptic curve over $\mathbb{Z}$ for all p up to some bound. Kedlaya's original algorithm has been subsequently generalized several times, for example to superelliptic curves [GG01], $C_{a,b}$ curves [DV06], even-degree hyperelliptic curves [Har12], and nondegenerate curves [CDV06]. More recently, Gonçalves [Gon15] extended Kedlaya's algorithm to cyclic covers of $\mathbb{P}^1$ and Tuitman [Tui16; Tui17] to general covers. All these generalizations kept the quasilinear time dependence in p. Minzlaff [Min10] improved Gaudry and Gürel's algorithm for superelliptic curves by incorporating Harvey's work, giving a $p^{1/2+o(1)}$ time algorithm. The algorithms described above are efficient in practice, and have been integrated into the current versions of Magma [BCP97] and SageMath [Sag].

In this paper, we build upon the work of Gonçalves, Harvey, and Minzlaff to obtain a practical $p^{1/2+o(1)}$ algorithm for cyclic covers of $\mathbb{P}^1$. We are aware of the existence of theoretical algorithms with such a time dependence on p (and their average polynomial time versions) for arbitrary schemes (see [Har15]), but these have never been implemented, and it is unclear if they can be made to work in practice. Our algorithm improves the run-time with respect to the other parameters over these very general algorithms and provides a step towards a practical average polynomial time in higher genus, analogous to the progression from $p^{1/2+o(1)}$ to average polynomial time for odd hyperelliptic curves by Harvey.

More recently, Tuitman [Tui19] combined Harvey's ideas with a deformation approach to give a $p^{1/2+o(1)}$ algorithm for computing zeta functions of generic projective hypersurfaces of higher dimension. Tuitman's algorithm has a similar theoretical dependence on the degree of the curve and the degree of the field (over $\mathbb{F}_p$) as our algorithm.

Throughout, we use a bit complexity model for computation and the notation $\widetilde{O}(x) = \bigcup_k O(x \log^k(x))$. Our main result is then as follows:

Theorem 1.1. *Let $\mathcal{C}$ be a cyclic cover of $\mathbb{P}^1$, of genus g, defined by*

$$\mathcal{C} : y^r = \bar{F}(x),$$

where $\bar{F} \in \mathbb{F}_q[x]$ is a squarefree polynomial of degree d. Let $\widetilde{\mathcal{C}}$ be the curve obtained from $\mathcal{C}$ by removing the δ points at infinity and the d points on the x-axis corresponding to the zeros of $\bar{F}(x)$. Let M_ϵ be the matrix of Frobenius acting on B_ϵ, where B_ϵ is a basis of the Monsky–Washnitzer cohomology of $\widetilde{\mathcal{C}}$ defined in (2.6).

Let $N \geq 1$, and assume

$$p > d(N + \epsilon)r \quad \text{and} \quad r + d \geq 5. \tag{1.2}$$

Then the entries of M are in $\mathbb{Z}_q$ and we may compute M modulo p^N in time

$$\widetilde{O}(p^{1/2} N^{5/2} d^{\omega} rn + N^4 r d^4 n \log p + N n^2 \log p)$$

and space

$$O((p^{1/2}N^{3/2}+rN^2)d^2n\log p),$$

where ω is a real number such that the matrix arithmetic operations on matrices of size $m \times m$ take $\widetilde{O}(m^{\omega})$ ring operations.

With the goal of computing $Z(\mathcal{C},t)$ we may apply Theorem 1.1 with $N = O(nrd)$, for example as in (6.1), and this gives the following result:

Theorem 1.3. *In the same setup as Theorem 1.1, assume $p > dr\left(\frac{1}{2}gn+\log_p(g)+2\right)$. We can compute the numerator of the zeta function of $\mathcal{C}$ in time*

$$\widetilde{O}(p^{1/2}n^{7/2}r^{7/2}d^{5/2+\omega}+n^5r^5d^8\log p)$$

and space $O((p^{1/2}+n^{1/2}r^{3/2}d^{1/2})n^{5/2}r^{3/2}d^{7/2}\log p)$.

We also provide the following $O(\log p)$ space alternative to Theorem 1.1; see Remark 5.3 for more details.

Theorem 1.4. *In the same setup as Theorem 1.1, we may we may compute M modulo p^N in time $\widetilde{O}(prd^3N^3n+n^2N\log p)$ and space $O(rd^2Nn\log p)$.*

In contrast with with Minzlaff's work, in all the theorems above we do not put any restrictions on r and $\deg(\overline{F})$ sharing a common factor. Theorem 1.4 reduces the space complexity of [Gon15, Proposition 5.1] from quasilinear to logarithmic. Theorem 1.3 reduces both time and space complexity of [Gon15, Proposition 5.1] from quasilinear in p to $p^{1/2+o(1)}$. Moreover, we provide a SageMath implementation of our algorithm for computing zeta functions [ACMT16].

As with all adaptations of Kedlaya's algorithm, the heart of our algorithm is a procedure for computing a p-adic approximation to the action of Frobenius on a well-chosen basis for (a slight modification of) the Monsky–Washnitzer cohomology of $\mathcal{C}$. This is described in Lemma 3.1.

The remainder of the paper is organized as follows. In Section 2, we recall the relevant definitions for Monsky–Washnitzer cohomology. In Section 3, we compute a "sparse" formula for the action of Frobenius on the basis B_{ϵ}. The formula from Section 3 includes terms of large positive x-degree and large negative y-degree. Sections 4A and 4B show how to replace terms with cohomologous terms with x- and y-degree closer to zero by "horizontal" and "vertical" reductions. Section 5 collects the full algorithms, including complexity statements. We close by demonstrating the practicality of our implementation in Section 6.

2. Setup and notation

Let p be a prime and let $q = p^n$ for some $n \geq 1$. Let $\mathbb{F}_q$ and $\mathbb{F}_p$ be the finite fields with q elements and p elements. We write $\mathbb{Q}_q$ for the unramified extension of degree n of $\mathbb{Q}_p$, and $\mathbb{Z}_q$ for its ring of integers.

We will work under the assumption that (1.2) holds.

Let $\bar{F}(x) \in \mathbb{F}_q[x]$ be a polynomial of degree d with no multiple roots. To $\bar{F}(x)$ we can associate an r-cyclic cover of the projective line $\mathcal{C}$ defined by

$$\mathcal{C}\colon y^r = \bar{F}(x). \tag{2.1}$$

Write $\delta := \gcd(r, d)$. Then the genus of $\mathcal{C}$ is $g = \frac{1}{2}((d-1)(r-1) - (\delta - 1))$. The curve $\mathcal{C}$ is naturally equipped with an automorphism of order r defined by

$$\rho_r\colon (x, y) \longmapsto (x, \zeta_r y), \tag{2.2}$$

where ζ_r is a primitive r-th root of unity in a fixed algebraic closure of $\mathbb{F}_q$.

As in Kedlaya's original algorithm [Ked01] we pick an arbitrary lift $F(x) \in \mathbb{Z}_q[x]$ of $\bar{F}(x)$, also of degree d. Let $\widetilde{\mathcal{C}}$ be the curve obtained from $\mathcal{C}$ by removing the δ points at infinity and the d points on the x-axis corresponding to the zeros of $\bar{F}(x)$. Let $\bar{A} = \mathbb{F}_q[x, y, y^{-1}]/(y^r - \bar{F}(x))$ denote the coordinate ring of $\widetilde{\mathcal{C}}$, and write

$$A = \mathbb{Z}_q[x, y, y^{-1}]/(y^r - F(x)) \tag{2.3}$$

for the lift of $\bar{A}$ associated to $F(x)$. Let $A^\dagger$ be the weak completion of A, i.e.,

$$A^\dagger = \mathbb{Z}_q^\dagger[[x, y, y^{-1}]]/(y^r - F(x)), \tag{2.4}$$

where $\mathbb{Z}_q^\dagger[[x, y, y^{-1}]]$ is the ring of power series whose radius of convergence is greater than one. We lift the p-power Frobenius on $\mathbb{F}_q$ to $A^\dagger$ as follows. On $\mathbb{Z}_q$, we take the canonical Witt vector Frobenius and set $\sigma(x) := x^p$. We then extend σ to $A^\dagger$ by the formula

$$\sigma(y^{-j}) := y^{-jp} \sum_{k=0}^{+\infty} \binom{-j/r}{k} \big(\sigma(F(x)) - F(x)^p\big)^k y^{-kpr}. \tag{2.5}$$

The above series converges (because p divides $\sigma(F(x)) - F(x)^p$) and the definitions ensure that σ is a semilinear (with respect to the Witt vector Frobenius) endomorphism of $A^\dagger$. We extend it to differential forms by $\sigma(f \mathrm{d}g) := \sigma(f)\mathrm{d}(\sigma(g))$.

In the spirit of Kedlaya's algorithm, we determine the zeta function of $\mathcal{C}$ by computing the Frobenius action on the subspace of $H^1_{\mathrm{MW}}(\widetilde{\mathcal{C}})$ spanned by the set

$$B_\epsilon = \left\{ x^i \frac{\mathrm{d}x}{y^{j+\epsilon r}} : i \in \{0, \dots, d-2\}, j \in \{1, \dots, r-1\} \right\}, \qquad \text{where } \epsilon = \begin{cases} 0 & \text{if } \delta = 1, \\ 1 & \text{if } \delta > 1. \end{cases} \tag{2.6}$$

This subspace is Frobenius stable and 0 is the only element fixed by the induced automorphism ρ_r. When $\delta > 1$, using the basis B_1 allows us to avoid divisions by zero while reducing differentials (cf. Lemma 4.6). This is critical for generalizing Harvey's work to this setting.

If $\eta\colon \langle B_\epsilon \rangle \to H^1_{\mathrm{MW}}(\mathcal{C})$ is the projection map, then we have

$$\langle B_\epsilon \rangle = H^1_{\mathrm{MW}}(\mathcal{C}) \oplus \ker(\eta), \tag{2.7}$$

where $\ker(\eta)$ is a $\delta - 1$ dimensional vector space stable under Frobenius. Thanks to Gonçalves's work [Gon15, proof of Theorem 7.5], we have an explicit description for the characteristic polynomial $U(t) := \det(t \cdot \mathrm{id} - \mathrm{Frob}_q \mid \ker(\eta))$ of Frobenius acting on $\ker(\eta)$:

$$U(t) := \det(t \cdot \mathrm{id} - \mathrm{Frob}_q \mid \ker(\eta)) = \det(t \cdot \mathrm{id} - P) \cdot (t-1)^{-1}, \tag{2.8}$$

where the matrix P represents the permutation induced by q-th power Frobenius action on the roots of $T^\delta - f_d$, where f_d is the leading term of $\bar{F}(x)$. In the case that $\bar{F}(x)$ is monic the expression above simplifies to $U(t) = \prod_{i \mid \delta, i > 1} \left(t^{k_i} - 1\right)^{\varphi(i)/k_i}$, where k_i is the order of q in $\left(\mathbb{Z}/i\mathbb{Z}\right)^\times$. Thus our goal is to compute a p-adic approximation of the matrix M_ϵ representing σ with respect to B_ϵ.

3. The Frobenius action on differentials

We now rewrite the Frobenius expansion of a basis element in a sparse way where the number of terms does not depend on p. This is a generalization of [Har07, Proposition 4.1] and [Min10, Proposition 4.1], which is made possible due to the analysis performed by Gonçalves in [Gon15, §6].

Lemma 3.1. *Let $N > 0$ be a positive integer, $0 \le i \le d-2$ and $\epsilon r + 1 \le j \le (1+\epsilon)r - 1$. Suppose $p > d(N+\epsilon)r$ and $x^i y^{-j} \mathrm{d}x \in B_\epsilon$. For $0 \le \ell < N$, write*

$$D_{j,\ell} := \sum_{k=\ell}^{N-1} (-1)^{k-\ell} \binom{-j/r}{k} \binom{k}{\ell} \quad \text{and} \quad \mu_{j,\ell,b} := p D_{j,\ell} \sigma(F)_b^\ell, \tag{3.2}$$

where $\sigma(F)_b^\ell$ is the coefficient of x^{pb} in $\sigma(F(x))^\ell$. The differentials $\sigma(x^i y^{-j} \mathrm{d}x)$ and

$$T_{(i,j)} := x^{p(i+1)-1} y^{-jp} \sum_{\ell=0}^{N-1} \sum_{b=0}^{d\ell} \mu_{j,\ell,b} x^{pb} y^{-\ell pr} \mathrm{d}x \tag{3.3}$$

differ in cohomology by an element of $p^N \operatorname{span}_{\mathbb{Z}_q}(B_\epsilon)$.

Proof. From (2.5) we obtain

$$\sigma(x^i y^{-j} \mathrm{d}x) = \sum_{k=0}^{+\infty} p x^{p(i+1)-1} \binom{-j/r}{k} \left(\sigma(F(x)) - F(x)^p\right)^k y^{-p(j+kr)} \mathrm{d}x \tag{3.4}$$

Let U_k be the k-th summand of the above sum. We claim that for $k \ge N$ the reductions of U_k lie in $p^N \operatorname{span}_{\mathbb{Z}_q}(B_\epsilon)$.

To show this we start by rewriting U_k. Since p divides $\sigma(F(x)) - F(x)^p$, we have

$$U_k = p^{k+1} H(x) y^{-p(j+kr)} \mathrm{d}x, \tag{3.5}$$

where $H(x) \in \mathbb{Z}_q[x]$ is of degree at most $pi + p - 1 + dkp < pd(k+1)$. Define

$$L = \begin{cases} p(k+1) - 1 & \text{if } \epsilon = 0 \\ \left\lfloor \frac{p(j+kr)}{r} \right\rfloor - \epsilon & \text{if } \epsilon > 0. \end{cases} \tag{3.6}$$

Now we will expand $H(x)$ F-adically to L terms. Taking $j' \in [1, r]$ congruent to $pj \mod r$, and applying the relation $F(x) = y^r$, we have

$$U_k = p^{k+1}\left(G(x)y^{-\epsilon r - j'} + \sum_{\ell=0}^{L} G_\ell(x) y^{r\ell - p(j+kr)}\right)\mathrm{d}x, \tag{3.7}$$

where each $G_\ell(x) \in \mathbb{Z}_q[x]$ has degree at most $d-1$ and $G(x)$ has degree at most

$$pd(k+1) - 1 - dL \le \begin{cases} d-1 & \text{if } \epsilon = 0, \\ 0 & \text{if } \epsilon > 0. \end{cases} \tag{3.8}$$

Taking $\nu = \lfloor \log_p p(j+kr) - r\ell \rfloor \le 1 + \lfloor \log_p (k+1+\epsilon)r \rfloor$, Gonçalves [Gon15, Proposition 6.1] shows that the reduction of $p^\nu G_\ell(x) y^{r\ell - p(j+kr)}\mathrm{d}x$ lies in $\operatorname{span}_{\mathbb{Z}_q}(B_\epsilon)$.

Similarly, [Gon15, Proposition 6.2] says that taking

$$\mu = \lfloor \log_p((r(\deg(G)+1) - (\epsilon r + j')d)/\delta) \rfloor \le 1 + \lfloor \log_p(rd) \rfloor, \tag{3.9}$$

the reduction of $p^\mu G(x) y^{-\epsilon r - j'}\mathrm{d}x$ lies in $\operatorname{span}_{\mathbb{Z}_q}(B_\epsilon)$.

Since $p > d(N+\epsilon)r$, both $\mu = 1$ and $\nu \le 1 + k - N$, so the reductions of U_k for $k \ge N$ lie in $p^N \operatorname{span}_{\mathbb{Z}_q}(B_\epsilon)$.

The lemma follows by the rearranging the truncated series as follows:

$$\begin{aligned} \sum_{k=0}^{N-1} \binom{-j/r}{k} \left(\sigma(F(x)) - y^{pr}\right)^k y^{-kpr} &= \sum_{k=0}^{N-1}\sum_{\ell=0}^{k} (-1)^{k-\ell} \binom{-j/r}{k}\binom{k}{\ell} \sigma\left(F(x)\right)^\ell y^{pr(k-\ell)} y^{-prk} \\ &= \sum_{\ell=0}^{N-1}\sum_{b=0}^{d\ell} D_{j,\ell}\sigma(F)^\ell_b x^{pb} y^{-\ell pr}. \end{aligned}$$ □

4. Reducing differentials

The powers of x and y appearing in $T_{(i,j)}$ (as in Lemma 3.1) are much larger than those appearing in our choice of representatives for the basis B_ϵ. We use relations (coboundaries) coming from the differentials of functions on our curve to "reduce" the terms from $T_{(i,j)}$ to linear combinations of elements of B_ϵ. We proceed in two-stages. Horizontal reduction reduces the x-degree while leaving the y-pole order constant. Vertical reduction decreases the y-pole order without increasing the x-degree. Given a differential ω, we call the unique cohomologous differential $\omega' \in \operatorname{span}(B_\epsilon)$ the *reduction of* ω. We may also abuse notation and call intermediate products of the vertical/horizontal reduction process *reductions* of ω.

Organizing our work carefully, we can compute the reduction of ω modulo p^N by performing intermediate steps modulo p^{N+1}.

4A. *Horizontal reductions.*

We follow the steps of Harvey and Minzlaff. Decompose $F(x)$ as $F(x) = f_d x^d + P(x)$, where $P(x)$ has degree at most $d-1$.

Definition 4.1. For $s \in \mathbb{Z}_{\ge -1}$ and $t \in \mathbb{Z}_{\ge 0}$ define the vector space

$$W_{s,t} = \{G(x)x^s y^{-t}\mathrm{d}x : \deg G \le d-1\} \tag{4.2}$$

equipped with the standard monomial basis.

Let $M_H^t(s)\colon W_{s,t} \to W_{s-1,t}$ be the linear map given by the matrix

$$M_H^t(s) = \begin{pmatrix} 0 & 0 & \cdots & 0 & C_0^t(s) \\ D_H^t(s) & 0 & \cdots & 0 & C_1^t(s) \\ 0 & D_H^t(s) & \cdots & 0 & C_2^t(s) \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & D_H^t(s) & C_{d-1}^t(s) \end{pmatrix}, \tag{4.3}$$

where $D_H^t(s) = (d(t-r) - rs)f_d$ and where $C_h^t(s)$ is the coefficient of x^h in the polynomial $C^t(x,s) = rsP(x) - (t-r)xP'(x)$. Moreover, for $s_0 < s_1$ we write

$$\begin{aligned} D_H^t(s_0, s_1) &:= D_H^t(s_0+1)D_H^t(s_0+2)\cdots D_H^t(s_1); \\ M_H^t(s_0, s_1) &:= M_H^t(s_0+1)M_H^t(s_0+2)\cdots M_H^t(s_1). \end{aligned} \tag{4.4}$$

Lemma 4.5. *For $s \in \mathbb{Z}_{\geq 0}$, $t \in \mathbb{Z}_{\geq 0}$, and $\omega \in W_{s,t}$, we have $D_H^t(s)\omega \sim M_H^t(s)\omega$ in cohomology.*

Proof. See [Har07, Proposition 5.4] or [Min10, Proposition 5.1]. The same algebraic manipulations hold in the cyclic cover setting, as long we do not divide by $D_H^t(s)$, as this might be zero. □

In the case that d and r share a common factor, i.e., $\delta > 1$ and $\epsilon = 1$, then $D_H^t(s)$ might be identically zero. The next lemma ensures this cannot happen due to our choice of basis B_ϵ.

Lemma 4.6. *We have $D_H^t(s) \neq 0$, while applying horizontal reductions to $T_{(i,j)}$, for $0 \leq i \leq d-2$ and $1 + \epsilon r \leq j \leq (1+\epsilon)r - 1$.*

Proof. By inspecting the Frobenius formula (3.3) for a fixed value of ℓ, we see the pole order of y is $t = p(j + r\ell)$, where $1 + \epsilon r \leq j \leq (1+\epsilon)r - 1$, and the largest power of x is at most $p(d\ell + i + 1) - 1 \leq pd(\ell+1) - 1$. Since the largest power of x in $W_{s,t}$ is $s + d - 1$, we need only consider the case $s + d - 1 \leq pd(\ell+1) - 1$.

If $\delta = 1$, then $\epsilon = 0$ and $d(t-r) - rs \equiv djp \not\equiv 0 \bmod r$.

If $\delta > 1$, then $\epsilon = 1$, so $j \geq 1 + r$ and $t \geq p(1 + r(\ell+1))$. Using $s + d < pd(\ell+1)$,

$$d(t-r) - rs = dt - r(s+d) \geq dp(1 + r(\ell+1)) - r(pd(\ell+1)) = dp > 0. \tag{4.7}$$

This concludes the proof. □

Corollary 4.8. *In the same setting as Lemma 4.6, $D_H^t(s) \equiv 0 \bmod p$ if and only if $s \equiv -d \bmod p$.*

Proof. As in Lemma 4.6, the pole order of y is $t = p(j + r\ell)$, thus

$$D_H^t(s) := (d(t-r) - rs)f_d \equiv -r(d+s)f_d \bmod p. \tag{4.9}$$

By assumption, neither r nor f_d is divisible by p, so we only divide by p exactly when $s \equiv -d \bmod p$. □

Lemma 4.10. *Suppose $p > d(N+\epsilon)r$ and $s \equiv -1 \bmod p$. Then $D_H^t(s - (d-1))$ is divisible by p, but it is not divisible by p^2.*

Proof. As $s-(d-1)\equiv -d \bmod p$, we know this denominator is divisible by p. It equals

$$f_d(d(t-r)-r(s-(d-1))) = f_d(dt-rs-r).$$

Since f_d is coprime to p, we analyze the piece $dt-r(s+1)$. Inspecting the Frobenius formula (3.3) and considering that horizontal reduction decreases the exponent of x, we see

$$\begin{aligned} p-1 \le s \le p(i+1)-1+pd(N-1), &\qquad 0\le i\le d-2, \\ 0\le t\le jp+(N-1)pr, &\qquad \epsilon r+1\le j\le (1+\epsilon)r-1, \end{aligned} \tag{4.11}$$

where $\epsilon\in\{0,1\}$. From these inequalities we obtain

$$|dt-r(s+1)| \le \max\{dt, r(s+1)\} < dp(N+\epsilon)r < p^2, \tag{4.12}$$

thus the denominator has p-valuation exactly 1. □

Now we describe the horizontal reduction procedure in a fashion similar to that in [Har07, §7.2]. Following the notation of (3.3), let v_ℓ be a vector representing a differential form in $W_{p\ell-1,t}$ that is cohomologous to

$$\sum_{b\ge\ell}^{dk} \mu_{j,k,b-i-1} x^{pb-1} y^{-t}\mathrm{d}x, \quad \text{where } t=p(kr+j). \tag{4.13}$$

As in [Har07, §7.2], we say a vector is 1-*correct* if the first coordinate (corresponding to the highest power of x) is both 0 modulo p and correct modulo p^{N+1}, and the other coordinates are correct modulo p^N.

Given v_ℓ which is 1-correct, we show how to compute $v_{\ell-1}$ which is also 1-correct. First we get down to $W_{\ell p-d-1,t}$, by doing the first d reductions modulo p^{N+1}, as follows:

$$\begin{aligned} v_\ell^{(1)} &= v_\ell && \in W_{\ell p-1,t} \\ v_\ell^{(2)} &= D_H^t(\ell p-1)^{-1}M_H^t(\ell p-1)v_\ell^{(1)} && \in W_{\ell p-2,t} \\ &\ \vdots && \ \vdots \\ v_\ell^{(d+1)} &= D_H^t(\ell p-d)^{-1}M_H^t(\ell p-d)v_\ell^{(d)} && \in W_{\ell p-d-1,t}. \end{aligned} \tag{4.14}$$

Then we get down to $W_{(\ell-1)p,t}$ via

$$v'_\ell = D_H^t((\ell-1)p, \ell p-d-1)^{-1}M_H^t((\ell-1)p,\ell p-d-1)v_\ell^{(d+1)}, \tag{4.15}$$

and then finally

$$v_{\ell-1} = \mu_{j,\ell,(\ell-1)-i-1}x^{p(\ell-1)-1}y^{-t}\mathrm{d}x + D_H^t((\ell-1)p)^{-1}M_H^t((\ell-1)p)v'_\ell. \tag{4.16}$$

An analysis similar to [Har07, §7.2.2] shows that all coefficients of $M_H^t(\ell p-d)v_\ell^{(d)}$ are divisible by p and correct modulo p^{N+1}. Then, Lemma 4.10 implies that $v_\ell^{(d+1)}$ is correct modulo p^N. By Corollary 4.8, v'_ℓ is correct modulo p^N. Since the first row of $M_H^t((\ell-1)p)$ is zero modulo p, the vector $v_{\ell-1}$ is 1-correct.

We may also speed up the evaluation of $M^t_H((\ell-1)p, \ell p-d-1)$ and $D^t_H((\ell-1)p, \ell p-d-1)$ by p-adically interpolating the remaining values from the first N values. See [Har07, §7.2.1] and Section 5 for more details.

4B. ***Vertical reductions.*** Vertical reduction replaces differentials with cohomologous differentials with smaller pole order in y. While we performed horizontal reductions by working with d-dimensional vector spaces of differential forms, vertical reductions arise most naturally on $(d-1)$-dimensional vector spaces.

Definition 4.17. For $t \in \mathbb{Z}_{\geq 0}$ and $j \in \{1, \dots, r-1\}$, define the vector space

$$V_t^j := W_{-1,rt+j} \cap W_{0,rt+j}, \tag{4.18}$$

equipped with the standard monomial basis.

Vertical reduction operates via a series of maps $V_t^j \to V_{t-1}^j$ which are identity maps in cohomology. To define the maps, we need a lemma.

Lemma 4.19. *Let $A \in \mathbb{Z}_q[x]$ be a polynomial with* $\deg(A) < 2d-1$. *Then, there exist unique polynomials $R, S \in \mathbb{Z}_q[x]$ such that* $\deg(R) < d-1$, $\deg(S) < d$, *and* $A(x) = R(x)F(x) + S(x)F'(X)$.

Proof. Since F is separable and $\bar{F}$ is squarefree, we can find R_0 and S_0 such that $1 = R_0F + S_0F'$ by the Euclidean algorithm. Then $A = (AR_0)F + (AS_0)F'$. There is a unique S and T satisfying $AS_0 = TF + S$ and $\deg(S) < d$. Set $R = AR_0 - TF'$. Since $\deg(A) < 2d-1$ and $\deg(SF') < 2d-1$, it follows that $\deg(RF) < 2d-1$, so $\deg(R) < d-1$.

Uniqueness follows immediately, since the vector spaces of polynomials of degree less than $2d-1$ and of pairs of polynomials of degrees less than $d-1$ and less than d both have dimension $2d-1$. □

We may now define the vertical reduction maps.

Definition 4.20. For each $i \in \{0, \dots, d-2\}$, let R_i and S_i in $\mathbb{Z}_q[x]$ be the unique polynomials with $\deg(R_i) < d-1$ and $\deg(S_i) < d$, respectively, such that

$$x^i = R_i(x)F(x) + S_i(x)F'(x). \tag{4.21}$$

Write $(rt-r+j)R_i(x) + rS_i'(x) = \gamma_{i,0} + \gamma_{i,1}x + \dots + \gamma_{i,d-2}x^{d-2}$. Define $M_V^j(t)$ and $D_V^j(t)$ by

$$\begin{aligned} M_V^j(t) &:= \begin{pmatrix} \gamma_{0,0} & \gamma_{1,0} & \cdots & \gamma_{d-2,0} \\ \gamma_{0,1} & \gamma_{1,1} & \cdots & \gamma_{d-2,1} \\ \vdots & & \ddots & \vdots \\ \gamma_{0,d-2} & \gamma_{1,d-2} & \cdots & \gamma_{d-2,d-2} \end{pmatrix}, \\ D_V^j(t) &:= rt - r + j. \end{aligned} \tag{4.22}$$

Further define

$$\begin{aligned} M_V^j(t_1, t_2) &:= M_V^j(t_1+1) \cdot M_V^j(t_1+2) \cdots M_V^j(t_2), \\ D_V^j(t_1, t_2) &:= D_V^j(t_1+1) \cdot D_V^j(t_1+2) \cdots D_V^j(t_2). \end{aligned} \tag{4.23}$$

Lemma 4.24. *Consider $M_V^j(t)$ as a linear map from V_t^j to V_{t-1}^j with respect to their standard bases. Then, for any $\omega \in V_t^j$,*

$$D_V^j(t)\omega \sim M_V^j(t)\omega \tag{4.25}$$

in cohomology. More generally, considering $M_V^j(t_1, t_2)$ as a linear map from $V_{t_2}^j$ to $V_{t_1}^j$ with respect to their standard bases, for any $\omega \in V_{t_2}^j$,

$$D_V^j(t_1, t_2)\omega \sim M_V^j(t_1, t_2)\omega. \tag{4.26}$$

Proof. For any $S(x) \in \mathbb{Q}_q[x]$,

$$0 \sim \mathrm{d}\left(\frac{-r}{rt-r+j} S(x) y^{-(rt-r+j)}\right) = S(x)F'(x)y^{-(rt+j)}\mathrm{d}x + \frac{-r}{rt-r+j} S'(x)y^{-(rt-r+j)}\mathrm{d}x. \tag{4.27}$$

So, writing $x^i = R_i(x)F(x) + S_i(x)F'(x)$ as in (4.21), we have

$$\begin{aligned} x^i y^{-(rt+j)}\mathrm{d}x &= R_i(x)F(x)y^{-(rt+j)}\mathrm{d}x + S_i(x)F'(x)y^{-(rt+j)}\mathrm{d}x \\ &\qquad\qquad \sim R_i(x)y^{-(rt-r+j)}\mathrm{d}x + \frac{r}{rt-r+j} S_i'(x)y^{-(rt-r+j)}\mathrm{d}x \\ &= \frac{(r(t-1)+j)R_i(x) + rS_i'(x)}{r(t-1)+j} y^{-(r(t-1)+j)}\mathrm{d}x \\ &= (D_V^j(t_1, t_2))^{-1}(\gamma_{i,0} + \gamma_{i,1}x + \cdots + \gamma_{i,d-2}x^{d-2})y^{-(r(t-1)+j)}\mathrm{d}x. \end{aligned}$$

From this, (4.25) follows by linearity. Then (4.26) is immediate from (4.25). □

Remark 4.28. If we could work at infinite (or even very large) precision without it costing us computation time, this would be sufficient. However, in practice (and in theory), working with fewer extra bits results in significant time savings. Fortunately, we will see that when p is sufficiently large, the valuations of the coefficients of $D_V^j(t_1, t_2)^{-1}M_V^j(t_1, t_2)$ are never less than -1. As a result, given any element of V_t^j, we will be able to compute a cohomologous element of V_0^j while only losing a single digit of p-adic absolute precision.

Now, we follow Harvey's lead and study the coefficients of the matrices $M_V^j(t_1, t_2)$ and scalars $D_V^j(t_1, t_2)$. Lemma 4.29 will be our main technical tool.

Lemma 4.29. *Suppose $A \in \mathbb{Z}_q[x]$ and $B, G_{-t_2+1}, \ldots, G_{-t_1} \in \mathbb{Q}_q[x]$ satisfy*

$$A(x)y^{-rt_2-j}\mathrm{d}x = B(x)y^{-rt_1-j}\mathrm{d}x + \mathrm{d}\left(\sum_{t=-t_2+1}^{-t_1} G_t(x)y^{rt-j}\right). \tag{4.30}$$

Fix $C \in \mathbb{Z}_q$. If

$$\frac{C}{rt_1+j}, \frac{C}{r(t_1+1)+j}, \ldots, \frac{C}{r(t_2-1)+j} \in \mathbb{Z}_q$$

then $C \cdot B(x) \in \mathbb{Z}_q[x]$.

Remark 4.31. In our setting, $rt_1 + j \leq rt_2 + j < p^2$, so we may take $C = p$. Applying Lemma 4.29 with $A(x) = 1, x, \ldots, x^{d-1}$, the coefficients of $pD_V^j(t_1, t_2)^{-1}M_V^j(t_1, t_2)$ all belong to $\mathbb{Z}_q$.

We defer the proof of Lemma 4.29 to the end of the section, and collect the consequences needed for our main algorithm.

Lemma 4.32. *If $r(t-1) \equiv -j \mod p$, then $M_V^j(t)^{-1}$ is integral.*

The proof is identical to the proof of [Har07, Lemma 7.7] after replacing each occurrence of $2g$ with $d-1$. Indeed, the matrices are the same, up to multiplication by a unit.

Lemma 4.33. *If $rt_1 \equiv -j \mod p$, then $M_V^j(t_1, t_1+p)$ is zero modulo p.*

Proof. Here, the proof generalizes [Har07, Lemma 7.9]. By Lemma 4.29,

$$X := pD_V^j(t_1, t_1+p+1)^{-1} M_V^j(t_1, t_1+p+1) \tag{4.34}$$

has integral coefficients. By a computation similar to Lemma 4.10, $D_V^j(t_1, t_1+p+1) = p^2 \cdot u$ for some unit $u \in \mathbb{Z}_q^\times$, since the first and last terms contribute exactly one power of p and no other terms contribute. Then,

$$M_V^j(t_1, t_1+p) = p^{-1} D_V^j(t_1, t_1+p+1) X M_V^j(t_1+p+1)^{-1} = puXM_V^j(t_1+p+1)^{-1}.$$

Lemma 4.32 implies $M_V^j(t_1+p+1)^{-1}$ is integral, so $M_V^j(t_1, t_1+p) \equiv 0 \mod p$. □

Lemma 4.33 implies that the matrix $Y := D_V^j(t_1, t_1+p)^{-1} M_V^j(t_1, t_1+p)$ is integral when $rt_1 \equiv -j \mod p$. Hence the denominators of "vertically reductions" of differentials do not grow, at least if we reduce in appropriate batches of p steps.

Unfortunately, we may not start with t_1 satisfying $rt_1 \equiv -j \mod p$. Reducing to this case involves dividing by p at most once. To compensate, we must compute Y to one extra digit of p-adic precision.

Having collected our results, we now prove Lemma 4.29. Much like Kedlaya's proof of [Ked01, Lemma 2], we compare power series expansions of differentials in the uniformizer y near $(\theta_i, 0)$ for all roots θ_i of F. We give a full proof for clarity. The argument relies heavily on the following lemma:

Lemma 4.35. *Let $G \in \mathbb{Q}_q[x]$ be a polynomial which has $\deg(G) < d$. View G as an element of $\mathbb{Q}_q[x, y]/(y^r - F(x))$. Let $\theta_1, \ldots, \theta_d$ be the roots of F. Let $K_i \cong \mathbb{Q}_q((y))$ be the fraction field of the completion of the local ring at $(\theta_i, 0)$. The following are equivalent:*

(i) *G has integral coefficients as a polynomial.*

(ii) *G has integral coefficients as a power series in K_i for all i.*

(iii) *The coefficient of y^0 of G as a power series in K_i is integral for all i.*

Proof. It is trivial that (ii) implies (iii).

(iii) ⇒ (i) This follows immediately from the observation that the coefficient of y^0 of G as a power series in K_i is equal to $G(\theta_i)$. Since $\deg(G) < d$ and the roots of F are distinct mod p, the Lagrange interpolation formula shows that $G \in \mathbb{Z}_q[x]$.

(i) ⇒ (ii) This follows immediately from the fact that F has distinct roots mod p, so expanding x as a power series in y in K_i never requires division by a nonunit. □

With Lemma 4.35, the proof of Lemma 4.29 follows from the observation that the map d commutes with passage to the local ring.

Proof of Lemma 4.29. Note that $F'(\theta_i) \in \mathbb{Z}_q^\times$ for all roots θ_i of F, since $\bar{F}$ is separable. Then, as power series in y (near $(\theta_i, 0)$),

$$\begin{aligned} A(x) y^{r(-t_2)-j} \mathrm{d}x = r A(x) y^{r(-t_2+1)-j-1} F'(x)^{-1} \mathrm{d}y &= \sum_{t=-t_2+1}^{\infty} a_{i,t} y^{rt-j-1} \mathrm{d}y, \\ B(x) y^{r(-t_1)-j} \mathrm{d}x &= \sum_{t=-t_1+1}^{\infty} b_{i,t} y^{rt-j-1} \mathrm{d}y, \end{aligned}$$

where the $a_{i,t}$ are integral by Lemma 4.35, but we have no bounds (yet) on the $b_{i,t}$. Then,

$$\mathrm{d}\Bigg(\sum_{t=-t_2+1}^{-t_1} G_t(x) y^{rt-j} \Bigg) = \sum_{t=-t_2+1}^{-t_1} a_{i,t} y^{rt-j-1} \mathrm{d}y + \sum_{t=-t_1+1}^{\infty} (a_{i,t} - b_{i,t}) y^{rt-j-1} \mathrm{d}y.$$

Integrating term by term,

$$\sum_{t=-t_2+1}^{-t_1} G_t(x) y^{rt-j} = \sum_{t=-t_2+1}^{-t_1} \frac{a_{i,t}}{rt-j} y^{rt-j} + \sum_{t=-t_1+1}^{\infty} \frac{a_{i,t} - b_{i,t}}{rt-j} y^{rt-j}. \tag{4.36}$$

In particular, if C satisfies $C/(r \cdot t + r - j) \in \mathbb{Z}_q$, for all $t \in \{-t_2, \ldots, -t_1 - 1\}$, then the coefficients of $y^{r(-t_2+1)-j}, y^{r(-t_2+2)-j}, \ldots, y^{r(-t_1-1)-j}, y^{r(-t_1)-j}$ in all of the power series expansions at points $(\theta_i, 0)$ of $\sum_{t=-t_2+1}^{-t_1} C \cdot G_t(x) y^{rt-j}$ are integral.

In particular, $C \cdot G_{-t_2+1}$ satisfies (iii) of Lemma 4.35. Then the series expansions of $C \cdot G_{-t_2+1}(x)$ are all integral by condition (ii). Subtracting off $C \cdot G_{-t_2+1}$, we see $C \cdot G_{-t_2+2}$ satisfies (iii) of Lemma 4.35, hence condition (ii) and so on, so that all of the coefficients in all of the expansions of $\sum_{t=-t_2+1}^{-t_1} G_t(x) y^{rt-j}$ are integral. They remain integral upon differentiating.

Rearranging (4.30), the expansions of $C \cdot B(x) y^{-rt_1+j} \mathrm{d}x$ at each $(\theta_i, 0)$ as Laurent series in $\mathbb{Q}_q((y))\mathrm{d}y$ are integral. Replacing $\mathrm{d}y$ with $F'(x) y^{1-r}/r \mathrm{d}x$ preserves integrality. A final application of Lemma 4.35 shows that $C \cdot B(x)$ is integral. □

5. Main algorithm

We now combine the techniques of the previous sections to compute the matrix representing the p-th power Frobenius action with respect to $\langle B_\epsilon \rangle \subset H^1_{\mathrm{MW}}(\widetilde{C})$ modulo p^N. We summarize the procedure in Algorithm 1, where we take all intervals to be discrete, i.e., intersected with $\mathbb{Z}$.

We now analyze the time and space complexity of Algorithm 1. First, we recall that all our underlying ring operations are done in $\mathbb{Z}_q/p^N$ or $\mathbb{Z}_q/p^{N+1}$. Using bitstrings of length $O(Nn \log p)$ to represent elements of these rings, the basic ring operations (addition, multiplication, and inversion) have bit complexity $\widetilde{O}(Nn \log p)$, the matrix arithmetic operations on matrices of size $m \times m$ have bit complexity $\widetilde{O}(m^\omega Nn \log p)$, and polynomial multiplication of polynomials of degree m has bit complexity

Algorithm 1: computes the matrix representing the p-th power Frobenius action with respect to $\langle B_\epsilon \rangle \subset H^1_{\mathrm{MW}}(\widetilde{\mathcal{C}})$ modulo p^N

```
for $k \in [0, N-1],\ i \in [0, d-2],\ j \in [1+\epsilon r, (1+\epsilon)r - 1],\ \ell \in [0, dk+i+1]$ do
    $T_{(i,j),k,\ell} \leftarrow \mu_{j,k,\ell-i-1} x^{p\ell-1} y^{-p(kr+j)}$    // see Lemma 3.1
// Horizontal reductions:
for $k \in [0, N-1],\ j \in [1+\epsilon r, (1+\epsilon)r - 1]$ do
    $t \leftarrow p(kr+j)$
    $L \leftarrow \min(N-1, dk+d-2)$
    // Horizontal reductions modulo $p^N$, by linear recurrences:
    for $\ell \in [0, L]$ do
        $D(\ell), M(\ell) \leftarrow D_H(p\ell, p(\ell+1)-d-1), M_H(p\ell, p(\ell+1)-d-1)$
    // Deduce the remaining $M(\ell)$ modulo $p^N$, by interpolation:
    for $\ell \in [L+1, dk+d-2]$ do
        $D(\ell), M(\ell) \leftarrow D_H(p\ell, p(\ell+1)-d-1), M_H(p\ell, p(\ell+1)-d-1)$
    // Reduce $T_{(i,j),k}$ horizontally:
    for $i \in [0, d-2]$ do
        $v \leftarrow T_{(i,j),k,dk+i+1}$    // $v \in W_{p(dk+i+1)-1,t}$
        for $\ell = dk+i$ to $0$ do
            for $e \in [1, d]$ do    // $W_{p(\ell+1)-1,t} \to W_{p\ell-1,t}$
                $v \leftarrow D^t_H(p(\ell+1)-e)^{-1}(M^t_H(p(\ell+1)-e) \cdot v)$
            $v \leftarrow T_{(i,j),k,l} + (D^t_H(p\ell)^{-1} M^t_H(p\ell)) \cdot (D(\ell)^{-1} M(\ell)) \cdot v$
        $w_{(i,j),k} \leftarrow v$    // $w_{(i,j),k} \in W_{-1,t}$
// Vertical reductions:
for $j \in [1+\epsilon r, (1+\epsilon)r - 1]$ do
    // $p(kr+j) = r(pk+\alpha) + \beta = pr(k+\lambda) + r\gamma + r\epsilon + \beta$
    $\alpha, \beta \leftarrow \lfloor pj/r \rfloor, pj \bmod r$
    $\lambda, \gamma \leftarrow \lfloor (\alpha-\epsilon)/p \rfloor, (\alpha-\epsilon) \bmod r$
    $\delta \leftarrow \gamma + \epsilon$
    // Vertical reductions modulo $p^{N+1}$, by linear recurrences:
    $M(0) \leftarrow D^\beta_V(\epsilon, \delta)^{-1} M^\beta_V(\epsilon, \delta)$
    for $\ell \in [1, \lambda+N-1]$ do
        $M(\ell) \leftarrow D^\beta_V(\delta + p(\ell-1), \delta + p\ell)^{-1} M^\beta_V(\delta + p(\ell-1), \delta + p\ell)$
    for $i \in [0, d-2]$ do
        $v \leftarrow w_{(i,j),N-1+\lambda}$    // $v \in V^\beta_{p(N-1+\lambda)+\delta}$
        for $k = N-1+\lambda$ to $1$ do    // $V^\beta_{pk+\delta} \to V^\beta_{p(k-1)+\delta}$
            if $k \geq \lambda$ then
                $v \leftarrow w_{(i,j),k-\lambda} + M(k)v$
            else
                $v \leftarrow M(k)v$
        $w_{(i,j)} \leftarrow M(0) \cdot v$
return $w_{(i,j)},\ i \in [0, d-2],\ j \in [1+\epsilon r, (1+\epsilon)r - 1]$
```

$\widetilde{O}(mNn\log p)$. Applying Frobenius to such an element has complexity $\widetilde{O}(n\log^2 p + nN\log p)$ [Hub10, Corollary 3].

For p sufficiently large, the dominant steps are the horizontal and vertical reductions, i.e., lines 7 and 23 in Algorithm 1. In either case, we apply a modification of [BGS07, Theorem 15] to achieve the $p^{1/2+o(1)}$ time dependence.

Proposition 5.1 (linear recurrences method, [Har07, Theorem 6.1]). *Let $R = \mathbb{Z}_q/p^N$ or $\mathbb{Z}_q/p^{N+1}$, and $M(x) := M_0 + xM_1 \in R[x]^{m\times m}$. Let $0 \le \alpha_1 < \beta_1 \le \alpha_2 < \beta_2 \le \cdots \le \alpha_h < \beta_h \le K$ be integers. Assume $h < \sqrt{K} < p-1$ and write $M(\alpha, \beta) := M(\alpha+1)\cdots M(\beta)$. Then $M(\alpha_i, \beta_i)$ for $i = 1, \ldots, h$ can be computed using $\widetilde{O}(m^\omega\sqrt{K})$ ring operations in space $O(m^2\sqrt{K})$.*

For the horizontal reductions, we apply Proposition 5.1 once for each pair

$$(k, j) \in [0, N-1] \times [1+\epsilon r, (1+\epsilon)r - 1],$$

with $K = O(pN)$ and $m = O(d)$. For the vertical reductions, we apply Proposition 5.1 once for each j, again with $K = O(pN)$ and $m = O(d)$. This adds up to $\widetilde{O}(p^{1/2}N^{3/2}rd^\omega)$ ring operations in space $O(p^{1/2}N^{1/2}d^2)$.

Now we bound the time for the remaining steps. We will see that the number of ring operations for the remaining steps is independent of p, so that they contribute at most a $\log p$ term to the bit complexity.

To compute $\mu_{j,\ell,b}$ we start by replacing the coefficients of $F(x)$ by their images under σ. We then calculate all $\sigma(F)_b^\ell$ in $O(d^2N^2)$ ring operations. Evaluating all the binomial coefficients and finding the $D_{j,\ell}$ uses $O(rN^2)$ ring operations. In total, we compute all the $\mu_{j,\ell,b}$ in $O(rd^2N^2)$ ring operations plus $O(d)$ Frobenius substitutions.

We also use the p-adic interpolation method introduced by Harvey [Har07, §7.2.1] and attributed to Kedlaya. This allows us to reduce the number of matrix products that must be computed using the linear recurrence algorithm. The rest can then be obtained by solving a linear system involving a Vandermonde matrix. In our setting, an analogous complexity analysis holds, and the total number of ring operations required is $O(rd^3N^3)$, where the extra r factor is due to the j loop.

The matrix $M_H^t(s)$ is sparse; for each t, it requires $O(d)$ ring operations to compute. We need to do this $O(rN)$ times, thus the total is $O(rdN)$.

During the horizontal reduction, for each ℓ we do $O(d)$ sparse vector-matrix multiplications and one dense vector-matrix multiplication. This requires $O(d^2)$ ring operations per ℓ. Hence, lines 10–16 add up to $O(rd^4N^2)$ ring operations. The number of vector-matrix multiplications during the vertical reduction is $O(dN)$, thus negligible in comparison with the horizontal phase.

Computing all the R_i and S_i requires $O(d^3)$ total ring operations. Then for each $j \in [r\epsilon+1, (1+\epsilon)r-1]$, the matrix $M_V^j(t)$ can be computed in $O(d^2)$ ring operations. The total number of ring operations for these steps is $O(rd^2 + d^3)$.

The total number of operations is $O(p^{1/2}N^{3/2}rd^\omega + rd^4N^3)$ plus $O(d)$ Frobenius substitutions. Converting this to bit complexity, our algorithm runs in time

$$\widetilde{O}(p^{1/2}N^{5/2}rd^\omega n + N^4rd^4n\log p + Ndn^2\log p). \tag{5.2}$$

In addition to the space required by Proposition 5.1, we use $O(rd^2N)$ space for the interpolation, to store $w_{(i,j),k}$ and to do the vector-matrix multiplications. This adds up to $O((p^{1/2}N^{3/2} + rN^2)d^2 n \log p)$ space, and Theorem 1.1 follows.

Remark 5.3. Under certain conditions, the time-space tradeoff provided by Proposition 5.1 might not be ideal or possible. In those cases, one can instead do the reductions one step at a time with naive vector-matrix multiplications. The horizontal phase amounts to $O(prd^2N^2)$ sparse matrix-vector multiplications of size $O(d)$ in space $O(rd^2N)$. The vertical phase amounts to $O(prdN)$ dense matrix-vector multiplications of the same size, and no extra space is required. With the single exception of the $O(d)$ Frobenius substitutions, all the other steps are negligible in comparison. In terms of bit complexity, this amounts to $\widetilde{O}(prd^3N^3n + n^2N \log p)$ time and $O(rd^2Nn \log p)$ space, and Theorem 1.4 follows.

6. Sample computations

We have implemented both versions of our method using SageMath. However, the $p^{1/2+o(1)}$ version, i.e., Theorem 1.3 and Algorithm 1, is only implemented for the case $n = 1$, as we rely on Harvey's implementation of Proposition 5.1 in C++. Our implementation is on track to be integrated in one of the upcoming versions SageMath [ACMT16]. An example session:

```
sage: x = PolynomialRing(GF(10007),"x").gen();
sage: CyclicCover(5, x^5 + 1).frobenius_polynomial()
x^12 + 300420147*x^8 + 30084088241167203*x^4 + 1004207356863602508537649
```

Our examples were computed on one core of a desktop machine with an `Intel Core i5-4590 3.30 GHz processor`. In all the examples, we took

$$N = \max\left\{\left\lceil \log_p(4g/i) + ni/2 \right\rceil : i = 1, \ldots, g\right\}, \tag{6.1}$$

and thus by employing Newton identities we can pinpoint the numerator of $Z(\mathcal{C}, t)$; see, for example, [Ked13, slide 8]. In practice, we may even work with lower N, and then hopefully verify that there is only one possible lift that satisfies the Riemann hypothesis and the functional equation in the Weil conjectures; see [Ked08].

In Table 1 we present the running times for computing $Z(\mathcal{C}, t)$ for three examples where $(g, d, r) =$ (6, 5, 5), (25, 6, 12), and (45, 11, 11), over a range of p values. This sample of running times confirms the practicality and effectiveness of our method for a wide range of p and tuples (d, r). We are not aware of any other alternative method that can handle p and g in these ranges.

Our implementation is also favorable when compared with Minzlaff's implementation (in Magma 2.24-1), which deals only with superelliptic curves, rather than arbitrary cyclic covers. For example, consider the superelliptic curve,

$$C\colon y^7 = x^3 + 4x^2 + 3x - 1.$$

p	time	p	time	p	time
$2^{14}-3$	1.21s	$2^{22}-3$	21.7s	$2^{30}-35$	5m 58s
$2^{16}-15$	3.05s	$2^{24}-3$	40.9s	$2^{32}-5$	11m 36s
$2^{18}-5$	5.74s	$2^{26}-5$	1m 23s	$2^{34}-41$	32m 59s
$2^{20}-3$	10.9s	$2^{28}-57$	2m 54s	$2^{36}-5$	1h 7m

Genus 6 curve $\mathcal{C}\colon\ y^5 = x^5 - x^4 + x^3 - 2x^2 + 2x + 1$ with $N = 4$

p	time	p	time	p	time
$2^{10}+45$	4m 37s	$2^{18}-5$	12m 2s	$2^{26}-5$	2h 38m
$2^{12}-3$	5m 31s	$2^{20}-3$	21m 34s	$2^{28}-57$	5h 24m
$2^{14}-3$	6m 20s	$2^{22}-3$	37m 21s	$2^{30}-35$	12h 12m
$2^{16}-15$	8m 15s	$2^{24}-3$	1h 13m	$2^{32}-5$	23h 35m

Genus 25 curve $\mathcal{C}\colon\ y^6 = x^{12} + 10x^{11} + x^{10} + 2x^9 - x^7 - x^5 - 4x^4 + 31x$ with $N = 13$

p	time	p	time	p	time
$2^{12}-3$	24m 1s	$2^{18}-5$	1h 2m	$2^{24}-3$	7h 21m
$2^{14}-3$	29m 50s	$2^{20}-3$	1h 52m	$2^{26}-5$	16h 24m
$2^{16}-15$	37m 14s	$2^{22}-3$	3h 22m	$2^{28}-57$	33h 17m

Genus 45 $\mathcal{C}\colon\ y^{11} = x^{11} + 21x^9 + 22x^8 + 12x^7 + 5x^4 + 15x^3 + 6x^2 + 99x + 11$ with $N = 23$

Table 1. Running times for three curves, for various p. Each row represents a (roughly) four-fold increase in p and a doubling in running time, compared to the row preceding it, confirming that our implementation has a $p^{1/2+o(1)}$ running time.

If we wish to compute all L polynomials of C for $p < 2^{24}$ using our implementation we estimate that this will take about 6 months on one core (on the same desktop mentioned above), whereas with Minzlaff's it would take around 3 years. The curve C has some interesting properties and it arose recently in some in progress work of D. Roberts, F. Rodriguez-Villegas, and J. Voight.

References

[ACMT16] Vishal Arul, Edgar Costa, Richard Magner, and Nicholas Triantafillou, *Hasse–Weil zeta function of a cyclic cover of $\mathbb{P}^1$ over finite fields*, 2016, https://trac.sagemath.org/ticket/20264.

[BCP97] Wieb Bosma, John Cannon, and Catherine Playoust, *The Magma algebra system, I: The user language*, J. Symbolic Comput. **24** (1997), no. 3-4, 235–265. MR 1484478

[BGS07] Alin Bostan, Pierrick Gaudry, and Éric Schost, *Linear recurrences with polynomial coefficients and application to integer factorization and Cartier–Manin operator*, SIAM J. Comput. **36** (2007), no. 6, 1777–1806. MR 2299425

[CDV06] W. Castryck, J. Denef, and F. Vercauteren, *Computing zeta functions of nondegenerate curves*, Int. Math. Res. Pap. **2006** (2006), art. id. 72017. MR 2268492

[DV06] Jan Denef and Frederik Vercauteren, *Counting points on C_{ab} curves using Monsky–Washnitzer cohomology*, Finite Fields Appl. **12** (2006), no. 1, 78–102. MR 2190188

[GG01] Pierrick Gaudry and Nicolas Gürel, *An extension of Kedlaya's point-counting algorithm to superelliptic curves*, Advances in cryptology—ASIACRYPT 2001, Lecture Notes in Comput. Sci., no. 2248, Springer, 2001, pp. 480–494. MR 1934859

[GKS11] Pierrick Gaudry, David Kohel, and Benjamin Smith, *Counting points on genus 2 curves with real multiplication*, Advances in cryptology—ASIACRYPT 2011, Lecture Notes in Comput. Sci., no. 7073, Springer, 2011, pp. 504–519. MR 2935020

[Gon15] Cécile Gonçalves, *A point counting algorithm for cyclic covers of the projective line*, Algorithmic arithmetic, geometry, and coding theory, Contemp. Math., no. 637, Amer. Math. Soc., Providence, RI, 2015, pp. 145–172. MR 3364447

[GS04] Pierrick Gaudry and Éric Schost, *Construction of secure random curves of genus 2 over prime fields*, Advances in cryptology—EUROCRYPT 2004, Lecture Notes in Comput. Sci., no. 3027, Springer, 2004, pp. 239–256. MR 2153176

[GS12] ———, *Genus 2 point counting over prime fields*, J. Symbolic Comput. **47** (2012), no. 4, 368–400. MR 2890878

[Har07] David Harvey, *Kedlaya's algorithm in larger characteristic*, Int. Math. Res. Not. **2007** (2007), no. 22, art. id. rnm095. MR 2376210

[Har12] Michael C. Harrison, *An extension of Kedlaya's algorithm for hyperelliptic curves*, J. Symbolic Comput. **47** (2012), no. 1, 89–101. MR 2854849

[Har14] David Harvey, *Counting points on hyperelliptic curves in average polynomial time*, Ann. of Math. (2) **179** (2014), no. 2, 783–803. MR 3152945

[Har15] ———, *Computing zeta functions of arithmetic schemes*, Proc. Lond. Math. Soc. (3) **111** (2015), no. 6, 1379–1401. MR 3447797

[Hub10] Hendrik Hubrechts, *Fast arithmetic in unramified p-adic fields*, Finite Fields Appl. **16** (2010), no. 3, 155–162. MR 2610706

[Ked01] Kiran S. Kedlaya, *Counting points on hyperelliptic curves using Monsky–Washnitzer cohomology*, J. Ramanujan Math. Soc. **16** (2001), no. 4, 323–338. MR 1877805 arXiv math/0105031

[Ked08] ———, *Search techniques for root-unitary polynomials*, Computational arithmetic geometry, Contemp. Math., no. 463, Amer. Math. Soc., Providence, RI, 2008, pp. 71–81. MR 2459990

[Ked13] ———, *Computing zeta functions of nondegenerate toric hypersurfaces via controlled reduction*, lecture slides, Univ. Oxford, 2013.

[Min10] Moritz Minzlaff, *Computing zeta functions of superelliptic curves in larger characteristic*, Math. Comput. Sci. **3** (2010), no. 2, 209–224. MR 2608297

[Pil90] J. Pila, *Frobenius maps of abelian varieties and finding roots of unity in finite fields*, Math. Comp. **55** (1990), no. 192, 745–763. MR 1035941

[Sag] The Sage developers, *Sagemath, the Sage mathematics software system*, http://www.sagemath.org.

[Sch85] René Schoof, *Elliptic curves over finite fields and the computation of square roots mod p*, Math. Comp. **44** (1985), no. 170, 483–494. MR 777280

[Tui16] Jan Tuitman, *Counting points on curves using a map to $\mathbb{P}^1$*, Math. Comp. **85** (2016), no. 298, 961–981. MR 3434890

[Tui17] ———, *Counting points on curves using a map to $\mathbb{P}^1$, II*, Finite Fields Appl. **45** (2017), 301–322. MR 3631366

[Tui19] ———, *Computing zeta functions of generic projective hypersurfaces in larger characteristic*, Math. Comp. **88** (2019), no. 315, 439–451. MR 3854065

Received 2 Mar 2018. Revised 17 Sep 2018.

VISHAL ARUL: varul@mit.edu
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, United States

ALEX J. BEST: alexjbest@gmail.com
Department of Mathematics, Boston University, Boston, MA, United States

EDGAR COSTA: edgarc@mit.edu
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, United States

RICHARD MAGNER: rmagner@bu.edu
Department of Mathematics, Boston University, Boston, MA, United States

NICHOLAS TRIANTAFILLOU: ngtriant@mit.edu
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, United States

VOLUME EDITORS

Renate Scheidler
University of Calgary
Calgary, AB T2N 1N4
Canada

Jonathan Sorenson
Butler University
Indianapolis, IN 46208
United States

The cover image is based on a design by Linh Chi Bui.

The contents of this work are copyrighted by MSP or the respective authors. All rights reserved.

Electronic copies can be obtained free of charge from http://msp.org/obs/2 and printed copies can be ordered from MSP (contact@msp.org).

The Open Book Series is a trademark of Mathematical Sciences Publishers.

ISSN: 2329-9061 (print), 2329-907X (electronic)

ISBN: 978-1-935107-02-6 (print), 978-1-935107-03-3 (electronic)

First published 2019.

msp

MATHEMATICAL SCIENCES PUBLISHERS
798 Evans Hall #3840, c/o University of California, Berkeley CA 94720-3840
contact@msp.org http://msp.org

THE OPEN BOOK SERIES 2

Thirteenth Algorithmic Number Theory Symposium

The Algorithmic Number Theory Symposium (ANTS), held biennially since 1994, is the premier international forum for research in computational number theory. ANTS is devoted to algorithmic aspects of number theory, including elementary, algebraic, and analytic number theory, the geometry of numbers, arithmetic algebraic geometry, the theory of finite fields, and cryptography.

This volume is the proceedings of the thirteenth ANTS meeting, held July 16-20, 2018, at the University of Wisconsin-Madison. It includes revised and edited versions of 28 refereed papers presented at the conference.

Edited by Renate Scheidler and Jonathan Sorenson

CONTRIBUTORS

Simon Abelard
Sonny Arora
Vishal Arul
Angelica Babei
Jens-Dietrich Bauch
Alex J. Best
Jean-François Biasse
Alin Bostan
Reinier Bröker
Nils Bruin
Xavier Caruso
Stephanie Chan
Qi Cheng
Gilles Christol
Owen Colman
Edgar Costa
Philippe Dumas
Kirsten Eisenträger
Claus Fieker
Shuhong Gao
Pierrick Gaudry
Alexandre Gélin
Alexandru Ghitza
Laurent Grémy
Jeroen Hanselman
David Harvey
Tommy Hofmann
Everett W. Howe
David Hubbard
Kiran S. Kedlaya
Thorsten Kleinjung
David Kohel
Wanlin Li
Richard Magner
Anna Medvedovsky
Michael Musty
Ha Thanh Nguyen Tran
Christophe Ritzenthaler
David Roe
J. Maurice Rojas
Nathan C. Ryan
Renate Scheidler
Sam Schiavone
Andrew Shallue
Jeroen Sijsling
Carlo Sircana
Jonathan Sorenson
Pierre-Jean Spaenlehauer
Andrew V. Sutherland
Nicholas Triantafillou
Joris van der Hoeven
Christine Van Vredendaal
John Voight
Daqing Wan
Lawrence C. Washington
Jonathan Webster
Benjamin Wesolowski
Yinan Zhang
Alexandre Zotine