Counting solutions to equations in many variables over finite fields

Alan G.B. Lauder *
Mathematical Institute
Oxford University

September 30, 2003

Abstract

We present a polynomial-time algorithm for computing the zeta function of a smooth projective hypersurface of degree d over a finite field of characteristic p, under the assumption that p is a suitably small odd prime and does not divide d. This improves significantly upon an earlier algorithm of the author and Wan which is only polynomial-time when the dimension is fixed.

1 Introduction

Let f be a polynomial in n variables with coefficients in a finite field with q elements of characteristic p. A compelling problem in algorithmic number theory is to count in an efficient manner the number of solutions to the equation $f = 0$ over the base field, and also over all finite extensions of the base field. This information can be encoded in a generating function, the zeta function of the affine hypersurface defined by f. Dwork's rationality theorem asserts that such a zeta function is always the quotient of two polynomials with integer coefficients [6]. Thus the zeta function is a finite object, and a sensible question to ask is: can one compute it, and if so how quickly? These two problems subsume our original problem on counting solutions. Bombieri has given an explicit bound on the total degree of the zeta function [4], and using this one may easily show that the zeta function is effectively computable [24, Corollary 2.7]. This answers the first question affirmatively. To address the second question, one must first specify a sensible measure of the size of the input and output. A natural measure of size for the input is $d^n \log(q)$ bits, where d is the total degree of the polynomial f. (This is the measure appropriate for densely represented polynomials. It is unlikely that one can say too much about sparsely represented polynomials.) By Bombieri's bound, the output has size $(d^n \log(q))^{\mathcal{O}(1)}$ bits, where $\mathcal{O}(1)$ indicates a constant. (When $d = 0$ or 1 the problem is easy, but the input/output size should be taken as $(d+2)^n \log(q)$ if one wishes to include these trivial cases.) The adjective "quickly" usually means in polynomial-time in the input/output size, and so our second question can be refined as follows, c.f. [24, Problem 4.2].

*The author is supported by the EPSRC (Grant GR/N35366/01) and St John's College, Oxford, and thanks his colleagues in Oxford for their help and encouragement. He also wishes to thank the American Institute of Mathematics for their hospitality and support during a visit in March 2003, and Pierre Berthelot, Don Coppersmith, David Cox, Kiran Kedlaya, Bernard Le Stum and Daqing Wan for their advice on specific questions. The anonymous referees provided many helpful comments, for which the author is very grateful. *Mathematics Subject Classification 2000*: 11Y99, 11M38, 11T99.

Question: Does there exist an explicit deterministic algorithm which takes as input a polynomial f in n variables of degree $d \geq 2$ over the field with q elements, gives as output the zeta function of the affine hypersurface defined by the equation $f = 0$, and requires $(d^n \log(q))^{\mathcal{O}(1)}$ bit operations?

Considerable work has been done on this problem over the last few decades, motivated in part by applications in cryptography. However, to the author's knowledge, only three general "qualitative" results have been obtained (that is, disregarding constants in the exponents). The first is for the case $n = 1$, which is relatively easy. Here one can obtain the optimal $(d \log(q))^{\mathcal{O}(1)}$ bit operations. This was proved by both Schwarz and Butler, and later by Berlekamp as part of his pioneering work on univariate polynomial factorisation [11, Note 14.8]. The second is the theorem of Schoof-Pila for curves, which essentially says that in the case $n = 2$ a complexity of $\log(q)^{C_d}$ can be attained, where the constant C_d depends exponentially upon d [19, 22]. The third is the general result of the author and Wan, which gives a complexity of $(pd^n \log(q))^{\mathcal{O}(n)}$ bit operations [15]. (It should also be possible to obtain an estimate of this form for smooth affine hypersurfaces using the method of Kedlaya [13], with the exponent improved by a constant factor. This is certainly the case for smooth projective hypersurfaces, using a direct application of Dwork's cohomological theory.) In the present paper we develop in full the deformation method introduced by the author in [14, Section 2]. The principal result is a complexity of $(pd^n \log(q))^{\mathcal{O}(1)}$ bit operations, for suitably generic homogeneous polynomials and under mild restrictions on p and d. We now introduce the notation and definitions necessary to fully explain this result.

In this paper we shall actually be concerned with projective hypersurfaces, rather than the affine hypersurfaces described above. Let $\mathbb{F}_q$ denote the finite field with q elements, where q is a power of a prime p. Fix an algebraic closure $\bar{\mathbb{F}}_q$ of $\mathbb{F}_q$, and let $\mathbb{F}_{q^k}$ denote the unique subfield of $\bar{\mathbb{F}}_q$ of order q^k, for k a positive integer. Let $f \in \mathbb{F}_q[X_1, \ldots, X_n]$ be homogeneous of degree d, where $n \geq 2$. For $k \geq 1$, let N_k denote the number of $\mathbb{F}_{q^k}$-rational points on the projective hypersurface defined by the equation $f = 0$. Specifically, the integer N_k can be defined via the equation

$$(q^k - 1)N_k + 1 = \#\{(x_1, \ldots, x_n) \in \mathbb{F}_{q^k}^n \mid f(x_1, \ldots, x_n) = 0\},$$

where $\#$ denotes the cardinality of a set. The zeta function of the projective hypersurface defined by f is

$$Z(f/\mathbb{F}_q, T) = \exp\left(\sum_{k=1}^{\infty} N_k \frac{T^k}{k}\right).$$

This is a rational function by Dwork's theorem. The main result of this paper is as follows. (Here the term "suitably generic" means in an explicit Zariski dense open subset of the space of all homogeneous polynomials of degree d in n variables over $\mathbb{F}_q$.)

Theorem 1 *There exists an explicit deterministic algorithm with the following input, output and bit complexity. The input is any suitably generic homogeneous polynomial $f \in \mathbb{F}_q[X_1, \ldots, X_n]$ of degree $d \geq 2$, where q is a power of a prime p. We assume that $p \neq 2$ and p does not divide d. The output is the zeta function $Z(f/\mathbb{F}_q, T)$ of the projective hypersurface defined by f. The algorithm requires $(pd^n \log(q))^{\mathcal{O}(1)}$ bit operations.*

Our generic condition implies in particular that the hypersurface is smooth. In fact, smoothness is the only restriction that one really needs to insist upon, and we sketch a proof of precisely how one relaxes the requirements in Note 21. With regard to the exact running time, the exponents in

the complexity estimate are given explicitly in Theorem 10. For example, the dependence on $\log(q)$ is essentially cubic, and the author believes the algorithm should be practicable in many interesting cases. The algorithm should out-perform all previous approaches when $n \geq 4$. When $n = 3$, the case of curves, the running time is broadly comparable to the approach of Kedlaya.

The theorem is proved using a very indirect method, which avoids the difficulties inherent in the more straightforward approaches of the author and Wan, and of Kedlaya. All three approaches use, in some form, the p-adic theory developed by Dwork in his study of the zeta function of a hypersurface. By contrast, the theorem of Schoof-Pila uses Weil's l-adic construction, where l is a prime distinct from p. The latter theory is constructive but restricted to curves (and abelian varieties). Even for curves Weil's theory does not appear to yield the desired complexity estimate of $(d \log(q))^{\mathcal{O}(1)}$ bit operations, at least when applied in a direct manner. One would like to use the general l-adic cohomology theory of Grothendieck to compute zeta functions, but this is apparently non-constructive in its present form. So it seems that Dwork's theory is the only one at present which is constructive for general hypersurfaces. Since it is a p-adic theory it is difficult to see how one might remove the factor $p^{\mathcal{O}(1)}$ from the complexity estimates. Thus it is probable that $(pd^n \log(q))^{\mathcal{O}(1)}$ bit operations is the optimal qualitative result that can be obtained using p-adic cohomology. Of course, there is always scope for improvement in the constant in the exponent, and much work remains to be done on obtaining good implementations of such algorithms.

It is of interest to observe that Theorem 10 even gives some improvement in the extreme case in which q is prime. We have the estimate of $\mathcal{O}(q^{2+\varepsilon})$ bit operations for computing the zeta function, when all other parameters are fixed, for any $\varepsilon > 0$. A nice corollary along these lines, of theoretical interest, is the following.

Theorem 2 *Let $f \in \mathbb{Z}[X_1, \ldots, X_n]$ be a suitably generic polynomial and $\varepsilon > 0$. There exists an explicit deterministic algorithm which takes as input a prime p, gives as output the number of solutions to the equation $f = 0 \bmod p$, and takes $\mathcal{O}(p^{2+\varepsilon})$ bit operations.*

Here "suitably generic" means that some explicit integer polynomial in the coefficients of f has a non-zero value. This theorem improves upon the elementary estimate of $\mathcal{O}(p^{n-1+\varepsilon})$ bit operations which can be obtained using Berlekamp's root counting algorithm. (One could also presumably prove an estimate of $\mathcal{O}(p^{n-2+\varepsilon})$ bit operations, for $n \geq 2$, by slicing the hypersurface into curves and using the Schoof-Pila algorithm. Results of this nature are described by Elkies in [10, Pages 34-35].)

The algorithm underlying Theorem 1 can be roughly described in the following manner. Given a hypersurface of the type described in the theorem, it is sufficient to compute the action of the qth power Frobenius map on the (primitive) middle-dimensional p-adic cohomology space constructed by Dwork. The qth power Frobenius map factors as a product of $\log_p(q)$ copies of the pth power Frobenius map, and it is a matrix for this action, the absolute Frobenius matrix, which we will compute. The middle-dimensional cohomology space of Dwork is the quotient of an $(n+1)$-variate p-adic power series ring by the sum of the images of n differential operators. Using the direct approach of Kedlaya, one computes the action of the pth power Frobenius by first lifting to the power series ring, computing the Frobenius action in this ring, and then performing cohomological reduction via the differential operators. In the method of the author and Wan, one does all computations in the power series ring itself. The drawback of both approaches is that, because of the decay rate of the $(n+1)$-variate p-adic power series, one requires $(pd^n \log(q))^{\mathcal{O}(n)}$ bits to represent them modulo the required power of p. This dominates the complexity of both approaches.

One desires a completely different approach which avoids all computations with the $(n+1)$-variate power series, and works solely at the level of the cohomology space itself. This is achieved in the present paper. The new algorithm uses the relative cohomology theory of Dwork, which allows one to study hypersurfaces in a family, and relate the absolute Frobenius matrices of different members of the family. The idea is that given a suitably generic polynomial f, one can vary the non-diagonal terms in the polynomial by premultiplying them by another variable Y. When Y is set to one, we just get our original polynomial back. When Y is set to zero, the non-diagonal terms disappear, and we are left with a diagonal form. Thus we can "deform" our hypersurface through a one-dimensional family to get a diagonal hypersurface. The absolute Frobenius matrix for a diagonal form is easy to compute, as there is an explicit formula for its entries which can be quickly evaluated. The variation of the absolute Frobenius matrices in the family is controlled by a differential equation. Solving this differential equation around the origin, and using the absolute Frobenius matrix for the diagonal form as a starting point, one can construct a p-adic power series expansion in Y for the "generic" absolute Frobenius matrix. Unfortunately, this expansion does not necessarily converge at $Y = 1$. However, using knowledge of the domain of holomorphy of the "generic" absolute Frobenius matrix, one can use the local expansion indirectly to recover the value of this matrix at the specialisation $Y = 1$. This is just the absolute Frobenius matrix of the original polynomial, and the zeta function can now be easily computed. The key point is that since the deformation is one-dimensional, the complexity of this indirect approach is largely independent of n, and we get the bound of $(pd^n \log(q))^{\mathcal{O}(1)}$ bit operations on the overall running time. A more detailed sketch of this strategy is given in [14, Section 2] and the theoretical ideas underlying it are concisely summarised by Katz in [12].

The method is extremely flexible, and when cast in the language of relative rigid cohomology [2, Section 4.3], should allow the efficient computation of zeta functions and L-functions in great generality. (Note that the author has not studied the general problem in any detail, and this paragraph just contains speculation on what he believes should be possible.) Here is a very brief description of the method in this setting: Given an overconvergent F-isocrystal on a smooth variety, embed it in a family whose base is a smooth affine curve through the origin. Moreover, do so in such a way that the action of Frobenius on the cohomology of the fibre at the origin is easy to compute. (The author does not claim that one can always find such an embedding.) The commutativity of relative Frobenius and the connection, on each piece of the cohomology of the family, gives a factorisation of the relative Frobenius in terms of a basis of horizontal sections of the connection around the origin and the Frobenius at the origin. This allows one to compute a local expansion around the origin of the relative Frobenius on each piece of cohomology, and thus by analytic continuation the Frobenius action, and hence L-function, of the original F-isocrystal. One could apply this strategy to, for example, Artin-Schreier crystals [14], Kummer crystals or a suitable hybrid to compute L-functions of additive, multiplicative or mixed character sums, respectively. (Note that the zeta function of a variety is the L-function of the trivial rank one crystal on the variety.) One expects that the complexity will be uniformly bounded in terms of the characteristic, field extension degree and "Betti numbers". Although the required equations are readily available when working with rigid cohomology, proving such algorithmic results in this setting is not an entirely straightforward task. For example, one requires a precise estimate on the domain of holomorphy of the relative Frobenius. Rigid cohomology just tells one that it is overconvergent, and the required result is not presently available in the literature. It would be very interesting to fully develop the method in this setting.

The remainder of the paper is structured in the following manner. In Section 2 we collect

together the main theoretical results we shall need in our algorithm. The results themselves are proved in Section 4.1 and Appendix A. The algorithm is presented in Section 3, and precise estimates on the running time of this algorithm are given in Theorem 9. Our algorithm actually tackles a slightly more general problem than that stated in this section, and in Theorem 10 we specialise our results to the case of primary interest. Both Theorems 1 and 2 are corollaries of Theorem 10. The contents of Sections 4, 5, 6, 7, 8, 9 and 10 are explained in Section 3. Essentially, they contain the proof of correctness and complexity analysis of the algorithm, as well as a detailed discussion of how to perform the various steps. Finally, in Section 11 we complete the proofs of the results in the introduction. The paper is organised in a similar manner to [14], although the author has tried to avoid repetition by referring to [14] when only minor modifications are required in an argument.

Numerous papers on the computation of zeta functions have been written over the last few decades. Motivated by applications in cryptography, they have mainly focused on the case of curves (see the bibliography in [3] and also the paper of Poonen [20]). To the author's knowledge, the general problem was first discussed in detail by Elkies [10, Page 33-35] and Wan [23]. The seminal paper of Schoof [22] attacked the problem for elliptic curves with Weil's l-adic theory, while Satoh [21] and Wan pioneered the use of p-adic methods. The author heartily recommends Wan's expository paper [24] as a good introduction to the subject. The method used in [14] and the present paper is significantly different from all previous approaches.

2 Background Theory

In Appendix A we describe the relative p-adic cohomology theory of Dwork for one-dimensional families of smooth projective hypersurfaces [8, 9]. Our presentation and proofs are somewhat different to that in Dwork's paper, but the essential content is the same. In this section we gather together the main results from this theory which will be required in our algorithm.

We first make some conventions regarding multi-index notation. For any vector $u \in \mathbb{Z}^n$ define $X^u = X_1^{u_1} \dots X_n^{u_n}$ where $u = (u_1, \dots, u_n)$. For any vector $u \in \mathbb{Z}^{n+1}$ define $X^u = X_0^{u_0} X_1^{u_1} \dots X_n^{u_n}$ where $u = (u_0, \dots, u_n)$ and X_0 is a new indeterminate. Thus the notation X^u can mean a monomial in either n or $n+1$ variables depending upon the domain in which u lies. This will always be clear from the context. The symbol X unadorned will usually be used to denote the list of variables $X_0, X_1, \dots, X_n$; however, at times it is useful to use X for the list $X_1, \dots, X_n$ and we will explicitly state when this is done.

We define several subsets of $\mathbb{Z}_{\geq 0}^{n+1}$.

$$\begin{aligned} \mathcal{I} &= \{(u_0, u_1, \dots, u_n) \,|\, du_0 = u_1 + \dots + u_n\}, \\ \mathcal{I}^{(i)} &= \{(u_0, \dots, u_n) \in \mathcal{I} \,|\, u_j > 0 \text{ for } j \neq i\}, \text{ for } 1 \leq i \leq n, \\ \mathcal{I}^o &= \mathcal{I} \cap \mathbb{Z}_{>0}^{n+1}. \end{aligned}$$

These sets will be used to label monomials in polynomial and power series rings.

Next, we introduce some notation related to p-adic rings. Let $\mathbb{Q}_p$ be the field of p-adic numbers, and $\mathbb{Z}_p$ the ring of p-adic integers. Let $\mathbb{C}_p$ denote the completion of an algebraic closure of $\mathbb{Q}_p$. We denote by $\mathbb{Q}_q$ the unique unramified extension of $\mathbb{Q}_p$ in $\mathbb{C}_p$ of degree $\log_p(q)$, and by $\mathbb{Z}_q$ the ring of integers of $\mathbb{Q}_q$. Let $\pi \in \mathbb{C}_p$ be an element such that $\pi^{p-1} = -p$. Let ord be the p-adic order function on $\mathbb{C}_p$ normalised so that $\mathrm{ord}(p) = 1$, and $\mathcal{O}_{\mathbb{C}_p}$ the ring of integers of $\mathbb{C}_p$. Observe that $\mathrm{ord}(\pi) = 1/(p-1)$. For $y \in \mathcal{O}_{\mathbb{C}_p}$ in the unique unramified extension of $\mathbb{Z}_q$ of degree r for

some $r \geq 1$, let τ denote the automorphism of $\mathbb{Q}_q(\pi, y)$ which reduces to the pth power map on its residue class field and fixes π. Let

$$\bar{f} = \sum_{i=1}^{n} \bar{a}_i X_i^d + Y\bar{h}(X_1, \ldots, X_n) \in \mathbb{F}_q[Y][X_1, \ldots, X_n],$$

where the polynomial $\bar{h}$ is homogeneous of degree d with no diagonal terms and $\bar{a}_1 \ldots \bar{a}_n \neq 0$. For $\bar{y} \in \mathbb{F}_{q^r}$, we shall denote by $\bar{f}(X, \bar{y})$ the polynomial obtained by setting $Y = \bar{y}$ in $\bar{f}$ (so here the X means $X_1, \ldots, X_n$). Let f denote the polynomial obtained by taking the Teichmüller liftings of the coefficients of $\bar{f}$. Specifically, writing $\bar{h} = \sum_{j \in J} \bar{b}_j X^j$ we have

$$f = \sum_{i=1}^{n} a_i X_i^d + Y \sum_{j \in J} b_j X^j.$$

Here the removal of bars indicates we have taken Teichmüller liftings.

Let $\mathcal{R}$ denote the subring of $\mathbb{Q}_q(\pi)(Y)$ which consists of rational functions which have no pole at the origin. Let $\mathcal{L}_Y$ denote the $\mathcal{R}$-module spanned by the monomials X^u with $u \in \mathcal{I}$. This module is also a ring under the usual multiplication of polynomials. Let $\mathcal{L}_Y^o$ and $\mathcal{L}_Y^{(i)}$ be the $\mathcal{R}$-submodules spanned by monomials with $u \in \mathcal{I}^o$ and $u \in \mathcal{I}^{(i)}$, respectively. (In Appendix A we actually work with a bigger ring $R \supset \mathcal{R}$ and bigger modules $L_Y \supset \mathcal{L}_Y$ etc., but these are more complicated to describe.) For $1 \leq i \leq n$ let the first-order differential operators

$$D_{i,Y} := X_i \frac{\partial}{\partial X_i} + \pi X_0 X_i \frac{\partial f}{\partial X_i}$$

act on the ring $\mathcal{L}_Y$. The set

$$\{\pi^{u_0} X^u \,|\, u \in \mathcal{B}\} \text{ where } \mathcal{B} := \{u = (u_0, u_1, \ldots, u_n) \in \mathcal{I}^o \,|\, u_1, \ldots, u_n < d\}$$

is an $\mathcal{R}$-basis for the quotient $\mathcal{R}$-module

$$\mathcal{L}_Y^o / \sum_{i=1}^{n} D_{i,Y}(\mathcal{L}_Y^{(i)}).$$

(See the explicit reduction formulae in Section 5.2.) By a simple argument one sees that the size of $\mathcal{B}$ is $\frac{1}{d}\sum_{i=0}^{d-1} g(\eta^i)$, where $g(z) = (z + z^2 + \ldots + z^{d-1})^n$ and η is a primitive dth root of unity [18, Page 89]. Thus

$$\#\mathcal{B} = \frac{1}{d}\{(d-1)^n + (-1)^n(d-1)\}.$$

We can now describe the matrix which defines the "differential equation of the deformation".

Definition 3 *For $u, v \in \mathcal{B}$, let $B_{u,v}(Y)$ denote the coefficient of $\pi^{u_0} X^u$ in the reduction of*

$$\pi X_0 h(X_1, \ldots, X_n) \times \pi^{v_0} X^v$$

modulo the images of the operators $D_{i,Y}$, where $1 \leq i \leq n$. Let $B(Y) = (B_{u,v})$ be the corresponding square matrix of size $\#\mathcal{B}$ over $\mathcal{R}$.

Note that $\mathcal{R}$ can be embedded in $\mathbb{Q}_q(\pi)[[Y]]$, via expansions of rational functions at the origin, and so we can consider the matrix $B(Y)$ as having entries in $\mathbb{Q}_q(\pi)[[Y]]$. The deformation matrix itself is as follows.

Definition 4 *Let $C(Y)$ be the matrix over $\mathbb{Q}_q(\pi)[[Y]]$ which is the unique solution around the origin to the differential equation and initial condition*

$$\frac{\partial C}{\partial Y} = C(Y)B(Y),\ C(Y) \equiv I \bmod (Y). \tag{1}$$

The idea is that the matrix $C(Y)$ controls the change in the "generic" zeta function "$Z(\bar{f}(X,Y),T)$" as one moves from $Y = 0$ to a "generic" choice of Y. We give an explicit formula for the zeta function at the specialisation $Y = 0$; more precisely, for the matrix which represents the action of the absolute Frobenius map on the cohomology space constructed by Dwork.

Proposition 5 *Let $\alpha(0)$ be the matrix for the action of the absolute Frobenius map on the cohomology space associated to the zeta function $Z(f(X,0),T)$. (This is defined in Section A.2.) The entry in the uth row and vth column of $\alpha(0)$ for $u, v \in \mathcal{B}$ is*

$$\pi^{v_0-u_0} \prod_{i=1}^{n} \sum_{m,r\geq 0} \lambda_m (u_i/d)_r (-1)^r \pi^{-r} \tau^{-1}(a_i^m) a_i^{-r}. \tag{2}$$

Here the sum is over all $m, r \geq 0$ such that $pu_i - v_i = d(m - pr)$, where $u = (u_0, u_1, \ldots, u_n)$ and $v = (v_0, v_1, \ldots, v_n)$.

The element λ_m is the coefficient of z^m in the power series $\exp(\pi(z - z^p))$. The p-adic integer $(u_i/d)_r$ is defined to be 1 when $r = 0$, and for $r > 0$ to be

$$(u_i/d)_r := (u_i/d)((u_i/d) + 1) \ldots ((u_i/d) + (r - 1)).$$

This proposition is proved in Section 4.1.

The main result of the deformation theory is as follows.

Proposition 6 *Let $\alpha(Y)$ denote the absolute Frobenius matrix for the "generic" Y. (This matrix is defined explicitly in Section A.2.) Then we have the following identity of matrices over $\mathbb{Q}_q(\pi)[[Y]]$:*

$$\alpha(Y) = C(Y^p)^{-1} \alpha(0) C^{\tau^{-1}}(Y).$$

Here τ^{-1} acts on the coefficients of power series, fixing the variable Y.

This proposition is proved in Appendix A. Specifically, see Equation (30) and Section A.3.

The generic Frobenius matrix as defined in Section A.2 contains entries which are power series in Y. These local expansions around the origin will not in general converge at points on the p-adic unit disk. However, it can be continued to a bigger domain, which includes the Teichmüller liftings of all but finitely many points in $\bar{\mathbb{F}}_q$. In the next proposition, the matrix $\alpha(y^{\tau^{-1}})$ means the value taken by the p-adic holomorphic function which continues the local expansion of α at the point $\tau^{-1}(y)$. It is closely related to the zeta functions we wish to compute.

Proposition 7 *Write*

$$Z(\bar{f}(X,\bar{y})/\mathbb{F}_{q^r},T) = \frac{P(T)^{(-1)^{n+1}}}{(1-T)(1-q^rT)\dots(1-q^{r(n-2)}T)}$$

for the zeta function of the smooth projective hypersurface $f(X,\bar{y})=0$ defined by specialising $Y=\bar{y}$ where $\bar{y}\in\mathbb{F}_{q^r}$ with $\mathfrak{g}(y)\neq 0$ and $R(\bar{y})\neq 0 \bmod p$. Then

$$P(qT) = \det(I - \alpha(y^{\tau^{-1}})^{\tau^{r\log_p(q)}}\alpha(y^{\tau^{-1}})^{\tau^{r\log_p(q)-1}}\dots\alpha(y^{\tau^{-1}})^{\tau}T).$$

Here y is the Teichmüller lifting of the field element $\bar{y}$.

The above proposition is proved in Section A.4. The product $\mathfrak{g}(Y)R(Y)$ specifies the generic condition we shall need in our algorithm. These two polynomials are defined in (6) and (7).

3 The Algorithm and Theorems

We now put the results in the previous section together in an appropriate manner to give our algorithm for computing the zeta function of a projective hypersurface. Note that we have made no effort whatsoever to minimise the constant factor in the running time, and our p-adic and Y-adic accuracies can certainly be taken to be much smaller.

Algorithm 8
Input: A homogeneous polynomial $f(X,\bar{y}) = \sum_{i=1}^{n}\bar{a}_iX_i^d + \bar{y}\bar{h}$, where $\bar{h}$ has no diagonal terms and $\bar{a}_1\dots\bar{a}_n\neq 0$. Also $\bar{y}\in\mathbb{F}_{q^r}$ for some $r\geq 1$ with $\mathfrak{g}(\bar{y})R(\bar{y})\neq 0 \bmod p$. We insist $p>2$ and $d\geq 2$, with d not divisible by p. (See Equations (6) and (7) for the definitions of $\mathfrak{g}(Y)$ and $R(Y)$.)
Output: The zeta function of the smooth projective hypersurface defined by equation $f(X,\bar{y})=0$.

Step 0: Set-up
We use the notation defined in Section 2. All computations below are performed with p-adic numbers in $\mathbb{Q}_q(\pi)$ and $\mathbb{Q}_q(\pi,y)$ working "modulo" some power of p. Define

$$\begin{aligned} N &= 2d^{n-1}rn\log_p(q) \\ N_Y &= 12p(ed)^{n-1}(N+(d-1)^{n-1}rn\log_p(q)+n) \\ \tilde{N} &= 171d^{n-1}(ed)^{n-1}rn\log_p(q). \end{aligned}$$

Step 1: Teichmüller liftings
Compute modulo $p^{\tilde{N}}$ the Teichmüller liftings of the coefficients of $\bar{h}(X)$.

Step 2: Compute the differential system
Let $B(Y)$ be the matrix of the differential system, as in Definition 3. Using the method of Section 6.1 compute the matrix $B(Y)$ with coefficients modulo $p^{\tilde{N}}$.

Step 3: Solve the differential system at the origin
Let $C(Y)$ be the unique solution matrix around the origin to the differential system, as in Definition

4. Working modulo $(p^{\tilde{N}}, Y^{N_Y})$, compute $C(Y)$ using the method in Section 6.2.

Step 4: Matrix inversion
Working modulo $(p^{\tilde{N}}, Y^{N_Y})$ compute the inverse matrix $C(Y^p)^{-1}$ using the Newton iteration method in [14, Section 5.2.2].

Step 5: Find the absolute Frobenius matrix for the diagonal case
Let $\alpha(0)$ be the matrix for the absolute Frobenius map in the case $\bar{y} = 0$, as defined in Proposition 5. Compute $\alpha(0)$ modulo $p^{\tilde{N}}$ using the summation bounds in Section 4.2.

Step 6: Compute the local expansion of the generic absolute Frobenius matrix around the origin
Working modulo $(p^{\tilde{N}}, Y^{N_Y})$ compute the the matrix product $\alpha(Y) := C(Y^p)^{-1}\alpha(0)C^{\tau^{-1}}(Y)$.

Step 7: Evaluate the generic absolute Frobenius matrix at a Teichmüller point
Compute the Teichmüller lifting $\tau^{-1}(y)$ of the element $\bar{y}^{1/p}$ modulo $p^{\tilde{N}}$. Compute $\alpha(y^{\tau^{-1}})$ modulo $p^{\tilde{N}}$ using the "analytic continuation" method in Section 8.

Step 8: Exponentiate and take the characteristic polynomial
Let $R(T)$ be the rational function over $\mathbb{Z}_p$ defined as

$$R(T)^{(-1)^{n+1}} = \det(I - \alpha(y^{\tau^{-1}})^{\tau^{r\log_p(q)}}\alpha(y^{\tau^{-1}})^{\tau^{r\log_p(q)-1}}\ldots\alpha(y^{\tau^{-1}})^{\tau}T).$$

Compute $R(T)$ modulo $p^{\tilde{N}}$ using fast exponentiation and the algorithm for characteristic polynomials from [14, Section 9].

Step 9: The zeta function
Let $P(T)$ with $P(T)^{(-1)^{n+1}} \in 1 + T\mathbb{Z}[T]$ be the unique rational function with coefficients in the range $(-p^{N-1}, p^{N-1}]$ such that $P(qT) \equiv R(T) \bmod p^N$. Output the rational function

$$Z(f(X, \bar{y}), T) = \frac{P(T)}{(1-T)(1-q^rT)\ldots(1-q^{r(n-2)}T)}.$$

In Section 6 we explain how to perform the non-trivial tasks in Steps 2, 3 and 4. Sections 4.2 and 8.2 explain how to perform Step 5 and Step 7, respectively. The tasks required in Steps 1, 5, 8 and 9 are relatively straightforward, and are discussed during the complexity analysis in Section 10.

The theorems which underlie the algorithm are located in Section 2. The proofs of these theorems are given in Section 4.1 and Appendix A. These theorems show that the algorithm would perform correctly if all computations could be performed to infinite p-adic and Y-adic accuracy. We must justify that the various p-adic and Y-adic accuracies at which power series are truncated does not compromise the correctness of the final answer. This is done in Section 9. For this, one must have reasonably good bounds on the domain of holomorphy of the generic Frobenius matrix $\alpha(Y)$; these are found in Section 7.1. This will prove the correctness of the algorithm. The analysis of the complexity of the algorithm is given in Section 10. This gives a proof of the following theorem,

in which we use Soft-Oh notation to hide logarithmic factors, as defined in [15, Section 6.3]. (Note that we have dropped the bars in the statement of the theorems to simplify notation.)

Theorem 9 *There exists an explicit deterministic algorithm (namely, Algorithm 8) with the following input, output and bit complexity. The input is any homogeneous polynomial*

$$f(X_1, \ldots, X_n, Y) = \sum_{i=1}^{n} a_i X_i^d + Yh(X_1, \ldots, X_n) \in \mathbb{F}_q[Y][X_1, \ldots, X_n]$$

of degree $d \geq 2$, where h has no diagonal terms and $a_1 \ldots a_n \neq 0$, along with $y \in \mathbb{F}_{q^r}$ for some $r \geq 1$. Here q is a power of p, with $p \neq 2$ and d not divisible by p. We assume that $\mathfrak{g}(y)R(y) \neq 0 \bmod p$, where $\mathfrak{g}(Y)R(Y) \bmod p$ is a non-zero univariate polynomial over $\mathbb{F}_q$ constructed from the coefficients of f. The output is the zeta function of the smooth projective hypersurface defined by the polynomial $f(X_1, \ldots, X_n, y)$. The complexity is

$$\tilde{\mathcal{O}}((\max\{d^{5+\omega}e^3, d^6e^5\})^{n-1}n^3p^2\log(q)^3r^3)$$

bit operations, and the algorithm requires $\tilde{\mathcal{O}}(d^{6(n-1)}e^{4(n-1)}n^2p^2\log(q)^3r^2)$ bits of space. Here e is the base of the natural logarithms and ω the exponent of matrix multiplication.

We note that the polynomial $\mathfrak{g}(Y)R(Y) \bmod p$ can be evaluated at any point $\bar{y} \in \mathbb{F}_{q^r}$ in $\mathcal{O}(n(ed)^{3n-3})$ operations in $\mathbb{F}_{q^r}$, where e is the base of the natural logarithms. The coefficients of $\mathfrak{g}(Y)R(Y)$ are integer polynomial expressions in the coefficients of f, and by Lemma 13 the constant term in Y is not zero. By Lemma 20, treating the coefficients of f as independent variables, the expression $a_1 \ldots a_n\mathfrak{g}(1)R(1)$ is a non-zero integer polynomial in these independent variables. Denote this polynomial by $\Delta_{n,d}$, since it depends only on d and n. Let $\Delta_{n,d}(\bar{f}) \in \mathbb{F}_q$ denote the value taken modulo p by this polynomial when evaluated at the coefficients of f. The condition $\Delta_{n,d}(\bar{f}) \neq 0$ is our generic requirement. When this is satisfied, $\bar{f}$ defines a smooth projective hypersurface, although the converse is not quite true, see Note 21. Setting $y = 1$ in the above theorem gives the next result.

Theorem 10 *The exists an explicit deterministic algorithm with the following input, output and bit complexity. The input is any polynomial $f \in \mathbb{F}_q[X_1, \ldots, X_n]$ which is homogeneous of degree $d \geq 2$ such that $\Delta_{n,d}(f) \neq 0$. We assume q is a power of p where $p \neq 2$ and p does not divide d. The output is the zeta function of the smooth projective hypersurface defined by f. The running time is*

$$\tilde{\mathcal{O}}((\max\{d^{5+\omega}e^3, d^6e^5\})^{n-1}n^3p^2\log(q)^3)$$

bit operations, and the algorithm requires $\tilde{\mathcal{O}}(d^{6(n-1)}e^{4(n-1)}n^2p^2\log(q)^3)$ bits of space.

Theorem 1 is an immediate corollary of this result. Theorem 2 follows with a few lines of work. This is described in Section 11, where we also discuss the forgotten cases with $p|d$ or $p = 2$, and explain how our generic condition can be relaxed so that we just require that the projective hypersurface defined by $\bar{f}$ is smooth.

4 The absolute Frobenius matrix for the diagonal form

4.1 Proof of Proposition 5

In this section we prove Proposition 5 which gives an explicit formula for the absolute Frobenius matrix of a diagonal form. For $u, v \in \mathcal{B}$, this matrix is defined to have (u, v)th entry the coefficient of $\pi^{u_0} X^u$ in the reduction modulo the differential operators $D_{i,0} (1 \leq i \leq n)$ of

$$\psi_p \circ F(X, 0)(\pi^{v_0} X^v). \tag{3}$$

(See Section A.2.) Here $F(X, 0) = \prod_{i=1}^{n} \theta(a_i X_0 X_i^d)$ where $\theta(z) = \exp(\pi(z - z^p)) = \sum_{m=0}^{\infty} \lambda_m z^m$. Also ψ_p acts on power series as

$$\psi_p(\sum_r A_r X^r) = \sum_r \tau^{-1}(A_{pr}) X^r$$

where the sum here is over $(n+1)$-tuples of non-negative integers. It is convenient to introduce a dth root $X_0^{1/d}$ of the indeterminate X_0. Extend the operator ψ_p to act as $\psi_p(X_0^{r/d}) = X_0^{r/pd}$ if p divides $r \geq 0$, and zero otherwise. Since p does not divide d this is consistent with the action of ψ_p on X_0. We may now rewrite (3) as

$$\pi^{v_0} \prod_{i=1}^{n} \sum_{m \geq 0,\, p | dm + v_i} \lambda_m \tau^{-1}(a_i^m)(X_0^{1/d} X_i)^{(dm+v_i)/p}.$$

Here we use the fact $\psi_p(\prod_{i=1}^{n} B_i(X_0^{1/d} X_i)) = \prod_{i=1}^{n} \psi_p(B_i(X_0^{1/d} X_i))$ for $B_i(X_0^{1/d} X_i) \in \mathbb{Q}_q(\pi)[[X_0^{1/d} X_i]]$. Let the operators

$$D_{i,0} = X_i \frac{\partial}{\partial X_i} + \pi d a_i X_0 X_i^d,\ 1 \leq i \leq n$$

act in the obvious manner on the ring obtained by adjoining the dth root $X_0^{1/d}$; so $D_{i,0}(X_0^{r/d} *) = X_0^{r/d} D_{i,0}(*)$. One may check that

$$D_{i,0}(B_i(X_0^{1/d} X_i)) D_{j,0}(B_j(X_0^{1/d} X_j)) = D_{i,0}(B_i(X_0^{1/d} X_i) D_{j,0}(B_j(X_0^{1/d} X_j)))$$

for $B_i(X_0^{1/d} X_i) \in \mathbb{Q}_q(\pi)[[X_0^{1/d} X_i]]$ and $B_j(X_0^{1/d} X_j) \in \mathbb{Q}_q(\pi)[[X_0^{1/d} X_j]]$ with $i \neq j$. It follows that it is enough to determine the coefficient of $X_0^{u_0/d} X_i^{u_i}$ in the reduction modulo $D_{i,0}$ of each summation

$$\sum_{m \geq 0,\, p | dm + v_i} \lambda_m \tau^{-1}(a_i^m)(X_0^{1/d} X_i)^{(dm+v_i)/p}$$

and take the product for $i = 1, \ldots, n$. One checks directly that

$$(X_0^{1/d} X_i)^{u_i + dr} \equiv \pi^{-r} (-1)^r (u_i/d)_r a_i^{-r} (X_0^{1/d} X_i)^{u_i} \bmod D_{i,0}.$$

Hence for $(dm + v_i)/p = u_i + dr$ we have that

$$\lambda_m \tau^{-1}(a_i^m)(X_0^{1/d} X_i)^{(dm+v_i)/p} \equiv \lambda_m \tau^{-1}(a_i^m) \pi^{-r} (-1)^r (u_i/d)_r a_i^{-r} (X_0^{1/d} X_i)^{u_i} \bmod D_{i,0}.$$

Thus the coefficient of $(X_0^{1/d}X)^{u_i}$ in the reduction of the ith univariate power series modulo $D_{i,0}$ is

$$\sum_{m,r\geq 0,\, dm+v_i=p(u_i+dr)} \lambda_m \tau^{-1}(a_i^m)\pi^{-r}(-1)^r(u_i/d)_r a_i^{-r}.$$

Hence the coefficient of $\pi^{u_0}X^u$ in the reduction of (3) is

$$\pi^{v_0-u_0}\prod_{i=1}^{n}\sum_{m,r\geq 0,\, dm+v_i=p(u_i+dr)} \lambda_m \tau^{-1}(a_i^m)a_i^{-r}\pi^{-r}(-1)^r(u_i/d)_r,$$

as claimed in Equation (2).

4.2 Auxiliary Routines: Step 5

In Step 5 we need to evaluate the sums which occur in (2) modulo $p^{\tilde{N}}$. The bounds in [14, Section 6.2] show that it is enough to compute each sum for m, r such that $m \leq 2p^2(\tilde{N}+(2n-3))/(p-1)$.

5 Explicit Reduction Formulae

In this section we examine in detail the quotient space $\mathcal{L}_Y^o/\sum_{i=1}^n D_{i,Y}(\mathcal{L}_Y^{(i)})$. We give explicit formulae which may be used to reduce elements in $\mathcal{L}_Y^o$ modulo the subspace $\sum_{i=1}^n D_{i,Y}(\mathcal{L}_Y^{(i)})$ to a sum of monomials in the basis $\pi^{u_0}X^u$ for $u \in \mathcal{B}$. These formulae are required in three distinct contexts. First, in the construction of the matrix $B(Y)$ in Section 6.1. Second, in the study of the domain of holomorphy of the matrix $\alpha(Y)$ in Section 7.1. Third, in Appendix A when we study the analytically defined space $L_Y/\sum_{i=1}^n D_{i,Y}(L_Y^{(i)})$ upon which the Frobenius map will act. Throughout this section, the notation f_i will be used to denote $X_i\frac{\partial f}{\partial X_i}$ for $1 \leq i \leq n$.

5.1 Standard forms in an n-variate quotient space

We begin with a lemma which will be used later in this section.

Lemma 11 *Let $\mathcal{D}(Y)$ be an $m \times m$ matrix over the ring $\mathbb{Z}_q[Y]$ with entries which are polynomials in Y of degree at most 1. Define $D(Y) = \det(\mathcal{D}(Y))$, and assume that $D(Y) \bmod p$ is not the zero polynomial. Then $\mathcal{D}(Y)^{-1}$ has entries of the form $A(Y)/D(Y)$ where $\deg_Y(A(Y)) < m$ and $\deg_Y(D(Y)) \leq m$. Moreover, both $D(Y)$ and the entries in $\mathcal{D}(Y)^{-1}$ may be computed deterministically in $\tilde{\mathcal{O}}(m^4)$ operations in $\mathbb{Z}_q$.*

Proof: The claim on the degree of $D(Y)$ is immediate from the explicit definition of the determinant. Likewise, the claim on the form of the entries in $\mathcal{D}(Y)^{-1}$ follows directly from the description of this matrix as the "adjugate" matrix times the inverse of the determinant $D(Y)$. We now discuss how these may be computed. First, suppose $c(Y) \in \mathbb{Z}_q[Y]$ is an irreducible polynomial, which remains irreducible of the same degree upon reduction modulo p. Reducing the entries in $\mathcal{D}(Y)$ modulo $c(Y)$, we get a matrix over an unramified extension of $\mathbb{Z}_q$. We use Gaussian elimination to put this matrix into "row-reduced form". By the latter, we mean simply that the leading entry in the jth row lies to the right of that in the ith row for $i < j$. Note that when performing row-reduction in a p-adic ring, one must choose each "pivot-row" to have a non-zero leading coefficient of minimal

p-adic order among the column being "cleared". Also, let us assume that the determinant of the matrix does not change during the reduction, except possibly by sign, i.e., we perform only two types of the usual three row operations. The determinant of the matrix modulo $c(Y)$ is the product of diagonal entries, up to a sign. If this product is not a unit, it means that $c(Y) \bmod p$ is a factor of $D(Y) \bmod p$. Let us assume that indeed the determinant is a unit. The row-reduced matrix can then be completely reduced to give a diagonal matrix, and keeping track of the row-operations, we can assume that the inverse matrix has also now been computed modulo $c(Y)$. The inverse matrix has entries in the ring $\mathbb{Z}_q[Y, 1/D(Y)]$, and we have computed it modulo the ideal $(c(Y))$ in this ring. The determinant lies in the ring $\mathbb{Z}_q[Y]$, and we have computed it modulo the ideal $(c(Y))$ in this ring. Now let $\{c(Y)\}$ be a set of "small degree" irreducible polynomials in $\mathbb{Z}_q[Y]$. We assume that all $c(Y)$ remain irreducible modulo p, are not factors of $D(Y) \bmod p$, and have degrees summing to an integer greater than m and less than $2m$. Such a set can be found deterministically within the required complexity bounds (some initially chosen $c(Y)$ may need to be discarded if the determinant turns out to a non-unit modulo $(c(Y), p)$). Using the Chinese Remainder Theorem (CRT), along with the above algorithm for a single $c(Y)$, one can recover the determinant $D(Y)$ exactly. Once $D(Y)$ is known, the unknown entries $A(Y)/D(Y)$ in the inverse can be recovered using the CRT from their reduction modulo $c(Y)$. The complexity of this approach is dominated by the Gaussian elimination. For each $c(Y)$, this requires $\tilde{\mathcal{O}}(m^3 \deg_Y c(Y))$ operations in $\mathbb{Z}_q$, using a soft-Oh linear time algorithm (in $\mathbb{Z}_q$-operations) for computing in $\mathbb{Z}_q[Y]/(c(Y))$. This gives a total complexity of $\tilde{\mathcal{O}}(m^4)$ operations in $\mathbb{Z}_q$, since $\sum \deg_Y(c(Y)) \leq 2m$. □

In this section, the notation X^m will be used to denote a monomial $X_1^{m_1} \ldots X_n^{m_n}$ where $m = (m_1, \ldots, m_n) \in \mathbb{Z}_{\geq 0}^n$. Moreover, we define $|m| = m_1 + \ldots + m_n$ for such a vector. For any i with $1 \leq i \leq n$, we say that a monomial X^m is reduced in $\{X_1^d, \ldots, X_{i-1}^d\}$ if it is not divisible by X_j^d for any $j < i$. Our convention for $i = 1$ is that all monomials are "reduced in $\{X_1^d, \ldots, X_{i-1}^d\}$".

For $\delta \in \mathbb{Z}_{\geq 0}$, let T_δ denote the set of monomials of total degree δ in $X_1, \ldots, X_n$ which are divisible by $X_1 \ldots X_n$. Thus $\#T_\delta = \binom{\delta-1}{n-1}$. For each $1 \leq i \leq n$, let $T_{\delta,i}$ be the set of monomials of degree δ which are reduced in $\{X_1^d, \ldots, X_{i-1}^d\}$ and which are divisible by $X_1 \ldots X_{i-1} X_{i+1} \ldots X_n$. For $1 \leq i \leq n$, let $T_\delta^{(i)} \subseteq T_\delta$ be the subset of monomials which are divisible by X_i^d, and are reduced in $\{X_1^d, \ldots, X_{i-1}^d\}$. Let $T_\delta^* \subseteq T_\delta$ be the subset of monomials not divisible by any X_i^d. Thus

$$T_\delta - T_\delta^* = \cup_{i=1}^n T_\delta^{(i)},\ T_\delta^{(i)} \cap T_\delta^{(j)} = \emptyset \text{ for } i \neq j.$$

Define $\bar{T}_\delta = T_\delta - T_\delta^*$, where the minus is set-theoretic difference. So $\bar{T}_\delta$ contains monomials of degree δ divisible by some X_i^d and by $X_1 \ldots X_n$. Denote by $X_i^d T_{\delta-d,i}$ the set of products $X_i^d t$ for $t \in T_{\delta-d,i}$. Then it is easily seen that $T_\delta^{(i)} = X_i^d T_{\delta-d,i}$. Thus

$$\sum_{i=1}^n \#T_{\delta-d,i} = \sum_{i=1}^n \#T_\delta^{(i)} = \#\bar{T}_\delta.$$

Definition 12 *For any $\delta \in \mathbb{Z}_{\geq 0}$, let $\mathcal{D}_\delta(Y)$ be the $\#\bar{T}_\delta \times \#\bar{T}_\delta$ matrix over $\mathbb{Z}_q[Y]$ defined as follows. The rows are labelled by the monomials $X^u \in \bar{T}_\delta$ and the columns by pairs (X^v, f_i) where $X^v \in T_{\delta-d,i}$. The entry in the row labelled by X^u and column by (X^v, f_i) is the coefficient of the monomial X^u in the product $X^v f_i$. Let $D_\delta(Y)$ be the determinant of $\mathcal{D}_\delta(Y)$. (When $\delta < d$ this matrix is empty, and we take the determinant to be 1.)*

Note that each f_i is divisible by X_i, and so multiplication by f_i maps monomials divisible by $X_1 \dots X_{i-1} X_{i+1} \dots X_n$ to those divisible by $X_1 \dots X_n$. For $\delta \geq n(d-1)+1$, every monomial of degree δ is divisible by some X_i^d, and thus $T_\delta^* = \emptyset$. So $\bar{T}_\delta = T_\delta$, and in this case the column of $\mathcal{D}_\delta(Y)$ labelled by (X^v, f_i) gives all the coefficients in the product $X^v f_i$, since all the monomials in this product have degree δ and are divisible by $X_1 \dots X_n$, and T_δ contains all such monomials. For $\delta \leq n(d-1)$, the set T_δ^* is non-empty. It contains monomials $X_1^{m_1} \dots X_n^{m_n}$ with $m_1 + \ldots + m_n = \delta$ and all $0 < m_i < d$. Thus in this case the column of $\mathcal{D}_\delta(Y)$ labelled by (X^v, f_i) just contains those coefficients of monomials appearing in $X^v f_i$ which are not in T_δ^*.

Lemma 13 *We have $D_\delta(0) \neq 0 \bmod p$. Thus the determinant $D_\delta(Y)$ is not identically zero (even modulo p).*

Proof: Each column contains exactly one non-zero constant, namely da_i, where the column is labelled by $(*, f_i)$ for some $*$. That each row also contains exactly one non-zero constant follows since $T_\delta^{(i)} = X_i^d T_{\delta-d,i}$, and the $T_\delta^{(i)}$ partition $\bar{T}_\delta$. Specifically, the row labelled by monomial $X^u \in T_\delta^{(i)}$ contains da_i in the column labelled by $(X_i^{-d} X^u, f_i)$. Thus the constant term of $D_\delta(Y)$ is $\pm \prod_{i=1}^n (da_i)^{\# T_{\delta-d,i}} \neq 0 \bmod p$, since by assumption $da_1 \dots a_n \neq 0 \bmod p$. □

Define $\mathbb{Z}_q[\pi, Y][X_1, \dots, X_n]^o = X_1 \dots X_n \mathbb{Z}_q[\pi, Y][X_1, \dots, X_n]$, and for $1 \leq i \leq n$ define

$$\mathbb{Z}_q[\pi, Y][X_1, \dots, X_n]^{(i)} = X_1 \dots X_{i-1} X_{i+1} \dots X_n \mathbb{Z}_q[\pi, Y][X_1, \dots, X_n].$$

The matrices $\mathcal{D}_\delta(Y)$ can be used to reduce homogeneous polynomials in $\mathbb{Z}_q[\pi, Y][X_1, \dots, X_n]^o$ to a standard form modulo the space $\sum_{i=1}^n f_i \mathbb{Z}_q[\pi, Y][X_1, \dots, X_n]^{(i)}$, up to some annihilating factor.

Proposition 14 *Let $U \in \mathbb{Z}_q[\pi, Y][X_1, \dots, X_n]^o$ be homogeneous of degree δ in the variables $X_1, \dots, X_n$. Let $\deg_Y(U)$ denote the degree in Y of U. If $\delta \geq n(d-1)+1$ then one may write*

$$U = \frac{1}{D_{n(d-1)+1}(Y)} \sum_{i=1}^n \left(\sum_{|j|=\delta-d} A_{ij}(Y) X^j \right) f_i(X_1, \dots, X_n, Y)$$

where $A_{ij}(Y) \in \mathbb{Z}_q[\pi, Y]$ with $\deg_Y(A_{ij}) < \binom{n(d-1)}{n-1} + \deg_Y(U)$. (When $\delta = n(d-1)+1$ the inner sum can in fact be taken over j with $X^j \in T_{\delta-d,i}$.) If $\delta \leq n(d-1)$ then one may write

$$U = \frac{1}{D_\delta(Y)} \sum_{i=1}^n \left(\sum_{X^j \in T_{\delta-d,i}} A_{ij}(Y) X^j \right) f_i(X_1, \dots, X_n, Y) + \sum_k \left(\frac{B_k(Y)}{D_\delta(Y)} + c_k(Y) \right) X^k.$$

Here $A_{ij}(Y), B_k(Y) \in \mathbb{Z}_q[\pi, Y]$ with $\deg_Y(A_{ij}), \deg_Y(B_k) < \binom{\delta-1}{n-1} + \deg_Y(U)$. The latter summation is over exponents $k = (k_1, \dots, k_n)$ with $k_1 + \ldots + k_n = \delta$ and $0 < k_1, \dots, k_n < d$, and $c_k(Y)$ is the coefficient of X^k in U. In both cases, writing $A_i(X) := \sum_j A_{ij}(Y) X^j$ we have $A_i(X) \in \mathbb{Z}_q[\pi][Y][X_1, \dots, X_n]^{(i)}$. (Here the summation range is as above for the two cases.) Moreover, when $\delta \leq n(d-1)+1$ each one of these representations may be computed deterministically in $\tilde{\mathcal{O}}\left(\binom{n(d-1)}{n-1}^4 + \binom{n(d-1)}{n-1}^2 \max\{\binom{n(d-1)}{n-1} - 1, \deg_Y(U)\} \right)$ operations in the ring $\mathbb{Z}_q[\pi]$.

Proof: First suppose $\delta \geq n(d-1)+1$. Write $U = \sum_r X^r U_r$ where each U_r is homogeneous of degree $n(d-1)+1$ and divisible by $X_1 \ldots X_n$, and the sum is finite. (This may be done in several ways, and we just choose one according to some simple rule.) Thus it suffices in this case to "reduce" a homogeneous polynomial of degree exactly $n(d-1)+1$ which is divisible by $X_1 \ldots X_n$. Given such a polynomial U, write its coefficients as an $\#T_{n(d-1)+1}$ column vector u, say. Then we need to solve the equation $\mathcal{D}_{n(d-1)+1}(Y)v = u$ for some column vector v. Since $D_{n(d-1)+1}(Y) \neq 0$ this can be done uniquely, and the claimed degree bounds follow immediately from the degree bounds in Lemma 11. This proves the existence of such a representation in the case $\delta \geq n(d-1)+1$. Suppose now that $\delta = n(d-1)+1$, and so we may assume the finite sum $\sum_r X^r U_r$ has only one term, U itself. One may compute $\mathcal{D}_{n(d-1)+1}(Y)^{-1}$ in $\tilde{\mathcal{O}}(\binom{n(d-1)}{n-1}^4)$ operations in $\mathbb{Z}_q$, by Lemma 11. The product $\mathcal{D}_{n(d-1)+1}(Y)^{-1}u$ can then be computed in $\mathcal{O}(\binom{n(d-1)}{n-1}^2)$ operations with polynomials of degree bounded by $\max(\binom{n(d-1)}{n-1}-1, \deg_Y(U))$ in the ring $\mathbb{Z}_q[\pi]$. This gives the claimed complexity bounds for $\delta = n(d-1)+1$.

Assume now that $\delta \leq n(d-1)$. A similar approach allows one to compute polynomials $A_{ij}(Y)$ such that

$$\frac{1}{D_\delta(Y)} \sum_{i=1}^{n} \left(\sum_{X^j \in T_{\delta-d,i}} A_{ij}(Y) X^j \right) f_i(X_1, \ldots, X_n, Y)$$

and U agree in monomials $X^u \in \bar{T}_\delta$. This gives the claimed expression for U, where $c_k(Y)$ is the coefficient of X^k in U itself, and the other coefficient of X^k the negative of that which appears in the above polynomial. Again, this can be computed using Lemma 11 and the complexity and degree bounds are straightforward. (Note that the binomial coefficient which appears in the proposition is just $\#T_\delta = \binom{\delta-1}{n-1}$, which is a bound on the size of $\#\bar{T}_\delta$.) □

5.2 Reduction by differential operators

In this section, we continue to use the notation X^m to denote a monomial $X_1^{m_1} \ldots X_n^{m_n}$ for $m = (m_1, \ldots, m_n) \in \mathbb{Z}^n_{\geq 0}$. We explain how to compute standard representatives in the quotient module

$$\mathcal{L}_Y^o / \sum_{i=1}^{n} D_{i,Y}(\mathcal{L}_Y^{(i)}).$$

Recall that the spaces $\mathcal{L}_Y^o$, $\mathcal{L}_Y^{(i)}$ and the quotient above are modules over $\mathcal{R}$, the ring of rational functions in $\mathbb{Q}_q(\pi)(Y)$ regular at the origin. For the purposes of the computations in Section 6.1 and theory in Section 7.1, it is enough to consider the case of coefficients which are polynomials in Y. (For the purposes of the theory in the proof of Proposition 24, it is enough to just consider monomials — the operators $D_{i,Y}$ are additive and linear over $\mathcal{R}$ anyway.)

Let $V \in \mathcal{L}_Y^o$ be of the form $V = (\pi X_0)^\delta U(X_1, \ldots, X_n)$, where $U \in \mathbb{Z}_q[\pi, Y][X_1, \ldots, X_n]^o$ is homogeneous of degree $d\delta$ in the variables $X_1, \ldots, X_n$. First, consider the case $d\delta \geq n(d-1)+1$. By Proposition 14 we have that

$$\begin{aligned} (\pi X_0)^\delta U &= \frac{(\pi X_0)^\delta}{D_{n(d-1)+1}(Y)} \sum_{i=1}^{n} \left(\sum_{|j|=d(\delta-1)} A_{ij}(Y) X^j \right) f_i(X_1, \ldots, X_n, Y) \\ &= \frac{(\pi X_0)^{\delta-1}}{D_{n(d-1)+1}(Y)} \left\{ \sum_{i=1}^{n} D_{i,Y}(A_i(X_1, \ldots, X_n, Y)) - X_i \frac{\partial A_i(X_1, \ldots, X_n, Y)}{\partial X_i} \right\}, \end{aligned}$$

where $A_i(X_1, \ldots, X_n, Y) = \sum_{|j|=d(\delta-1)} A_{ij}(Y) X^j$. We have that $(\pi X_0)^{\delta-1} A_i \in \mathcal{L}_Y^{(i)}$. Thus modulo $\sum_{i=1}^n D_{i,Y}(\mathcal{L}_Y^{(i)})$, we see that V is equivalent to

$$-\frac{(\pi X_0)^{\delta-1}}{D_{n(d-1)+1}(Y)} \sum_{i=1}^n X_i \frac{\partial A_i(X_1, \ldots, X_n, Y)}{\partial X_i}.$$

This is homogeneous of degree $d(\delta-1)$ in the variables $X_1, \ldots, X_n$ and lies in $\mathcal{L}_Y^o$. The case $d\delta \leq n(d-1)$ is similar. Specifically, we find that V is equivalent modulo the subspace to

$$-\frac{(\pi X_0)^{\delta-1}}{D_{d\delta}(Y)} \sum_{i=1}^n X_i \frac{\partial A_i(X_1, \ldots, X_n, Y)}{\partial X_i} + (\pi X_0)^{\delta} \sum_k \left(\frac{B_k(Y)}{D_{d\delta}(Y)} + c_k(Y) \right) X^k.$$

Here the first sum is homogeneous of degree $d(\delta-1)$ and lies in $\mathcal{L}_Y^o$. The second sum is over elements in the basis set. Specifically, the sum k runs over vectors $(k_1, \ldots, k_n)$ with $k_1 + \ldots + k_n = d\delta$ and $0 < k_1, \ldots, k_n < d$.

These reduction steps may be iterated until we obtain a sum of basis elements. The complexity estimates in Proposition 14 yield identical estimates for the complexity of performing each reduction step.

6 Auxiliary Routines: Steps 2, 3 and 4

In this section we explain how to compute the matrix $B(Y)$, numerically solve the system (1), and invert the matrix $C(Y^p)$. We shall estimate the complexity of these steps in p-adic ring operations, ignoring for the time being that we are actually working modulo certain powers of p.

6.1 Construction of the differential system

The (u, v)th entry of $B(Y)$ is defined to be the coefficient of $\pi^{u_0} X^u$ in the reduction of $\pi X_0 h \times \pi^{v_0} X^v$ modulo the operators $D_{i,Y}$ for $1 \leq i \leq n$. Using the reduction method from Section 5.2 this matrix may be found, as we now explain. The product $\pi X_0 h \times \pi^{v_0} X^v$ has degree in $X_1, \ldots, X_n$ at most $n(d-1)+d$. One first precomputes $\mathcal{D}_\delta(Y)^{-1}$ for $\delta = n(d-1)+1$, and all $\delta \leq n(d-1)$ with $d|\delta$. There are $\lfloor n(d-1)/d \rfloor + 1 \leq n$ such matrices, and finding each one requires $\tilde{\mathcal{O}}(\binom{n(d-1)}{n-1}^4)$ operations in $\mathbb{Z}_q$, by Lemma 11. The degree in Y of $\pi X_0 \times \pi^{v_0} X^v$ is zero. On each reduction step, a factor $(\pi X_0 D_\delta(Y))^{-1}$ is introduced on the denominator. Here δ is the minimum of $n(d-1)+1$ and the degree in $X_1, \ldots, X_n$ of the polynomial being reduced. On each reduction step the degree in Y is increased by less than $\binom{\delta-1}{n-1} \leq \binom{n(d-1)}{n-1}$. One needs $\lfloor (n(d-1)+d)/d \rfloor \leq n$ reduction steps, and so the degree in Y which occurs in any polynomial is less than $n\binom{n(d-1)}{n-1}$. Thus each reduction step can be computed in $\tilde{\mathcal{O}}(n\binom{n(d-1)}{n-1}^3)$ operations in $\mathbb{Z}_q$, by the second term in the complexity estimate in Proposition 14. (Note that in the first step one needs to reduce a polynomial of degree at most $n(d-1)+d$, rather than $n(d-1)+1$ as in Proposition 14; however, because of the factor X^v, one can easily reduce the degree to $n(d-1)+1$ on the first step by taking a single monomial out as a common factor of all terms.) This gives a complexity of $\tilde{\mathcal{O}}(n^2\binom{n(d-1)}{n-1}^3)$ operations in $\mathbb{Z}_q$ to find one column in $B(Y)$, ignoring the precomputation of matrix inverses. There are $\#\mathcal{B} < \binom{n(d-1)}{n-1}$ columns, and so one obtains an overall complexity of $\tilde{\mathcal{O}}(n^2\binom{n(d-1)}{n-1}^4)$ operations in $\mathbb{Z}_q$ to find $B(Y)$, including the precomputation.

Note that each coefficient in the matrix $B(Y)$ is of the form

$$\frac{A(Y)}{D_{n(d-1)+1}(Y)\prod_{\delta\le n(d-1),\, d|\delta} D_\delta(Y)} \tag{4}$$

where $A(Y) \in \mathbb{Z}_q[Y]$ has degree less than

$$\binom{n(d-1)}{n-1} + \sum_{\delta\le n(d-1),\, d|\delta} \binom{\delta-1}{n-1} \le n\binom{n(d-1)}{n-1}.$$

The entries in $A(Y)$ lie in $\mathbb{Z}_q$, rather than just $\mathbb{Q}_q(\pi)$, since exactly the correct power of π is divided out during the reduction. Thus all calculations involve polynomials over $\mathbb{Z}_q$ rather than $\mathbb{Z}_q[\pi]$.

6.2 Solution of the system around the origin

In this section we define $\Lambda = n\binom{n(d-1)}{n-1}$, to simplify the notation. To solve the differential system (1), we write $B(Y)$ in the form

$$B(Y) = \frac{\sum_{k=0}^{\Lambda-1} B_k Y^k}{\sum_{j=0}^{\Lambda} e_j Y^j}.$$

Here B_k are matrices of size $\#\mathcal{B} \times \#\mathcal{B}$ over $\mathbb{Z}_q$, and $e_j \in \mathbb{Z}_q$ are the coefficients of the polynomial on the denominator of (4). We write $C(Y) = \sum_{\ell=0}^{\infty} C_\ell Y^\ell$ with the C_ℓ matrices. The equation $dC/dY = C(Y)B(Y)$ can then be written in the form

$$\left(\sum_{j=0}^{\Lambda} e_j Y^j\right)\left(\sum_{\ell=1}^{\infty} \ell C_\ell Y^{\ell-1}\right) = \left(\sum_{\ell=0}^{\infty} C_\ell Y^\ell\right)\left(\sum_{k=0}^{\Lambda-1} B_k Y^k\right).$$

From Lemma 13 the constant term $e_0 \neq 0 \bmod p$. Equating coefficients of $Y^{\ell-1}$ on both sides for $\ell \ge 1$, and using the fact $C_0 = I$, allows us to solve for each ℓC_ℓ as

$$\ell C_\ell = -e_0^{-1}\left\{\left(\sum_{j=1}^{\min(\Lambda,\,\ell-1)} e_j(\ell-j)C_{\ell-j}\right) + \left(\sum_{k=0}^{\min(\Lambda-1,\,\ell-1)} C_{\ell-1-k}B_k\right)\right\}.$$

It follows that the matrix coefficient C_ℓ of Y^ℓ in this expansion of $C(Y)$ is such that $\ell! C_\ell$ has entries in $\mathbb{Z}_q$. This gives a linear lower bound on the p-adic order of the coefficients in $C(Y)$. The complexity of computing the expansion of $C(Y)$ modulo (Y^{N_Y}) is $\mathcal{O}(\Lambda N_Y)$ operations in the ring of $\#\mathcal{B} \times \#\mathcal{B}$ matrices over $\mathbb{Q}_q$.

6.3 Matrix Inversion

The matrix $C(Y^p)^{-1}$ can be computed by inverting $C(Y)$ and substituting Y for Y^p. This can be done using Newton iteration with quadratic convergence, with respect to the ideal (Y) in $\mathbb{Q}_q[Y]$. See [14, Section 5.2.2] for details.

6.4 An Example

We now look at a simple but non-trivial example, to illustrate the method of computing $B(Y)$ which was described in Section 6.1. Consider the Hesse family of elliptic curves described by the polynomial $f = X_1^3 + X_2^3 + X_3^3 - 3YX_1X_2X_3$. We have $f_i = 3(X_i^3 - YX_1X_2X_3)$ for $i = 1, 2, 3$. The polynomial f is defined over the integers. All elliptic curves over finite fields of characteristic not equal to three can be defined by reducing the equation $f = 0$ modulo the ideal $(p, Y - y)$ for some $p \neq 3$ and algebraic integer y with $y^3 \neq 1$. We describe how to compute the matrix $B(Y)$ in this example. First note that according to our presentation, one should first fix a prime p, and then take Teichmüller liftings of the coefficients of f in $\mathbb{Z}_p$. Lifting the coefficient 3 in the final monomial would complicate the presentation, and so we shall just leave it as it is. (This can be justified by replacing the variable Y by $Z = 3Y$ throughout the calculations, and switching back for Y at the end.) All of our calculations will take place with rational numbers, and will in fact be independent of the choice of prime p.

We have $d = n = 3$ and so the basis $\mathcal{B}$ of the quotient space $\mathcal{L}_Y^o / \sum_{i=1}^3 D_{i,Y}(\mathcal{L}_Y^{(i)})$ is the set $\{\pi X_0X_1X_2X_3, (\pi X_0X_1X_2X_3)^2\}$. We have $h(X) = -3X_1X_2X_3$ and so the first column of $B(Y)$ contains $-3\pi X_0X_1X_2X_3 \times \pi X_0X_1X_2X_3$, written in terms of the basis elements. This product is a multiple of a basis element, so the first column is just the transpose of $(0, -3)$. To compute the second column we need to reduce $-3\pi X_0X_1X_2X_3 \times (\pi X_0X_1X_2X_3)^2$ modulo the differential operators $D_{i,Y}$. The special form of the polynomial f makes it quite easy to do this by hand, resulting in the following matrix c.f. [9, Page 287].

$$\begin{pmatrix} 0 & -\frac{Y}{3(1-Y^3)} \\ -3 & \frac{3Y^2}{1-Y^3} \end{pmatrix}.$$

We now explain how one can calculate the second column using the systematic approach of Section 5.2. To reduce $-3(\pi X_0X_1X_2X_3)^3$ we first consider the monomial $X_1^3X_2^3X_3^3$. We need to write this monomial as a linear combination of the polynomials f_1, f_2 and f_3, with coefficients from the appropriate ring. The total degree of this monomial is 9 which exceeds $n(d-1)+1 = 7$. Thus we should first remove a monomial of degree 2, say X_3^2. The monomial $X_1^3X_2^3X_3$ can be written as a linear combination of the polynomials f_i, since we know from the theory that any monomial of degree 7 can be. To do this, we construct the matrix $\mathcal{D}_7(Y)$. The rows of this matrix are labelled by the 15 monomials of degree 7 which are divisible by $X_1X_2X_3$. The indices of the columns are pairs (X^v, f_i) where $X^v \in T_{7-3,i}$ and $i = 1, 2, 3$. Ignoring for the moment the second element in each labelling pair we have that: The first six columns are labelled by monomials of degree 4 which are divisible by X_2X_3; the second five columns by monomials of degree 4 which are not divisible by X_1^3 but are divisible by X_1X_3; and the final four columns by monomials of degree 4 which are not divisible by X_1^3 or X_2^3, but are divisible by X_1X_2. Writing the exponents of the monomials as vectors, the indices of the columns are

$$(2,1,1),(1,2,1),(0,3,1),(0,2,2),(0,1,3),(1,1,2);$$

$$(2,1,1),(2,0,2),(1,2,1),(1,1,2),(1,0,3);$$

$$(2,2,0),(2,1,1),(1,2,1),(1,1,2)$$

The column with label (X^v, f_i) for v an exponent in the ith row immediately above, gives the coefficients of the polynomial $X^v f_i$. For example, $X_1^2X_2X_3(3X_1^3 - 3YX_1X_2X_3) = 3X_1^5X_2X_3 -$

$3YX_1^3X_2^2X_3^2$ and so the first column contains 3 in the row labelled by $X_1^5X_2X_3$ and $-3Y$ in the row labelled by $X_1^3X_2^2X_3^2$. Indeed, every column just contains one 3 and one $-3Y$ with the remaining entries zero. The determinant of this matrix is $3^{15}(Y^3-1)^3$ (this was an easy calculation for the computer algebra package Magma). However, every non-zero entry in the adjugate matrix is of one of the following forms: $3^{14}(Y^3-1)^3$ or $-3^{14}Y^i(Y^3-1)^2$ for $i=0,1,2,3$. Thus the non-zero entries of the inverse matrix are of the form $1/3$ or $-Y^i/3(Y^3-1)$ for $i=0,1,2,3$. Multiplying the inverse matrix by the column representing $X_1^3X_3^3X_3$, we get a column with three non-zero entries. This leads us to the equation

$$3(1-Y^3)X_1^3X_2^3X_3 = f_1(X_2^3X_3) + f_2(YX_1X_2X_3^2) + f_3(Y^2X_1^2X_2^2).$$

We can then write $-3(\pi X_0X_1X_2X_3)^3$ as

$$\frac{-1}{(1-Y^3)}\Big\{D_{1,Y}((\pi X_0)^2X_2^3X_3^3) + D_{2,Y}(Y(\pi X_0)^2X_1X_2X_3^4) + D_{3,Y}(Y^2(\pi X_0X_1X_2X_3)^2)$$

$$-Y(\pi X_0)^2X_1X_2X_3^4 - 2Y^2(\pi X_0X_1X_2X_3)^2\Big\}. \tag{5}$$

The last term is a multiple of a basis element, and so we need only reduce the second last term. We do this using the matrix $\mathcal{D}_6(Y)$. This 9×9 matrix has rows labelled by all monomials in three variables of degree 6 which are divisible by $X_1X_2X_3$, excluding $X_1^2X_2^2X_3^2$. Specifically, we shall order the exponents of such monomials in the following way:

$$(4,1,1),(3,1,2),(3,2,1),(1,4,1),(1,3,2),(2,3,1),(1,1,4),(1,2,3),(2,1,3).$$

The columns are labelled by pairs (X^v, f_i) for all v in the ith row below:

$$(1,1,1),(0,1,2),(0,2,1);$$

$$(1,1,1),(1,0,2),(2,0,1);$$

$$(1,1,1),(1,2,0),(2,1,0).$$

With these orderings of rows and columns the matrix $\mathcal{D}_6(Y)$ is

$$\begin{pmatrix}
3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0\\
0 & 3 & 0 & 0 & 0 & -3Y & 0 & 0 & 0\\
0 & 0 & 3 & 0 & 0 & 0 & 0 & 0 & -3Y\\
0 & 0 & 0 & 3 & 0 & 0 & 0 & 0 & 0\\
0 & 0 & -3Y & 0 & 3 & 0 & 0 & 0 & 0\\
0 & 0 & 0 & 0 & 0 & 3 & 0 & -3Y & 0\\
0 & 0 & 0 & 0 & 0 & 0 & 3 & 0 & 0\\
0 & -3Y & 0 & 0 & 0 & 0 & 0 & 3 & 0\\
0 & 0 & 0 & 0 & -3Y & 0 & 0 & 0 & 3
\end{pmatrix}.$$

This matrix has determinant $3^9(Y^3-1)^2$, and all non-zero terms in the adjugate are of the form $3^8(Y^3-1)^2$ or $-3^8Y^i(Y^3-1)$ for $i=0,1,2$. The required expression for $X_1X_2X_3^4$ is obtained by picking out the seventh column from the inverse matrix. This column has only one non-zero

entry, 1/3 in the seventh position. This tells us that $f_3(\frac{1}{3}X_1X_2X_3)$ and $X_1X_2X_3^4$ have the same coefficients in all monomials except possibly $X_1^2X_2^2X_3^2$. Indeed

$$X_1X_2X_3^4 = \frac{1}{3}f_3(X_1X_2X_3) + YX_1^2X_2^2X_3^2.$$

We can now reduce $(\pi X_0)^2X_1X_2X_3^4$ as

$$\frac{1}{3}D_{3,Y}(\pi X_0X_1X_2X_3) + Y(\pi X_0X_1X_2X_3)^2 - \frac{1}{3}\pi X_0X_1X_2X_3.$$

Combining this equation with (5) shows that modulo the differential operators $-3(\pi X_0X_1X_2X_3)^3$ and $-\frac{Y}{3(1-Y^3)}(\pi X_0X_1X_2X_3) - \frac{3Y^2}{1-Y^3}(\pi X_0X_1X_2X_3)^2$ agree, which gives us the second column in $B(Y)$. This completes our example illustrating the general reduction method for computing the matrix $B(Y)$.

7 p-adic estimates

In this section we derive various lower bounds on the p-adic orders of the expressions which arise in our algorithm.

7.1 Domain of holomorphy of the generic absolute Frobenius matrix

The purpose of this section is to give a "lower bound" on the domain of holomorphy of the entries in the generic absolute Frobenius matrix. (We recall the precise meaning of this phrase in the context of p-adic analysis in Section A.4.) This lower bound will be of crucial importance when we try to reconstruct the function $\alpha(Y)$ modulo a power of p from its local expansion around the origin.

Define

$$\mathfrak{g}(Y) := \prod_{1\leq\delta\leq n(d-1),\, d|\delta} D_\delta(Y) \tag{6}$$

$$R(Y) := D_{n(d-1)+1}(Y). \tag{7}$$

The polynomials $D_\delta(Y)$ are described in Definition 12. The notation $\mathfrak{g}(Y)$ is chosen to reflect that used by Dwork in the dual setting [9, Pages 256-257] [8, Equation (8)]. One presumes that the two uses are related, although Dwork's $\mathfrak{g}(Y)$ is not defined in an explicit manner. The notation "$R(Y)$" is used by Dwork for the resultant with respect to Y of $X_i\frac{\partial f}{\partial X_i}$, for $1 \leq i \leq n$. This resultant characterises exactly those hypersurfaces whose intersections with all coordinate subspaces are smooth. This polynomial is certainly different from $D_{n(d-1)+1}(Y)$, but the two are related: The matrix $\mathcal{D}_{n(d-1)+1}(Y)$ is precisely the one which occurs in Macaulay's construction of the resultant which respect to Y of $\frac{\partial f}{\partial X_i}$, for $1 \leq i \leq n$ [17, Sections 6-8]. (We worked with $X_i\frac{\partial f}{\partial X_i}$, and insisted upon certain divisibility properties by $X_1 \ldots X_n$ and $X_1 \ldots X_{i-1}X_{i+1} \ldots X_n$, but the matrix obtained is the same as Macaulay's "D", taking "$F_i = \frac{\partial f}{\partial X_i}$" and "$l_i = d-1$".) The resultant of $\frac{\partial f}{\partial X_i}$ for $1 \leq i \leq n$ is the ratio $D_{n(d-1)+1}(Y)/\Delta(Y)$, where $\Delta(Y)$ is an explicit minor of the matrix $\mathcal{D}_{n(d-1)+1}(Y)$ [17, Section 8]. The resultant does not vanish modulo p precisely for those $\bar{y}$ such that $\frac{\partial f(\bar{y})}{\partial X_i}$, $1 \leq i \leq n$, have no common projective zero. Since $d \neq 0 \bmod p$, Euler's relation $df = \sum_{i=1}^n X_i\frac{\partial f}{\partial X_i}$, tells us that this resultant does not vanish precisely when $\bar{f}(X, \bar{y})$ defines a

smooth projective hypersurface. Thus if our $R(\bar{y}) \neq 0 \bmod p$, we know that $\bar{f}(X, \bar{y}) = 0$ is smooth, but the converse is not necessarily true because of the presence of the extraneous factor $\Delta(Y)$.

We first examine the reduction formulae in Section 5.2. Note that we will use the operators D_{i,Y^p} rather than $D_{i,Y}$, and so one must replace Y by Y^p in the formulae. Consider a term $\mathfrak{g}(Y^p)a_m(Y)X^m$ where $m \in \mathcal{I}^o$ and $a_m(Y) \in \mathbb{Q}_q(\pi)[Y]$. The coefficients of the monomials X^u, for $u \in \mathcal{B}$, in the reduction of the term $\mathfrak{g}(Y^p)a_m(Y)X^m$ modulo the operators D_{i,Y^p}, for $1 \leq i \leq n$, are rational functions in $\mathbb{Q}_q(\pi)[Y, 1/R(Y^p)]$. Here the factor $\mathfrak{g}(Y^p)$ cancels any factor $1/\mathfrak{g}(Y^p)$ which may be introduced in the last few steps of the reduction process. (The actual basis elements are $\pi^{u_0}X^u$, but we shall just ignore the factor π^{u_0} initially.) These coefficients have p-adic order

$$\geq \mathrm{ord}(a_m(Y)) - \frac{m_0}{p-1}, \tag{8}$$

since at most a power of π^{-1} can be introduced on each reduction step. By the p-adic order of $a_m(Y)$ we mean the minimum order among the coefficients of non-zero terms. The degree in Y of the numerator is

$$\leq \deg_Y(a_m(Y)) + p\deg_Y(\mathfrak{g}(Y)) + m_0 p\left\{\binom{n(d-1)}{n-1} - 1\right\}, \tag{9}$$

since at most a polynomial of degree less than $\binom{n(d-1)}{n-1}$ in Y^p is introduced on each reduction step. Similarly, the power of $R(Y^p)$ in the denominator is certainly

$$\leq m_0. \tag{10}$$

The next proposition gives our estimate for the domain of holomorphy of the generic Frobenius matrix $\alpha(Y)$. For our purposes, it is convenient to state this via lower bounds on the decay rate of coefficients in a natural expansion of $\mathfrak{g}(Y^p)\alpha(Y)$.

Proposition 15 *Let $\alpha(Y) = (\alpha_{u,v}(Y))$ where u and v run over the set $\mathcal{B}$. One may write*

$$\mathfrak{g}(Y^p)\alpha_{u,v}(Y) = \sum_{i=0}^{\infty} A_{u,v}^{(i)} Y^i + \sum_{j=1}^{\infty} \frac{B_{u,v}^{(j)}(Y)}{R^{\tau^{-1}}(Y)^j}$$

where $A_{u,v}^{(i)} \in \mathbb{Q}_q(\pi)$ and $B_{u,v}^{(j)}(Y) \in \mathbb{Q}_q(\pi)[Y]$ with $\deg_Y(B_{u,v}^{(j)}) < \deg_Y(R)$. Moreover, we have the following lower bounds on the p-adic orders.

$$\mathrm{ord}(A_{u,v}^{(i)}) \geq \frac{1}{6p(ed)^{n-1}}i - n,\ \mathrm{ord}(B_{u,v}^{(j)}) \geq \frac{1}{6p}j - n.$$

Here e is the base of the natural logarithms. Thus the entries in $\mathfrak{g}(Y^p)\alpha(Y)$ have p-adic order at least $-n$.

Let $M \geq 1$ be an integer. Then "modulo" p^M the matrix $\mathfrak{g}(Y^p)\alpha(Y)$ contains rational functions in $\mathbb{Q}_q(\pi)[Y, 1/R(Y)]$ of p-adic order at least $-n$, with numerator of degree less than $12p(ed)^{n-1}(M+n)$ and denominator a power of $R^{\tau^{-1}}(Y)$ less than $6p(M+n)$.

Here by "modulo" p^M we mean that one truncates the p-adic expansions of the entries in $\mathfrak{g}(Y^p)\alpha(Y)$ after the power p^{M-1}.

Proof: The matrix $\alpha(Y)$ is defined in Section A.2 and we use some of the notation from that section. Write $F(X,Y) = \sum_m G_m(Y)X^m$ where the sum is over vectors $m \in \mathcal{I}$. We first need to understand the degrees and p-adic orders of the polynomials $G_m(Y)$. By definition we have

$$F = \prod_{i=1}^{n} \theta(a_i X_0 X_i^d) \prod_{j \in J} \theta(Y b_j X_0 X^j).$$

Here $\bar{h} = \sum_{j\in J} \bar{b}_j X^j$, and $h = \sum_{j\in J} b_j X^j$ with b_j the Teichmüller lifting of $\bar{b}_j$. Recall that $\theta(z) = \exp(\pi(z - z^p))$ and write $\theta(z) = \sum_{r=0}^{\infty} \lambda_r z^r$. Write $\theta(Y b_j X_0 X^j) = \sum_m a_{j,m}(Y) X^m$ with the sum over $m = (m_0, \ldots, m_n) \in \mathcal{I}$. We see that $a_{j,m}(Y) = \lambda_{m_0}(Y b_j)^{m_0}$ for $m = m_0(1, j_1, \ldots, j_n)$ where $j = (j_1, \ldots, j_n)$, and $a_{j,m}(Y) = 0$ when $m \neq m_0(1, j_1, \ldots, j_n)$. Thus $\deg_Y(a_{j,m}) \leq m_0$ for all $m \in \mathcal{I}$. Hence writing $\prod_{j\in J} \theta(Y b_j X^j) = \sum_m c_m(Y) X^m$, by [14, Lemma 16] we have $\deg_Y(c_m) \leq m_0$. Certainly writing $\prod_{i=1}^{n} \theta(a_i X_i^d) = \sum_m d_m(Y) X^m$ we have $\deg_Y(d_m) = 0$. Thus by [14, Lemma 16]

$$\deg_Y(G_m) \leq m_0. \tag{11}$$

We have the estimate $\mathrm{ord}(\lambda_r) \geq \frac{p-1}{p^2} r$ [7, Pages 55-57]. Thus the p-adic order of each $a_{j,m}(Y)$ is at least $\frac{(p-1)m_0}{p^2}$. Similarly, the coefficients in $\prod_{i=1}^{n} \theta(a_i X_i^d)$ satisfy the bound $\mathrm{ord}(d_m(Y)) \geq \frac{(p-1)m_0}{p^2}$. By [14, Lemma 16] this gives a lower bound of

$$\mathrm{ord}(G_m) \geq \frac{(p-1)m_0}{p^2}. \tag{12}$$

We now have all the results required on the degrees and p-adic orders of the coefficients of $F(X,Y)$ itself.

The vth column of the matrix $\mathfrak{g}(Y^p)\alpha(Y)$ has uth entry the coefficient of $\pi^{u_0} X^u$ in the reduction modulo $\sum_{i=1}^{n} D_{i,Y^p}(L_Y^{(i)})$ of $\mathfrak{g}(Y^p)\psi_p(F\pi^{v_0}X^v)$. Now $F\pi^{v_0}X^v = \pi^{v_0}\sum_m G_{m-v}X^m$. Hence $\mathfrak{g}(Y^p)\psi_p(F\pi^{v_0}X^v) = \mathfrak{g}(Y^p)\pi^{v_0}\sum_m \tau^{-1}(G_{pm-v})X^m$. We need to understand the reduction of each term $\mathfrak{g}(Y^p)\tau^{-1}(G_{pm-v})X^m$ in this series. (Note that the action of τ^{-1} is inconsequential since it fixes Y and does not change p-adic estimates.) With regard to the degree, by estimates (9) and (11) the coefficient of each basis monomial X^u in the reduction of this term has degree in Y at most

$$\deg_Y(G_{pm-v}) + p\deg_Y(\mathfrak{g}(Y)) + m_0 p\left\{\binom{n(d-1)}{n-1} - 1\right\}$$

$$\leq (pm_0 - v_0) + pn\binom{n(d-1)}{n-1} + m_0 p\left\{\binom{n(d-1)}{n-1} - 1\right\}.$$

By estimates (8) and (12) the p-adic order of the coefficient of each basis monomial in the reduction is at least

$$\mathrm{ord}(G_{pm-v}) - \frac{m_0}{(p-1)} \geq \frac{(p-1)(pm_0 - v_0)}{p^2} - \frac{m_0}{(p-1)}.$$

By estimate (10) the power of $R(Y^p)$ which occurs in the denominator of the coefficient of each basis monomial in the reduction is at most m_0. Now since $v \in \mathcal{B}$ we have $1 \leq v_0 \leq \lfloor n(d-1)/d \rfloor \leq n-1$. Thus $pm_0 - (n-1) \leq pm_0 - v_0 \leq pm_0 - 1$. From the estimate $n^n/n! \leq e^{n-1}$ we see that $\binom{n(d-1)}{n-1} \leq (e(d-1))^{n-1} < (ed)^{n-1}$, where e is the base of the natural logarithms. (Write $\binom{n(d-1)}{n-1} = (n(d-1))(n(d-1)-1)\ldots(n(d-1)-(n-2))/(n-1)!$ and forcibly extract a factor n

from each of the $n-1$ terms on the numerator.) For the degree upper bound we can therefore take the estimate $p(m_0+n)(ed)^{n-1}$. For the p-adic order lower bound we have the estimate

$$\geq m_0\left(\frac{p-1}{p}-\frac{1}{p-1}\right)-\frac{(n-1)(p-1)}{p^2}. \tag{13}$$

We take the lower bound of $-2n/9$ for the second term, and $1/6$ for the coefficient of m_0 in the first (recall that $p\geq 3$). Thus (13) is

$$\geq \frac{m_0}{6}-\frac{2n}{9}. \tag{14}$$

We now have the necessary inequalities on the reduction of each term in $\mathfrak{g}(Y^p)\psi_p(FX^v)$. It only remains to put these results together to obtain the desired lower bounds.

The coefficient of X^u in the reduction of $\mathfrak{g}(Y^p)\tau^{-1}(G_{pm-v})X^m$ modulo the operators D_{i,Y^p} is an element of $\mathbb{Q}_q(\pi)[Y,1/R(Y^p)]$. We write it in the form

$$A_{u,v,m}(Y)+\frac{B_{u,v,m}(Y)}{R(Y^p)^{m_0}}.$$

Here $A_{u,v,m}(Y)$ is a polynomial of degree at most

$$p(m_0+n)(ed)^{n-1}-pm_0\deg(R(Y))\leq p(m_0+n)(ed)^{n-1}. \tag{15}$$

By (14) and (15) the Newton polygon of the polynomial $A_{u,v,m}(Y)$ lies on or above the graph

$$y=\frac{1}{6p(ed)^{n-1}}x-\frac{7n}{18}. \tag{16}$$

(Specifically, we put together the inequalities $y\geq(m_0/6)-(2n/9)$ and $m_0\geq(x/p(ed)^{n-1})-n$.) The polynomial $B_{u,v,m}(Y)$ has degree strictly less than that of the denominator $R(Y^p)^{m_0}$. This fact, along with (14), shows that the Newton polygon of $B_{u,v,m}(Y)$ lies on or above the graph

$$y=\frac{1}{6p\deg_Y(R(Y))}x-\frac{2n}{9}. \tag{17}$$

Write

$$\pi^{u_0-v_0}\mathfrak{g}(Y^p)\alpha_{u,v}(Y)=\sum_{i=0}^{\infty}\tilde{A}_{u,v}^{(i)}Y^i+\sum_{j=1}^{\infty}\frac{\tilde{B}_{u,v}^{(j)}(Y)}{R(Y^p)^j} \tag{18}$$

where $\tilde{A}_{u.v}^{(i)}\in\mathbb{Q}_q(\pi)$ and $\tilde{B}_{u,v}^{(j)}\in\mathbb{Q}_q(\pi)[Y]$ with $\deg_Y(\tilde{B}_{u,v}^{(j)})<p\deg_Y(R)$. From (16) and (17) it follows that

$$\mathrm{ord}(\tilde{A}_{u,v}^{(i)})\geq\frac{1}{6p(ed)^{n-1}}i-\frac{7n}{18},\ \mathrm{ord}(\tilde{B}_{u,v}^{(j)})\geq\frac{1}{6}j-\frac{2n}{9}. \tag{19}$$

The inequalities in Proposition 15 for $\mathrm{ord}(A_{u,v}^{(i)})$ and $\mathrm{ord}(B_{u,v}^{(j)})$ now follow easily, once we have rewritten our expression for $\mathfrak{g}(Y^p)\alpha_{u,v}(Y)$ in the correct form. To do this, write $R(Y^p)=R^{\tau^{-1}}(Y)^p-pS(Y)$ with $S(Y)\in\mathbb{Z}_q[Y]$, and observe that

$$\frac{1}{R(Y^p)}=\frac{1}{R^{\tau^{-1}}(Y)^p}\sum_{j=0}^{\infty}p^j\frac{S(Y)^j}{R^{\tau^{-1}}(Y)^{pj}}. \tag{20}$$

The power series on the righthand side decays quickly. Substitute in (18) for $1/R(Y^p)$ using equation (20). This gives an expansion of the "fractional part" of $\pi^{u_0-v_0}\mathfrak{g}(Y^p)\alpha_{u,v}(Y)$ in powers of $1/R^{\tau^{-1}}(Y)^p$ with the same decay rate as the expansion in $1/R(Y^p)$. Rewriting this new expansion in terms of powers of $1/R^{\tau^{-1}}(Y)$, an extra factor of p arises in the denominator of the cofactor of j in the decay rate estimate from (19). We also need to adjust for the factor $\pi^{u_0-v_0}$, and do this using the fact $u_0 - v_0 \leq n-2$, and so $-2n/9-(n-2)/(p-1) \geq -n$ and $-7n/18-(n-2)/(p-1) \geq -n$. This finishes the proof of the first part of the proposition.

The claim on the size of entries in the matrix "modulo" p^M follows easily. The proof is complete. □

7.2 Lower bounds on p-adic orders of intermediate results

We also need lower bounds on the p-adic order of the entries in the matrices $C(Y) \bmod Y^{N_Y}$, $\alpha(0)$, and the $\log_p(q)$th "power" of $\alpha(y^{\tau^{-1}})$ which occurs in the equation in Step 8. (We do not believe that these bounds are optimal.)

Lemma 16 *The polynomial entries in $C(Y) \bmod Y^{N_Y}$ have p-adic order at least $-(N_Y-1)/(p-1)$.*

Proof: Let $\ell > 0$. From Section 6.2 we know the matrix $\ell! C_\ell$ has entries in $\mathbb{Z}_q$. Now $\mathrm{ord}(\ell!) < \ell/(p-1)$, and so C_ℓ itself has entries of order at least $-\ell/(p-1)$. Since $C(Y) \bmod Y^{N_Y} = \sum_{\ell=0}^{N_Y-1} C_\ell Y^\ell$ the lemma follows. □

Note 17 Using the equation $C(Y) = \tau(\alpha(0)^{-1}C(Y^p)\alpha(Y))$ and "Dwork's trick", one can prove the much better lower bound

$$\mathrm{ord}(C_\ell) \geq -\lceil \log_p(\ell) \rceil \#\mathcal{B}\frac{4n+3}{2} - 1.$$

This shows $C(Y)$ actually converges on the open unit disk. Given a similar bound for the coefficients in $C(Y)^{-1}$, one can improve the complexity in Theorem 9 by a factor d^{n-1}; however, this requires certain information on $\alpha(Y)^{-1}$ which we have been unable to derive.

Lemma 18 *The entries in the matrix $\alpha(0)$ have p-adic order at least $-3n/2$.*

Proof: Following the proof of [14, Lemma 20], we first need a lower bound on

$$m\frac{(p-1)}{p^2} - S_r \tag{21}$$

where $m, r \geq 0$ are integers, with S_r the sum of the p-adic digits in r, and $d(m-pr) = pu_i - v_i$. Now $pu_i - v_i \geq p \times 1 - (d-1) = p - d + 1$ and so $m - pr \geq (p/d) - 1 + (1/d) > (p/d) - 1$. Therefore $r < (m/p) - (1/d) + (1/p)$ and so since r is an integer we must have $r \leq (m/p)$. So we find (21) is at least

$$r\frac{(p-1)}{p} - S_r \geq (p-1)\left(\frac{r}{p} - (\log_p(r)+1)\right).$$

For $r \geq 3p$ the second expression is non-negative. For $p \leq r < 3p$ one can show directly that the first expression is always at least zero (using the fact $p \geq 3$). For $r < p$ the first expression equals

$-r/p > -1$. It now follows from Equation (2) and [14, Inequality (11)] that $\pi^{u_0-v_0}\alpha(0)$ has entries of p-adic order $> -n$. Finally $u_0 - v_0 \leq n-2$ and $-n-(n-2)/(p-1) \geq -3n/2$, since $p \geq 3$. □

Lemma 19 *The entries in the matrix $\alpha(y^{\tau^{-1}})$ have p-adic order at least $-n$. Those in the product $\alpha(y^{\tau^{-1}})^{\tau^{r\log_p(q)-1}} \ldots \alpha(y^{\tau^{-1}})$ have p-adic order at least $-nr\log_p(q)$.*

Proof: The first claim follows from Proposition 15 and the second from the first, using the fact $\mathrm{ord}(y) = 0$ and that the map τ does not change p-adic orders. □

8 Analytic continuation of power series

8.1 The underlying problem

In Step 7 of the algorithm we encounter the problem of recovering the value of the matrix function $\alpha(Y)$ at a non-singular point on the unit disk, from its local expansion about the origin. The underlying problem here seems of some independent interest, and in this section we will present a simple solution to it.

Consider the following scenario. Let $T(Y) \in \mathbb{Z}_q[Y]$ with $T(0) \neq 0 \bmod p$. Let $\beta(Y)$ be a p-adic holomorphic function on some domain (see Section A.4 for a definition of this term). Assume that $\beta(Y)$ has a "Mittag-Leffler" expansion of the form

$$\beta(Y) = \sum_{j=0}^{\infty} b_j Y^j + \sum_{j=1}^{\infty} \frac{c_j(Y)}{T(Y)^j},\ b_j \in \mathbb{Z}_q[\pi],\ c_j \in \mathbb{Z}_q[\pi][Y],\ \deg_Y(c_j) < \deg_Y(T).$$

Suppose that we know $\mathrm{ord}(b_j) \geq rj$ and $\mathrm{ord}(c_j) \geq sj$ for some positive real numbers r and s, but the coefficients b_j and c_j themselves are unknown. Assume further that we have computed, modulo some power of Y and power of p, a formal expansion of $\beta(Y)$ around the origin

$$\beta(Y) \equiv \sum_{k=0}^{M_1-1} e_k Y^k \bmod (Y^{M_1}, p^{M_2}).$$

This equation is in the ring of formal power series $\mathbb{Z}_q[\pi][[Y]]$ modulo the ideal (Y^{M_1}, p^{M_2}). The question we wish to address is, to compute $\beta(y) \bmod p^M$ for some $M \geq 1$ and element $y \in \mathcal{O}_{\mathbb{C}_p}$ with $\mathrm{ord}(T(y)) = 0$, what values do we need to choose for M_1 and M_2? More specifically, we wish to compute the rational function

$$\beta_M(Y) := \sum_{j=0}^{\lfloor M/r \rfloor} b_j Y^j + \sum_{j=1}^{\lfloor M/s \rfloor} \frac{c_j(Y)}{T(Y)^j} \bmod p^M. \tag{22}$$

For then $\beta_M(y) \equiv \beta(y) \bmod p^M$ for any $y \in \mathcal{O}_{\mathbb{C}_p}$ with $\mathrm{ord}(T(y)) = 0$. Taking $M_2 = M$, and putting the righthand-side of (22) over a common denominator we have

$$\frac{S(Y)}{T(Y)^{\lfloor M/s \rfloor}} \equiv \sum_{k=0}^{M_1-1} e_k Y^k \bmod (Y^{M_1}, p^M). \tag{23}$$

Here $S(Y)$ is an unknown polynomial of degree at most $\deg_Y(T)\lfloor M/s \rfloor + \lfloor M/r \rfloor$. Recall the $T(0) \neq 0 \bmod p$ and so indeed $1/T(Y) \in \mathbb{Z}_q[\pi][[Y]]$. Essentially we have the local expansion of a rational function at the origin, where the poles are known with bounds on their orders. It is now a simple matter to recover the rational function: Multiplying both sides of (23) by $T(Y)^{\lfloor M/s \rfloor}$ we get

$$S(Y) \equiv T^{\lfloor M/s \rfloor} \sum_{k=0}^{M_1-1} e_k Y^k \bmod (Y^{M_1}, p^M). \tag{24}$$

Set $M_1 = \deg_Y(T)\lfloor M/s \rfloor + (\lfloor M/r \rfloor + 1)$. Recall that we know the coefficients e_k, and so we can compute a Y-adic expansion of the righthand-side of (24). This can be done in $\tilde{\mathcal{O}}(M_1)$ operations in $\mathbb{Z}_q[\pi]/(p^M)$, using a square-and-multiply algorithm for expanding the power of $T(Y)$, and fast polynomial multiplication, see [5]. Thus we can recover the unknown coefficients in $S(Y)$, as required.

8.2 Auxiliary Routines: Step 7

In Step 7 we are given as input the expansion of $\mathfrak{g}(Y^p)\alpha(Y)$ around the origin up to some p-adic and Y-adic accuracy, and we need to evaluate this at some Teichmüller point. The method in Section 8.1 applies directly. We take "$T(Y)$" as our $R^{\tau^{-1}}(Y)$ and "$\beta(Y)$" an entry in the matrix $p^n\mathfrak{g}(Y^p)\alpha(Y)$. (The factor p^n arises to account for the fact that the matrix entries might have small denominators.) Note that $R^{\tau^{-1}}(0) \neq 0 \bmod p$, by Lemma 13. Once we have the required expansion for all entries in $p^n\mathfrak{g}(Y^p)\alpha(Y)$, we can evaluate at a Teichmüller point $Y = \tau^{-1}(y)$ provided $\mathrm{ord}(R^{\tau^{-1}}(y^{\tau^{-1}})) = 0$, i.e., $R(\bar{y}) \neq 0 \bmod p$. Furthermore, if $\mathfrak{g}((\tau^{-1}(y))^p) \neq 0$, we can then recover $\alpha(\tau^{-1}(y))$ itself modulo the required power of p. Notice that if $\mathfrak{g}(y) = 0 \bmod p$, we will lose some accuracy here. As such, we just assume the stronger condition that $\mathfrak{g}(y) \neq 0 \bmod p$.

9 Error Analysis

In this section we use the bounds from Section 7 to show that the final output is correct.

9.1 Choice of final p-adic accuracy

We first justify the final p-adic accuracy modulo p^N to which the zeta function is computed. Our generic condition implies that $\bar{f}(X, \bar{y})$ defines a smooth hypersurface, as explained in Section 7.1. The Riemann hypothesis for smooth projective hypersurfaces then tells us that if $(1 - \alpha T)$ is a complex factor of $P(T)^{(-1)^{n+1}}$ then $||\alpha|| = q^{r(n-2)/2}$. Here $P(T)^{(-1)^{n+1}}$ is the unknown polynomial in the zeta function of $\bar{f}(X, \bar{y})$ where $\bar{y} \in \mathbb{F}_{q^r}$. (One could also use the more naive bound $||\alpha|| \leq q^{r(n+1)}$ on the complex absolute values.) The polynomial $P(T)^{(-1)^{n+1}}$ has degree $\#\mathcal{B} \leq (d-1)^{n-1}$. Hence its coefficients have absolute value at most $q^{r(n-2)(d-1)^{n-1}/2}2^{(d-1)^{n-1}}$. Thus computing $P(qT)^{(-1)^{n+1}}$ modulo p to the power

$$\lceil (d-1)^{n-1}((rn/2)\log_p(q) + \log_p(2)) \rceil + 1$$

is sufficient to recover the integer polynomial $P(T)$ exactly. We define $N = 2d^{n-1}rn\log_p(q)$, a simple upper bound on this number.

9.2 Reverse analysis of error propagation

Even though the final answer is a rational function with integer coefficients, the calculations followed in finding this involve matrices with possibly p-adically non-integral entries. We showed in Section 9.1 that a final "absolute error" of order p^N, in the p-adic sense, was sufficient to recover the zeta function exactly. At each step of the algorithm, the absolute error can increase by an amount depending upon the p-adic order of the matrices with which we are computing. In this section we analyse the "propagation of errors" through the algorithm. We determine an overall "p-adic accuracy" $p^{\tilde{N}}$ to which one can compute throughout the algorithm and be sure that the final error is of the correct magnitude. The analysis in this section closely follows the approach taken in [14, Section 8].

Our starting point is the observation that the polynomial $R(T)^{(-1)^{n+1}}$ computed in Step 8 has coefficients in $\mathbb{Z}_p$. We need to compute it modulo p^N. Writing $R(T)^{(-1)^{n+1}} = \det(1 - *T)$, as in [14, Section 8] we see that it is enough to compute $*$ modulo p^{N_1} where

$$N_1 = N + ((d-1)^{n-1} - 1)rn \log_p(q).$$

Here we use Lemma 19 to bound the order of the entries in $*$. Now $*$ is obtained by essentially raising $\alpha(y^{\tau^{-1}})$ to the "power" $r \log_p(q)$. Thus to determine $*$ modulo p^{N_1} it is enough by the first part of Lemma 19 to know $\alpha(y^{\tau^{-1}})$ modulo p^{N_2} where

$$N_2 = N_1 + nr \log_p(q) (= N + (d-1)^{n-1} rn \log_p(q)),$$

compare with [14, Section 8]. The matrix $\alpha(y^{\tau^{-1}})$ is obtained by evaluating $\alpha(Y)$ at an integral point, and so we need to find $\alpha(Y)$ modulo p^{N_2}. Putting $M = N_2$ in Proposition 15 we see that $\alpha(Y)$ modulo p^{N_2} is a matrix of rational functions whose numerators have degree less than

$$N_Y = 12p(ed)^{n-1}(N_2 + n).$$

Thus in Steps 3, 4, 6, and 7, it suffices to work modulo (Y^{N_Y}), as coefficients of any higher powers of Y cannot contribute to the unknown numerators in $\alpha(Y)$ modulo p^{N_2}.

We wish to compute $\alpha(Y)$ modulo p^{N_2}. A similar analysis to that in the relevant paragraphs of [14, Section 8] shows the following: It is enough to compute $C(Y)$ mod (Y^{N_Y}) with coefficients modulo p^{N_4} where

$$N_4 := N_3 + \lfloor 4N_Y/p(p-1) \rfloor,$$

with $N_3 := N_2 + \lfloor N_Y/(p-1) \rfloor + (3n/2)$. (It is convenient here to retain the numbering of the "N_*" in [14, Section 8]. The only difference with the analysis in [14, Section 8] is that we can drop a factor n in several places since we know $B(Y)$ has integral coefficients, and we replace $n(p+1)$ by $3n/2$.) This part of the analysis requires Lemma 18 and the bounds on the growth rates of the coefficients in $C(Y)$ alluded to in Section 6.2. Moreover, it suffices in Step 5 to compute the elements a_i^{-1} and $\tau^{-1}(a_i)$ modulo p to the power

$$N_2 + \lfloor N_Y/(p-1) \rfloor + \lfloor N_Y/p(p-1) \rfloor + 1 < N_4.$$

Next, we need to determine the p-adic accuracy required for the matrix $B(Y)$ to compute $C(Y)$ mod (Y^{N_Y}) with coefficients modulo p^{N_4}. We need to find C_ℓ for $0 \leq \ell < N_Y$ with coefficients modulo p^{N_4}. Since $B(Y)$ has p-adic integral elements, it follows immediately that it is enough to

compute the entries in $B(Y)$ modulo p^{N_4}. (This contrasts with [14, Section 8], where the matrix $B(Y)$ did not have integral elements.)

Finally, we determine the p-adic accuracy required in Step 2. Computing $B(Y)$ from the coefficients of $h(X)$ involves multiplication by a power of π at most $n+1$, followed by the removal of those powers of π. Thus we need to compute the coefficients of $h(X)$ modulo p^{N_5} where

$$N_5 := N_4 + \lceil (n-1)/(p-1) \rceil.$$

For simplicity we shall work to a common accuracy modulo $p^{\tilde{N}}$ throughout Steps 1-9 of the algorithm, for some $\tilde{N} \geq N_5$. Specifically, define

$$\tilde{N} := 171 d^{n-1} (ed)^{n-1} rn \log_p(q).$$

Observe also that

$$N_Y = 12p(ed)^{n-1}(N + (d-1)^{n-1} rn \log_p(q) + n).$$

We have shown that working to these accuracies is enough to determine the coefficients of the zeta function modulo p^N, and hence recover the zeta function exactly.

10 Complexity Analysis

We count the number of bit operations required in each of the steps of Algorithm 8. The analysis in this section is mostly parallel to [14, Section 9] except that n is often replaced by $n-1$ here. Let ω be the exponent of deterministic matrix multiplication over arbitrary rings, see [11, Section 12.1]. We use fast polynomial multiplication over matrix rings to get a time of $\tilde{\mathcal{O}}(\ell)$ ring operations to multiply two matrix polynomials of degree bounded by ℓ [5]. The calculations involve addition and multiplication in the p-adic rings $\mathbb{Q}_q(\pi)$ and $\mathbb{Q}_q(\pi, y)$ "modulo" $p^{\tilde{N}}$ during Steps 2-9 of the algorithm. The numbers we manipulate have explicit lower bounds on their p-adic orders, as given in Proposition 15 and Lemmas 16, 18 and 19. In particular, the p-adic order is certainly always at least $-\tilde{N}$. Addition and multiplication of such numbers in $\mathbb{Q}_q(\pi)$ and $\mathbb{Q}_q(\pi, y)$ "modulo" $p^{\tilde{N}}$ have the same complexity as in the rings $\mathbb{Z}_q[\pi]/(p^{\tilde{N}})$ and $\mathbb{Z}_q[\pi, y]/(p^{\tilde{N}})$, respectively, up to a constant factor. We take a Soft-Oh linear time bound for multiplication, division by units and addition in p-adic rings, see [11, Theorem 8.23]. Thus computing in $\mathbb{Q}_q(\pi)$ and $\mathbb{Q}_q(\pi, y)$ "modulo" $p^{\tilde{N}}$ requires $\tilde{\mathcal{O}}(\tilde{N} p \log(q))$ and $\tilde{\mathcal{O}}(\tilde{N} p \log(q) r)$ bits of time/space, respectively. Throughout this section e will denote the base of the natural logarithms, and we shall use the estimate $\binom{n(d-1)}{n-1} \leq (e(d-1))^{n-1} < (ed)^{n-1}$.

Step 1: This needs $\tilde{\mathcal{O}}(d^{n-1} \log(q))$ operations in $\mathbb{Z}_q/(p^{\tilde{N}})$, see [14, Section 9: Step 1].

Step 2: By the complexity estimate in Section 6.1, this requires $\tilde{\mathcal{O}}(n^2 (ed)^{4(n-1)})$ operations in $\mathbb{Z}_q/(p^{\tilde{N}})$.

Step 3: By the complexity estimate in Section 6.2, this requires $\tilde{\mathcal{O}}(n(ed)^{n-1} N_Y)$ operations in the ring of $\#\mathcal{B} \times \#\mathcal{B}$ matrices over $\mathbb{Q}_q$ "modulo" $p^{\tilde{N}}$. Since $\#\mathcal{B} < d^{n-1}$, this amounts to $\tilde{\mathcal{O}}(n d^{\omega(n-1)} (ed)^{n-1} N_Y)$ operations in $\mathbb{Q}_q$ "modulo" $p^{\tilde{N}}$.

STEP 4: As in [14, Section 9, STEP 4], this requires $\tilde{\mathcal{O}}(d^{\omega(n-1)}N_Y)$ operations in $\mathbb{Q}_q$ "modulo" $p^{\tilde{N}}$.

STEP 5: This is almost identical to [14, Section 9, STEP 5], the only difference being that the sum is over slightly different pairs (u, v). The complexity is $\tilde{\mathcal{O}}(d^{2(n-1)}n\tilde{N})$ operations in $\mathbb{Q}_q(\pi)$ "modulo" $p^{\tilde{N}}$.

STEP 6: Again, this is very similar to [14, Section 9, STEP 6], and the complexity is $\tilde{\mathcal{O}}(d^{\omega(n-1)}N_Y)$ operations in $\mathbb{Q}_q(\pi)$ "modulo" $p^{\tilde{N}}$.

STEP 7: The element $\tau^{-1}(y)$ is in the ring $\mathbb{Z}_q[y]$, and it can be found modulo $p^{\tilde{N}}$ in $\tilde{\mathcal{O}}(r\log(q))$ operations in $\mathbb{Z}_q[y]/(p^{\tilde{N}})$ by taking the Teichmüller lifting of $\bar{y}^{1/p}$. The polynomial $R(Y)$ can be found using Lemma 11 in $\mathcal{O}((ed)^{4(n-1)})$ operations in $\mathbb{Z}_q/(p^{\tilde{N}})$, and $R^{\tau^{-1}}(Y)$ can then be computed in a further $\tilde{\mathcal{O}}(\log(q)(ed)^{2(n-1)})$ operations in this ring. We need to recover the rational functions entries in $p^n\mathfrak{g}(Y^p)\alpha(Y)$ from their local expansions around the origin. These entries have the form $S(Y)/R^{\tau^{-1}}(Y)^{6p(\tilde{N}+n)}$, where $\deg_Y(R^{\tau^{-1}}(Y)) < (ed)^{n-1}$ and $\deg_Y(S(Y)) < 12pd(ed)^{n-1}(\tilde{N}+n)$. Using the soft-linear time method of Section 8, the recovery of all $(\#\mathcal{B})^2 = \mathcal{O}(d^{2(n-1)})$ entries can be done in $\tilde{\mathcal{O}}(N_Y d^{2(n-1)})$ operations in $\mathbb{Z}_q[\pi]$ modulo $p^{\tilde{N}+n}$. Evaluation at $\tau^{-1}(y)$ now requires $\mathcal{O}(d^{2(n-1)} \times p(ed)^{n-1}d^{n-1}rn\log_p(q))$ operations in $\mathbb{Z}_q[\pi, y]$ modulo $p^{\tilde{N}+n}$.

STEP 8: This is essentially the same as [14, Section 9, STEP 8], and requires $\mathcal{O}(d^{(\omega+1)(n-1)})$ operations in $\mathbb{Q}_q(\pi, y)$ "modulo" $p^{\tilde{N}}$.

STEP 9: Here one is just required to take the negative of $\mathcal{O}(d^{n-1})$ integers, and do some exact division by q, which is absorbed in the above estimates.

Now $N = \mathcal{O}(d^{n-1}rn\log_p(q))$, $\tilde{N} = \mathcal{O}(d^{n-1}(ed)^{n-1}rn\log_p(q))$ and $N_Y = \mathcal{O}(p(ed)^{n-1}d^{n-1}rn\log_p(q))$. Substituting these values, and using our estimates for computations in truncated p-adic fields, we get a total complexity for Algorithm 8 of

$$\tilde{\mathcal{O}}((\max\{d^{5+\omega}e^3, d^6e^5\})^{n-1}n^3p^2\log(q)^3r^3)$$

bit operations. The space complexity in bits is

$$\tilde{\mathcal{O}}(d^{6(n-1)}e^{4(n-1)}n^2p^2\log(q)^3r^2).$$

Here we have quadratic rather than cubic dependence on r: Step 7 requires less space than time in r. We recall that e is the base of the natural logarithms and ω is the exponent for deterministic matrix multiplication. Thus one has $2.376 \leq \omega \leq 3$, the precise value depending upon the method used [11, Section 12.1].

11 Completion of the proofs

Theorem 10 follows immediately from our complexity estimate and proof of correctness of Algorithm 8, and the following result. (Recall that $\Delta_{n,d}$ is defined in the paragraph immediately preceding Theorem 10.)

Lemma 20 *The polynomial $\Delta_{n,d}$ is not identically zero modulo p.*

Proof: In $\mathfrak{g}(Y)R(Y)$ the constant term in Y is non-zero and it contains the highest powers of $a_1, \ldots, a_n$, viewing these elements as independent variables. Precisely, by Lemma 13 it equals $\pm(da_1)^{t_1} \ldots (da_n)^{t_n}$ where $t_i = \sum_{1 \leq \delta \leq n(d-1),\, d|\delta} \#T_{\delta-d,i} + \#T_{n(d-1)+1-d,i}$. Any other term is of the form $a_1^{s_1} \ldots a_n^{s_n} b Y^j$ where b is a polynomial expression in the other variables, and $s_i \leq t_i$ for $1 \leq i \leq n$ with the inequality strict for at least one i. This follows from the fact that each column and each row of $\mathcal{D}_\delta(Y)$ contains exactly one entry da_i, as explained in the proof of Lemma 13. Hence setting $Y = 1$ in this polynomial, this "leading term" cannot be cancelled. Since p does not divide d we see $\mathfrak{g}(1)R(1)$ is not identically zero modulo p, and the lemma follows. □

Now since $d \geq 2$ we have that $(\max\{d^{5+\omega}e^3, d^6e^5\})^{n-1}n^3 = d^{\mathcal{O}(n)}$. Theorem 1 follows from Theorem 10 and this estimate. Here is the proof of Theorem 2: let X_f denote the affine hypersurface defined by f, $\tilde{X}_f$ its projective completion, and H_f denote the intersection of $\tilde{X}_f$ with the hyperplane at infinity. Assume that both $\tilde{X}_f$ and H_f, which are both projective hypersurfaces, satisfy the generic condition in Theorem 10 over $\mathbb{Z}$. Precisely, the homogenised version of f is not a zero of $\Delta_{n+1,d}$, with a similar condition holding for the homogeneous polynomial in n variables (the leading form of f) which defines H_f. This is just a new generic condition on f itself. Then f satisfies this new generic condition modulo p for all but finitely many primes. These exceptional primes, as well as those dividing the degree d and the prime $p = 2$, may be dealt with in constant time. So assume that f satisfies the new generic condition modulo p, and $p \neq 2$ with d not divisible by p. By Theorem 10, we can compute the number of rational points on these two hypersurfaces modulo p, say $\#\tilde{X}_f(\mathbb{F}_p)$ and $\#H_f(\mathbb{F}_p)$, in $\mathcal{O}(p^{2+\varepsilon})$ bit operations, since d and n are fixed. One can now recover the required number of rational points on X_f via the formula $\#\tilde{X}_f(\mathbb{F}_p) = \#X_f(\mathbb{F}_p) + \#H_f(\mathbb{F}_p)$. This completes the proofs of the results in the introduction.

Note 21 We conclude by making some comments on the the restrictions on p and d, and the generic requirement in Theorem 1.

The theorem excludes the cases $p = 2$ or $p|d$. There are two possible approaches to the case $p = 2$. One is to observe that it is likely the theory we have presented is true in the case $p = 2$; however, to prove this a much more careful analysis of the denominators introduced during the cohomological reduction process is required, in the manner of [16, Section 7.2]. It would be very interesting to carry out this analysis, and implement the algorithm in the case $p = 2$. The second approach for $p = 2$ is to use a more complicated "splitting function", in the manner of [16, Note 33]. The author has not studied the situation when p divides d. Monsky studies this case in detail in [18, Chapter 8], and one expects that his techniques may be of use.

Regarding our generic condition, it would be preferable if one could just insist that the projective hypersurface is smooth. There are a few small obstacles to this, all of which can be overcome.

First, we require that $\bar{a}_1 \ldots \bar{a}_n \neq 0$, so that the "diagonal fibre" in our family is smooth. Provided $q > d + 1$, one can enforce this condition in deterministic polynomial-time by making a non-singular linear change of variable. We omit the details of this simple algorithm, as there is an alternative approach. The assumption that $\bar{h}(X_1, \ldots, X_n)$ has no diagonal terms is never used in any essential manner, only to simplify certain arguments, e.g., in the proof of Lemma 13. This means that one can actually choose $\bar{a}_1, \ldots, \bar{a}_n$ completely freely all non-zero, and then define $\bar{h} = \bar{f} - \sum_{i=1}^n \bar{a}_i X_i^d$. This gives some flexibility in the choice of family through which one deforms $\bar{f}$ as well as removing the diagonal coefficients condition — thanks to Laurent Moret-Bailly for

raising this point.

The second restriction is that our algorithm requires $D_{n(d-1)+1}(1) \neq 0 \bmod p$, whereas we would like to work with just the condition that the resultant of $\frac{\partial f(Y)}{\partial X_i}$, $1 \leq i \leq n$, does not vanish modulo p at $Y = 1$. Call this resultant $r(Y)$. The resultant $r(Y)$ is an explicit factor of $R(Y) = D_{n(d-1)+1}(Y)$, as explained in Section 7.1. Write $R(Y) = \Delta(Y)r(Y)$. By, for example, [17, Page 8], in the first equation in Proposition 14 one can replace $D_{n(d-1)+1}(Y)$ by the resultant $r(Y)$, provided one allows the jth index in the ith summand to run over all monomials of degree $\delta - d$ which are divisible by all X_k for $k \neq i$. The author does not know of any explicit degree bounds on the polynomials A_{ij}, or a method for computing this expression efficiently, in this case. However, this shows elements in $\mathcal{L}_Y^o$ can be reduced modulo $\sum_{i=1}^n D_{i,Y}(\mathcal{L}_Y^{(i)})$ to a sum of the basis elements with coefficients which are rational functions of the form $E(Y)/\mathfrak{g}(Y)r(Y)^m$, for some positive integer m and polynomial $E(Y)$. Using our reduction method the corresponding basis elements have coefficients of the form $A(Y)/\mathfrak{g}(Y)R(Y)^m$ for some polynomial $A(Y)$ (compare with (4)). For any specialisation $Y = y$ with $r(y) \neq 0 \bmod p$ we know that this representation is unique, since $L((p-1)/p)^o \otimes \mathbb{Q}_q(y)/\sum_{i=1}^n D_{i,y}(L((p-1)/p)^{(i)} \otimes \mathbb{Q}_q(y))$ is a $\mathbb{Q}_q(\pi, y)$-vector space with basis $\pi^{u_0}X^u$ for $u \in \mathcal{B}$ (c.f. Section A.4). It follows that the two rational functions must be equal, i.e., the factor $\Delta(Y)^m$ cancels in $A(Y)/\mathfrak{g}(Y)R(Y)^m$. This shows that in fact $\mathfrak{g}(Y)B(Y)$ has rational functions entries with denominators a power of $r(Y)$. Likewise, the entries in the matrix $\mathfrak{g}(Y^p)\alpha(Y)$ can be written as a power series in X plus an expansion in $1/r^{\tau^{-1}}(Y)$. So modulo a power of p, it is just a rational function with denominator a power of $r^{\tau^{-1}}(Y)$. It can therefore be evaluated at any Teichmüller point $\tau^{-1}(y)$ with $r(y) \neq 0 \bmod p$. If the hypersurface defined by f is smooth, then indeed $r(1) \neq 0 \bmod p$ and the theory applies.

Finally, we imposed the generic requirement that $\mathfrak{g}(y) \neq 0 \bmod p$, although in the theory one only requires $\mathfrak{g}(y) \neq 0$. This restriction can be removed by following an idea of Dwork [8, Pages 251 Note (4), 255]. Let $\eta \in \mathbb{Z}_q$ with $\eta \equiv 1 \bmod p$. One can take different liftings ηa_i of the diagonal coefficients of $\bar{f}$ and the theory still works. In doing this, we get a new polynomial $\mathfrak{g}_0(Y) := \eta^{\deg_Y(\mathfrak{g})}\mathfrak{g}(\eta^{-1}Y)$ the non-vanishing of which now defines a new generic condition. Now choose δ minimal so that $q^{\delta-1} > n\binom{n(d-1)}{d-1} \geq \deg_Y(\mathfrak{g}(Y))$. By simply trying sufficiently many $\eta \in \mathbb{Z}_q$, one can find $\eta \equiv 1 \bmod p$ such that η^{-1} is different modulo p^δ from all roots of $\mathfrak{g}(Y)$. For such a choice of η, setting $Y = 1$ we get that $\mathfrak{g}_0(1) \neq 0 \bmod p^\delta$. This condition is sufficient, although we will lose a little p-adic accuracy in Step 7.

In summary, using the ideas above our generic condition in Theorem 1 can be relaxed to just the requirement that the hypersurface is smooth. Likewise, in Theorem 2 one can replace a "suitably generic polynomial" by a "smooth homogeneous polynomial".

A Appendix

In this appendix we develop a relative p-adic cohomology theory for one-dimensional families of smooth projective hypersurfaces. More precisely, we define a "generic absolute Frobenius matrix" which acts on a "generic cohomology space". We show that this generic matrix may be recovered around the origin from the solution matrix of a system of differential equations. We also explain how a suitable analytic continuation of this generic Frobenius matrix is related to zeta functions. Our theory is inspired by that of Dwork [9]. However, to the author's knowledge, Dwork does not actually develop the relative theory as we present it. Rather, he develops an analogous theory in

his "dual spaces", and states but does not prove any results for the types of spaces we consider. Dwork's dual theory seems more complicated, and we prefer to remain in the usual "Dwork space".

A.1 Generic spaces

Let $p > 2$ and q be a power of p. Let R denote the ring of power series in Y over the unramified extension $\mathbb{Q}_q(\pi)$ of $\mathbb{Q}_p(\pi)$ of degree $\log_p(q)$ which converge on some closed disk of non-zero radius containing the origin. Precisely, elements in R have an expansion of the form $\sum_{i=0}^{\infty} a_i Y^i$ where $a_i \in \mathbb{Q}_q(\pi)$ with $\mathrm{ord}(a_i) - \epsilon i \to \infty$ for some real number ϵ, not necessarily positive. The ring R contains power series expansions of holomorphic functions at the origin.

We first describe the space which is used by Dwork in his cohomology theory for a single smooth projective hypersurface [7]. For $b > 0$ and c real numbers, let $L(b, c)$ denote the $\mathbb{Z}_q[\pi]$-module of power series over $\mathbb{Q}_q(\pi)$ of the form

$$\sum_{m \in \mathcal{I}} a_m X^m, \ \mathrm{ord}(a_m) \geq b m_0 + c$$

where

$$\mathcal{I} = \{m = (m_0, m_1, \ldots, m_n) \in \mathbb{Z}_{\geq 0}^{n+1} \mid d m_0 = m_1 + \ldots + m_n\}.$$

Let $L(b) = \cup_c L(b, c)$, a ring. It is the space $L((p-1)/p)$ which occurs in Dwork's theory. (This is not actually a p-adic Banach space, and one often also works in a slightly smaller space which is, although we shall not need this.) The choice $b = (p-1)/p$ ensures that the space is large enough to be stable under the Frobenius map, and yet small enough for a certain factor space to be finite dimensional. Dwork also works with certain subspaces. For any space $*$ whose elements are sums of monomials X^m with $m \in \mathcal{I}$, denote by $*^o$ the space obtained by restricting the summation to $m \in \mathcal{I}^o$, and for each $1 \leq i \leq n$ denote by $*^{(i)}$ the space obtained by restricting summation to $m \in \mathcal{I}^{(i)}$. So, for example, $L(b)^o = L(b) \cap X_0 X_1 \ldots X_n \mathbb{Q}_q(\pi)[[X]]$, and $L(b)^{(i)} = L(b) \cap X_0 X_1 \ldots X_{i-1} X_{i+1} \ldots X_n \mathbb{Q}_q(\pi)[[X]]$, the set of power series in $L(b)$ in which every term is divisible by all X_k for $k \neq i$. We will also use the spaces $L(b, c)^o$ and $L(b, c)^{(i)}$.

We wish to "extend the scalars" in $L((p-1)/p)$ to obtain a module over the ring R. We need to do this very carefully, to ensure we get all the necessary properties.

Definition 22 *The space L_Y is defined to be the set of power series of the form $\sum_{j=0}^{\infty} a_j(X) Y^j$ where $a_j(X) \in L((p-1)/p, c_j)$ with*

$$c_j - \epsilon j \to \infty \ \textit{as} \ j \to \infty \tag{25}$$

for some real number ϵ. Here ϵ depends upon the power series, and we do not assume any lower bound on the values required for ϵ over the elements in L_Y.

Condition (25) on the c_j ensures that any series $a(X, Y) \in L_Y$ will converge to an element of $L((p-1)/p)$ on substitution of $Y = y$ for sufficiently small y, i.e., $\mathrm{ord}(y) \geq -\epsilon$. This is useful in the proof of Proposition 24. It is easy to prove that L_Y is a module over R. We need some R-submodules of L_Y. Following our convention, we define $L_Y^o \subset L_Y$ to be the submodule of power series $\sum_{j=0}^{\infty} a_j(X) Y^j$ such that each $a_j(X) \in X_0 X_1 \ldots X_n \mathbb{Q}_q(\pi)[[X]]$. For $1 \leq i \leq n$ let $L_Y^{(i)} \subset L_Y$ be those power series with each $a_j(X) \in X_0 X_1 \ldots X_{i-1} X_{i+1} \ldots X_n \mathbb{Q}_q(\pi)[[X]]$. It is not difficult to show that L_Y is a ring, and both L_Y^o and $L_Y^{(i)}$ are modules over L_Y.

We now introduce the "first-order differential operators" which will act on our generic space L_Y. Recall that

$$f(X_1, \ldots, X_n, Y) = \sum_{i=1}^{n} a_i X_i^d + Y h(X_1, \ldots, X_n)$$

where $a_1 \ldots a_n \neq 0 \bmod p$ and p does not divide d. Let

$$D_{i,0} = \exp(-\pi X_0 \sum_{i=1}^{d} a_i X_i^d) \circ X_i \frac{\partial}{\partial X_i} \circ \exp(\pi X_0 \sum_{i=1}^{d} a_i X_i^d) = X_i \frac{\partial}{\partial X_i} + \pi d a_i X_0 X_i^d,$$

and

$$\begin{aligned} D_{i,Y} &= \exp(-\pi Y X_0 h) \circ D_{i,0} \circ \exp(\pi Y X_0 h) \\ &= X_i \frac{\partial}{\partial X_i} + \pi X_0 X_i \frac{\partial f}{\partial X_i} \end{aligned} \tag{26}$$

act on $\mathbb{Q}_q(\pi)[[X, Y]]$, for $1 \leq i \leq n$. Let D_{i,Y^p} be as $D_{i,Y}$ with Y replaced by Y^p.

Lemma 23 *Each of the three operators $D_{i,0}$, $D_{i,Y}$ and D_{i,Y^p} maps the space $L_Y^{(i)}$ to the space L_Y^o.*

Proof: Similar to the proof of [14, Lemma 24]. □

Proposition 24 *The quotient R-modules*

$$L_Y^o / \sum_{i=1}^{n} D_{i,Y}(L_Y^{(i)}),\ L_Y^o / \sum_{i=1}^{n} D_{i,Y^p}(L_Y^{(i)}),\ L_Y^o / \sum_{i=1}^{n} D_{i,0}(L_Y^{(i)})$$

are free R-modules spanned by the set

$$\{\pi^{u_0} X^u \,|\, u = (u_0, u_1, \ldots, u_n) \in \mathcal{I},\ 0 < u_i < d\}.$$

Proof: We will first of all prove the claims for the space $L_Y^o / \sum_{i=1}^{n} D_{i,Y}(L_Y^{(i)})$. Consider a series $\sum_m A_m X^m \in L((p-1)/p, c)^o$ for some real number c. After one reduction step a term $A_m X^m$ is shown by the formulae in Section 5.2 to be equivalent modulo $\sum_{i=1}^{n} D_{i,Y}(L_Y^{(i)})$ to a sum of monomials X^r, with $r_0 = m_0 - 1$, whose coefficients are series of the form

$$A_m \pi^{-1} \frac{E(Y)}{D_{\min(m_0, n(d-1)+1)}(Y)}.$$

Here $E(Y) \in \mathbb{Z}_q[\pi][Y]$ is the coefficient of some monomial in the function "$X_i(\partial A_i / \partial X_i)$" which occurs in the formulae in Section 5.2. The quotient of polynomials in this expression can be written as a series with p-adic integral coefficients, since $D_*(0) \neq 0 \bmod p$ by Lemma 13. Thus after at most m_0 reduction steps, the term $A_m X^m$ can be written as a linear combination of the basis monomials with coefficients of the form $\sum_{k=0}^{\infty} d_{m,k} Y^k$ where $\mathrm{ord}(d_{m,k}) \geq ((p-1)/p - (1/p-1))m_0 + c$. Therefore the reduction of $\sum_{m \in \mathcal{I}^o} A_m X^m$ is equal to $\sum_{k=0}^{\infty} (\sum_m d_{m,k}) Y^k$. For k fixed, since $(p-1)/p - 1/(p-1) > 0$ (recall $p \neq 2$), the sum $\sum_m d_{m,k}$ is defined and has p-adic order $\geq c$. Thus the coefficient of a basis monomial in the reduction of $\sum_m A_m X^m \in L((p-1)/p, c)^o$ is of the form $\sum_k e_k Y^k$ where $\mathrm{ord}(e_k) \geq c$. Consider now an element $\sum_j a_j(X) Y^j$ where $a_j(X) = \sum_m A_m^{(j)} X^m$. Adding the

superscript notation (j) to the series above, and replacing c by c_j, we see that the coefficient of a basis monomial in the reduction of $\sum_j a_j(X)Y^j$ is of the form $\sum_j(\sum_k e_k^{(j)})Y^{j+k}$ where $\mathrm{ord}(e_k^{(j)}) \geq c_j$. This series is certainly defined, since the the coefficient of each power of Y is a finite sum. We need to show that it lies in R. Rewrite this as $\sum_\ell f_\ell Y^\ell$ where $f_\ell = \sum_{j,k:\, j+k=\ell} e_k^{(j)}$, a finite sum. We have

$$\mathrm{ord}(f_\ell) \geq \min(\mathrm{ord}(e_k^{(j)}) \,|\, j+k=\ell) \geq \min(c_j \,|\, j \leq \ell).$$

Since $c_j - \epsilon j \to \infty$ as $j \to \infty$ it follows that $\mathrm{ord}(f_\ell) - \varepsilon\ell \to \infty$ as $\ell \to \infty$. Here ε can be chosen arbitrarily less than zero when $\epsilon \geq 0$. When $\epsilon < 0$ we choose ε arbitrarily less than ϵ. Thus $\sum_\ell f_\ell Y^\ell \in R$, as we wished to show.

This shows that the reduction process "converges" in an appropriate sense. We must also check that the operands of $D_{i,Y}$ which arise when reducing the series in L_Y^o also lie in $L_Y^{(i)}$. Reducing a series $\sum_m A_m X^m \in L((p-1)/p, c)^o$ gives an R-linear combination of the basis set plus $\sum_{i=1}^n D_{i,Y}(\sum_{m\in\mathcal{I}^{(i)}} b_{i;m}(Y)X^m)$, say. Let us focus on the ith operand, and write it simply as $\sum_m b_m(Y)X^m$. Writing $b_m(Y) = \sum_k b_{m,k}Y^k$ we see from the formulae in Section 5.2 that $\mathrm{ord}(b_{m,k}) \geq ((p-1)/p)(m_0+1) - (1/(p-1)) + c$. Once again adding the superscript (j) etc, and focusing still on one operand, we see that when reducing $\sum_j(\sum_m A_m^{(j)}X^m)Y^j \in L_Y^o$, we must check that $\sum_m(\sum_j\sum_k b_{m,k}^{(j)}Y^{j+k})X^m \in L_Y^{(i)}$ where $\mathrm{ord}(b_{m,k}^{(j)}) \geq ((p-1)/p)(m_0+1) - (1/(p-1)) + c_j$. Rewrite this sum as $\sum_\ell(\sum_m d_{m,\ell}X^m)Y^\ell$ where $d_{m,\ell} = \sum_{j,k:\, j+k=\ell} b_{m,k}^{(j)}$. Now

$$\begin{array}{rcl} \mathrm{ord}(d_{m,\ell}) & \geq & \min\left(\frac{p-1}{p}(m_0+1) - \frac{1}{p-1} + c_j \,|\, j+k=\ell\right) \\ & \geq & \frac{p-1}{p}m_0 + \min(c_j \,|\, j \leq \ell). \end{array}$$

Define $C_j = \min(c_j \,|\, j \leq \ell)$. Then $\sum_m d_{m,\ell}X^m \in L((p-1)/p, C_j)^{(i)}$. Moreover, we have $C_j - \varepsilon j \to \infty$, where ε is defined from ϵ as in the previous paragraph. Hence $\sum_j(\sum_m d_{m,\ell}X^m)Y^j \in L_Y^{(i)}$, as we wished to show.

The proof that the spanning set is linearly independent is done via a specialisation and modular reduction argument: Consider again reduction via the operators $D_{i,Y}$. Suppose that the claimed basis set is not linearly independent. Then we have an equation of the form

$$\sum_u a_u(Y)(\pi^{u_0}X^u) = \sum_{i=1}^n D_{i,Y}(b_i(X,Y))$$

where u runs over a non-empty subset of $\mathcal{B}$, and $a_u(Y) \in R - \{0\}$ with $b_i \in L_Y^{(i)}$. Choose $y \in \mathbb{Q}_q$ small enough that each $a_u(Y)$ converges at $Y = y$ to a non-zero element of $\mathbb{Q}_q(\pi)$, and each b_i converges at $Y = y$ to an element of $L((p-1)/p)^{(i)}$. Then $\sum_u a_u(y)\pi^{u_0}X^u = \sum_{i=1}^n D_{i,y}(b_i(X,y))$ with $a_u(y) \in \mathbb{Q}_q(\pi) - \{0\}$ and $b_i(X,y) \in L((p-1)/p)^{(i)}$. Here $D_{i,y}$ is the operator with Y specialised to y acting on $L((p-1)/p)$. Multiply both sides of the equation by a power π^m so that $\pi^m b_i(X,y) \in L((p-1)/p, 0)^{(i)}$ and $\pi^m a_u(y) \in \mathbb{Z}_q[\pi]$, for all i and u, to get a new equation. Now let B denote the p-adic Banach module over $\mathbb{Z}_q[\pi]$ consisting of power series of the form $\sum_{m\in\mathcal{I}} a_m\pi^{m_0}X^m$ where $a_m \in \mathbb{Z}_q[\pi]$ with $\mathrm{ord}(a_m) \to \infty$. Then $L((p-1)/p, 0)^o \subset B$ and B is stable under $D_{i,y}$. Monsky's result [18, Theorem 8.5 (2)] can now be applied to show that $B/\sum_{i=1}^n D_{i,y}(B)$ is a free $\mathbb{Z}_q[\pi]$-module with basis $\{\pi^{u_0}X^u \,|\, u \in \mathcal{I},\, u_1, \ldots, u_n < d\}$. (Specifically, in Monsky's notation, "$H_0(\bar{C})$" is the space $S/(X_0X_1^d, \ldots, X_0X_n^d)$ where S is the $\mathbb{F}_q$-vector space spanned by monomials X^u for $u \in \mathcal{I}$. A basis for this reduced space is $\{X^u \,|\, u \in \mathcal{I},\, u_1, \ldots, u_n < d\}$,

and Monsky's result tells us this is also a basis for "$H_0(C)$", which is just $B/\sum_{i=1}^n D_{i,y}(B)$.) But our new equation contradicts the linear independence of this set. This completes the proof for the operators $D_{i,Y}$.

Similar arguments work for D_{i,Y^p}, just replacing Y by Y^p in the formulae in Section 5.2. The argument for $D_{i,0}$ is similar, but much simpler. □

We will use the shorthand notation $L_Y^o/\sum_{i=1}^n D_{i,Y}$ for $L_Y^o/\sum_{i=1}^n D_{i,Y}(L_Y^{(i)})$, and likewise for the other two modules.

A.2 Frobenius maps and deformations

Recall that τ is the map on $\mathbb{Q}_q(\pi)$ which reduces to the pth power map on its residue field and fixes π. Extend τ to act on $\mathbb{Q}_q(\pi)[[Y]]$ by fixing Y (and being linear and continuous under the Y-adic norm). Let the operator ψ_p act on the space of formal power series $\mathbb{Q}_q(\pi)[[X,Y]]$ in the following way. For a monomial X^m where $m=(m_0,m_1,\ldots,m_n)$, the image $\psi_p(X^m)$ is $X_0^{m_0/p}X_1^{m_1/p}\ldots X_n^{m_n/p}$ when p divides each m_i, and zero otherwise. Also ψ_p is τ^{-1}-linear over $\mathbb{Q}_q(\pi)[[Y]]$ (and continuous under the X-adic norm); in particular, ψ_p fixes Y. Explicitly,

$$\psi_p : \sum_{r=0}^{\infty}\sum_{m\in\mathbb{Z}_{\geq 0}^{n+1}} a_{m,r}X^mY^r \mapsto \sum_{r=0}^{\infty}\sum_{m\in\mathbb{Z}_{\geq 0}^{n+1}} \tau^{-1}(a_{pm,r})X^mY^r,\ a_m\in\mathbb{Q}_q(\pi).$$

Define $\alpha_Y : \mathbb{Q}_q(\pi)[[X,Y]]\to\mathbb{Q}_q(\pi)[[X,Y]]$ by

$$\begin{aligned}\alpha_Y &= \exp(-\pi X_0 f(X_1,\ldots,X_n,Y^p))\circ\psi_p\circ\exp(\pi X_0 f(X_1,\ldots,X_n,Y)) \qquad (27)\\ &= \psi_p\circ F(X,Y)\end{aligned}$$

where

$$\begin{aligned}F(X,Y) &= \exp(\pi(X_0 f(X_1,\ldots,X_n,Y) - X_0^p f^\tau(X_1^p,\ldots,X_n^p,Y^p)))\\ &= \prod_{i=1}^n \theta(a_i X_0 X_i^d)\prod_{j\in J}\theta(Y b_j X_0 X^j)\end{aligned}$$

Here $\theta(z)=\exp(\pi(z-z^p))$, and τ fixes the variable Y and in f^τ acts only on the coefficients of f. The set J contains those vectors $(j_1,\ldots,j_n)$ with $|j|=d$ which do not belong to the "diagonal" set $\{(d,0,\ldots,0),\ldots,(0,\ldots,0,d)\}$. The coefficients b_j are defined via the equation $h=\sum_{j\in J} b_j X^j$, that is, they are the non-diagonal coefficients in f.

Lemma 25 *The space L_Y^o is stable under the map α_Y.*

Proof: Similar to the proof of [14, Lemma 26]. □

Lemma 26 *The action $\alpha_Y : L_Y^o\to L_Y^o$ induces a map*

$$\alpha_Y : L_Y^o/\sum_{i=1}^n D_{i,Y} \to L_Y^o/\sum_{i=1}^n D_{i,Y^p}.$$

Proof: Similar to the proof of [14, Lemma 27]. □

Define

$$\alpha_0 = \exp(-\pi X_0 f(X_1, \ldots, X_n, 0)) \circ \psi_p \circ \exp(\pi X_0 f(X_1, \ldots, X_n, 0)) \tag{28}$$
$$= \psi_p \circ F(X, 0),$$

a map on $\mathbb{Q}_q(\pi)[[X, Y]]$. Then L_Y^o is stable under α_0, and α_0 induces a map

$$\alpha_0 : L_Y^o / \sum_{i=1}^n D_{i,0} \to L_Y^o / \sum_{i=1}^n D_{i,0}.$$

(These facts are proved in the same manner as Lemmas 25 and 26.) Let $T_{Y,0}$ denote the bijective map "multiplication by $\exp(\pi Y X_0 h)$" from $\mathbb{Q}_q(\pi)[[X, Y]]$ to itself (the inverse is multiplication by $\exp(-\pi Y X_0 h)$).

Lemma 27 *The space L_Y^o is stable under the map $T_{Y,0}$. Moreover, $T_{Y,0}$ is a bijection on L_Y^o.*

Proof: Similar to the proof of [14, Lemma 28]. □

The map $T_{Y,0}$ induces a bijection

$$T_{Y,0} : L_Y^o / \sum_{i=1}^n D_{i,Y} \to L_Y^o / \sum_{i=1}^n D_{i,0}.$$

On sees this by considering the factorisations of the operators $D_{i,Y}$ and $D_{i,0}$. Define $T_{Y^p,0}$ to be the map "multiplication by $\exp(\pi Y^p X_0 h)$" on L_Y^o. Then $T_{Y^p,0}$ also induces a bijection, as in the righthand vertical arrow in the next diagram.

Proposition 28 *The following diagram commutes:*

$$\begin{array}{ccc} L_Y^o / \sum_{i=1}^n D_{i,Y} & \xrightarrow{\alpha_Y} & L_Y^o / \sum_{i=1}^n D_{i,Y^p} \\ \downarrow T_{Y,0} & & \downarrow T_{Y^p,0} \\ L_Y^o / \sum_{i=1}^n D_{i,0} & \xrightarrow{\alpha_0} & L_Y^o / \sum_{i=1}^n D_{i,0}. \end{array}$$

Thus we have

$$\alpha_Y = T_{Y^p,0}^{-1} \circ \alpha_0 \circ T_{Y,0}. \tag{29}$$

Proof: Similar to the proof of [14, Proposition 29]. □

Let $C(Y)$ denote the matrix for the map multiplication by $\exp(\pi Y X_0 h)$ from $L_Y^o / \sum_{i=1}^n D_{i,Y}$ to $L_Y^o / \sum_{i=1}^n D_{i,0}$ with respect to the basis of monomials of the two spaces. Write $\alpha(Y)$ and $\alpha(0)$ for the matrices of the maps α_Y and α_0, respectively. (Our matrix convention is that the entry in the uth row and vth column, the (u, v)th entry, of the matrix for a map gives the coefficient of $\pi^{u_0} X^u$ in the image of $\pi^{v_0} X^v$ under the map.)

Lemma 29 *The matrix of the map $T_{Y^p,0}^{-1} : L_Y^o / \sum_{i=1}^n D_{i,0} \to L_Y^o / \sum_{i=1}^n D_{i,Y^p}$ is $C(Y^p)^{-1}$.*

Proof: Similar to the proof of [14, Lemma 30]. □

From (29) and Lemma 29 we get

$$\alpha(Y) = C(Y^p)^{-1}\alpha(0)C(Y)^{\tau^{-1}}. \tag{30}$$

Here the τ^{-1} arises since α_0 is τ^{-1} linear. This is an identity in the ring R of power series expansions convergent on some closed disk of non-zero radius around the origin.

Note 30 One can present the results in this section in an alternative, and perhaps more natural, fashion. Instead of defining ψ_p to fix Y, one can introduce a formal root $Y^{1/p}$ and define ψ_p to map Y to $Y^{1/p}$. Then $\alpha_Y : L^o_Y/\sum_{i=1}^n D_{i,Y} \to L^o_{Y^{1/p}}/\sum_{i=1}^n D_{i,Y}$ where $L^o_{Y^{1/p}}$ is defined exactly as before with $Y^{1/p}$ replacing Y. The resulting factorisation is "$\alpha(Y) = C(Y)^{-1}\alpha(0)C(Y^{1/p})^{\tau^{-1}}$" where $C(Y)$ is the same matrix as before. Thus our new $\alpha(Y)$ is just the old $\alpha(Y)$ with Y replaced by $Y^{1/p}$. Given a suitable Teichmüller point y, one now evaluates the new $\alpha(Y)$ directly at y and takes $y^{1/p}$ to equal $\tau^{-1}(y)$ to get the semi-linear Frobenius matrix c.f. the second paragraph in Section A.4.

A.3 The differential equation satisfied by the deformation matrix

The proof that the matrix $C(Y)$ from the previous section is the unique solution matrix of the differential system (1) is very similar to [14, Appendix A.3] and is omitted. This fact, combined with Equation (30), proves Proposition 6.

A.4 Zeta functions and the specialisation of the generic Frobenius matrix

By definition, the matrix $\alpha(Y)$ has entries which are power series expansions around the origin. In general, these expansions will not converge on the closed unit disk. However, the matrix $\alpha(Y)$ can be analytically continued to a much larger region. Precisely, there is a matrix of p-adic holomorphic functions $\alpha(Y)^*$, say, on a set S, say, with the following properties: The set $S \subset \mathbb{C}_p$ is a closed disk of radius greater than one around the origin, with open disks of radius less than one around finitely many points of norm one removed. The matrix $\alpha(Y)^*$ has entries holomorphic on S, and the functions $\alpha(Y)$ and $\alpha(Y)^*$ agree on some closed disk of radius less than one around the origin. This matrix of holomorphic functions $\alpha(Y)^*$ is uniquely determined by the matrix of local expansions $\alpha(Y)$. (A holomorphic function on S is defined as the uniform limit of rational functions with no poles in S. Such functions can be described explicitly via the p-adic Mittag-Leffler theorem, see for example [1, Page 286]. Any such function is defined uniquely by its local expansion around the origin, according to a theorem of Krasner [8, Page 256 Lines 7-9].) Bounds on the region of holomorphy of $\alpha(Y)^*$ are calculated in Section 7.1. Specifically, the entries in $\mathfrak{g}(Y^p)\alpha(Y)^*$ are expanded as a power series in Y and $1/R(Y)$; such a representation of a holomorphic function is a kind of "Mittag-Leffler theorem over $\mathbb{Q}_q(\pi)$". In Section 7.1 and the present section, the notation $\alpha(Y)$ is used to denote the matrix of p-adic holomorphic function $\alpha(Y)^*$, rather than just the matrix of local expansions around the origin.

The significance of the (analytically continued) matrix $\alpha(Y)$ from the point of view of zeta functions is that evaluating this matrix at $\tau^{-1}(y)$ for some Teichmüller point y with $\mathfrak{g}(y) \neq 0$ and $R(y) \neq 0 \bmod p$ gives the semi-linear Frobenius matrix. Specifically, let $y \in \mathcal{O}_{\mathbb{C}_p}$ with $y^{q^r} = y$ such

that $\mathfrak{g}(y) \neq 0$ and $R(y) \neq 0 \bmod p$. The quotient $\mathbb{Q}_q(\pi, y)$-module

$$L((p-1)/p)^o \otimes \mathbb{Q}_q(\pi, y) / \sum_{i=1}^{n} D_{i,y}(L((p-1)/p)^{(i)} \otimes \mathbb{Q}_q(\pi, y)) \tag{31}$$

is free on the set $\{\pi^{u_0} X^u \,|\, u \in \mathcal{B}\}$. Here $D_{i,y}$ is just $D_{i,Y}$ with the variable Y replaced by the Teichmüller point y. Also, the tensor product notation indicates that we extend the scalars to the larger field, but retain the same decay conditions. Let $\alpha_y = \psi_p \circ F(X, y)$ with $F(X, Y)$ exactly as above, with ψ_p extended to act on $\mathbb{Q}_q(\pi, y)$ by $\psi_p(y) = \tau^{-1}(y)$. As before α_y induces a map on the quotient space (31). Let (α_y) be the matrix for this map with respect to the basis of monomials. (Note that we now have two maps α_0: one as in (28) and that defined immediately above. The former acts on the R-module spanned by the basis set, and the latter on the $\mathbb{Q}_q(\pi)$-module. However, the matrices for these two maps are the same.) We have that (α_y) is equal to $\alpha(Y)$ evaluated at $\tau^{-1}(y)$, as can be seen by examining the action of α_y and α_Y and the operators $D_{i,y}$ and "D_{i,Y^p} evaluated at $Y = \tau^{-1}(y)$". The conditions $\mathfrak{g}(y) \neq 0$ and $R(y) \neq 0 \bmod p$ are needed to ensure that $\alpha(Y)$ converges at y, and also that the the quotient space (31) is finite dimensional. Now $\alpha_y^{r \log_p(q)}$ is the Frobenius map on Dwork's cohomology space. Write $Z(\bar{f}(X, \bar{y})/\mathbb{F}_{q^r}, T)$ for the zeta function of the smooth projective hypersurface $\bar{f}(X, \bar{y}) = 0$. Define

$$P(qT) = \det(1 - T\alpha_y^{r \log_p(q)}).$$

Because of τ^{-1}-linearity, the matrix for the linear map $\alpha_y^{r \log_p(q)}$ is equal to

$$(\alpha_y)(\alpha_y)^{\tau^{-1}} \ldots (\alpha_y)^{\tau^{-r \log_p(q)+1}}.$$

We have from [18, Theorem 8.8 (3),(4)] that

$$Z(\bar{f}(X, \bar{y})/\mathbb{F}_{q^r}, T) = \frac{P(T)^{(-1)^{n+1}}}{(1-T)(1-q^r T)\ldots(1-q^{r(n-2)}T)}.$$

Note that $\tau^{-i} = \tau^{r \log_p(q) - i}$ on $\mathbb{Q}_q(\pi, y)$, which gives the precise formulation in Proposition 7.

References

[1] A. Adolphson, An index theorem for p-adic differential operators, Trans. Amer. Math. Soc., Vol. 216, (1976), 279-293.

[2] P. Berthelot, Géométrie rigide et cohomologie des variétés algebriques de caractéristique p, in Introductions aux cohomologies p-adiques (Luminy, 1984), Mém. Soc. Math. France 23 (1986), 7-32.

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

[4] E. Bombieri, On exponential sums in finite fields II, Invent. Math. 47, (1978), 29-39.

[5] D.G. Cantor and E. Kaltofen, On fast multiplication of polynomials over arbitrary algebras, Acta Inform., 28, 693-701, 1991.

[6] B. Dwork, On the rationality of the zeta function of an algebraic variety, Amer. J. Math., 82, (1960), 631-648.

[7] B. Dwork, On the zeta function of a hypersurface, Pub. Math. IHES No. 12, (1962), 5-68.

[8] B. Dwork, A deformation theory for the zeta function of a hypersurface, Proc. Internat. Congr. Mathematicians (Stockholm 1962), 247-259.

[9] B. Dwork, On the zeta function of a hypersurface II, Ann. Math. (2) 80, (1964), 227-299.

[10] N. Elkies, Elliptic and modular curves over finite fields and related computational issues, in "*Computational perspectives in number theory: Proceedings of a conference in honour of A.O.L. Atkin*" , (D.A. Buell and J.T. Teitelbaum), American Mathematical Society International Press 7, 1998, 21-76.

[11] J. von zur Gathen and J. Gerhard, Modern Computer Algebra, Cambridge University Press, 1999.

[12] N.M. Katz, On the differential equations satisfied by period matrices, Pub. Math. IHES 35, (1968), 71-106.

[13] K. Kedlaya, Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology, Journal of the Ramanujan Mathematical Society, 16 (2001), 323-338.

[14] A.G.B. Lauder, Deformation theory and the computation of zeta functions, to appear in the Proceedings of the London Mathematical Society. Preprint available at: http://web.comlab.ox.ac.uk/oucl/work/alan.lauder/

[15] A.G.B. Lauder and D. Wan, Counting points on varieties over finite fields of small characteristic, to appear in Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography (Mathematical Sciences Research Institute Publications), J.P. Buhler and P. Stevenhagen (eds), Cambridge University Press.

[16] A.G.B. Lauder and D. Wan, Computing zeta functions of Artin-Schreier curves over finite fields, London Math. Soc. JCM, Vol. 5, (2002), 34-55.

[17] F.S. Macaulay, The Algebraic Theory of Modular Systems, Cambridge University Press, 1916.

[18] P. Monsky, p-adic Analysis and Zeta Functions, Lectures in Mathematics, Kyoto University, Tokyo, 1970.

[19] J. Pila, Frobenius maps of abelian varieties and finding roots of unity in finite fields, Math. Comp. 55, 745-763, 1990 .

[20] B. Poonen, Computational aspects of curves of genus at least 2, in "*Algorithmic Number Theory II*" (H. Cohen), Lecture Notes in Computer Science 1122, Springer, 1996, 283-306.

[21] T. Satoh, The canonical lift of an ordinary elliptic curve over a finite field and its points counting, J. Ramanujan Math. Soc. 15, (2000), 247-270.

[22] R. Schoof, Elliptic curves over finite fields and the computation of square roots mod p, Math. Comp. 44, no. 170, (1985), 483-494.

[23] D. Wan, Computing zeta functions over finite fields, Contemporary Mathematics, 225 (1999), 131-141.

[24] D. Wan, Algorithmic theory of zeta functions, to appear in Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography (Mathematical Sciences Research Institute Publications), J.P. Buhler and P. Stevenhagen (eds), Cambridge University Press.
Available at: http://www.math.uci.edu/∼dwan/preprint.html