

COUNTING POINTS ON VARIETIES OVER FINITE FIELDS OF SMALL CHARACTERISTIC

ALAN G.B. LAUDER AND DAQING WAN

ABSTRACT. We present a deterministic polynomial time algorithm for computing the zeta function of an arbitrary variety of fixed dimension over a finite field of small characteristic. One consequence of this result is an efficient method for computing the order of the group of rational points on the Jacobian of a smooth geometrically connected projective curve over a finite field of small characteristic.

1. INTRODUCTION

The purpose of this paper is to give an elementary and self-contained proof that one may efficiently compute zeta functions of arbitrary varieties of fixed dimension over finite fields of suitably small characteristic. This is achieved via the p -adic methods developed by Dwork in his proof of the rationality of the zeta function of a variety over a finite field [4, 5]. Dwork's theorem shows that it is in principle possible to compute the zeta function. Our main contribution is to show how Dwork's trace formula, Bombieri's degree bound [3] and a semi-linear reduction argument yield an efficient algorithm for doing so. That p -adic methods may be used to efficiently compute zeta functions for small characteristic was first suggested in [20, 22], where Wan gives a simpler algorithm for counting the number of solutions to an equation over a finite field modulo small powers of the characteristic.

We now give more details of our results. For $q = p^a$, where p is a prime number and a a positive integer, let \mathbb{F}_q denote a finite field with q elements. Let $\bar{\mathbb{F}}_q$ denote an algebraic closure of \mathbb{F}_q , and \mathbb{F}_{q^k} the subfield of $\bar{\mathbb{F}}_q$ of order q^k . Denote by $\mathbb{F}_q[X_1, \dots, X_n]$ the ring of all polynomials in n variables over \mathbb{F}_q .

Date: October 27, 2002.

Alan Lauder gratefully acknowledges the support of the EPSRC (Grant GR/N35366/01) and St John's College, Oxford, and thanks Richard Brent. Daqing Wan is partially supported by the NSF and the NSFC. The authors are pleased to thank Colin McDiarmid and Bernd Sturmfels for answering some questions on convex geometry. Many excellent suggestions were made for improving the paper by the anonymous referee, and these were incorporated by the authors during the revision. They are especially grateful for this help.

Mathematics Subject Classification 2000: 11Y16, 11T99, 14Q15 *Key words and phrases:* variety, finite field, zeta function, algorithm.

For a polynomial $f \in \mathbb{F}_q[X_1, \dots, X_n]$, we denote by N_k the number of solutions to the equation $f = 0$ with coordinates in \mathbb{F}_{q^k} . The zeta function of the variety defined by f is the formal power series in T with non-negative integer coefficients

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

Dwork's theorem asserts that $Z(f/\mathbb{F}_q)$ is a rational function $r(T)/s(T)$ with integer coefficients. From this it follows that knowledge of explicit bounds $\deg(r) \leq D_1$ and $\deg(s) \leq D_2$, and of the values N_k for $k = 1, 2, \dots, D_1 + D_2$ is enough to efficiently determine $Z(f/\mathbb{F}_q)$, see [22]. The Bombieri degree bound tells us that $\deg(r) + \deg(s) \leq (4d + 9)^{n+1}$, see [3, 22], and so in particular we may take $D_1 = D_2 = (4d + 9)^{n+1}$. Each number N_k can be computed in a naive fashion by straightforward counting, using q^{nk} evaluations of the polynomial f . Thus one may compute the zeta function of a variety, but the naive method described requires a number of steps that is exponential in the parameters d^n and $\log q$, where d is the total degree of f . Note that the dense input size of f is $\mathcal{O}((d+1)^n \log q)$, and the size of the zeta function is polynomial in $\mathcal{O}((d+1)^n \log q)$, by the Bombieri degree bound. We prove the following theorem.

Theorem 1. *There exist an explicit deterministic algorithm and an explicit polynomial P such that for any $f \in \mathbb{F}_q[X_1, \dots, X_n]$ of total degree d , where $q = p^a$ and p is prime, the algorithm computes the zeta function $Z(f/\mathbb{F}_q)(T)$ of f in a number of bit operations which is bounded by $P(p^n d^2 a^n)$.*

In particular, this computes the zeta function of a polynomial in a fixed number of variables over a finite field of “small characteristic” in deterministic polynomial time. Note that our result makes no assumption of non-singularity on the variety defined by the polynomial. Also, we shall explicitly describe all the algorithms in this paper, rather than just prove their existence, and we assume that the finite field \mathbb{F}_q itself is presented as input via an irreducible polynomial of degree a over the prime field \mathbb{F}_p , as explained in Section 3.

With regard to the exponents in the algorithm, we state these precisely in Theorem 37. For now we observe that if p , d and n are fixed then the time required to compute the number of points in \mathbb{F}_q is $\mathcal{O}(a^{3n+7})$, with space complexity $\mathcal{O}(a^{2n+4})$. Here we are ignoring logarithmic factors (see also Proposition 36). Note that all our complexity estimates are made using standard methods for multiplication in various rings, and can be modestly reduced with faster methods.

We also present refinements to this result based upon the ideas of Adolphson and Sperber [2], and indeed from the outset will follow their approach, as it involves little extra complication. This refinement takes into account the terms which actually occur in the polynomial f rather than working solely with the total degree. We shall need more definitions: the support

of a polynomial f is the set of exponents $r = (r_1, \dots, r_n)$ of non-zero terms $a_r X_1^{r_1} \dots X_n^{r_n}$ which occur in f , thought of as points in \mathbb{R}^n . The Newton polytope of f is defined to be the convex hull in \mathbb{R}^n of the support of f . Our refined version of Theorem 1 essentially replaces the parameter d^n with the normalised volume of the Newton polytope of f (see Proposition 35 and Section 6.3.3).

Zeta functions may also be defined for a finite collection of polynomials, and we next describe how our results may be extended to this case. An affine variety V over \mathbb{F}_q is the set of common zeros in $(\overline{\mathbb{F}}_q)^n$ of a set of polynomials f_1, f_2, \dots, f_r . An analogous zeta function $Z(V/\mathbb{F}_q)$ may then be defined in terms of the number of solutions in each finite extension field of \mathbb{F}_q . These numbers may be computed using an inclusion-exclusion argument involving the polynomials $\prod_{i \in S} f_i$ where S is a subset of $\{1, 2, \dots, r\}$, see [22]. Theorem 1 then easily yields the following.

Corollary 2. *There exist an explicit deterministic algorithm and an explicit polynomial Q with the following property. Let $f_1, \dots, f_r \in \mathbb{F}_q[X_1, \dots, X_n]$ have total degrees d_1, \dots, d_r respectively, where $q = p^a$ and p is prime. Denote by V the affine variety defined by the common vanishing of these polynomials, and define $d = \sum_{i=1}^r d_i$. The algorithm computes the zeta function $Z(V/\mathbb{F}_q)(T)$ in a number of bit operations which is bounded by $Q(p^n d^{n^2} a^n 2^r)$.*

Thus one has an efficient algorithm for computing the zeta function of an arbitrary affine variety over \mathbb{F}_q assuming the characteristic, dimension and the number of defining polynomials are fixed. More generally still, an arbitrary variety is defined through patching together suitable affine varieties. Zeta functions for such general varieties may be defined. These zeta functions may be computed using the above ideas provided explicit data are given on how to construct them from affine patches. As an example, the zeta functions of arbitrary projective varieties or toric varieties may be computed in this way.

Conceptually our algorithm is rather straightforward. The zeta function can be expressed in terms of the “characteristic power series” (Fredholm determinant) of a certain “lifting of Frobenius” which acts on an infinite dimensional p -adic Banach space constructed by Dwork. Under modular reduction, we obtain an operator acting on a finite dimensional vector space. Unfortunately, this operator cannot be computed efficiently directly from its definition; however, it can be expressed as a product of certain semi-linear operators each of which can be computed efficiently if the characteristic p is small. This last step is thus of crucial importance in deriving an efficient algorithm. Note that the same idea is used in [20, 22] in a simpler situation. In more concrete language, the algorithm requires one to construct a certain “semi-linear” finite matrix and compute the “linear” matrix which is the product of the Galois conjugates of the semi-linear matrix. The number of points is then read off from the trace of the final linear matrix. The zeta

function can be computed from the characteristic polynomial of the final linear matrix. The semi-linear matrix itself is defined over a certain finite “ p -adic lifting” of the original finite field.

In the literature algorithms have already been developed for computing zeta functions of curves and abelian varieties [1, 6, 14, 17, 18]. These utilise the theory originally developed by Weil for abelian varieties, whereas we use Dwork’s more general and simpler p -adic theory. For example, in the case of a smooth geometrically irreducible projective plane curve of degree d over a field of size $q = p^a$ these algorithms have a time complexity which grows as $(\log q)^{C_d}$, where the constant C_d grows exponentially in the degree d . Given an absolutely irreducible bivariate polynomial f of degree d over \mathbb{F}_q , the zeta function of the unique smooth projective curve birational to the affine curve defined by f may be computed in time polynomial in d , p and a using our approach. Thus our more general method is far better in terms of the degree d if the characteristic p is small, since the running time has polynomial growth in d (but much worse if p is large). The zeta function immediately gives the order of the group of rational points on the Jacobian, see [22].

Corollary 3. *There exist an explicit deterministic algorithm and an explicit polynomial R with the following property. Let V be a geometrically irreducible affine curve defined by the vanishing of polynomials $f_1, \dots, f_r \in \mathbb{F}_q[X_1, \dots, X_n]$ of total degrees d_1, \dots, d_r respectively, where $q = p^a$ and p is prime. Denote by \tilde{V} the unique smooth projective curve birational to the affine curve V , and let $d = \sum_{i=1}^r d_i$. The algorithm computes the order of the group of rational points on the Jacobian of \tilde{V} in a number of bit operations bounded by $R(p^n d^{n^2} a^n 2^r)$.*

Thus one may compute the order of the group of rational points on the Jacobian of a smooth geometrically irreducible projective curve over a finite field of small characteristic in deterministic polynomial time, provided the number of variables and the number of defining equations are fixed. (Notice that the case $r = 1$ and $n = 2$ corresponds to that of being given a possibly singular plane model of the curve.) In particular, this answers a question posed in [15], which Poonen attributes to Katz and Sarnak. We note that recently a similar result for special classes of plane curves was independently obtained in [8, 10], using Monsky-Washnitzer’s method. Also, a different p -adic approach for elliptic curves has been developed in [16].

This paper is written primarily for theoretical computational interest, in obtaining a deterministic polynomial time algorithm for computing the zeta function in full generality if p is small. It can certainly be improved in many ways for practical computations. What we have done in this paper is to work on the easier but more flexible “chain level”. A general improvement in the smooth case, is to work on the cohomology level. Then there are several related p -adic cohomology theories available, each leading to a somewhat different version of the algorithm. Again, as indicated in [20, 22], these p -adic

methods are expected to be practical only for small p . (The authors have recently applied Dwork's cohomology theory to derive a practical algorithm in a special case [12].)

2. ADDITIVE CHARACTER SUMS OVER FINITE FIELDS

The most natural objects of study in Dwork's theory are certain additive character sums over finite fields. In this section we introduce these sums, and explain their connection to varieties.

An additive character Ψ is a mapping from \mathbb{F}_{q^k} to the group of units of some commutative ring S with identity 1 such that $\Psi(x+y) = \Psi(x)\Psi(y)$ for $x, y \in \mathbb{F}_{q^k}$. We say that it is non-trivial if $\Psi(x) \neq 1$ for some $x \in \mathbb{F}_{q^k}$. Let $\{\Psi_k\}_{k \geq 1}$ be any family of mappings with each Ψ_k a non-trivial additive character from \mathbb{F}_{q^k} to some extension ring of the integers whose image is a group of order p with elements summing to zero. We assume that the family $\{\Psi_k\}$ forms a tower of characters in the sense that for each $k \geq 2$,

$$\Psi_k = \Psi_1 \circ \text{Tr}_{\mathbb{F}_{q^k}/\mathbb{F}_q}.$$

In our application, the ring S will be taken to be a certain p -adic ring.

For the remainder of the paper we will use multi-index notation. Specifically, we let X^u represent the monomial $X_0^{u_0} X_1^{u_1} \dots X_n^{u_n}$ for an integer vector $u = (u_0, u_1, \dots, u_n)$; let x be an $(n+1)$ -tuple (x_0, x_1, \dots, x_n) of field elements; and X the list of indeterminates X_0, X_1, \dots, X_n . Observe here that we have introduced an extra indeterminate X_0 . Even though the polynomial f whose zeta function we wish to compute is in the n variables X_1, \dots, X_n , in Dwork's theory the extra indeterminate arises naturally, as we are about to see.

Lemma 4. *Let $f \in \mathbb{F}_q[X_1, \dots, X_n]$ and N_k^* denote the number of solutions to the equation $f = 0$ in the affine torus $(\mathbb{F}_{q^k}^*)^n$. Then*

$$\sum_{x \in (\mathbb{F}_{q^k}^*)^{n+1}} \Psi_k(x_0 f(x_1, \dots, x_n)) = q^k N_k^* - (q^k - 1)^n$$

where $x = (x_0, x_1, \dots, x_n)$.

Proof. We have that for any $u \in \mathbb{F}_{q^k}$

$$\sum_{x_0 \in \mathbb{F}_{q^k}} \Psi_k(x_0 u) = \begin{cases} 0 & \text{if } u \in \mathbb{F}_{q^k}^* \\ q^k & \text{if } u = 0. \end{cases}$$

This is a standard result from the theory of additive character sums [13, page 168]. Thus $\sum \Psi_k(x_0 f(x_1, \dots, x_n))$, where the sum is taken over points $x \in \mathbb{F}_{q^k} \times (\mathbb{F}_{q^k}^*)^n$, equals $q^k N_k^*$. Removing the contribution of $(q^k - 1)^n$ from the terms with $x_0 = 0$ in this sum gives the required result. \square

Our next step will be to find an alternative formula for the lefthand side of the equation in Lemma 4. This is achieved in Proposition 11, which leads us eventually to Dwork's Trace Formula (Theorem 26).

3. p-ADIC THEORY

3.1. p-adic rings. We first introduce notation for the p -adic rings we shall need, before explaining how to construct them and compute in them. Let \mathbb{Q}_p be the field of p -adic rationals, and \mathbb{Z}_p the ring of p -adic integers (see [11]). Denote by Ω the completion of an algebraic closure of \mathbb{Q}_p . Select $\pi \in \Omega$ with $\pi^{p-1} = -p$ and define $R_1 = \mathbb{Z}_p[\pi]$, a totally ramified extension of \mathbb{Z}_p of degree $p-1$. By binomial expansion and Hensel's lifting lemma, one sees that the equation $(1 + \pi t)^p = 1$ has exactly p distinct solutions t in R_1 . In particular, R_1 contains all p -th roots of unity. The motivation behind the introduction of R_1 is that to define an additive character of order p , we need a small p -adic ring which contains a primitive p th root of unity.

Let R_0 denote the ring of integers of the unique unramified extension of \mathbb{Q}_p in Ω of degree a , where $q = p^a$. Finally, let R be the compositum ring of R_1 and R_0 . We have the diagram of ring extensions

$$\begin{array}{ccc}
 & \Omega & \\
 & | & \\
 & R & \\
 R_0 & \swarrow \quad \searrow & R_1 \\
 & \mathbb{Z}_p &
 \end{array}$$

The residue class ring of R_0 is \mathbb{F}_q , and we shall “lift” the coefficients of the polynomial $f \in \mathbb{F}_q[X_1, \dots, X_n]$ to this characteristic zero ring, using the Teichmüller lifting. Precisely, the Teichmüller lift $\omega(x)$ of a non-zero element $x \in \mathbb{F}_q$ is defined as the unique root of unity in the maximal unramified extension of \mathbb{Q}_p which is congruent to x modulo p and has order coprime to p . We define $\omega(0) = 0$. Thus the compositum R contains both the lifting of the coefficients of f and the image of an additive character we shall construct.

3.2. Algorithmic aspects.

3.2.1. Construction and lifting Frobenius. We assume that \mathbb{F}_q is presented as the quotient $\mathbb{F}_p[y]/(h)$ where $h(y)$ is a monic, irreducible polynomial of degree a over the prime field \mathbb{F}_p . For any positive integer N we describe how the quotient ring $R/(p^N)$ may be constructed: Lift the polynomial h to an integer polynomial \hat{h} whose coefficients lie in the open interval $(-p/2, (p+1)/2)$. Take the unramified extension R_0 to be $\mathbb{Z}_p[\mu] = \mathbb{Z}_p[y]/(\hat{h})$. Elements in $R_0/(p^N)$ can now be represented as linear combinations over $\mathbb{Z}_p/(p^N)$ of the basis elements $1, \mu, \dots, \mu^{a-1}$. (Recall that $\mathbb{Z}_p/(p^N)$ can be identified with $\mathbb{Z}/(p^N)$.) The extension $R/(p^N) = R_0[\pi]/(p^N)$ is easily constructed by adjoining an element π and specifying the relation $\pi^{p-1} = -p$.

We define a lifting of the Frobenius automorphism on \mathbb{F}_q to an automorphism of R which is the identity on R_1 . Define the map $\tau : R \rightarrow R$ by setting $\tau(\mu)$ to be the unique root of \hat{h} which is congruent to μ^p modulo p . Define $\tau(\pi) = \pi$, and extend to the whole of R by insisting τ is an automorphism. We extend τ to act on $R[[X]]$ coefficient-wise, fixing monomials. Here $R[[X]]$ is the ring of all power series in the indeterminates $X = X_0, \dots, X_n$ with coefficients from R .

3.2.2. Complexity of arithmetic. In this section we bound the complexity of the basic arithmetic operations in the ring $R/(p^N)$, along with that of computing the map τ and Teichmüller lifts. Note that these estimates are all the simplest possible, and can be improved using more advanced methods. The reader may wish to skip the proof of the next lemma, and refer back when required in Section 6.3.2.

Lemma 5. *Elements in $R/(p^N)$ can be represented using $\mathcal{O}(paN \log p)$ bits. Addition and subtraction can be performed in $\mathcal{O}(paN \log p)$ bit operations, and multiplication and inversion of units in $\mathcal{O}((paN \log p)^2)$ bit operations. The Teichmüller lifting to $R/(p^N)$ of a finite field element can be computed in $\mathcal{O}((a \log p)^3 N^2)$ bit operations. For any $1 \leq i \leq a-1$, the map τ^i on $R/(p^N)$ may be evaluated using $\mathcal{O}(p(aN \log p)^2)$ bit operations. (For the powers of τ we require a total of $\mathcal{O}(a^4 N^2 (\log p)^3)$ bits of precomputation.)*

Proof. Elements in $R/(p^N)$ can be written as

$$(1) \quad \sum_{i=0}^{a-1} \sum_{j=0}^{p-2} c_{ij} \pi^j \mu^i$$

where the coefficients c_{ij} belong to the ring $\mathbb{Z}_p/(p^N)$ of size p^N . The bit size of such an expression is $a(p-1) \log(p^N) = \mathcal{O}(paN \log p)$. Addition and subtraction are straightforward, just involving the addition of integers and reduction modulo p^N . Likewise, multiplication of two expansions of the form (1) is straightforward, using the reduction relations $\pi^{p-1} = -p$ and $\hat{h}(\mu) = 0$.

For Teichmüller lifting and inversion we shall use Newton iteration with quadratic convergence. Specifically, we shall use Newton lifting with respect to the prime π in R , and define $l = (p-1)N$ so that $\pi^l = (-p)^N$. Given a polynomial $\phi(Y) \in R[Y]$ and an element $g_0 \in R$ such that $\phi(g_0) \equiv 0 \pmod{\pi}$ and $\phi'(g_0)$ is invertible (modulo π) with inverse s_0 modulo π , this algorithm computes an element $g \in R/(\pi^l)$ such that $\phi(g) \equiv 0 \pmod{\pi^l}$ and $g \equiv g_0 \pmod{\pi}$ (compare with [7, Algorithm 9.22]). For $i \geq 1$, assuming that g_{i-1} and s_{i-1} have been found, we define $g_i = g_{i-1} - \phi(g_{i-1})s_{i-1} \pmod{\pi^{2^i}}$ and $s_i = 2s_{i-1} - \phi'(g_i)s_{i-1}^2 \pmod{\pi^{2^i}}$. As in [7, Theorem 9.23] one checks that $g_i \equiv g_0 \pmod{\pi}$, $\phi(g_i) \equiv 0 \pmod{\pi^{2^i}}$ and $s_i \equiv \phi'(g_i)^{-1} \pmod{\pi^{2^i}}$ at the i th step. Thus after $\lceil \log_2(l) \rceil$ steps we shall have found the required approximate root $g = g_{\lceil \log_2(l) \rceil}$. The i th step involves three additions/subtractions and three

multiplications in the ring $R/(\pi^{2^i})$, which requires $\mathcal{O}((2^i a \log p)^2)$ bit operations, along with evaluation of the polynomials ϕ and ϕ' modulo π^{2^i} . Let $c(\phi, i)$ be the complexity of these latter operations. Thus the total complexity is $\mathcal{O}((paN \log p)^2 + c(\phi))$ bit operations where $c(\phi) = \sum_{i=1}^{\lceil \log_2(b) \rceil} c(\phi, i)$. Note that to compute approximate roots in the ring R_0 rather than R , one can lift using the prime p rather than π and obtain a complexity of $\mathcal{O}((aN \log p)^2 + c(\phi))$.

Suppose now we are given a unit $x \in R/(p^N)$. We compute a Newton lifting starting from an element $g_0 \in R/(\pi)$ such that $xg_0 \equiv 1 \pmod{\pi}$; that is, we use the equation $\phi(Y) = xY - 1$. Now $R/(\pi)$ is just the finite field \mathbb{F}_q and so an inverse g_0 of x modulo π can be computed in $\mathcal{O}((a \log p)^2)$ bit operations [7, Corollary 4.6]. (Here $\phi'(Y) = x$ and $s_0 = g_0$, and so in fact we only need to iterate the formula for s_i .) In this case $c(\phi, i)$ is just one multiplication and a subtraction modulo π^{2^i} . By the above paragraph the total complexity is then $\mathcal{O}((paN \log p)^2)$ for inversion. For Teichmüller lifts we use the same approach, only with the polynomial $\phi(Y) = Y^{q-1} - 1$ and lifting in R_0 via the element p . Here using a fast exponentiation routine we find that $c(\phi, i)$ involves $\mathcal{O}(\log q)$ multiplications and a subtraction in $R_0/(p^{2^i})$. Thus $c(\phi) = \mathcal{O}((aN \log p)^2 \log q)$ which gives the Teichmüller lifting estimate.

We precompute τ on the basis elements $1, \mu, \dots, \mu^{a-1}$. To do this, recall that $\tau(\mu) \in R_0$ is defined as the unique root of \hat{h} which is congruent modulo p to μ^p . This may be approximated modulo p^N by Newton lifting in R_0 with respect to p using the polynomial $\phi(Y) = \hat{h}(Y)$ and the initial value $g_0 = \mu^p$. (Finding μ^p takes $\mathcal{O}(a^2(\log p)^3)$ bit operations and this is absorbed in the stated precomputation estimate.) Here $c(\phi, i)$ is $\mathcal{O}(a)$ additions and multiplications in $R_0/(p^{2^i})$, using Horner's method for polynomial evaluation [7, Page 93]. Thus the Newton lifting estimate gives a complexity of $\mathcal{O}(a(aN \log p)^2)$ bit operations. Using $\tau(\mu^i) = (\tau(\mu))^i$ one can now find the image of all basis elements in a further $\mathcal{O}(a(aN \log p)^2)$ bit operations. One stores this information as a matrix for τ acting on $R_0/(p^N)$ as an $\mathbb{Z}_p/(p^N)$ -module with basis the powers of μ . The map τ can now be computed on any element in $R_0/(p^N)$ in $\mathcal{O}(a^2(N \log p)^2)$ bit operations using linear algebra. By taking powers of the matrix, matrices for the maps τ^i for $1 \leq i \leq a-1$ can also be found in $\mathcal{O}(a^3 a(N \log p)^2)$ bit operations. Thus the total precomputation is bounded by $\mathcal{O}(a^4 N^2 (\log p)^3)$, and each evaluation of τ^i on $R_0/(p^N)$ takes $\mathcal{O}((aN \log p)^2)$ bit operations. Finally, to evaluate τ on $R/(p^N)$ one writes elements of R on the R_0 -basis $1, \pi, \dots, \pi^{p-2}$ and applies τ component-wise. \square

3.3. p -adic valuations and convergence of power series. Denote by ord the additive valuation on Ω normalised so that $\text{ord}(p) = 1$. Thus $\text{ord}(\pi) = 1/(p-1)$. Define a p -adic norm $|\cdot|_p$ on Ω by $|x|_p = p^{-\text{ord}(x)}$. The set of all $x \in \Omega$ with $|x|_p \leq 1$ (equivalently $\text{ord}(x) \geq 0$) is called the closed

unit disk. Given any formal power series $\sum_r A_r X^r$ where $A_r \in \Omega$, we say it converges at a point $x = (x_0, \dots, x_n) \in \Omega^{n+1}$ if the sequence of partial sums $\sum_{r, |r| < e} A_r x^r$ tends to a limit under the p -adic norm. (Here $|r| = \sum_{i=0}^n r_i$.) This sequence of partial sums will converge if and only if the summands $A_r x^r$ tend to zero p -adically (that is, are divisible in the ring of integers of Ω by increasingly large powers of p), as $|r|$ goes to infinity. In particular, if there is a real number $c > 0$ such that $\text{ord}(A_r) \geq c|r|$ for all r , then the series will certainly converge for all points x which are Teichmüller liftings of points over \mathbb{F}_q . Note that throughout the paper we shall use the additive valuation ord rather than the p -adic norm $|\cdot|_p$ itself.

4. ANALYTIC REPRESENTATION OF CHARACTERS

4.1. Dwork's splitting functions. We now need to find a suitable p -adic expression for a non-trivial additive character from \mathbb{F}_q to R_1 . In the case of complex characters, this is done via the exponential function. However, the radius of convergence of the exponential function in Ω is not large enough; in particular, it does not converge on the Teichmüller lifting of all the points in \mathbb{F}_q . Instead we use the power series constructed by Dwork using the exponential function. (The reader may find the discussion of the related Artin-Hasse function on [11, Pages 92-93] helpful.)

Let μ be the Möbius function. Taking the logarithmic derivative, one checks that the exponential function has the following product expansion

$$\exp(z) = \sum_{k=0}^{\infty} \frac{z^k}{k!} = \prod_{k=1}^{\infty} (1 - z^k)^{-\mu(k)/k}.$$

This can be rewritten as

$$\exp(z) = \prod_{(k,p)=1}^{\infty} (1 - z^k)^{-\mu(k)/k} (1 - z^{kp})^{\mu(k)/(kp)}.$$

It follows that

$$(2) \quad \exp\left(z + \frac{z^p}{p}\right) = \prod_{(k,p)=1}^{\infty} (1 - z^k)^{-\mu(k)/k} (1 - z^{kp^2})^{\mu(k)/(kp^2)}.$$

Replacing z by πz in the above relation and noting that $\pi^p = -p\pi$, we define a power series in z by

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

Writing $\theta(z) = \sum_{r=0}^{\infty} \lambda_r z^r$ we see that $\lambda_r = \pi^r/r!$ for $r < p$, and we shall shortly show that all λ_r lie in R_1 . From (2) we get the product expansion

$$(3) \quad \theta(z) = \prod_{(k,p)=1}^{\infty} (1 - \pi^k z^k)^{-\mu(k)/k} (1 - \pi^{kp^2} z^{kp^2})^{\mu(k)/(kp^2)}.$$

By the binomial expansion, the first factor

$$(1 - \pi^k z^k)^{-\mu(k)/k} = \sum_{j=0}^{\infty} (-1)^j \binom{-\mu(k)/k}{j} \pi^{jk} z^{jk} = \sum_{j=0}^{\infty} b_j(k) z^{jk}$$

is a power series in z^k . Now for $(k, p) = 1$, we have [11, Page 82]

$$-\mu(k)/k \in \mathbb{Z}_p, \quad \binom{-\mu(k)/k}{j} \in \mathbb{Z}_p.$$

Thus for $j > 0$ the coefficient of z^{jk} satisfies

$$\text{ord}(b_j(k)) \geq \frac{jk}{p-1} > \frac{p-1}{p^2} jk.$$

Similarly, for the second factor, we write

$$(1 - \pi^{kp^2} z^{kp^2})^{\mu(k)/(kp^2)} = \sum_{j=0}^{\infty} (-1)^j \binom{\mu(k)/(kp^2)}{j} (\pi z)^{jkp^2} = \sum_{j=0}^{\infty} c_j(kp^2) z^{jkp^2}.$$

Now for $j > 0$ we have $\text{ord}(j!) < \frac{j}{p-1}$, see [11, Page 79]. Thus for $j > 0$ and k positive and coprime to p , the coefficient of z^{jkp^2} satisfies

$$\text{ord}(c_j(kp^2)) > \frac{jkp^2}{p-1} - 2j - \frac{j}{p-1} \geq \frac{p-1}{p^2} jkp^2.$$

Putting the above two inequalities together, we conclude that for $r > 0$ the coefficients λ_r of $\theta(z)$ satisfy

$$(4) \quad \text{ord}(\lambda_r) > \frac{(p-1)r}{p^2}, \quad \lambda_r \in R_1.$$

This shows that the power series $\theta(z)$ is convergent in the disk $|z|_p < 1 + \epsilon$ for some $\epsilon > 0$. In particular, $\theta(z)$ converges on the closed unit disk, and Definition 6 makes sense.

In the proof of the next lemma we shall use the fact that $\text{ord}(\lambda_r) \geq 2/(p-1)$ for $r \geq 2$, and so

$$(5) \quad \theta(1) \equiv 1 + (\pi z) \pmod{(\pi z)^2}.$$

This can be seen as follows: The Artin-Hasse exponential function [11, Page 93]

$$E(z) := \prod_{(k,p)=1}^{\infty} (1 - z^k)^{-\mu(k)/k} = \exp \left(z + \frac{z^p}{p} + \frac{z^{p^2}}{p^2} + \dots \right)$$

has coefficients in \mathbb{Z}_p , because each factor in the product expansion does. Since $E(\pi z) \equiv \theta(z) \pmod{z^{p^2}}$ we see

$$\text{ord}(\lambda_r) \geq \frac{r}{p-1}$$

for $0 \leq r < p^2$. This estimate combined with (4) gives (5).

Definition 6. [Dwork's splitting function] Let

$$\Phi_k(z) = \prod_{i=0}^{ak-1} \theta(z^{p^i}) \in R_1[[z]],$$

and

$$\Psi_k = \Phi_k \circ \omega : \mathbb{F}_{q^k} \rightarrow R_1,$$

where ω is the Teichmüller map. (Recall that $q = p^a$.)

(That Ψ_k has image in R_1 can be seen as follows: Let $R_1[\omega(\mathbb{F}_{q^k})]$ be R_1 adjoined the image of ω on \mathbb{F}_{q^k} . Then $R_1[\omega(\mathbb{F}_{q^k})]$ is an unramified extension of R_1 of degree k . The Galois group of the corresponding quotient field extension is generated by τ . The map τ acts on a Teichmüller point $\omega(x)$ as $\tau(\omega(x)) = \omega(x)^p$. Hence it fixes the element $\Psi_k(x)$ for $x \in \mathbb{F}_{q^k}$, and so $\Psi_k(x) \in R_1$.)

Lemma 7. The maps Ψ_k form a tower of non-trivial additive characters from the fields \mathbb{F}_{q^k} to the ring R_1 .

Proof. (This is the case “ $s = 1$ ” on [5, Pages 55-57].) We first show that $\theta(1)$ is a primitive p th root of unity. By (5) we see $\theta(1) \neq 1$. As a formal power series in z ,

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

Now, $\theta(z)$, $\exp(p\pi z)$ and $\exp(-p\pi z^p)$ are all convergent in $|z|_p < 1 + \varepsilon$ for some $\varepsilon > 0$. We can thus substitute $z = 1$ and find that $\theta(1)^p = \exp(p\pi) \exp(-p\pi) = 1$. Thus $\theta(1)$ is a primitive p -th root of unity in R_1 .

Next, for $\gamma \in R_0$ with $\gamma^{p^{ak}} = \gamma$, we claim that

$$\prod_{i=0}^{ak-1} \theta(\gamma^{p^i}) = \theta(1)^{\gamma + \gamma^p + \dots + \gamma^{p^{ak-1}}}.$$

Using (5) it is clear that both sides are congruent to

$$1 + \pi(\gamma + \gamma^p + \dots + \gamma^{p^{ak-1}})$$

modulo π^2 . To prove the claim, it remains to prove that both sides are p -th roots of unity. The right side is a p -th root of unity since $\theta(1)$ is a p -th root of unity. The p -th power of the left side is

$$\prod_{i=0}^{ak-1} \theta(\gamma^{p^i})^p = \exp(p\pi \sum_{i=0}^{ak-1} (\gamma^{p^i} - \gamma^{p^{i+1}})) = \exp(p\pi\gamma) \exp(-p\pi\gamma^{p^{ak}}) = 1.$$

Thus, the left side is also a p -th root of unity. The claim is proved. Note that the individual factor $\theta(\gamma)$ is not necessarily a p -th root of unity. We conclude that

$$\Psi_k(x) = \theta(1)^{\text{Tr}_k(x)}$$

for any $x \in \mathbb{F}_{q^k}$, where Tr_k is the trace function from \mathbb{F}_{q^k} to \mathbb{F}_p , and the exponent is thought of as an integer. \square

Note 8 The infinite sum

$$\exp(\pi(z - z^p)) = \sum_{k=0}^{\infty} \frac{(\pi(z - z^p))^k}{k!}$$

is convergent for $|z|_p < 1$, but not necessarily convergent for $|z|_p = 1$. If $|z|_p = 1$, then it is possible that $|z - z^p|_p = 1$ and for such z the above infinite sum does not converge. Since the infinite sum does not converge everywhere on the disk $|z|_p \leq 1$, one cannot simply substitute $z = 1$ into the above infinite sum and get the contradiction that $\theta(1) = 1$. There is no contradiction here!

We now define another power series related to our original polynomial f whose relevance will become apparent in Proposition 11.

Definition 9. Let f be the polynomial whose zeta function we wish to compute, and write

$$X_0 f = \sum_{j \in J} \bar{a}_j X^j,$$

where J is the support of $X_0 f$. Let a_j be the Teichmüller lifting of \bar{a}_j . Let F be the formal power series in the indeterminates X with coefficients in R given by

$$F(X) = \prod_{j \in J} \theta(a_j X^j).$$

Let $F^{(a)}(X)$ be the formal power series in the indeterminates X with coefficients in R given by

$$F^{(a)}(X) = \prod_{j \in J} \prod_{s=0}^{a-1} \theta((a_j X^j)^{p^s}).$$

The relation between $F(X)$ and $F^{(a)}(X)$ is clear.

Lemma 10. Let the power series $F^{(a)}$ and F be as in Definition 9. Then

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

where the map τ acts coefficient-wise on the power series F .

The power series $F^{(a)}(X)$ relates to rational point counting in the following way.

Proposition 11. Let $f \in \mathbb{F}_q[X_1, \dots, X_n]$ and $F^{(a)} \in R[[X]]$ be as in Definition 9. Denoting by N_k^* the number of solutions to the equation $f = 0$ in the affine torus $(\mathbb{F}_{q^k}^*)^n$, we have

$$q^k N_k^* - (q^k - 1)^n = \sum_{x^{q^k-1}=1} F^{(a)}(x) F^{(a)}(x^q) \dots F^{(a)}(x^{q^{k-1}})$$

where the summation is over the Teichmüller lifting of points on the torus $(\mathbb{F}_{q^k}^*)^{n+1}$.

Proof. For any point \bar{x} in $(\mathbb{F}_{q^k})^{n+1}$ with Teichmüller lifting x we have that

$$\begin{aligned}
\Psi_k(\bar{x}_0 f(\bar{x}_1, \dots, \bar{x}_n)) &= \Psi_k(\sum_{j \in J} \bar{a}_j \bar{x}^j) \\
&= \prod_{j \in J} \Psi_k(\bar{a}_j \bar{x}^j) \\
&= \prod_{j \in J} \Phi_k(a_j x^j) \\
&= \prod_{j \in J} \{\prod_{i=0}^{ak-1} \theta((a_j x^j)^{p^i})\} \\
&= \prod_{j \in J} \{\prod_{i=0}^{k-1} \prod_{s=0}^{a-1} \theta((a_j x^j)^{q^i p^s})\} \\
&= \prod_{i=0}^{k-1} \{\prod_{j \in J} (\prod_{s=0}^{a-1} \theta((a_j^{q^i} x^{jq^i})^{p^s}))\} \\
&= \prod_{i=0}^{k-1} \{\prod_{j \in J} (\prod_{s=0}^{a-1} \theta((a_j (x^{q^i})^j)^{p^s}))\} \\
&= F^{(a)}(x) F^{(a)}(x^q) \dots F^{(a)}(x^{q^{k-1}}),
\end{aligned}$$

where $F^{(a)}(X)$ is given by

$$F^{(a)}(X) = \prod_{j \in J} \prod_{s=0}^{a-1} \theta((a_j X^j)^{p^s}).$$

(We pause to justify the steps above: the first four equalities follow straight from definitions and from the homomorphic property of Ψ_k ; the fifth and sixth by rearrangement; and the seventh since a_j satisfies $a_j^q = a_j$.)

Thus we have

$$\sum_{\bar{x} \in (\mathbb{F}_{q^k}^*)^{n+1}} \Psi_k(\bar{x}_0 f(\bar{x}_1, \dots, \bar{x}_n)) = \sum_{x^{q^k-1}=1} F^{(a)}(x) F^{(a)}(x^q) \dots F^{(a)}(x^{q^{k-1}})$$

where the latter sum is over the Teichmüller lifting in Ω^{n+1} of points in $(\mathbb{F}_{q^k}^*)^{n+1}$. Combining this with Lemma 4 gives us the result. \square

4.2. Decay rates and weight functions. We now describe the decay rates of the coefficients of the power series $F^{(a)}$ and F . Specifically, we obtain lower bounds for the p -adic order of the coefficients of the power series F expressed in terms of a certain weight function on integer vectors.

Write $F = \sum_r F_r X^r$ where the sum is over non-negative integer vectors in $\mathbb{Z}_{\geq 0}^{n+1}$. Let A be the $(n+1) \times |J|$ matrix whose columns are $j = (j_0, j_1, \dots, j_n) \in J$. Then from Definition 9 one sees

$$(6) \quad F_r = \sum_u \left(\prod_{j \in J} \lambda_{u_j} a_j^{u_j} \right)$$

where the outer sum is over all $|J|$ -tuples $u = (u_j)$ of non-negative integers such that

$$(7) \quad Au = r,$$

thinking of u and r as column vectors. Since $j_0 = 1$ for all $(j_0, \dots, j_n) \in J$, the first row of the matrix A is the vector $(1, 1, \dots, 1)$. The first equation

in the above linear system is then

$$(8) \quad \sum_{j \in J} u_j = r_0.$$

Now F_r is zero if (7) has no solutions. Otherwise, since $\text{ord}(\lambda_{u_j} a_j^{u_j}) = \text{ord}(\lambda_{u_j})$ from (4),(6) and (8) we get

$$(9) \quad \text{ord}(F_r) \geq \inf_u \left\{ \sum_{j \in J} \frac{(p-1)u_j}{p^2} \right\} = \frac{p-1}{p^2} r_0,$$

where the inf is over all non-negative integer vector solutions u of (7). We now define a weight function w such that (9) gives estimates on $\text{ord}(F_r)$ in terms of this weight function.

Let $\delta_1 \subset \mathbb{R}^n$ denote the convex hull of the support of f (the set of exponents of non-zero terms). Let $\delta_2 \subset \mathbb{R}^n$ be the convex hull of the origin and the n points $(d, 0, \dots, 0), (0, d, \dots, 0), \dots, (0, \dots, 0, d)$ where d is the total degree of f . We call δ_1 the Newton polytope of f ; the polytope δ_2 is just a simplex containing δ_1 .

Definition 12. Let δ be any convex polytope with integer vertices such that $\delta_1 \subseteq \delta \subseteq \delta_2$. Denote by Δ the convex polytope in \mathbb{R}^{n+1} obtained by embedding δ in \mathbb{R}^{n+1} via the map $x \mapsto (1, x)$ for $x \in \mathbb{R}^n$, and taking the convex hull with the origin. Denote by $C(\Delta)$ the cone generated in \mathbb{R}^{n+1} as the positive hull of Δ . So $C(\Delta)$ is the union of all rays emanating from the origin and passing through Δ .

Ultimately, in Sections 6.3.3 and 6.4, we shall only be interested in the simplest choice of polytope $\delta = \delta_2$ (although the choice $\delta = \delta_1$ leads to the most refined algorithm). Letting Δ_1 denote the polytope in \mathbb{R}^{n+1} obtained by choosing $\delta = \delta_1$ we see that $C(\Delta_1)$ is the cone generated by the exponents of non-zero terms in $X_0 f$. Equation (7) has no non-negative integer (or even real) solutions when r does not lie in $C(\Delta_1)$. Thus for any choice of $\delta (\supseteq \delta_1)$ and corresponding $\Delta (\supseteq \Delta_1)$, all exponents of F lie in the cone $C(\Delta)$.

Definition 13. Define a weight function w from \mathbb{R}^{n+1} to $\mathbb{R} \cup \{\infty\}$ in the following way: For $r = (r_0, r_1, \dots, r_n) \in \mathbb{R}^{n+1}$ define

$$w(r) = \begin{cases} r_0 & \text{if } r \in C(\Delta), \\ \infty & \text{otherwise.} \end{cases}$$

In particular $w(r)$ is a non-negative integer for any $r \in C(\Delta) \cap \mathbb{Z}^{n+1}$.

Note 14 Choosing $\delta = \delta_1$ corresponds to working with the weight function of Adolphson and Sperber [2], and taking $\delta = \delta_2$ to Dwork's original weight function [4]. Note that Dwork's weight function can equivalently be defined as $w(r) = r_0$ if $r_1 + \dots + r_n \leq r_0 d$ and ∞ otherwise, where d is the total degree of f .

The weight function has a simple geometric interpretation: For a real number c define $c\Delta = \{cx \mid x \in \Delta\}$. The next lemma is straightforward.

Lemma 15. *For any point $r \in C(\Delta)$ we have that $w(r)$ is the smallest non-negative number c such that $r \in c\Delta$. If $r \notin C(\Delta)$ then $w(r) = \infty$.*

By (9) and the sentence preceding Definition 13, we have the following result.

Lemma 16. *We have the inequality*

$$\text{ord}(F_r) \geq w(r) \frac{p-1}{p^2}.$$

We also note the following simple property whose proof is straightforward.

Lemma 17. *Let $r, r' \in \mathbb{R}^{n+1}$ and k a non-negative integer. Then $w(kr) = kw(r)$ and $w(r + r') \leq w(r) + w(r')$. In particular when $w(r') \neq \infty$,*

$$w(kr - r') \geq kw(r) - w(r').$$

We shall work in certain subrings of $R[[X]]$ defined in terms of the weight function.

Definition 18. *Define L_Δ to be the subring of $R[[X]]$ given by*

$$L_\Delta = \left\{ \sum_{r \in C(\Delta) \cap \mathbb{Z}^{n+1}} A_r X^r \mid A_r \in R \right\}.$$

Thus, L_Δ is just the ring of all power series over R whose terms have exponents in the cone $C(\Delta)$. Certainly $F \in L_\Delta$ and from Lemma 10 we see easily that $F^{(a)} \in L_\Delta$.

Lemma 19. *The power series F and $F^{(a)}$ both belong to L_Δ .*

This concludes all results in this section which shall be essential to the proof of our modular version of Dwork's Trace Formula. We conclude with a definition and some comments which we will refer to in the analysis of the running time of our algorithm.

Definition 20. *For any positive real number b , we define a set of power series by*

$$L_\Delta(b) = \left\{ \sum_r A_r X^r \in L_\Delta \mid \text{ord } A_r \geq bw(r) \right\}.$$

The set $L_\Delta(b)$ is easily seen to be a subring of L_Δ . Elements in $L_\Delta(b)$ for large b can be thought of as having fast decaying coefficients. Such rings will reduce to rings of small dimension modulo small powers of p .

We have that

$$(10) \quad F \in L_\Delta \left(\frac{p-1}{p^2} \right)$$

$$(11) \quad F^{(a)} \in L_\Delta \left(\frac{p-1}{qp} \right).$$

The first is immediate from the inequality in Lemma 16 and the second follows since $\tau^i(F(X^{p^i})) \in L_\Delta(\frac{p-1}{p^i p^2})$ for each $0 \leq i \leq a-1$. Thus for $a > 1$ the coefficients of $F^{(a)}$ decay more slowly than those of F itself.

5. DWORK'S TRACE FORMULA

5.1. Lifting Frobenius. We now introduce Dwork's "left inverse of Frobenius" mapping ψ_p on the ring $R[[X]]$.

Definition 21. Let ψ_p be defined on the monomials in $R[[X]]$ by

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

and extend ψ_p by τ^{-1} -linearity to all of $R[[X]]$. That is

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

Here $p|r$ means that p divides all of the entries in the integer vector r . This map is a left inverse of the "Frobenius" map on the ring $R[[X]]$ which takes a power series $\sum_r A_r X^r$ to $\sum_r \tau(A_r) X^{pr}$.

Definition 22. Let α_a be the map from $R[[X]]$ to itself defined as

$$\alpha_a = \psi_p^a \circ F^{(a)}.$$

Precisely, this is the map which is the composition of multiplication by the power series $F^{(a)}$ followed by the mapping ψ_p^a on the ring $R[[X]]$. (Notice that ψ_p^a just acts as

$$\psi_p^a\left(\sum_r A_r X^r\right) = \sum_r A_{ar} X^r$$

since $\tau^{-a}(A_r) = A_{ar}$.) Let the map α from $R[[X]]$ to itself be defined as

$$\alpha = \psi_p \circ F.$$

Thus α is multiplication by F followed by the mapping ψ_p .

We have the following result relating these two maps, which shall be of crucial importance in our derivation of an efficient algorithm.

Lemma 23. With α_a and α as in Definition 22 we have

$$\alpha_a = \alpha^a$$

where the second exponent is a power under composition.

Proof. Firstly let $H \in R[[X]]$ and denote by $\psi_p \circ H(X^p)$ the map composed of multiplication by $H(X^p)$ followed by ψ_p . We claim that

$$(12) \quad \psi_p \circ H(X^p) = \tau^{-1}(H(X)) \circ \psi_p$$

To see this write $H = \sum_r H_r X^{pr}$. Then we have that

$$\psi_p \circ H(X^p) = \sum_r \tau^{-1}(H_r) (\psi_p \circ X^{pr}).$$

Here the infinite series is interpreted as a mapping. Now $\psi_p \circ X^{p^r} = X^r \circ \psi_p$ as these two maps are τ^{-1} -linear and agree on monomials. Hence we have $\psi_p \circ H(X^p) = \sum_r \tau^{-1}(H_r)(X^r \circ \psi_p) = \tau^{-1}(H(X)) \circ \psi_p$.

Next we claim that for any $b \geq 1$ and power series H we have

$$\psi_p^b \circ \prod_{i=0}^{b-1} \tau^i(H(X^{p^i})) = (\psi_p \circ H(X))^b.$$

We prove this by induction, the result trivially holding if $b = 1$. For $b > 1$, by $b - 1$ applications of (12) we get

$$\psi_p^b \circ \prod_{i=0}^{b-1} \tau^i(H(X^{p^i})) = (\psi_p \circ H(X)) \circ (\psi_p^{b-1} \circ \prod_{i=0}^{b-2} \tau^i(H(X^{p^i}))).$$

The second claim then follows by induction. Putting $H = F$ and $b = a$ we get the required result. \square

The map α_a is linear and continuous, in the sense that

$$\alpha_a\left(\sum_r A_r X^r\right) = \sum_r A_r \alpha_a(X^r)$$

for any element $\sum_r A_r X^r \in R[[X]]$. The map α is τ^{-1} -linear and continuous, in the sense that

$$\alpha\left(\sum_r A_r X^r\right) = \sum_r \tau^{-1}(A_r) \alpha(X^r).$$

From Lemma 19 the next lemma follows easily.

Lemma 24. *The subring L_Δ is stable under both α and α_a .*

Both maps when restricted to the subring L_Δ are determined by their action on the monomials X^r (with $w(r) < \infty$), which we now consider.

5.2. Matrix representations of mappings. Recall that $C(\Delta)$ is the cone in \mathbb{R}^{n+1} from Definition 12. The set

$$\Gamma_\Delta = \{X^u \mid u \in C(\Delta) \cap \mathbb{Z}_{\geq 0}^{n+1}\}$$

written as a row vector, is a *formal basis* for the space L_Δ . Precisely, this means that any power series in L_Δ may be written in exactly one way as an infinite sum $\sum_u A_u X^u$ with X^u in the above set. Notice that this is different from the usual notion of a basis in linear algebra, since we allow infinite combinations of basis elements. It is also different from the notion of an “orthonormal basis” in the literature, where one requires that the coefficient A_u goes to zero as $w(u)$ goes to ∞ (see [21] for a more detailed discussion of these notions).

By Lemma 24, both α and α_a send the ring L_Δ to itself. We define certain matrices associated to the maps α and α_a restricted to L_Δ with regard to the formal row basis Γ_Δ of monomials.

Definition 25. Let the infinite matrices M and M_a have columns describing the images of the monomials $X^v \in L_\Delta$ under the maps α and α_a with respect to our formal row basis Γ_Δ :

$$\alpha(\Gamma_\Delta) = \Gamma_\Delta M, \quad \alpha_a(\Gamma_\Delta) = \Gamma_\Delta M_a.$$

Specifically, the (u, v) th entries of M and M_a for $u, v \in C(\Delta)$ are m_{uv} and $m_{uv}^{(a)}$ respectively where

$$\begin{aligned} m_{uv} &= \tau^{-1}(F_{pu-v}) \\ m_{uv}^{(a)} &= F_{qu-v}^{(a)}. \end{aligned}$$

Here $F^{(a)} = \sum_r F_r^{(a)} X^r$ and as before $F = \sum_r F_r X^r$, and we take the coefficients of exponents r with negative entries to be zero.

(Note that as yet we have not ordered the basis; however, we shall choose a convenient ordering in the proof of Theorem 28.) We have by Lemmas 16 and 17

$$\begin{aligned} \text{ord}(F_{pu-v}) &\geq \frac{p-1}{p^2} w(pu-v) \\ (13) \quad &\geq \frac{p-1}{p} \left(w(u) - \frac{1}{p} w(v) \right) \end{aligned}$$

and certainly $\text{ord}(m_{uv}) = \text{ord}(\tau^{-1}(F_{pu-v})) = \text{ord}(F_{pu-v})$.

Notice that the matrix powers M_a^k and M^k are defined for every positive integer k since the entries in M_a^k , say, are just finite sums of the entries in M_a and M_a^{k-1} . This follows for M_a , say, since all entries $m_{uv}^{(a)}$ in M_a are zero when the vector $qu-v$ contain negative entries. We define the trace of an infinite matrix to be the sum of its diagonal entries, when this sum converges, and ∞ when the sum does not converge. We shall see shortly that the trace of the infinite matrix M_a^k is finite. We write this as $\text{Tr}(M_a^k)$.

We can now prove the following result of Dwork.

Theorem 26 (Dwork's Trace Formula). *Let $f \in \mathbb{F}_q[X_1, \dots, X_n]$ and let α_a be the mapping on the ring $R[[X]]$ given as $\alpha_a = \psi_p^a \circ F^{(a)}$, where $F^{(a)}$ and ψ_p are described in Definitions 9 and 21. Let M_a denote the infinite matrix representing the map α_a restricted to the subring L_Δ as described in Definition 25. For $k \geq 1$, denote by N_k^* the number of solutions to the equation $f = 0$ in the torus $(\mathbb{F}_{q^k}^*)^n$. Then*

$$(q^k - 1)^{n+1} \text{Tr}(M_a^k) = q^k N_k^* - (q^k - 1)^n.$$

Proof. (The following is a hybrid of the matrix proof given in [19] and Dwork's original argument [4].)

By Proposition 11 we have that

$$q^k N_k^* - (q^k - 1)^n = \sum_{x^{q^k-1}=1} F^{(a)}(x) F^{(a)}(x^q) \dots F^{(a)}(x^{q^{k-1}})$$

where the sum is over all $(n+1)$ -tuples of $(q^k - 1)$ st roots of unity in Ω (namely the Teichmüller lifting in Ω^{n+1} of points on the torus $(\mathbb{F}_{q^k}^*)^{n+1}$).

We first consider the case $k = 1$. Since $F^{(a)} \in L_\Delta$, we can write $F^{(a)}(X) = \sum_r F_r^{(a)} X^r$ where the sum is over all lattice vectors r which belong to $C(\Delta)$. Then the latter sum is

$$\begin{aligned} \sum_{x^{q-1}=1} F^{(a)}(x) &= \sum_{x^{q-1}=1} (\sum_r F_r^{(a)} x^r) \\ &= \sum_r F_r^{(a)} (\sum_{x^{q-1}=1} x^r) \\ &= (q-1)^{n+1} \sum_{r, (q-1)|r} F_r^{(a)} \\ &= (q-1)^{n+1} \sum_s F_{(q-1)s}^{(a)} \\ &= (q-1)^{n+1} \text{Tr}(M_a). \end{aligned}$$

Here by $(q-1)|r$ we mean that the integer $q-1$ divides every entry in the vector r . Also, we use the fact that for any integer r_i [11, page 120]

$$\sum_{x_i^{q-1}=1} x_i^{r_i} = \begin{cases} q-1 & \text{if } (q-1)|r_i \\ 0 & \text{otherwise.} \end{cases}$$

For $k > 1$, define M_{ak} to be the matrix for the map $\psi_p^{ak} \circ \prod_{i=0}^{k-1} F^{(a)}(X^{q^i})$ with respect to the formal basis Γ_Δ . Then by an analogous argument to that in the case $k = 1$ we see

$$\sum_{x^{q^k-1}=1} \prod_{i=0}^{k-1} F^{(a)}(x^{q^i}) = (q^k - 1)^{n+1} \text{Tr}(M_{ak}).$$

Using (12) in the proof of Lemma 23 one sees that

$$\psi_p^{ak} \circ \prod_{i=0}^{k-1} F^{(a)}(X^{q^i}) = (\psi_p^a \circ F^{(a)}(X))^k.$$

Since both maps are linear it follows that $M_{ak} = M_a^k$, and the theorem is proved. \square

To compute the trace of the matrix in Dwork's formula we shall use the following matrix identity derived from Lemma 23.

Lemma 27. *Let M_a and M denote the matrices for the maps α_a and α described above. Then*

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

where the map τ^{-i} acts entry wise on the matrix M .

Proof. First suppose that N_1 and N_2 are matrices representing maps β_1 and β_2 on some subspace of $R[[X]]$. We assume that β_1 is τ^{-1} -linear and β_2 is τ^{-j} -linear for some $j \in \mathbb{Z}$. Then it is not difficult to prove that the matrix for the $\tau^{-(j+1)}$ -linear map $\beta_1 \circ \beta_2$ is just $N_1 \tau^{-1}(N_2)$.

By Lemma 23 we have $\alpha_a = \alpha^a$. We claim that for any positive integer b the matrix for the map α^b is $\prod_{i=0}^{b-1} \tau^{-i}(M)$. The result is trivially true if $b = 1$. For $b > 1$ we have $\alpha^b = \alpha \circ \alpha^{b-1}$. By induction the matrix for α^{b-1}

is $\prod_{i=0}^{b-2} \tau^{-i}(M)$. By τ^{-1} -linearity of α it follows from the observations in the preceding paragraph that the matrix for α^b is $M\tau^{-1}(\prod_{i=0}^{b-2} \tau^{-i}(M)) = \prod_{i=0}^{b-1} \tau^{-i}(M)$. The required result now follows by taking $b = a$. \square

We note in passing that M_a in Lemma 27 also equals

$$\tau^{a-1}(S)\tau^{a-2}(S)\dots S,$$

where $S = \tau(M)$ is the matrix with (u, v) th entry simply F_{pu-v} . Although this is slightly more desirable from a practical point of view we shall not use this expression for M_a .

5.3. Modular reduction of the trace formula. We now examine the reduction of the Dwork trace formula modulo a power p^N of p . Observe that both sides in the trace formula, and all the entries in the matrices M and M_a , are elements of R , and so this reduction is defined.

We first recall some notation: Let $C(\Delta)$ be the cone in \mathbb{R}^{n+1} from Definition 12, and L_Δ denote the ring of power series over R whose monomials have exponents lying in $C(\Delta)$ (Definition 18). Let M denote the matrix for the τ^{-1} -linear map $\alpha = \psi_p \circ F$ with respect to the formal row basis Γ_Δ (Definition 25). Here ψ_p is the “left inverse of Frobenius” given in Definition 21 and F is the power series obtained from the polynomial f as in Definition 9.

Theorem 28. *Let N denote any positive integer and A_N the finite square matrix over the finite ring $R/(p^N)$ obtained by reducing modulo p^N all those entries in M whose rows and columns are indexed by vectors $u \in C(\Delta) \cap \mathbb{Z}_{\geq 0}^{n+1}$ with $w(u) < (p/(p-1))^2 N$. Then*

$$(q^k - 1)^{n+1} \text{Tr}(\left(\prod_{i=0}^{a-1} \tau^{-i}(A_N)\right)^k) = q^k N_k^* - (q^k - 1)^n \pmod{p^N}$$

where N_k^* is the number of solutions to the equation $f = 0$ in the affine torus $(\mathbb{F}_{q^k}^*)^n$. Moreover, the size of A_N is $W = \#(t\Delta)$, where $t = \lceil p^2 N / (p-1)^2 \rceil - 1$ and $\#(t\Delta)$ is number of lattice points in a dilation by a factor t of the polytope Δ .

Proof. For any finite or infinite matrix L with coefficients in R , we define \overline{L} to be the matrix obtained by reducing all its entries modulo p^N . Thus \overline{L} has entries in $R/(p^N)$.

The theorem will follow from the Dwork Trace Formula (Theorem 26) once we find a suitable expression for the reduction modulo p^N of the trace of the matrix M_a^k . This is equal to the trace of the matrix $\overline{M_a^k} = (\overline{M_a})^k$. By Lemma 27 this matrix can be computed as a matrix product from the matrix \overline{M} .

By inequality (13) every entry $m_{uv} = \tau^{-1}(F_{pu-v})$ in M satisfies

$$\text{ord}(m_{uv}) \geq \frac{p-1}{p} \left(w(u) - \frac{w(v)}{p} \right).$$

Define t to be the greatest integer less than $(p/(p-1))^2 N$. When $w(u) \geq w(v)$ and $w(u) > t$ we have

$$(14) \quad \frac{p-1}{p} \left(w(u) - \frac{w(v)}{p} \right) \geq \frac{(p-1)^2}{p^2} w(u) \geq \frac{(p-1)^2}{p^2} (t+1) \geq N.$$

Recall that we have not yet ordered the basis. Now choose any total ordering on the basis set such that for distinct lattice points $u, v \in C(\Delta)$, the monomial X^v comes before X^u if $w(v) < w(u)$. Thus, by the inequalities in (14) and the choice of ordering of the basis, the matrix \overline{M} is of the form

$$(15) \quad \begin{pmatrix} A_N & B_N \\ 0 & C_N \end{pmatrix}$$

where C_N is a strictly upper triangular infinite matrix and A_N is the finite square matrix indexed by lattice points $u \in C(\Delta)$ such that $w(u) \leq t$. The size of A_N is the number of lattice points u with $w(u) \leq t$. By Lemma 15 this is exactly the number of lattice points in the polytope $t\Delta$.

By Lemma 27 and modular reduction, one has

$$(\overline{M}_a)^k = \left(\prod_{i=0}^{a-1} \tau^{-i}(\overline{M}) \right)^k.$$

By (15) we see that $(\overline{M}_a)^k$ is of the form

$$\begin{pmatrix} (\prod_{i=0}^{a-1} \tau^{-i}(A_N))^k & B'_N \\ 0 & C'_N \end{pmatrix}$$

where $C'_N = (\prod_{i=0}^{a-1} \tau^{-i}(C_N))^k$ is strictly upper triangular.

Hence the trace of $(\overline{M}_a)^k$ equals the trace of the finite matrix

$$(16) \quad \left(\prod_{i=0}^{a-1} \tau^{-i}(A_N) \right)^k.$$

The theorem now follows from (16) and Theorem 26. \square

6. ALGORITHMS

In this section we present an algorithm for counting points based upon Theorem 28, and complete the proofs of the results in the introduction. This is a relatively straightforward matter, although the precise complexity estimates require a little care.

6.1. Toric point counting algorithm. We first give the algorithm.

Algorithm 29 Toric point counting

Input: Positive integers a, k, n, d and a prime p ; a polynomial f ; a polytope Δ . (Here f is a polynomial in n variables of total degree d with coefficients in the field \mathbb{F}_q where $q = p^a$. The polytope Δ is as in Definition 12. We assume a model of \mathbb{F}_q is given as in Section 3.2.1.)

Output: The number of solutions N_k^* to the equation $f = 0$ in the torus

$$(\mathbb{F}_{q^k}^*)^n.$$

Step 0: Set $N = (n + 1)ak$, where $q = p^a$.

Step 1: Compute the polynomial $F \bmod p^N$ in the ring $(R/(p^N))[X]$ where F is the power series in Definition 9. Specifically, writing $X_0 f = \sum_{j \in J} \bar{a}_j X^j$ we have $F = \prod_{j \in J} \theta(a_j X^j) \bmod p^N$. Here a_j is the Teichmüller lifting of \bar{a}_j and $\theta(z) = \exp(\pi(z - z^p))$. (See Section 3 for a description of the ring $R/(p^N)$ and the element π .)

Step 2: Construct the matrix A_N which occurs in the statement of Theorem 28. Specifically, the matrix A_N is indexed by pairs (u, v) where u and v are lattice points in the dilation by a factor t of the polytope Δ , and $t = \lceil p^2 N / (p - 1)^2 \rceil - 1$. The (u, v) th entry of A_N is τ^{-1} of the coefficient of X^{pu-v} in the polynomial $F \bmod p^N$. The action of τ is as described in Section 3.

Step 3: Compute the product $(\prod_{i=0}^{a-1} \tau^{-i}(A_N))^k$. Let T denote the trace of this product.

Step 4: Output

$$N_k^* = q^{-k} [((q^k - 1)^{n+1} T + (q^k - 1)^n) \bmod p^N],$$

where the square brackets denote the smallest non-negative residue modulo p^N .

In the algorithm we assume that the polynomial is presented as input explicitly via its list of non-zero terms. The manner of presentation of Δ only affects the time required to find all lattice points in the dilated polytope $t\Delta$. For concreteness, let us say it is presented via its list of vertices, although any other reasonable presentation would suffice.

6.2. Proof of correctness of the algorithm. We know by Theorem 28 that

$$q^k N_k^* - (q^k - 1)^n = (q^k - 1)^{n+1} \text{Tr} \left(\left(\prod_{i=0}^{a-1} \tau^{-i}(A_N) \right)^k \right) \bmod p^N$$

in the ring $R/(p^N)$. Thus,

$$q^k N_k^* = (q^k - 1)^n + (q^k - 1)^{n+1} T \bmod p^N.$$

The lefthand side is a non-negative integer. The first term on the righthand side is an integer. The second term on the righthand side is the reduction modulo p^N of the trace of M_a^k , which is known to be an integer by Dwork's Trace Formula (Theorem 26), and thus it is also an integer. Since $N_k^* \leq$

$(q^k - 1)^n$ it follows that the lefthand side is smaller than $q^{k(n+1)}$. Hence in the case $N \geq ak(n+1)$ we must have

$$q^k N_k^* = [(q^k - 1)^{n+1} T + (q^k - 1)^n] \bmod p^N.$$

The proof is complete.

6.3. Complexity analysis. We shall use Big-Oh and Soft-Oh notation in our analysis of the complexity of the above algorithm. If C_1 and C_2 are real functions we write $C_1 = \mathcal{O}(C_2)$ if $|C_1| \leq c(|C_2| + 1)$ for some positive constant c . We write $C_1 = \tilde{\mathcal{O}}(C_2)$ if $C_1 = \mathcal{O}(C_2 \log(|C_2| + 1)^{c'})$ for some constant c' . Thus in the latter notation one ignores logarithmic factors.

6.3.1. Ring operations. We shall first of all count the number of operations in the ring $R/(p^N)$ required in Steps 1, 2, 3, ignoring for the time being any other auxiliary computations. More precisely, because of the complexity bounds in Lemma 5 it is convenient for our analysis to define a “ring operation” to be either arithmetic or the evaluation of the map τ^i , for $1 \leq i \leq a-1$, in the ring $R/(p^N)$ (excluding precomputation). In Section 6.3.2 we shall add back in the small contribution from computing Teichmüller liftings and also the precomputation required for the maps τ^i . Similarly, in Section 6.3.3 we shall account for the remaining operations in Steps 1, 2, 3, arising mainly from computations with the exponents of polynomials (at this stage we will restrict the input polytope Δ to avoid complications from convex geometry). The contributions from Steps 0 and 4 are easily seen to be absorbed into the other estimates, and we shall not mention them again.

Our running time will be in terms of the parameters $t, \tilde{t}, W, \tilde{W}$. Here

$$\begin{aligned} t &= \lceil \left(\frac{p}{p-1}\right)^2 N \rceil - 1, & W &= \#(t\Delta), \\ \tilde{t} &= \lceil \frac{p^2}{p-1} N \rceil - 1, & \tilde{W} &= \#(\tilde{t}\Delta), \end{aligned}$$

where $N = ak(n+1)$ and the $\#$ operator counts lattice points in convex sets. The sizes of the sets W and \tilde{W} are the number of lattice points in certain “truncated” cones. Since \tilde{t} is about p times as large as t , the integer \tilde{W} will be around p^{n+1} times as large as W , since we are working in $n+1$ dimensional space. The integer W is precisely the size of the matrix which occurs in Step 2. The integer \tilde{W} will turn out to be the maximum number of terms in the polynomial we compute in Step 1. Define $L_\Delta((p-1)/p^2) \bmod p^N$ to be the ring of polynomials obtained by reducing the coefficients of power series in $L_\Delta((p-1)/p^2) \subseteq R[[X]]$ modulo p^N . Then \tilde{W} is the number of monomials which occur in the finite ring $L_\Delta((p-1)/p^2) \bmod p^N$.

For Step 1 we have the following estimate.

Lemma 30. *Let F be the power series given in Definition 9 and N a positive integer. The polynomial $F \bmod p^N$ may be computed in*

$$\mathcal{O}(|J|\tilde{W}^2)$$

operations in the ring $R/(p^N)$.

Proof. By (4) we have that $\theta(z) \bmod p^N$ is a polynomial of degree not greater than $p^2N/(p-1)$. Thus we can obtain $\theta(z)$ via the formula

$$\theta(z) = \exp(\pi z) \exp(-\pi z^p)$$

by computing the first $\mathcal{O}(pN)$ terms in the expansion for $\exp(\pi z)$, substituting $z = -z^p$, and one multiplication of polynomials. Note that $\exp(\pi z)$ has p -adic integral coefficients. Thus $\theta(z)$ can be found in time $\mathcal{O}((pN)^2)$ operations in the ring $R/(p^N)$ using standard polynomial arithmetic.

Now by (10) we have that $F \in L_\Delta((p-1)/p^2)$. Thus $F \bmod p^N \in L_\Delta((p-1)/p^2) \bmod p^N$. One may then compute $F \bmod p^N$ directly from Definition 9 in $|J| - 1$ multiplications of polynomials of the form $\theta(a_j X^j) \bmod p^N$. Each such polynomial lies in the ring $L_\Delta((p-1)/p^2) \bmod p^N$, because $\text{ord}(a_j) = 0, w(j) = 1$ and the coefficients of θ decay at a suitable rate. Hence all computations required in computing $F \bmod p^N$ involve polynomials in this ring. Such polynomials have at most \tilde{W} terms. Exactly $|J| - 1$ multiplications are required. Thus the complexity is $\mathcal{O}(|J|\tilde{W}^2)$ ring operations. Noting that $pN = \mathcal{O}(\tilde{W})$ we have the result. \square

Note 31 It is crucial here that we only need to compute $F \bmod p^N$ and not $F^{(a)} \bmod p^N$, as one might attempt to do using a more naive approach. The latter polynomial has very high degree ($\mathcal{O}(q)$) because of the slow decay rate of the coefficients of $F^{(a)}$.

With regard to Step 2, given that the polynomial $F \bmod p^N$ has already been computed the only task required is to identify those pairs of points (u, v) such that $u, v \in t\Delta$, compute the integer point $pu - v$, and copy τ^{-1} of the term $F_{pu-v} \bmod p^N$ from $F \bmod p^N$ into the correct position in the matrix. Thus no arithmetic operations in the ring are required here, except W^2 computations of $\tau^{-1} = \tau^{a-1}$. These arithmetic operations can safely be ignored since in Lemma 30 we have already counted $\mathcal{O}(|J|\tilde{W}^2)$ ring operations. Computation of the appropriate indices (u, v) does require one to find all lattice points in certain polytopes and we return to that in Section 6.3.3.

Finally, for Step 3 we have the following estimate.

Lemma 32. *With the notation as in the statement of Theorem 28, the product*

$$\left(\prod_{i=0}^{a-1} \tau^{-i}(A_N) \right)^k$$

can be computed given the matrix A_N in

$$\mathcal{O}(W^3 \log(ak))$$

operations in the ring $R/(p^N)$.

Proof. A fast square-and-multiply style algorithm may be used to compute the power α^a in $\mathcal{O}(\log a)$ “matrix ring operations”. Specifically, working with the matrix representations one may compute α^{r+s} from α^r and α^s using

$$\prod_{i=0}^{r+s-1} \tau^{-i}(A_N) = \prod_{i=0}^{r-1} \tau^{-i}(A_N) \left(\tau^{-r} \left(\prod_{i=0}^{s-1} \tau^{-i}(A_N) \right) \right).$$

Now the case $r = s = 2^c$ for some c gives us the “square” step (computing $\alpha^{2^{c+1}}$ from α^{2^c}) and the case $r = 2^c$ and s arbitrary the “multiply” step. These two operations may be combined to give a fast exponentiation method in a straightforward way. The time required to compute a matrix for α^a from one for α is thus $\mathcal{O}(\log a)$ “matrix ring operations”. By matrix ring operations we mean multiplication of matrices of size W over $R/(p^N)$, and also computing $\tau^{-i}(B)$ for some $1 \leq i \leq a-1$ and matrix B of this form. The former requires $\mathcal{O}(W^3)$ operations in the ring $R/(p^N)$ using standard algorithms. Since $\tau^{-i} = \tau^{a-i}$ the latter may be computed in $\mathcal{O}(W^2)$ ring operations (that is, applications of a power of τ). Thus the time for computing a matrix for α^a from one for α is $\mathcal{O}(W^3 \log a)$.

Having obtained a matrix for α^a one may then compute a matrix for α^{ak} using the standard square-and-multiply algorithm. This requires $\mathcal{O}(\log k)$ matrix multiplications, that is $\mathcal{O}(W^3 \log k)$ operations in $R/(p^N)$. Thus the total time required is as claimed. \square

Note that the exponent 3 for multiplication of matrices can be improved to around 2.4 using faster methods [7, page 330].

Gathering these results we find the following.

Lemma 33. *The running time of Algorithm 29 is*

$$\mathcal{O}(\tilde{W}^2 |J| + W^3 \log(ak))$$

ring operations. Here $|J|$ is the number of non-zero terms in f , and \tilde{W} and W are defined as at the start of Section 6.3.1, with $N = ak(n+1)$.

6.3.2. Bit complexity arising from ring operations. Using Lemma 5 one may now calculate the number of bit operations in the algorithm which arise from operations in the ring $R/(p^N)$. In this section we shall also count the small contribution from computing the Teichmüller lifting of the coefficients of f , and also the precomputation required for powers of τ , which it was convenient to ignore in Section 6.3.1,

Definition 34. *Let the polytope Δ from Definition 12 have dimension $\tilde{n} \leq n+1$. Let $V(\Delta)$ be the \tilde{n} -dimensional volume of Δ . Denote by $v = \tilde{n}!V(\Delta)$ the “normalised” volume of Δ .*

Since f is non-zero we have $\tilde{n} \geq 1$ and certainly $v > 0$.

Proposition 35. *The running time of Algorithm 29 is*

$$\tilde{\mathcal{O}}(a^{3n+7} k^{3n+5} n^{3n+5} v^3 p^{2n+4})$$

bit operations, plus the contribution from operations outside of the ring $R/(p^N)$. The space complexity in bits is

$$\tilde{\mathcal{O}}(a^{2n+4}k^{2n+3}n^{2n+3}v^2p)$$

plus the contribution from operations outside of $R/(p^N)$.

Proof. To compute the complexity first observe that $t = \lceil p^2N/(p-1)^2 \rceil - 1 \leq 4N$. Thus $t, N = \mathcal{O}(akn)$ in the algorithm. Also $\tilde{t} = \mathcal{O}(pN) = \mathcal{O}(pt)$. Now

$$\begin{aligned} W &= \#(t\Delta) \\ &\leq \tilde{n}!V(t\Delta) + \tilde{n} \\ &= \tilde{n}!t^{\tilde{n}}V(\Delta) + \tilde{n} \\ &= \mathcal{O}(vt^{n+1}). \end{aligned}$$

Here we use the Blichfeldt bound $\#(P) \leq m!V(P) + m$ for any m -dimensional polytope P from [9, page 144], the fact $V(t\Delta) = t^{\tilde{n}}V(\Delta)$ since Δ is \tilde{n} -dimensional, and also that $\tilde{n} \leq n + 1$. Similarly $\tilde{W} = \#(\tilde{t}\Delta) = \mathcal{O}(v\tilde{t}^{n+1}) = \mathcal{O}(vt^{n+1}p^{n+1})$.

Thus from Lemma 33 the number of ring operations is

$$\mathcal{O}((vt^{n+1}p^{n+1})^2(v+n) + (t^{n+1}v)^3 \log(ak))$$

since $|J| \leq \#(\Delta) - 1 \leq v + n$. Thus the bit complexity which arises from ring operations is by Lemma 5

$$\mathcal{O}(((vt^{n+1}p^{n+1})^2(v+n) + (t^{n+1}v)^3 \log(ak))(paN \log p)^2).$$

Tidying up and ignoring logarithmic factors we get

$$\tilde{\mathcal{O}}(v^3t^{2n+2}p^{2n+4}a^2N^2 + v^3t^{3n+3}p^2a^2N^2).$$

The second term is dominant in all factors except p . For simplicity we take the estimate of

$$\tilde{\mathcal{O}}(v^3t^{3n+3}p^{2n+4}a^2N^2).$$

Now putting $t, N = \mathcal{O}(ank)$ we get

$$\tilde{\mathcal{O}}(v^3a^{3n+7}k^{3n+5}n^{3n+5}p^{2n+4}).$$

This is the total bit complexity which arises from “ring operations”, as defined at the start of Section 6.3.1. There remains the contribution from computing the Teichmüller liftings of the coefficients of f in Step 1, and also the precomputation required for the map τ . By Lemma 5, this is easily seen to be absorbed in the above estimate.

With regard to the space complexity, this is dominated by the space required to store the matrix A_N , which is $\mathcal{O}(W^2)$ ring elements. Putting $W = \mathcal{O}(v(ank)^{n+1})$ and using Lemma 5 gives us the result. \square

One may replace n in the exponents in Proposition 35 by $\tilde{n} - 1$; however, this only gives an improvement when the Newton polytope is not full-dimensional.

6.3.3. Bit complexity arising from auxiliary operations. It remains to bound the complexity which arises from operations outside of the ring $R/(p^N)$ in Steps 1 and 2. In Step 1 manipulation of exponents of polynomials will add an extra term $\mathcal{O}(\log(pNd))$ to the running time, which can safely be ignored.

In Step 2 one is required to find all lattice points which lie in $t\Delta$, for $t = \mathcal{O}(ank)$. The complexity of this step will depend upon the input polytope Δ . For a “general” Δ one requires methods from computational convex geometry which are not in the spirit of the present exposition. Thus our total bit complexity estimate for Algorithm 29 will just be as in Proposition 35 “plus the contribution from finding all lattice points in $t\Delta$ ”.

At this stage for simplicity we shall restrict to the choice of $\delta = \delta_2$ in Definition 12. Thus we take $\Delta = \Delta_2$ as the convex hull in \mathbb{R}^{n+1} of the origin and the $n+1$ points

$$(1, 0, \dots, 0), (1, d, 0, \dots, 0), \dots, (1, 0, \dots, 0, d).$$

For this case the required set of lattice points is

$$\{(r_0, r_1, \dots, r_n) \mid r_1 + \dots + r_n \leq dr_0 \leq dt\}$$

and so no computations are required here. Also, now v equals d^n and directly from Proposition 35 we get the following result.

Proposition 36. *Let Algorithm 29' be exactly as Algorithm 29 only with input restricted to the choice of polytope $\Delta = \Delta_2$ described in the preceding paragraph. The total running time of Algorithm 29' is*

$$\tilde{\mathcal{O}}(a^{3n+7}k^{3n+5}n^{3n+5}d^{3n}p^{2n+4})$$

bit operations. The total space complexity in bits is

$$\tilde{\mathcal{O}}(a^{2n+4}k^{2n+3}n^{2n+3}d^{2n}p).$$

We shall use this restricted version of Algorithm 29 in the proofs of our main results in the next section.

6.4. Proofs of the results in the Introduction. To compute the number of points on the affine variety defined by a polynomial f one simply uses the torus decomposition of $\mathbb{F}_{q^k}^n$. Specifically, for any subset $S \subseteq \{1, 2, \dots, n\}$ let G_k^S denote the set of points

$$\{(x_1, \dots, x_n) \mid x_i \in \mathbb{F}_{q^k}, x_i = 0 \Leftrightarrow i \in S\}.$$

Denote by f^S the polynomial obtained from f by setting to zero all indeterminates X_i which occur in f for $i \in S$. Denote by N_k^S the number of solutions of $f^S = 0$ in the torus G_k^S of dimension $n - |S|$. Then $N_k = \sum_S N_k^S$ where the sum is over all subsets of $\{1, 2, \dots, n\}$. Each number N_k^S can be computed using Algorithm 29'. (If some f^S is identically zero or has degree 0 then $N_k^S = (q^k - 1)^{n-|S|}$ or 0, respectively, and Algorithm 29' is not required!) Thus by 2^n applications of this algorithm we obtain N_k as desired.

Now to obtain the whole zeta function $Z(f/\mathbb{F}_q)$ it suffices to count N_k for all $k = 1, \dots, \deg(r) + \deg(s)$, where

$$Z(f/\mathbb{F}_q)(T) = \frac{r(T)}{s(T)}$$

with r and s coprime polynomials in $1 + T\mathbb{Z}[T]$. More precisely, it is enough to know upper bounds $\deg(r) \leq D_1$ and $\deg(s) \leq D_2$, and compute N_k for $k = 1, \dots, D_1 + D_2$. Then use the linear algebra method described prior to [22, Corollary 2.8], which we now supplement with further details.

Let $u(T) = 1 + \sum_{i=1}^{D_1} u_i T^i$ and $v(T) = 1 + \sum_{i=1}^{D_2} v_i T^i$ have indeterminate coefficients. Write $Z(f/\mathbb{F}_q)(T) = 1 + z_1 T + z_2 T^2 + \dots$. This power series has non-negative integer coefficients and it can easily be computed modulo $T^{D_1+D_2+1}$ given N_k for $k = 1, \dots, D_1 + D_2$. The equation $v(T)Z(f/\mathbb{F}_q) \equiv u(T) \pmod{T^{D_1+D_2+1}}$ defines a linear system $Ax = y$ where A is a known square $D_1 + D_2$ integer matrix and y a known integer column vector. The entries in A and y are just coefficients from the power series $Z(f/\mathbb{F}_q) \pmod{T^{D_1+D_2+1}}$. Let b be a bound on their bit length. The unknown entries in x are the coefficients of u and v . By [22] the set of all solutions to this system consists of precisely those vectors x derived by specialising the coefficients of $u(T)$ and $v(T)$ to equal those of $d(T)r(T)$ and $d(T)s(T)$, respectively, for some $d(T) \in 1 + T\mathbb{Z}[T]$ with degree at most $\min(D_1 - \deg(r), D_2 - \deg(s))$. In particular, the system has a unique solution (i.e. $\det(A) \neq 0$) if and only if either $\deg(r) = D_1$ or $\deg(s) = D_2$ (or both). The determinant $\det(A)$ can be computed using the small primes method in [7, Algorithm 5.10] in a number of bit operations bounded by $\tilde{O}((D_1 + D_2)^4 b^2)$ [7, Theorem 5.12]. Let us assume now that $\det(A) \neq 0$ and so the system has a unique solution, namely the unknown vector containing the integer coefficients of r and s . Let B be a bound on the bit length of these coefficients. Find the unique solution to the linear system modulo enough small primes which do not divide $\det(A)$, and recover this integer solution using the Chinese remainder theorem. Precisely, work modulo a collection of such primes whose product has bit length greater than B . This second step requires $\tilde{O}((D_1 + D_2)^4 B^2)$ bit operations using Gaussian elimination (this can be improved with a Padé Approximation algorithm [7, Section 5.9]). Values for b and B may be deduced from the bound $N_k \leq q^{nk}$. Specifically, one may show from this that the absolute values of the reciprocal zeros of r and s are all $\leq q^n$, and so we can take $B = \mathcal{O}((D_1 + D_2)n \log q)$. Also, we can take $b = \mathcal{O}((D_1 + D_2)^2 n \log q)$. If in the above we find $\det(A) = 0$ then we must have that $\deg(r) < D_1$ and $\deg(s) < D_2$. In this case one must first reduce D_2 , say, and compute determinants until the correct value $D_2 = \deg(s)$ is found (then $\det(A) \neq 0$ and the above method works).

By the refinement of Bombieri's degree bound [3] from [2, Equation (1.13)], the "total degree" $\deg(r) + \deg(s)$ is bounded by $2^{n+1}6^{n+1}(n+1)!V(\Delta_1)$, where Δ_1 is the polytope in \mathbb{R}^{n+1} derived from the Newton polytope of f (see the paragraph following Definition 12). Certainly $(n +$

$1)!V(\Delta_1) \leq d^n$. Hence we may take $D_1, D_2 = 2^{4n+4}d^n$ and so $D_1 + D_2 = 2^{4n+5}d^n$.

Theorem 37. *Let f be a polynomial in n variables of total degree $d > 0$ over \mathbb{F}_q , where $q = p^a$. The full zeta function $Z(f/\mathbb{F}_q)$ can be computed deterministically in*

$$\tilde{\mathcal{O}}(2^{13n^2} a^{3n+7} d^{3n^2+9n} p^{2n+4})$$

bit operations. (Here we use soft-Oh notation which ignores logarithmic factors, as defined at the start of Section 6.3.)

Proof. From Proposition 36 and the torus decomposition method, the bit complexity of computing N_k for $k = 1, \dots, 2^{4n+5}d^n$ is

$$\tilde{\mathcal{O}}\left(\sum_{k=1}^{2^{4n+5}d^n} \{a^{3n+7} k^{3n+5} n^{3n+5} d^{3n} p^{2n+4}\} 2^n\right).$$

(The contribution from recovering $Z(f/\mathbb{F}_q)$ from the N_k is absorbed in this estimate.) Tidying up the factor in n we get the claimed result. \square

Since we may assume that $d > 1$ we have $2^{n^2} = \mathcal{O}(d^{n^2})$, and Theorem 1 now follows.

The proof of Corollary 2 was explained in the introduction, and we finish with some comments on Corollary 3. By Weil's theorem, the zeta function of the smooth projective curve \tilde{V} from Corollary 3 is of the form

$$Z(\tilde{V})(T) = \frac{P(T)}{(1-T)(1-qT)}$$

for some polynomial $P(T)$ whose reciprocal roots have complex absolute value $q^{1/2}$. Since the (possibly singular) affine curve V and the smooth projective curve \tilde{V} differ in only finitely many closed points, we deduce that the zeta function of V is of the form

$$Z(V)(T) = \frac{P(T)Q(T)}{1-qT},$$

where $Q(T)$ is a rational function whose zeros and poles are roots of unity. This zeta function, and in particular the rational function $P(T)Q(T)$, may be computed within the time bound in Corollary 3 by Corollary 2. In terms of the pure weight decomposition [22], the polynomial $P(T)$ (respectively, $Q(T)$) is exactly the pure weight 1 (respectively, weight 0) part of the product $P(T)Q(T)$, and can be recovered quickly from $P(T)Q(T)$ via the LLL polynomial factorization algorithm. In our current special case, one can proceed directly without using the LLL-factorization algorithm. By repeatedly removing the common factor of the numerator of $P(T)Q(T)$ with $T^s - 1$ for $\phi(s)$ (Euler totient function) not greater than the total degree of $P(T)Q(T)$, the desired polynomial $P(T)$ can be recovered. The order of the group of rational points on \tilde{V} is simply $P(1)$, see [22]. Thus for fixed dimension and

finite field one can compute the order of the group of rational points on the Jacobian of a smooth projective curve in time polynomial in the degree d .

REFERENCES

- [1] L. Adleman and M.D. Huang, Counting rational points on curves and abelian varieties over finite fields, in “*Algorithmic Number Theory II*” (H. Cohen), Lecture Notes in Computer Science 1122, Springer-Verlag, 1996, 1-16.
- [2] A. Adolphson and S. Sperber, Newton polyhedra and the degree of the L-functions associated to an exponential sum, *Invent. Math.*, 88 (1987), 555-569.
- [3] E. Bombieri, On exponential sums in finite fields, II, *Invent. Math.* 47, (1978), 29-39.
- [4] B. Dwork, On the rationality of the zeta function of an algebraic variety, *Amer. J. Math.*, 82, (1960), 631-648.
- [5] B. Dwork, On the zeta function of a hypersurface, *Pub. IHES No. 12*, 1962.
- [6] 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.
- [7] J. von zur Gathen and J. Gerhard, *Modern Computer Algebra*, Cambridge University Press, 1999.
- [8] P. Gaudry and N. Gürel, An extension of Kedlaya’s algorithm for counting points on superelliptic curves, in “*Advances in Cryptology - ASIACRYPT 2001*” (C. Boyd), Lecture Notes in Computer Science 2248, Springer, (2001), 480-494.
- [9] J.E. Goodman and J. O’Rourke (Eds), *Handbook of Discrete and Computational Geometry*, CRC Press, 1997.
- [10] K.S. Kedlaya, Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology, *J. Ramanujan Math. Soc.* 16 (2001), 323-338.
- [11] N. Koblitz, *p -Adic Numbers, p -Adic Analysis and Zeta Functions*, 2nd Edition, Springer, 1984.
- [12] 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.
- [13] R. Lidl and H. Niederreiter, *Introduction to Finite Fields and Their Applications*, Cambridge University Press, 1986.
- [14] J. Pila, Frobenius maps of Abelian varieties and finding roots of unity in finite fields, *Math. Comp* 55 (1990), 745-763.
- [15] 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.
- [16] T. Satoh, The canonical lift of an ordinary elliptic curve over a finite field and its point counting, *J. Ramanujan Math. Soc.* 15, (2000), 247-270.
- [17] R. Schoof, Elliptic curves over finite fields and the computation of square roots mod p , *Math. Comp.* 44 (1985), no. 170, 483-494.
- [18] R. Schoof, Counting points on elliptic curves over finite fields, *J. Théorie des Nombres de Bordeaux* 7, 1998, 219-254.
- [19] D. Wan, Meromorphic continuation of L-functions of p -adic representations, *Ann. Math.*, 143 (1996), 469-498.
- [20] D. Wan, Computing zeta functions over finite fields, *Contemporary Mathematics*, 225 (1999), 131-141.
- [21] D. Wan, Rank one case of Dwork’s conjecture, *J. Amer. Math. Soc.*, 13(2000), 853-908.
- [22] D. Wan, Algorithmic theory of zeta functions over finite fields, preprint.

COMPUTING LABORATORY, OXFORD UNIVERSITY, OXFORD OX1 3QD, UK *E-mail*
address: `ALAN.LAUDER@COMLAB.OX.AC.UK`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, IRVINE, CA 92697,
USA *E-mail address:* `DWAN@MATH.UCI.EDU`