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 Jacobian of a smooth, projective plane curve over a small finite field.

1. Introduction

The purpose of this paper is to give an elementary and self-contained proof that one may efficiently compute zeta functions of varieties over finite fields given that the dimension and number of defining equations are assumed fixed, and the characteristic is suitably small. 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 efficient methods of doing so. That p-adic methods may be used to efficiently compute zeta functions for small characteristic was first suggested in [18, 20], where Wan gives a simpler algorithm for computing congruence properties of the zeta function.

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 the 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 $\mathbb{F}_q[X_1, \ldots, X_n]$ the ring of all polynomials in n variables over $\mathbb{F}_q$.

For a polynomial $f \in \mathbb{F}_q[X_1, \ldots, X_n]$ 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

Date: July 12, 2001.

Alan Lauder gratefully acknowledges the support of the EPSRC and Wolfson College, Oxford, and thanks Richard Brent. Daqing Wan is partially support by the NSF. The authors are pleased to thank Alexander Barvinok, Colin McDiarmid, and Bernd Sturmfels for answering some questions on convex geometry.

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

integer coefficients

$$Z(f/\mathbb{F}_q)(T) := \exp(\sum_{k\geq 0} \frac{N_k T^k}{k}).$$

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 N_k for the values $k = 1, 2, \ldots, 2\deg(r) + 2\deg(s)$ is enough to explicitly and efficiently determine $Z(f/\mathbb{F}_q)$, see [20]. The numbers 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 an exponential number of steps 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 $\mathcal{O}(d^n \log q)$, which is also the output size of the zeta function. We prove

Theorem 1. *Let f be a polynomial of total degree d in n variables over the finite field $\mathbb{F}_q$ of q elements, where $q = p^a$ and p is the characteristic. One may compute the zeta function of f deterministically using a number of bit operations which is bounded by a polynomial in the parameters p^n, d^{n^2} and a^n.*

In particular, this computes the zeta function in deterministic polynomial time if n is fixed, and p is bounded by a polynomial in d and $\log q$. Note that our result makes no assumption of non-singularity on the variety defined by the polynomial.

With regard to the exponents in the algorithm, we state these precisely in Corollary 39. For now we observe that if p and d 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 $\mathcal{O}(a^{2n+4})$. For a and p fixed, the running time for computing points in $\mathbb{F}_q$ is $\mathcal{O}(d^{3n})$ with space $\mathcal{O}(d^{2n})$. Here we are ignoring logarithmic factors. (See Corollary 35 and Note 37.) Note that all our complexity estimates are made using standard methods for multiplication in various rings, and can be modestly reduced using 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, in a similar manner to [7]. We shall need more definitions: the support set of a polynomial f is the set of exponents $r = (r_1, \ldots, r_n)$ of non-zero terms $a_r X_1^{r_1} \ldots X_n^{r_n}$ which occur in f, thought of as points in $\mathbb{R}^n$. The Newton polytope of f, denoted $\mathrm{Newt}(f)$, is defined to be the convex hull in $\mathbb{R}^n$ of the support set 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 Theorem 38).

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 $(\bar{\mathbb{F}}_q)^n$ of a set of polynomials $f_1, f_2, \ldots, f_r$. An analogous zeta function $Z(V/\mathbb{F}_q)$ may then be defined in terms of the number of solutions in each extension field of $\mathbb{F}_q$. These numbers may be computed using an inclusion-exclusion argument involving the polynomials $\prod_S f_s$ where S is a subset of $\{1, 2, \ldots, r\}$, see [20]. Theorem 1 then easily yields the following

Corollary 2. *Let f_i $(1 \leq i \leq r)$ be polynomials of total degree d_i in n variables over the finite field $\mathbb{F}_q$, with $q = p^a$ where p is the characteristic. Let V denote the affine variety defined by these polynomials. Then one may compute the zeta function of V using a number of bit operations which is bounded by a polynomial in the parameters d^{n^2}, p^n, a^n and 2^r, where $d = \sum_i d_i$.*

Thus one has an efficient algorithm for an arbitrary affine variety assuming the characteristic, dimension and the number of defining polynomials is 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 is 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 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 space. Unfortunately, this operator cannot be computed efficiently directly; 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 key in deriving an efficient algorithm. Note that the same idea is used in [18, 20] 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, 15, 16]. These utilise the theory originally developed by Weil for abelian varieties whereas we use Dwork's more general and simpler p-adic theory. For the case of a smooth projective plane curve of degree d over a field of size $q = p^a$ these algorithms have a time complexity which grows as $\log(q)^{C_d}$ using standard arithmetic, where the constant C_d grows exponentially in the degree d. 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. The zeta function immediately gives the order of the Jacobian of f, see [20]. We deduce

Corollary 3. *Let f define a projective, smooth plane curve of degree d over a finite field $\mathbb{F}_q$. Then, assuming p is fixed, one may compute deterministically the order of the Jacobian of f in time polynomial in $d\log(q)$.*

Thus one may compute the order of the Jacobian of a projective, smooth plane curve over finite fields of small characteristic p in deterministic polynomial time. In particular, this answers a question posed in [14], which Poonen attributes to Katz and Sarnak. We note that recently a similar result for special classes of curves was worked out in [9, 11], using Monsky-Washnitzer's method.

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 chain level. A general improvement, especially in the smooth case, is to work on the cohomology level, which we plan to carry out in a later paper. There are then several related p-adic cohomology theories available, each leading to a somewhat different version of the algorithm. Again, as indicated in [18, 20], these p-adic methods are expected to be efficient only for small p.

2. Exponential sums over finite fields

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

An additive character Ψ from $\mathbb{F}_{q^k}$ to some field F is a non-zero mapping (i.e., to F^*) such that $\Psi(a+b) = \Psi(a)\Psi(b)$. We say that it is non-trivial if $\Psi(1) \neq 1$. 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 F of the rationals which contains a primitive pth root of unity. We assume that the family $\{\Psi_k\}$ forms a tower of characters in the sense that for each $k \geq 2$,

$$\Psi_k = \Psi_{k-1} \circ \mathrm{Tr}_{\mathbb{F}_{q^k}/\mathbb{F}_{q^{k-1}}}.$$

In our application, the field F will be taken to be a certain p-adic field.

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 $u = (u_0, u_1, \dots, u_n)$; x 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.

For every $g \in \mathbb{F}_q[X]$ we define the exponential sum

$$S_k^*(g) := \sum_{x \in (\mathbb{F}_{q^k}^*)^{n+1}} \Psi_k(g(x)).$$

Here the sum is over all points in affine space which have non-zero coordinates. (In the literature the notation $S_k(g)$ is reserved for the sum over all points in affine space.) We have

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 $(\mathbb{F}_{q^k}^*)^n$. For any non-trivial additive character Ψ_k from $\mathbb{F}_{q^k}$ to an extension F of the rationals let*

$$S_k^*(X_0 f) = \sum_{x \in (\mathbb{F}_{q^k}^*)^{n+1}} \Psi_k(x_0 f(x_1, \dots, x_n)).$$

Then

$$S_k^*(X_0 f) = q^k N_k^* - (q^k - 1)^n$$

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

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

This is a standard result from the theory of exponential sums [13, page 168]. Thus the sum $\sum_{x_0 \in \mathbb{F}_{q^k}^*} \Psi_k(x_0 f(x_1, \dots, x_n))$ is (-1) or $(q^k - 1)$ according to whether $f(x_1, \dots, x_n)$ is non-zero or zero. Letting N_k^* denote the number of points $(x_1, \dots, x_n) \in (\mathbb{F}_{q^k}^*)^n$ for which $f(x_1, \dots, x_n) = 0$ we have that

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

This completes the proof ◻

Our next step will be to find an alternative formulation for the lefthand side of the equation in Lemma 4. This is achieved in Proposition 9, which leads us eventually to Dwork's Trace Formula (Theorem 25).

3. p-adic theory

3.1. **p-adic fields.** Let $\mathbb{Q}_p$ denote the p-adic rationals and $\mathbb{Z}_p$ the p-adic integers. Explicitly, $\mathbb{Q}_p$ may be identified with the set of all expansions

$$\sum_{i \geq m} a_i p^i$$

where m is an integer and a_i belongs to the set $\{0, 1, 2, \dots, p-1\}$. The subring $\mathbb{Z}_p$ consists of those expansions with no negative powers of p, that is $m \geq 0$. The ring $\mathbb{Z}_p$ is a local ring with maximal ideal generated by p (expansions with $m \geq 1$). The field $\mathbb{Q}_p$ is the quotient field of $\mathbb{Z}_p$ and is thus a local field. Its residue class field is by definition $\mathbb{Z}_p / p\mathbb{Z}_p$ which is isomorphic to the finite field $\mathbb{F}_p$ with p elements. The ring $\mathbb{Z}_p$ is as such a natural lifting to characteristic zero of the finite field $\mathbb{F}_p$.

By adjoining a suitable algebraic element to $\mathbb{Z}_p$ one constructs a p-adic lifting of the field $\mathbb{F}_q$. Specifically, let Ω be a fixed algebraic closure of $\mathbb{Q}_p$. We need to construct a primitive $(q-1)$st root of unity in Ω. To this end, let $h(z)$ be a monic polynomial with coefficients in the set $\{0,1,\dots,p-1\}$ whose reduction modulo p is primitive of degree a, where $q=p^a$. (Thus each root of $h(z)$ has multiplicative order exactly $q-1$ when reduced modulo p [13, page 82].) Then $K:=\mathbb{Q}_p(\mu)$, where $\mu\in\Omega$ with minimal polynomial h, is the unique unramified extension of $\mathbb{Q}_p$ of degree a. Starting with a solution $z=\mu$ of $z^{q-1}\equiv 1 \bmod p$ in $\mathbb{Z}_p[\mu]$, hensel lift in $\mathbb{Z}_p[\mu]$ to obtain (a suitable approximation of) some element ϵ. Now $\epsilon^{q-1}=1$ and also the lower powers of ϵ are distinct modulo p, since $\epsilon=\mu \mod p$ is primitive in $\mathbb{Z}_p/p\mathbb{Z}_p$. Hence ϵ is a primitive $(q-1)$st root of unity, and we have $K=\mathbb{Q}_p(\epsilon)$. We can construct the minimal polynomial of ϵ as $g(z):=(z-\epsilon)(z-\epsilon^p)\dots(z-\epsilon^{p^{a-1}})$, which is a polynomial with coefficients in $\mathbb{Z}_p$. Elements in K may be represented as expansions

$$\sum_{k=0}^{a-1}(\sum_{i\geq m} a_{ik}p^i)\epsilon^k$$

with $a_{ik}\in\{0,1,\dots,p-1\}$ and we treat ϵ "symbolically" with arithmetic subject to the rule $g(\epsilon)=0$.

We wish to define non-trivial additive characters from $\mathbb{F}_p$ to some extension field of $\mathbb{Q}$. Such an extension must necessarily contain a primitive pth root of unity, namely the image of the identity under the character. Let $\pi\in\Omega$ satisfy the equation $\pi^{p-1}=-p$. If ζ is a primitive p-th root of unity, then $\mathrm{ord}_p(1-\zeta)=\mathrm{ord}_p\pi=1/(p-1)$. Thus, both $\mathbb{Q}_p(\pi)$ and $\mathbb{Q}_p(\zeta)$ are the unique totally ramified extension over $\mathbb{Q}_p$ of degree $(p-1)$. It follows that $\mathbb{Q}_p(\pi)=\mathbb{Q}_p(\zeta)$. Hence the field $\Omega_1:=\mathbb{Q}_p(\pi)$ is a suitable image for our additive character. It has residue class field $\mathbb{F}_p$. Elements in Ω_1 may be written as

$$\sum_{j=0}^{p-2}(\sum_{i\geq m} a_{ij}p^i)\pi^j$$

with $a_{ij}\in\{0,1,2\dots,p-1\}$ and arithmetic subject to the rule $\pi^{p-1}=-p$.

We shall need to work in the compositum of the two fields K and Ω_1. We denote this compositum field Ω_0. It is obtained from $\mathbb{Q}_p$ by adjoining both ϵ and π and elements may be represented as

$$\sum_{k=0}^{a-1}\sum_{j=0}^{p-2}(\sum_{i\geq m} a_{ijk}p^i)\pi^j\epsilon^k \tag{1}$$

with $a_{ijk} \in \{0, 1, \ldots, p-1\}$ and arithmetic subject to the rules $g(\epsilon) = 0$ and $\pi^{p-1} = -p$. We have the diagram of field extensions

$$\begin{tikzcd}[column sep=small] & \Omega \arrow[d, no head] & \\ & \Omega_0 \arrow[dl, no head] \arrow[dr, no head] & \\ \Omega_1 \arrow[dr, no head] & & K \arrow[dl, no head] \\ & \mathbb{Q}_p & \end{tikzcd}$$

The ring of integers of Ω_0 is just those expansions in (1) with $m \geq 0$. We will denote this as R_0.

The Teichmuller lifting of a finite field element $\bar{a} \in \mathbb{F}_q$ is defined to be the unique element $a \in K$ such that $a^q = a$ and $a \bmod p = \bar{a}$. For non-zero elements, this may be approximated via hensel lifting starting with the solutions $z = \bar{a}$ of the equation $z^{q-1} = 1$ modulo p in $\mathbb{F}_q$.

We define a lifting of the Frobenius automorphism on $\mathbb{F}_q$ to an automorphism of Ω_0 which fixes Ω_1. Define the map τ by setting $\tau(\epsilon) = \epsilon^p$ and $\tau(\pi) = \pi$, and extend to the whole of Ω_0 by insisting τ is an automorphism. Thus τ simply acts on expansions like (1) by taking pth powers of ϵ which occur and fixing everything else. We extend τ to act on $\Omega_0[[X]]$ coefficient-wise, fixing monomials. Here $\Omega_0[[X]]$ is the ring of all power series in the indeterminates $X = X_0, \ldots, X_n$ with coefficients from Ω_0.

3.2. **Computing in modular reductions of p-adic rings.** In Section 5.3 we will look at the reduction modulo a power p^N of elements in R_0. (N will be chosen later but for now is assumed to be any positive integer.) Elements in this ring $R_0/(p^N)$ are finite expansions of the form

$$\sum_{k=0}^{a-1} \sum_{j=0}^{p-2} \Big(\sum_{i=0}^{N-1} a_{ijk} p^i\Big) \pi^j \epsilon^k \tag{2}$$

with $a_{ijk} \in \{0, 1, \ldots, p-1\}$ and arithmetic subject to the rules $g(\epsilon) = 0$, $\pi^{p-1} = -p$ and $p^N = 0$. We define the "Teichmuller lifting to $\mathbb{R}_0/(p^N)$" of an element of $\mathbb{F}_q$ to be the reduction of its Teichmuller lifting modulo p^N. The lifting of elements $a \in \mathbb{F}_q$ can be computed using the method in the previous section.

One may compute τ acting on the ring $R_0/(p^N)$ as follows: Given an element of the form (2), it is enough to separately compute τ acting on the $(p-1)$ elements $\sum_{k=0}^{a-1} (\sum_{i=0}^{N-1} a_{ijk} p^i) \epsilon^k$, and add the results together, multiplying by suitable powers of π. Now τ of such an element is just $\sum_{k=0}^{a-1} (\sum_{i=0}^{N-1} a_{ijk} p^i) \epsilon^{kp}$. To reduce this back to the required form we use the relation $\epsilon^a = \sum_{k=0}^{a-1} (\sum_{i=0}^{N-1} g_{ik} p^i) \epsilon^k$, where the $g_{ik} \in \{0, 1, \ldots, p-1\}$ are obtained from the coefficients of $g(z)$. This requires $\mathcal{O}(ap)$ subtractions of an element of size $\mathcal{O}(aN \log p)$ from one of size $\mathcal{O}(paN \log p)$. To compute

τ on a full expansion (2) we need to do this $\mathcal{O}(p)$ times, giving a total of $\mathcal{O}(a^2p^2N\log p)$ bit operations.

One may estimate the complexity of the other basic operations in $R_0/(p^N)$ in a similar way to get

Lemma 5. *Let $R_0/(p^N)$ be the reduction of R_0 modulo the power p^N. Elements in $R_0/(p^N)$ can be represented using $\mathcal{O}(paN\log p)$ bits. Aaddition, subtraction, multiplication and inversion (of units) in R_0/p^N can be performed in $\mathcal{O}((paN\log p)^2)$ bit operations. The map τ on $R_0/(p^N)$ may be computed using $\mathcal{O}((ap)^2N\log p)$ bit operations. The Teichmuller lifting to $R_0/(p^N)$ of a finite field element can be computed in $\mathcal{O}((aN\log p)^3)$ bit operations.*

For arithmetic the exponent 2 may be improved using faster methods, leading to a theoretical improvement in the final algorithm. However, we shall be content with the estimate $\mathcal{O}((paN\log p)^2)$ for both arithmetic operations and the computation of τ in the analysis of our point counting algorithm. (Teichmuller lifting is performed so infrequently that the weaker bound does not cause any harm.)

3.3. **p-adic valuations and convergence of power series.** Denote by ord the additive valuation on Ω_0 normalised so that $\mathrm{ord}\,(p) = 1$. Thus $\mathrm{ord}\,(\epsilon) = 0$ and $\mathrm{ord}\,(\pi) = 1/(p-1)$, and ord of a power series in (1) is $i + (j/(p-1))$ where (i,j) is the least index, ordered lexicographically, such that $a_{ijk} \neq 0$ for some $0 \leq k \leq a-1$. Elements in Ω_0 which have large p-adic valuation ord can be thought of as "small", and two p-adic numbers "close" if their expansions agree to a high power of p. More precisely, one can define a p-adic norm $|.|_p$ on Ω_0 by $|x| = p^{-\mathrm{ord}\,(x)}$. Given any formal power series $\sum_r A_rX^r$ where $A_r \in \Omega_0$, we say it converges at a point $x = (x_0, \ldots, x_n) \in \Omega_0^{n+1}$ if the sequences of partial sums $\sum_{r,\,|r|<e} A_rx^r$ tends to a limit under this 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_rx^r tend to zero p-adically (that is, are divisible in R_0 by increasing large powers of p), as $|r|$ goes to infinity. In particular, if there is a real number $c > 0$ such that $\mathrm{ord}\,(A_r) \geq c|r|$ for all r, then the series will certainly converge for all points x which are Teichmuller liftings of points over $\mathbb{F}_q$. Thus p-adic convergence is a rather simple matter. Note that throughout the paper we shall use the additive valuations ord rather than the p-adic norm $|.|_p$ 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 Ω_1. In the case of complex characters, this is done via the exponential function. Explicitly, the map Ψ from $\mathbb{F}_p$ to $\mathbb{C}$ which takes $\bar{b} \in \mathbb{F}_p$ to $\exp((2\pi i/p)b)$ is a non-trivial additive character. Here exp is the usual exponential function and π the usual one rather than that defined above. Also b denotes the "naive lifting" to $\mathbb{Z}$ of the finite field element $\bar{b}$.

Notice that the above character is of the form "$\bar{b} \mapsto \exp(*b)$" where $*$ is some element in $\mathbb{C}$, in this case $2\pi i/p$, for which $(\exp(*))^p = 1$. The same idea very nearly works in the p-adic field setting. We now describe precisely how this is done.

Let $\exp(t)$ be the formal power series with coefficients in $\mathbb{Q}$ given by

$$\exp(t) := \sum_{n\geq 0} \frac{t^n}{n!} = \prod_{n=1}^{\infty}(1-t^n)^{-\mu(n)/n}.$$

Here μ is the Mobius function and the latter equality is easily proved [12]. Unfortunately the exponential function $\exp(t)$ does not have good convergence on the field Ω_0. In particular, it does not converge at the p-adic lifting of all the non-zero points in $\mathbb{F}_q$. Instead we use the Artin-Hasse exponential, which is obtained by removing the "bad" terms in the infinite product for which $p|n$. Specifically, the Artin-Hasse exponential may be defined as

$$E(t) := \exp(\sum_{j=0}^{\infty} \frac{t^{p^j}}{p^j}) = \prod_{p\nmid n}(1-t^n)^{-\mu(n)/n}.$$

Now once again choosing a suitable $*$ in Ω so that $(E(*))^p = 1$, we find that the map $\bar{a} \mapsto E(*a)$, where a is the Teichmuller lifting of $\bar{a}$, is an additive character from $\mathbb{F}_p$ to $\mathbb{Q}_p(*)$. Specifically, if we take γ_∞ to be an element in Ω which satisfies $\sum_{j=0}^{\infty} \frac{\gamma_\infty^{p^j}}{p^j} = 0$ and $\text{ord}\,(\gamma_\infty) = 1/(p-1)$, then $E(\gamma_\infty t)$ has the properties we shall need. (Note that subtlties in convergence of composed power series ensure that $E(\gamma_\infty) \neq 1$, and so the character is non-trivial, even though $\sum_{j=0}^{\infty} \gamma^{p^j}/p^j = 0$ and $\exp(0) = 1$.)

In [5] Dwork constructs an infinite family of functions with appropriate properties, which he calls "splitting functions", indexed by s, a positive integer or infinity. These are described in the next lemma, the proof of which we omit and can be found in [5, Pages 55-57].

Lemma 6 (Dwork's Splitting Functions). *Let s denote a positive integer or infinity. Choose $\gamma_s \in \mathbb{Z}_p[\pi]$ with $\text{ord}\,(\gamma_s) = 1/(p-1)$ and*

$$\sum_{j=0}^{s} \frac{\gamma_s^{p^j}}{p^j} = 0.$$

Define

$$\theta_s(t) := \exp(\sum_{j=0}^{s} \frac{(\gamma_s t)^{p^j}}{p^j}) = \sum_r \lambda_r(s)t^r \in \mathbb{Z}_p[\pi][[t]].$$

Then the coefficients $\lambda_r(s)$ satisfy

$$\text{ord}\,(\lambda_r(s)) \geq rc_s,$$

where

$$c_s := \frac{1}{p-1} - \frac{(p-1)(s+1)+1}{(p-1)p^{s+1}} > 0$$

and thus θ_s *converges on the Teichmuller lifting of elements in* $\bar{\mathbb{F}}_q$.

Moreover for $k \geq 1$ *define*

$$\Phi_{k,s}(t) := \prod_{i=0}^{ak-1} \theta_s(t^{p^i}).$$

Then the maps $\Psi_{k,s} = \Phi_{k,s} \circ$ *Teich form a tower of non-trivial additive characters from* $\mathbb{F}_{q^k}$ *to* Ω_1, *where* $q = p^a$. *(Here Teich is the map which takes a finite field element* $\bar{x}$ *to its Teichmuller lifting* x*.)*

We shall use the simplest splitting function obtained from the case $s = 1$. Then γ_1 is a root of $t+(t^p/p)$ with p-adic order $1/(p-1)$. It is straightforward to check that π, where $\pi^{p-1} = -p$, is such a root. Precisely

Definition 7. *Let* $\theta(t)$ *be the power series in* t *with coefficients in* $\mathbb{Z}_p[\pi]$ *given by*

$$\theta(t) := \exp(\pi(t - t^p)) = \sum_{m=0}^{\infty} \lambda_m t^m \in \mathbb{Z}_p[\pi][[t]].$$

Let

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

and $\Psi_k = \Phi_k \circ$ *Teich.*

By Lemma 6 the coefficients λ_m satisfy

$$\text{ord}\,(\lambda_m) \geq \frac{(p-1)m}{p^2},\ \lambda_m \in \Omega_1. \tag{3}$$

Furthermore, the maps Ψ_k form a tower of non-trivial additive characters from $\mathbb{F}_{q^k}$ to $\mathbb{Q}_p(\pi)$.

We now define another power series related to our original polynomial f whose relevance will become apparent in Proposition 9.

Definition 8. *Let* $F^{(a)}(X)$ *be the formal power series in the indeterminates* X *with coefficients in* R_0 *given by*

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

Here f *is the polynomial whose zeta function we wish to compute, and*

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

where J *is the support set of* $X_0 f$, *and* a_j *the Teichmuller lifting of* $\bar{a}_j$.

Proposition 9. *Let $f \in \mathbb{F}_q[X_1, \dots, X_n]$ and $F^{(a)} \in R_0[[X]]$ be as in Definition 8. Denoting by N_k^* the number of solutions to the equation $f = 0$ in the affine torus $(\mathbb{F}_{q^k}^*)^n$ we have*

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

where the summation is over the Teichmuller 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 Teichmuller lifting x we have that

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

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 Teichmuller lifting in $(K^*)^{n+1}$ of points in $(\mathbb{F}_{q^k}^*)^{n+1}$. Combining this with Lemma 4 gives us the result. ∎

The next step we need to take is to find a convenient expression for the righthand-side of the equation in Proposition 9. This is achieved through Dwork's trace formula (Theorem 25). This formula itself does not immediately yield an efficient method of computing zeta functions in small characteristic. We shall also need an alternative expression for the power series $F^{(a)}$.

Definition 10. *Let F be the power series in the indeterminates X with coefficients in R_0 given by*

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

where θ is as in Definition 7.

It follows immediately that

Lemma 11. *Let the power series $F^{(a)}$ and F be as in Definitions 8 and 10. Then*

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

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

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 $\mathbb{N}^{n+1}$. Let A be the $(n+1) \times |J|$ matrix whose columns are $j = (j_0, j_1, \cdots, j_n) \in J$. Then from Definitions 7 and 10

$$F_r = \sum_u (\prod_{j \in J} \lambda_{u_j} a_j^{u_j}) \tag{4}$$

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

$$Au = r, \tag{5}$$

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

$$\sum_{j \in J} u_j = r_0. \tag{6}$$

Now F_r is zero if (5) has no solutions. Otherwise, since $\mathrm{ord}\,(\lambda_{u_j} a_j^{u_j}) = \mathrm{ord}\,(\lambda_{u_j})$ from (3),(4) and (6) we get

$$\mathrm{ord}\,(F_r) \geq \inf_u \left\{ \sum_{j \in J} \frac{(p-1)u_j}{p^2} \right\} = \frac{p-1}{p^2} r_0, \tag{7}$$

where the inf is over all non-negative integer vector solutions u of (5). We now describe a weight function "$w(*)$" which gives estimates on $\mathrm{ord}\,(F_r)$.

Denote by Δ the polytope obtained as the convex hull in $\mathbb{R}^{n+1}$ of the support set of $X_0 f$ and 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. This cone is just obtained by embedding the Newton polytope of f itself in the hyperplane "$X_0 = 1$" of $\mathbb{R}^{n+1}$ and taking the positive hull with the origin.

Definition 12. *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*

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

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

The weight function has a simple geometric interpretation: For a real number c define $c\Delta=\{cx \mid x\in\Delta\}$. It is clear that we have,

Lemma 13. *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* $w(r)=\infty$.

Moreover, we note for later that $(r_0,r_1,\dots,r_n)\in C(\Delta)$ if and only if $(r_1,\dots,r_n)\in r_0\mathrm{Newt}(f)$, where $\mathrm{Newt}(f)$ is the Newton polytope of f as defined in Section 1.

By (7), we have

Lemma 14. *We have the following inequality*

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

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

Lemma 15. *Let* $r,r'\in C(\Delta)$ *and* k *a non-negative integer. Then* $w(kr)=kw(r)$ *and* $w(r+r')\leq w(r)+w(r')$. *In particular,*

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

We shall work in certain subrings of $R_0[[X]]$ defined in terms of the weight function.

Definition 16. *Define* L_Δ *to be the subring of* $R_0[[X]]$ *given by*

$$L_\Delta:=\{\sum_{r\in C(\Delta)\cap\mathbb{Z}^{n+1}}A_rX^r \mid A_r\in R_0\}.$$

Thus, L_Δ is just the ring of all power series over R_0 whose terms have exponents in the cone $C(\Delta)$. Certainly $F\in L_\Delta$ and from Lemma 11 we see easily that $F^{(a)}\in L_\Delta$.

Lemma 17. *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 proof of the running time of our algorithm.

Definition 18. *For any positive real number* b, *we define a set of power series by*

$$L_\Delta(b):=\{\sum_r A_rX^r\in L_\Delta \mid ord\,A_r\geq bw(r)\}.$$

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

$$F \in L_\Delta\left(\frac{p-1}{p^2}\right) \tag{8}$$

$$F^{(a)} \in L_\Delta\left(\frac{p-1}{qp}\right). \tag{9}$$

The first is immediate from the inequality in Lemma 14 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 the coefficients of $F^{(a)}$ decay much more slowly than F itself.

Note 19 The weight function w in this section is that of Adolphson and Sperber [2]. Dwork's original work uses the weight function w_D defined as follows: $w_D(r) = r_0$ if $r_1 + \ldots + r_n \leq r_0 d$ and ∞ otherwise. We have that $w(r) \geq w_D(r)$ for all points r, and the two weight functions agree on $C(\Delta)$. The essential difference is that w_D attaches finite weights to some points outside of $C(\Delta)$, and as such it is not as refined. However, Dwork's weight function does have the key properties in Lemmas 14 and 15, and the main theoretical result (Theorem 27) holds if we replace w by w_D throughout the paper, with minor modifications. The disadvantage is that the matrix "A_N" in Theorem 27 is slightly larger. However, using Dwork's weight function one does not encounter any complications from computational convex geometry, as experienced in Section 6.3.3. Moreover Theorem 1 can still be proved by working with Dwork's weight function. We return to this aside in Note 37.

5. Dwork's trace formula

5.1. **Lifting Frobenius.** We now introduce Dwork's "one-sided inverse of Frobenius" mapping ψ_p on the ring $R_0[[X]]$.

Definition 20. *Let ψ_p be defined on the monomials in $R_0[[X]]$ 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 $R_0[[X]]$. That is

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

Here $p|r$ means that p divides all of the entries in the integer vector r. This map is a one-sided inverse of the "Frobenius" map on the ring $R_0[[X]]$ which takes a power series $\sum_r A_r X^r$ to $\sum_r \tau(A_r) X^{pr}$.

Definition 21. *Let α_a be 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 $R_0[[X]]$. (Notice that ψ_p^a just acts as

$$\psi_p^a(\sum_r A_r X^r) = \sum_r A_{qr} X^r$$

since $\tau^{-a}(A_r) = A_r$.) Let the map α on $R_0[[X]]$ be defined as

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

Thus α comprise 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 22. *Let α and α_a denote the maps defined on $R_0[[X]]$ as $\alpha = \psi_p \circ F$ and $\alpha_a = \psi_p^a \circ F^{(a)}$. Then*

$$\alpha_a = \alpha^a$$

where the second exponent is a power under composition.

Proof. Firstly let $H \in R_0[[X]]$ and denote by $\psi_p \circ H$ the map composed of multiplication by H followed by ψ_p. We claim that

$$\psi_p \circ H(X^p) = \tau^{-1}(H(X)) \circ \psi_p \tag{10}$$

To see this write $H = \sum_r H_r X^r$. 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 $\psi_p \circ X^{pr} = X^r \circ \psi_p$ 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 proof this by induction, the result trivially holding if $b = 1$. By $b-1$ applications of (10) 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, in the sense that

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

for any element $\sum_r A_r X^r \in R_0[[X]]$. The map α is τ^{-1}-linear, in the sense that

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

¿From Lemma 17 it follows easily that

Lemma 23. *The maps α and α_a are both stable on the ring L_Δ.*

Thus 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.** Let $\mathbb{N}$ denote the set of non-negative integers. Partition the vectors u in $\mathbb{N}^{n+1}$ such that $w(u) < \infty$ into sets B_i ($i \in \mathbb{N}$) as follows: Let B_i consist of all lattice points u such that $w(u) = i$. The union $\cup_i B_i$ is just the set of lattice points in the cone $C(\Delta)$.

We now choose a total ordering $\prec$ on $\cup_i B_i$ such that if $u \in B_i$ and $v \in B_j$ with $i < j$ then $u \prec v$. This defines a total ordering on the set of monomials

$$\{X^u \mid w(u) < \infty\}$$

in a straightforward way. This set of monomials forms what we call a *formal basis* for the ring L_Δ. Precisely, this means that any power series in L_Δ may be written in exactly one way as an infinite sum $\sum_u A_u X^u$ where X^u are in the set of monomials with $w(u) < \infty$ ordered as described. 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_r goes to zero as $w(r)$ goes to ∞ (see [19] for a more detailed discussion of these notions).

Recall from Lemma 23 that α and α_a are stable on the ring L_Δ. We define certain matrices associated to the maps α and α_a restricted to L_Δ with regard to this formal basis of monomials.

Definition 24. *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 basis of monomials in L_Δ. Specifically, the (u, v)th entry of M and M_a for $u, v \in \cup_i B_i$ 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.

We have by Lemmas 14 and 15

$$\begin{aligned} \operatorname{ord}(F_{pu-v}) &\geq \frac{p-1}{p^2} w(pu - v) \\ &\geq \frac{p-1}{p}\left(w(u) - \frac{1}{p}w(v)\right) \qquad (11) \end{aligned}$$

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

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. For a fixed $u \in B_i$ this will be true for any $v \in B_j$ for j sufficiently large.) 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 $Tr(M_a^k)$. It is easily seen that for $k > 1$

$$Tr(M_a^k) = \sum_i \left(\sum_{j_1,\dots,j_{k-1}} m_{ij_1}^{(a)} m_{j_1j_2}^{(a)} \dots m_{j_{k-1}i}^{(a)} \right), \tag{12}$$

Here the inner sum is over all sequences of length $k-1$ of vectors j in $\cup_i B_i$.

We can now prove

Theorem 25 (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 $R_0[[X]]$ given as $\alpha_a = \psi_p^a \circ F^{(a)}$, where $F^{(a)}$ and ψ_p are described in Definitions 8 and 20. Let M_a denote the infinite matrix representing the map α_a restricted to the subring L_Δ as described in Definition 24. 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*

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

Proof. We will follow the matrix proof given in [17], rather than the slicker version in Dwork's original paper [4].

By Proposition 9 we have that

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

where the sum is over all $(n+1)$-tuples of $(q^k - 1)$st roots of unity in R_0 (namely the Teichmuller lifting in $(K^*)^{n+1}$ 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 $F^{(a)}(X) = \sum_r b_r X^r$ where the sum is over all vectors r which belong to $\cup_i B_i$. (Thus, for simplicity, we just replace our usual notation $F_r^{(a)}$ for the coefficient of X^r by b_r.) Then the latter sum is just

$$\begin{array}{rcl} \sum_{x^{q-1}=1} F^{(a)}(x) & = & \sum_{x^{q-1}=1} (\sum_r b_r x^r) \\ & = & \sum_r b_r (\sum_{x^{q-1}=1} x^r) \\ & = & (q-1)^{n+1} \sum_{r,(q-1)|r} b_r \\ & = & (q-1)^{n+1} \sum_s b_{(q-1)s} \\ & = & (q-1)^{n+1}\, Tr(M_a). \end{array}$$

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 [12, 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 $k>1$ the proof is similar. We have

$$\begin{aligned} & \sum_{x^{q^k-1}=1} F^{(a)}(x)\ldots F^{(a)}(x^{q^{k-1}}) \\ = & \sum_{u^{(0)},\ldots,u^{(k-1)}} b_{u^{(0)}}\ldots b_{u^{(k-1)}} \sum_{x^{q^k-1}=1} x^{q^{k-1}u^{(k-1)}+\ldots+qu^{(1)}+u^{(0)}} \\ = & (q^k-1)^{n+1} \sum_{u^{(0)}+u^{(1)}q+\ldots+u^{(k-1)}q^{k-1}\equiv 0 \bmod (q^k-1)} b_{u^{(0)}}\ldots b_{u^{(k-1)}}. \end{aligned} \tag{13}$$

To show this is equal to the number $(q^k-1)^{n+1}\, Tr(M_a^k)$ one writes all lattice point solutions of the equation

$$u^{(0)}+qu^{(1)}+\ldots+q^{k-1}u^{(k-1)} = (q^k-1)v^{(k-1)}$$

where the $u^{(j)}$ lie in the set $\cup_i B_i$ in the form

$$\begin{aligned} u^{(0)} &= qv^{(0)}-v^{(k-1)} \\ u^{(1)} &= qv^{(1)}-v^{(0)} \\ &\vdots \\ u^{(k-1)} &= qv^{(k-1)}-v^{(k-2)}. \end{aligned}$$

where each $v^{(j)}\in\cup_i B_i$.

Then from (13) we find that $q^k N_k^* - (q^k-1)^n$ is

$$\begin{aligned} & (q^k-1)^{n+1}\textstyle\sum_{v^{(0)},\ldots,v^{(k-1)}} b_{qv^{(0)}-v^{(k-1)}} b_{qv^{(1)}-v^{(0)}}\ldots b_{qv^{(k-1)}-v^{(k-2)}} \\ = & (q^k-1)^{n+1}\, Tr(M_a^k) \end{aligned}$$

the equality following from (12). □

To compute the trace of the matrix in Dwork's formula we shall use the following matrix identity derived from Lemma 22.

Lemma 26. *Let M_a and M denote the matrices for the maps α_a and α described above. Then*

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

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 $R_0[[X]]$. We assume that β_1 is τ^{-1}-linear and β_2 is τ^{-j}-linear for some $0\le j\le a-1$. 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 22 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 $b = a$. □

We note in passing that M_a in Lemma 26 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 p-adic integers in R_0, and so this reduction is defined.

Theorem 27. *For $i \geq 0$, denote by B_i the set of all integer vectors $u \in \mathbb{N}^{n+1}$ such that $w(u) = i$ (see Definition 12). Let L_Δ denote all power series over R_0 whose monomials have exponents lying in the cone $\cup_i B_i$ (Definition 16). Fix an ordering on the monomials $X^u \in L_\Delta$ such that those in B_i come before those in B_j for $i < j$. Let M denote the matrix for the τ^{-1}-linear map $\alpha = \psi_p \circ F$ with respect to this formal basis of monomials of L_Δ (Definition 24). Here ψ_p is the "one-sided inverse of Frobenius" given in Definition 20 and F is the power series obtained from the polynomial f as in Definition 10.*

Let N denote any positive integer and A_N the finite square matrix over the finite ring $R_0/(p^N)$ obtained by reducing modulo p^N all those entries in M whose rows and columns are indexed by vectors in the set B_i for $0 \leq i \leq \lfloor \frac{p^2}{(p-1)^2} N \rfloor$. Then

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

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 $t = \lfloor p^2 N/(p-1)^2 \rfloor$ and $\#(t\Delta)$ is number of lattice points in a dilation by a factor t of the polytope Δ.

Proof. For any matrix L finite or infinite with coefficients in R_0, we define $\overline{L}$ to be the matrix obtained by reducing all its entries, which are p-adic integers, modulo p^N. Thus $\overline{L}$ has entries in $R_0/(p^N)$, whose elements can be represented as expansions of length $a(p-1)N$ with coefficients in the set $\{0, 1, \dots, p-1\}$.

The theorem will follow from the Dwork Trace Formula (Theorem 25) 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 26 this matrix can be computed as a matrix product from the matrix $\overline{M}$.

We have the following block form for M

$$M = \begin{bmatrix} M_{01} & M_{02} & M_{03} & \dots \\ M_{11} & M_{12} & \dots & \\ M_{21} & \vdots & & \\ \vdots & & & \end{bmatrix}$$

where the rows and columns in the finite block matrices M_{ij} are labelled by vectors $u \in B_i$ and $v \in B_j$. Thus the (i,j)th block is a finite matrix of size $|B_i| \times |B_j|$. Now for $u \in B_i$ and $v \in B_j$ we have that $w(u) = i$ and $w(v) = j$.

By inequality (11) every entry $m_{uv} := \tau^{-1}(F_{pu-v})$ in M_{ij} satisfies

$$\mathrm{ord}\,(m_{uv}) \geq \frac{p-1}{p}\left(i - \frac{j}{p}\right)$$

Thus the matrix $\overline{M}$ is of the form

$$\begin{pmatrix} A_N & B_N \\ 0 & C_N \end{pmatrix} \tag{14}$$

where C_N is a strictly upper triangular infinite matrix and A_N is as follows: the matrix A_N is the finite square matrix of the form

$$\begin{bmatrix} \overline{M_{01}} & \dots & \overline{M_{0t}} \\ \vdots & & \vdots \\ \overline{M_{t1}} & \dots & \overline{M_{tt}} \end{bmatrix}.$$

Here the $\overline{M_{ij}}$ is the reduction modulo p^N of the finite block M_{ij} and t is chosen to satisfy

$$\frac{p-1}{p}\left(t - \frac{t}{p}\right) < N \leq \frac{p-1}{p}\left((t+1) - \frac{t+1}{p}\right).$$

Thus we need

$$\left(\frac{p}{p-1}\right)^2 N - 1 \leq t < \left(\frac{p}{p-1}\right)^2 N$$

from which it follows $t = \lfloor p^2 N/(p-1)^2 \rfloor$. (Actually for $p = 2$ this value of t is not quite optimal, but certainly suffices.) Note that C_N is strictly upper triangular because any entry in a block B_{ij} with $i \geq j > t$ has p-adic order at least

$$\frac{p-1}{p}\left(i - \frac{j}{p}\right) \geq \frac{p-1}{p}\left(j - \frac{j}{p}\right) \geq (t+1)\left(\frac{p-1}{p}\right)^2 \geq N.$$

The size of A_N is $W = |\cup_{i=0}^t B_i|$ where $t = \lfloor p^2 N/(p-1)^2 \rfloor$. By Lemma 13 this is exactly the number of lattice points in the polytope $t\Delta$.

By Lemma 26 and modular reduction, one has

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

Thus by (14) 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 thus strictly upper triangular.

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

$$(\prod_{i=0}^{a-1} \tau^{-i}(A_N))^k \tag{15}$$

The theorem now follows from (15) and Theorem 25. □

6. Algorithms

In this section we present an algorithm for counting points based upon Theorem 27, 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, making it as self-contained as possible for the casual reader.

Algorithm 28 Toric point counting

Input: Positive integers a, k, d and a polynomial f. Here f is a polynomial in n variables of total degree d with coefficients in the field $\mathbb{F}_q$ where $q = p^a$.

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 $(R_0/(p^N))[X]$ where F is the power series in Definition 10. 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 Teichmuller lifting of $\bar{a}_j$ and $\theta(z) = \exp(\pi(z - z^p))$. (See Section 3 for a description of the ring $R_0/(p^N)$.)

Step 2: Construct the matrix A_N which occurs in the statement of Theorem 27. 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. The (u, v)th entry of A_N is τ^{-1} of the coefficient of X^{pu-v} in the polynomial $F \bmod p^N$. Here Δ is the polytope in $\mathbb{R}^{n+1}$ obtained as the convex hull of the origin and $X_0 f$, $t = \lfloor p^2 N/(p-1)^2 \rfloor$, and 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 bracket denotes the smallest non-negative residue modulo p^N.

6.2. **Proof of correctness of the algorithm.** We know by Theorem 27 that

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

in the ring $R_0/(p^N)$. 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 25), and thus it is also an integer. Since $N_k^* \le (q^k-1)^n$ it follows that the lefthand side is not greater than $q^{k(n+1)}$. Hence in the case $N \ge ak(n+1)$ 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 $|C_1| \le c(|C_2|+1)$ for some positive constant c. We write $C_1 = \tilde{\mathcal{O}}(C_2)$ if $C_1 = \mathcal{O}(C_2 \log(|C_2|+1)^d)$ for some constant d. 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 $R_0/(p^N)$ 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 computation of the map τ in the ring $R_0/(p^N)$. In Section 6.3.2 we shall add back in the small contribution from computing Teichmuller liftings. Our running time will be in terms of the parameters $t, \tilde{t}, W, \tilde{W}$. Here

$$\begin{array}{ll} t = \lfloor \left(\frac{p}{p-1}\right)^2 N \rfloor, & W = |\cup_{i=0}^{t} B_i|, \\ \tilde{t} = \lfloor \frac{p^2}{p-1} N \rfloor, & \tilde{W} = |\cup_{i=0}^{\tilde{t}} B_i|. \end{array}$$

where $N = ak(n+1)$. 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 larger than t, $\tilde{W}$ will be around p^{n+1} times larger than W, as 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 $L_\Delta((p-1)/p^2) \subseteq R_0[[X]]$ 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

Lemma 29. *Let F be the power series given in Definition 10 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 $R_0/(p^N)$.

Proof. By (3) we have that $\theta(t) \bmod p^N$ is a polynomial of degree not greater than $p^2N/(p-1)$. Thus we can obtain $\theta(t)$ via the formula

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

by computing the first $\mathcal{O}(pN)$ terms in the expansion for $\exp(\pi z)$, substituting $z = t$ and $z = -t^p$, and one multiplication of polynomials. Note that $\exp(\pi z)$ has p-adic integral coefficients. Thus $\theta(t)$ can be found in time $\mathcal{O}((pN)^2)$ operations in the ring $R_0/(p^N)$ using standard polynomial arithmetic.

Now by (8) 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 10 in $|J|$ multiplications of polynomials of the form $\theta(a_jX^j) \bmod p^N$. Each such polynomial lies in the ring $L_\Delta(p-1/p^2) \bmod p^N$, because $\operatorname{ord}(a_j) = 0, w(j) = 1$ 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 $|J|$ multiplications are required. Thus the complexity is $\mathcal{O}(|J|\tilde{W}^2)$ ring operations. Noting that $pN = O(\tilde{W})$ we have the result. □

Note 30 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 $u, v \in \cup_{i=0}^{t} B_i$, 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 $\tau^{-1} = \tau^{a-1}$, which can safely be ignored. 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

Lemma 31. *With the notation as in the statement of Theorem 27, the product*

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

can be computed given the matrix A_N in

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

operations in the ring $R_0/(p^N)$.

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

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

Now the case $r = s = 2^m$ for some m gives us the "square" step (computing $\alpha^{2^{m+1}}$ from α^{2^m}) and the case $r = 2^m$ 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 $\mathcal{O}(\log_2(a))$ "matrix ring" operations. By "matrix ring" operations we mean multiplication of matrices of size W over $R_0/(p^N)$, and also computing $\tau^{-i}(B)$ for some $1 \leq i \leq a-1$ and matrix B of this form. The former requires $O(W^3)$ operations in the ring $R_0/(p^N)$ using standard algorithms. Since $\tau^{-i} = \tau^{a-i}$ the latter may be computed in $\mathcal{O}(W^2 \log(a))$ ring operations using a fast "square-and-multiply" method in the ring $R_0/(p^N)$. 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 ring operations. Thus the total time required is $O(W^3 \log(ak))$ operations in the ring $R_0/(p^N)$. □

Note that the exponent 3 for multiplication of matrices can be improved to around 2.4 using faster methods [8, page 330].

Gathering these results we find

Lemma 32. *The running time of Algorithm 28 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 $N = ak(n+1)$.

Thus the complexity of the algorithm is closely related to the geometry of the Newton polytope of f. Also, if the polynomial is sparse one may be able to further refine this complexity, although we do not pursue this.

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 $R_0/(p^N)$. In this section we shall also count the small contribution from computing the Teichmuller lifting of the coefficients of f, which it was convenient to ignore in Section 6.3.1. Without loss of generality, we may assume that $\mathrm{Newt}(f)$ is n-dimensional.

Definition 33. *Let $V(Newt(f))$ be the volume of Newton polytope of f. Denote by $m := n!V(Newt(f))$ the "normalised" volume of $Newt(f)$.*

Proposition 34. *Ignoring auxiliary operations outside of the ring $R_0/(p^N)$, the running time of Algorithm 28 is*

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

bit operations. The space complexity in bits is

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

Proof. To compute the complexity first observe that $t = \lfloor p^2N/(p-1)^2 \rfloor \leq 4N$ and certainly $t \geq N$. Thus $t, N = \mathcal{O}(akn)$ in the algorithm. Also $N = \mathcal{O}(t)$ and we have $\tilde{t} = \mathcal{O}(pN) = \mathcal{O}(pt)$.

Now $W = \#(t\Delta) \leq (n+1)!V(t\Delta) + (n+1) = (n+1)!t^{n+1}V(\Delta) + (n+1) = t^{n+1}n!V(\mathrm{Newt}(f)) + (n+1) = \mathcal{O}(mt^{n+1})$. (Here we use the Blichfeldt bound $\#(P) \leq n!V(P) + n$ for any n-dimensional polytope P from [10, page 144], the fact $V(t\Delta) = t^{n+1}V(\Delta)$ since Δ is full-dimensional, and also that $(n+1)!V(\Delta) = n!V(\mathrm{Newt}(f))$.) Similarly $\tilde{W} = \#(\tilde{t}\Delta) = \mathcal{O}(m\tilde{t}^{n+1}) = \mathcal{O}(mt^{n+1}p^{n+1})$.

Thus from Lemma 32 the number of ring operations is

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

since $|J| \leq \#(\mathrm{Newt}(f)) = \mathcal{O}(m)$. Thus the bit complexity which arises from ring operations is by Lemma 5

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

Tidying up and ignoring logarithmic factors we get

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

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

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

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

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

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 Teichmuller liftings of the coefficients of f in Step 1. 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, which is $\mathcal{O}(W^2)$ ring elements. Putting $W = \mathcal{O}(m(ank)^{n+1})$ and using Lemma 5 gives us the result. □

Corollary 35. *Let f be a polynomial of degree d in n variables over $\mathbb{F}_q$. Ignoring auxiliary operations outside of the ring $R_0/(p^N)$, the running time of Algorithm 28 is*

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

bit operations. The space complexity in bits is

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

In this case, it is clear that m is at most d^n and hence the corollary follows.

6.3.3. *Bit complexity arising from auxiliary operations.* It remains to bound the complexity which arises from operations outside of the ring $R_0/(p^N)$. First note that 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 $t = \mathcal{O}(ank)$. By the description of Δ in terms of $\mathrm{Newt}(f)$ (following Lemma 13) it suffices to find all rational points i/r in $\mathrm{Newt}(f)$ where i is an integer vector and the denominator r a positive integer bounded by t. This may be done in several ways, and we describe one naive method: First observe that the problem can be reduced to that of listing rational points in the $\mathcal{O}(|J|^{n+1})$ simplices obtained by selecting all subsets of $n+1$ affinely independent points from J and taking their convex hull. (That the union of all such simplices equals $\mathrm{Newt}(f)$ is Carathéodory's theorem [10, page 69].) Given such a simplex we may assume its vertices are $0, v_1, \ldots, v_n$, where v_i are column vectors of length n. Let A denote the invertible integer matrix with columns v_i. We need to find all column vectors $x := (x_1, \ldots, x_n)^t$ such that $0 \le x_i \le 1$, and $\sum_i x_i \le 1$, with Ax a rational vector whose coefficients lie in $\mathbb{Z}/t := \{z/t \,|\, z \in \mathbb{Z}\}$. Now for any vector y, if $Ay = z$ where z has coefficients in $\mathbb{Z}/t$, then $y = (1/\det(A))A^{adj}z$, where A^{adj} is the integer adjoint matrix of A. Hence the coefficients of y must lie in $\mathbb{Z}/\det(A)t$. Thus to find all suitable rational points in the simplex, it suffices to take all points x with coefficients in $\mathbb{Z}/\det(A)t$ between 0 and 1 with sum bounded by 1, and compute Ax. This will be a point in the simplex, and it only remains to check whether its coordinates lie in $\mathbb{Z}/t$. Now $\det(A)$ is just $n!$ times the volume of our simplex, and so $\det(A) = \mathcal{O}(m)$. Thus the complexity of this method for a single simplex is $\tilde{\mathcal{O}}((mt)^n)$ arithmetic operations with integers bound by $\log d$. There are $\mathcal{O}(m^{n+1})$ such simplices, giving an overall complexity of $\tilde{\mathcal{O}}(t^n m^{2n+1})$.

Lemma 36. *The contribution to the bit complexity of Algorithm 28 from auxiliary operations (in convex geometry) outside the ring $R_0/(p^N)$ is*

$$\tilde{\mathcal{O}}((ank)^n m^{2n+1}).$$

Note that a similar approach using the Smith Normal Form of A reduces the complexity to $\tilde{\mathcal{O}}((tm)m^{n+1})$.

Note 37 If one is interested solely in counting points over the field $\mathbb{F}_q$, rather than computing the full zeta function, then the above estimate is inconveniently large. Specifically it contributes $d^{\mathcal{O}(n^2)}$. Using Dwork's weight function instead, as defined in Note 19, the labelling set for the matrix "A_N" is just the set of lattice points $(r_0, \dots, r_n)$ with $r_1 + \ldots + r_n \leq r_0 d$ and $0 \leq r_0 \leq t$. Such points are easily found and the complications above vanish. In particular, for counting points over the base field using Dwork's weight function the auxiliary operations are bound by $\mathcal{O}(d^n)$ (thus justifying the claims on the growth in terms of d in Section 1). Similar comments apply to families of varieties, such as hyperelliptic curves, for which the Newton polytope has a very simple description.

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

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

Denote by f^S the polynomial obtained from f be 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 28. Thus by 2^n applications of this algorithm we obtain N_k as desired. For n and p fixed the estimates in Corollary 35 and Lemma 36 show that this is a polynomial time algorithm.

Now to obtain the whole zeta function $Z(f/\mathbb{F}_q)$ it suffices to count N_k for all $k = 1, \dots, 2\deg(r) + 2\deg(s)$, where

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

with r and s coprime polynomials. By the refinement of Bombieri's degree bound [3] from [2, Equation (1.13)], the "total degree" $\deg(r) + \deg(s)$ is bounded by $2^{n+1}6^{n+1}(n+1)!V(\Delta)$. We take the estimate $\mathcal{O}(2^{4n+4}m)$ for this bound.

Theorem 38. *Let f be a polynomial in n variables over $\mathbb{F}_q$, where $q = p^a$. Denote by m the normalised volume $n!V(Newt(f))$ of the Newton polytope of f. The full zeta function $Z(f/\mathbb{F}_q)$ can be computed deterministically in*

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

bit operations. (Note that in this estimate we have ignored logarithmic factors in d, the degree of f.)

Proof. ¿From Proposition 34 and the torus decomposition method, the bit complexity arising from ring operations is

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

Tidying up the factor in n, and noting that, by Lemma 36, this dominates the complexity arising from auxiliary operations, we get the claimed result. □

Putting $m = \mathcal{O}(d^n)$ we get

Corollary 39. *Let f be a polynomial of total degree d in n variables over $\mathbb{F}_q$ where $q = p^a$. Then 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.

Since we may assume that $d > 1$ we have $2^{n^2} = \mathcal{O}(d^{n^2})$, and Theorem 1 now follows. (Note that the large factor in n in the first term of this estimate is an artefact of our general bound on the total degree of the zeta function. Using Bombieri's original estimate [3] it can be modestly improved, and it can be greatly reduced under the assumption of non-singularity of the variety.)

The proof of Corollary 2 was explained in the introduction, and we finish with some comments on Corollary 3. The order of the Jacobian follows easily once the zeta function has been computed, see [20]. Thus for fixed dimension and finite field one can compute the order of the Jacobian 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] S. Gao and A.G.B. Lauder, Factoring polynomials via polytopes, preprint.
[8] J. von zur Gathen and J. Gerhard, Modern Computer Algebra, Cambridge University Press, 1999.
[9] P. Gaudry and N. Gürel, An extension of Kedlaya's algorithm for counting points on superelliptic curves. preprint.
[10] J.E. Goodman and J. O'Rourke (Eds), Handbook of Discrete and Computational Geometry, CRC Press, 1997.
[11] K.S. Kedlaya, Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology, preprint.
[12] N. Koblitz, p-Adic Numbers, p-Adic Analysis and Zeta Functions, 2nd Edition, Springer, 1984.
[13] R. Lidl and H. Niederreiter, Introduction to Finite Fields and Their Applications, Cambridge University Press, 1986.
[14] 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.
[15] R. Schoof, Elliptic curves over finite fields and the computation of square roots mod p, Math. Comp. 44 (1985), no. 170, 483-494.
[16] R. Schoof, Counting points on elliptic curves over finite fields, J. Theorie des Nombres de Bordeuax 7, 1998, 219-254.
[17] D. Wan, Meromorphic continuation of L-functions of p-adic representations, Ann. Math., 143 (1996), 469-498.
[18] D. Wan, Computing zeta functions over finite fields, Contemporary Mathematics, 225 (1999), 131-141.
[19] D. Wan, Rank one case of Dwork's conjecture, J. Amer. Math. Soc., 13(2000), 853-908.
[20] 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