Algorithmic Theory of Zeta Functions over Finite Fields

Daqing Wan*

March 30, 2003

Department of Mathematics
University of California
Irvine, CA 92697-3875
Email: dwan@math.uci.edu

Abstract

This is an expanded version of my lecture given at the 2000 MSRI introductory workshop on computational number theory. It gives a brief introductory account of the general algorithmic theory of the zeta function of an algebraic set defined over a finite field.

1 Introduction

Let $\mathbf{F}_q$ be a finite field of q elements and p its characteristic. Let X be an algebraic set defined over $\mathbf{F}_q$. For each positive integer k, let N_k denote the number of $\mathbf{F}_{q^k}$-rational points on X. The zeta function $Z(X)$ of X is the generating function

$$Z(X) = Z(X,T) = \exp\Big(\sum_{k=1}^{\infty} \frac{N_k}{k} T^k\Big).$$

The zeta function contains important arithmetic and geometric information about X. It has been studied extensively in connection with the celebrated Weil conjectures [37].

Both practical applications and theoretical investigations make a good understanding of the zeta function from an algorithmic point of view increasingly important. The aim of this paper is to present a brief introductory

*Partially supported by the NSF and the NSFC

account of the various fundamental problems and results in the emerging algorithmic theory of zeta functions. We shall focus on general properties rather than on results that are restricted to special cases. In particular, in most of this paper we do not assume X to be smooth and projective, although in that case one can often say more.

The contents are organized as follows. In Section 2 we review general properties of zeta functions from an algorithmic point of view. A naive effective algorithm for computing the zeta function is given. If the characteristic p is small, one can use Dwork's p-adic method to obtain a polynomial time algorithm for computing the zeta function in the case that the numbers of variables and defining equations for X are fixed.

In Section 3, we show that the general case of algebraic sets can be reduced in various ways to the case that X is a hypersurface. A more detailed discussion of that crucial case is given in Section 4, with emphasis on the smooth projective case. In Section 5 we consider the complex pure weight decomposition. Using the LLL factorization algorithm and Deligne's main theorem, we show that, when the zeta function is given, one can compute in polynomial time how many zeros and poles with a given complex absolute value it has. In Section 6, which is devoted to the p-adic pure slope decomposition, we use the theory of Newton polygons to obtain a similar result for the number of zeros and poles with a given p-adic absolute value.

We conclude the paper by giving, in Section 7, an algorithm for the simpler problem of computing the zeta function modulo p. This algorithm shares several characteristic features with the general p-adic method for computing the full zeta function that is presented in [22]. Section 7 may thus serve as an introduction to [22].

All algorithms in this paper are deterministic. Probabilistic algorithms will not be discussed. Time is measured in bit operations.

Acknowledgment. It is a pleasure to thank S. Gao, H.W. Lenstra, Jr., and A.M. Odlyzko for several interesting discussions. Special thanks are due to the referee for many helpful suggestions.

2 Generalities on computing zeta functions

Let X be an algebraic set defined over a finite field $\mathbf{F}_q$ of q elements of characteristic p. For computational purposes, we may assume that X is affine, i.e., that it is the subset of affine n-space $\mathbf{A}^n$ defined by a system of polynomial equations:

$$\begin{cases} f_1(x_1, \ldots, x_n) = 0, \\ \quad \vdots \\ f_m(x_1, \ldots, x_n) = 0, \end{cases}$$

where $f_i \in \mathbf{F}_q[x_1, \ldots, x_n]$. Let

$$X(\mathbf{F}_q) = \{x = (x_1, \ldots, x_n) \in \mathbf{F}_q^n \mid f_1(x) = \ldots = f_m(x) = 0\}$$

be the finite set of $\mathbf{F}_q$-rational points on X. It is clear that the cardinality $\#X(\mathbf{F}_q)$ is effectively computable.

For algorithmic purposes, 'giving' $\mathbf{F}_q$ means specifying p as well as an irreducible polynomial h in one variable over $\mathbf{F}_p$ that defines $\mathbf{F}_q$, so that q equals $p^{\deg h}$ (see [30] for a method to find such a polynomial h when q is given); elements of $\mathbf{F}_q$ are then represented as polynomials of degree less than $\deg h$ in a formal zero of h, with coefficients from $\mathbf{F}_p$. 'Giving' X means specifying a system of m defining polynomials f_i in n variables with coefficients in $\mathbf{F}_q$. Let d be the maximum of the total degrees of the polynomials f_i. Then the dense input size for X is $O(m\binom{d+n}{n}\log q)$, which is $O(m(d+1)^n \log q)$. Our first fundamental problem is the following.

Problem 2.1 *Given $\mathbf{F}_q$ and X, compute the number $\#X(\mathbf{F}_q)$ in time polynomial in the dense input size $O(m(d+1)^n \log q)$.*

This problem is trivial if q is fixed, so we may assume that q is large. In theory, the problem of counting $X(\mathbf{F}_q)$ can be reduced to the zero-dimensional case. Namely, let Y be the zero-dimensional algebraic set defined by

$$\{f_1 = \ldots = f_m = 0,\ x_1^q - x_1 = \ldots = x_n^q - x_n = 0\}.$$

Then it is clear that

$$\#X(\mathbf{F}_q) = \#Y(\mathbf{F}_q).$$

Following a suggestion of Eisenbud and Sturmfels, one may now compute a Gröbner basis for Y, and its cardinality equals $\#Y(\mathbf{F}_q)$. However, as q gets large, the cases where the Gröbner basis computation can be done efficiently are likely to become increasingly exceptional. (See [12] for Gröbner bases.)

Let $\bar{\mathbf{F}}_q$ denote a fixed algebraic closure of $\mathbf{F}_q$. For each positive integer k, let $\mathbf{F}_{q^k}$ denote the unique subfield of $\bar{\mathbf{F}}_q$ with q^k elements. Let $\#X(\mathbf{F}_{q^k})$ denote the number of $\mathbf{F}_{q^k}$-rational points on X. The following problem is harder but more interesting than Problem 2.1.

Problem 2.2 *Given* $\mathbf{F}_q$ *and* X, *compute the sequence of numbers* $\#X(\mathbf{F}_{q^k})$ $(k = 1, 2, \ldots)$.

It may not be clear how a finite algorithm can compute an infinite sequence of numbers, but this will be clarified below. As we shall see, one can encode the entire sequence in a suitably defined generating function, which turns out to be a rational function. This so-called *zeta function* of X is finite in nature and can thus be written down in a finite amount of time. Actually doing this for given X is the content of Problem 2.2.

A *geometric point* of X is an $\bar{\mathbf{F}}_q$-rational point of X. From the equality

$$\bigcup_{k=1}^{\infty} X(\mathbf{F}_{q^k}) = X(\bar{\mathbf{F}}_q)$$

we see that each geometric point of X will be counted somewhere in the sequence of numbers $\#X(\mathbf{F}_{q^k})$. This may explain why many of the subtle geometric invariants associated with an algebraic variety X can be read off from its zeta function, in addition to a wealth of arithmetic information.

Definition 2.3 *The zeta function of* X *is the generating function*

$$Z(X) = Z(X, T) = \exp\Big(\sum_{k=1}^{\infty} \frac{T^k}{k} \#X(\mathbf{F}_{q^k})\Big).$$

The qth power Frobenius map Frob_q is the permutation of the set $X(\bar{\mathbf{F}}_q)$ of geometric points of X defined by:

$$\mathrm{Frob}_q\colon\ x = (x_1, \ldots, x_n) \mapsto x^q = (x_1^q, \ldots, x_n^q).$$

The *degree* of a geometric point x is defined to be the smallest positive integer d such that

$$\mathrm{Frob}_q^d(x) = x$$

or, equivalently, such that $x \in X(\mathbf{F}_{q^d})$. A *closed point* over $\mathbf{F}_q$ is the orbit of a geometric point under Frob_q. All geometric points belonging to a given closed point have the same degree, and this common degree is called the *degree* of the closed point. We denote by $|X|$ the set of closed points of X over $\mathbf{F}_q$, and, for each positive integer k, by $M_k(X)$ the number of closed points of X of degree k. Since each closed point of degree k consists of exactly k points in $X(\mathbf{F}_{q^k})$, one deduces

$$\#X(\mathbf{F}_{q^k}) = \sum_{d|k} d M_d(X).$$

Considering the logarithmic derivative of the zeta function, one finds the *Euler product* expansion

$$Z(X) = \prod_{k=1}^{\infty} \frac{1}{(1-T^k)^{M_k(X)}} = \prod_{x \in |X|} \frac{1}{1-T^{\deg(x)}} \in 1 + T\mathbf{Z}[[T]].$$

As the Weil conjectures [37] predict, the zeta function is a rational function. The first proof, given by Dwork [11], used p-adic analysis. The second proof, given by Grothendieck, used the theory of ℓ-adic cohomology, where ℓ is a prime number different from p. These two proofs pioneered the general p-adic and ℓ-adic study of zeta functions over finite fields.

Theorem 2.4 *The zeta function $Z(X)$ is a rational function, i.e., it belongs to $\mathbf{Q}(T)$. If we write*

$$Z(X,T) = \frac{R_1(X,T)}{R_2(X,T)}, \quad (R_1, R_2) = 1, \quad R_i \in 1 + T\mathbf{Q}[T],$$

then we have $R_i \in 1 + T\mathbf{Z}[T]$.

The rationality of $Z(X)$ has an interesting consequence for the numbers $\#X(\mathbf{F}_{q^k})$, as follows. Let β_i and γ_j denote the reciprocal zeros of $R_1(X)$ and $R_2(X)$, respectively, so that one has $R_1(X) = \prod_i (1-\beta_i T)$ and $R_2(X) = \prod_j (1-\gamma_j T)$. Then one finds

$$\sum_{k=1}^{\infty} \#X(\mathbf{F}_{q^k}) T^k = T \frac{d \log Z(X,T)}{dT} = \sum_j \frac{\gamma_j T}{1-\gamma_j T} - \sum_i \frac{\beta_i T}{1-\beta_i T}.$$

This implies

$$\#X(\mathbf{F}_{q^k}) = \sum_j \gamma_j^k - \sum_i \beta_i^k \quad \text{for all } k \geq 1.$$

As a corollary, one deduces that for each positive integer k one has

$$Z(X \otimes \mathbf{F}_{q^k}) = \frac{\prod_i (1-\beta_i^k T)}{\prod_j (1-\gamma_j^k T)}.$$

The integrality of the coefficients of R_i can be deduced from the rationality of $Z(X)$ and the following elementary result.

Lemma 2.5 *Let $f \in 1 + T\mathbf{Z}[[T]]$ be a rational function. Write*

$$f = \frac{f_1}{f_2}, \quad (f_1, f_2) = 1 \quad f_i \in 1 + T\mathbf{Q}[T].$$

Then one has $f_i \in 1 + T\mathbf{Z}[T]$.

Proof. This is usually derived from the so-called lemma of Fatou, see Katz [15]. Here we include two additional proofs.

One proof uses the Newton polygon or the Weierstrass factorization theorem. Suppose some prime number ℓ occurs in the common denominator of the coefficients of f_1. Then the theory of Newton polygons shows that f_1 has an ℓ-adic zero in the open unit disk $|T|_\ell < 1$. But the power series $f \in 1 + T\mathbf{Z}[[T]]$ is clearly analytic and non-zero in the open unit disk $|T|_\ell < 1$. This gives the desired contradiction.

A second proof, suggested by Hendrik Lenstra, uses Gauss's lemma for power series. The *content* $\mathrm{cont}(g)$ of a non-zero power series $g = \sum_i a_i T^i \in \mathbf{Z}[[T]]$ is defined to be the greatest common divisor of its coefficients a_i. Call g *primitive* if $\mathrm{cont}(g) = 1$ or, equivalently, if g is not in the kernel of the natural map $\mathbf{Z}[[T]] \to \mathbf{F}_\ell[[T]]$ for any prime number ℓ. Since the rings $\mathbf{F}_\ell[[T]]$ are domains, the product of any two primitive power series is primitive. One deduces

$$\mathrm{cont}(g_1 g_2) = \mathrm{cont}(g_1)\mathrm{cont}(g_2),$$

which is Gauss's lemma for power series.

Lenstra's proof of Lemma 2.5 then proceeds as follow. Write $f = g_1/g_2$, where $g_i \in \mathbf{Z}[T]$ and $(g_1, g_2) = 1$ in $\mathbf{Q}[T]$. It is clear that $\mathrm{cont}(f) = 1$. The relation $g_1 = g_2 f$ implies that $\mathrm{cont}(g_1) = \mathrm{cont}(g_2)$. Cancelling this common factor, we may assume $\mathrm{cont}(g_1) = \mathrm{cont}(g_2) = 1$. Since g_1 and g_2 are relatively prime over $\mathbf{Q}$, there is a positive integer n such that

$$n \in g_1\mathbf{Z}[T] + g_2\mathbf{Z}[T] \subset g_2\mathbf{Z}[[T]],$$

the last inclusion because $g_1 = g_2 f$. Write $n = hg_2$ with $h \in \mathbf{Z}[[T]]$. Then one has

$$n = \mathrm{cont}(n) = \mathrm{cont}(h)\mathrm{cont}(g_2) = \mathrm{cont}(h).$$

Hence h is divisible by n. We conclude that $g_2(0) = \pm 1$. This implies 2.5.

In order to actually compute the zeta function, it is useful to know an upper bound for the total degree $\deg R_1 + \deg R_2$ of the zeta function. The following explicit bound was proved by Bombieri [5].

Theorem 2.6 *The total degree of $Z(X)$ satisfies*

$$\deg R_1 + \deg R_2 < (4d + 9)^{n+m},$$

where

$$d = \max_{1\le j\le m} \deg(f_j).$$

Bombieri's bound is a general purpose bound. Its proof depends on Dwork's p-adic method. It can be improved in various ways, especially when one takes the Newton polytope of the defining polynomials f_i into account, as was done by Adolphson and Sperber [2]. The bound is reasonably good as a function of d, but the dependence on m can probably be significantly improved.

An easy consequence is the following result.

Corollary 2.7 *The zeta function $Z(X)$ is effectively computable.*

This corollary is obvious in the special case that the zeta function $Z(X)$ is known to be a polynomial, or known to be the reciprocal of a polynomial, up to some trivial known factors. In the general case, one can deduce Corollary 2.7 from Theorem 2.4 and Theorem 2.6 in several ways, for instance by using the results of Berlekamp-Massey on linear recurring sequences (see [3] for more detail). Here we use a simple linear algebra argument explained to me by Hendrik Lenstra.

Let D_1 and D_2 be upper bounds for the degree of the numerator and the denominator of $Z(X)$, respectively. For instance, we can take $D_i = D = (4d + 9)^{n+m}$, by Theorem 2.6. Compute the first $D_1 + D_2 + 1$ terms of the power series

$$Z(X) = 1 + z_1T + z_2T^2 + \ldots + z_{D_1+D_2}T^{D_1+D_2} + \ldots$$

by explicitly counting $X(\mathbf{F}_{q^k})$ for $k \le D_1 + D_2$. Write a_i, b_i for the coefficients of R_1 and R_2 to be determined:

$$R_1(X) = 1 + a_1T + \ldots + a_{D_1}T^{D_1},$$
$$R_2(X) = 1 + b_1T + \ldots + b_{D_2}T^{D_2}.$$

The congruence

$$R_2(X)Z(X) \equiv R_1(X) \pmod{T^{D_1+D_2+1}}$$

gives a system of linear equations in the a_i's and the b_i's. This system has at least one rational solution, and using linear algebra we can find one. Denote it by

$$(a'_1, \ldots, a'_{D_1}; b'_1, \ldots, b'_{D_2}).$$

Let

$$R'_1(X) = 1 + a'_1 T + \ldots + a'_{D_1} T^{D_1},$$
$$R'_2(X) = 1 + b'_1 T + \ldots + b'_{D_2} T^{D_2}.$$

Then, the congruence

$$R'_2(X) Z(X) \equiv R'_1(X) \pmod{T^{D_1+D_2+1}}$$

holds as well, and therefore

$$R_2(X) R'_1(X) \equiv R_2(X) R'_2(X) Z(X) \equiv R'_2(X) R_1(X) \pmod{T^{D_1+D_2+1}}.$$

Since each of $R_2 R'_1$ and $R'_2 R_1$ has degree at most $D_1 + D_2$, we deduce $R_2 R'_1 = R'_2 R_1$, so

$$\frac{R'_1(X)}{R'_2(X)} = \frac{R_1(X)}{R_2(X)} = Z(X).$$

Removing the greatest common factor of $R'_1(X)$ and $R'_2(X)$, one obtains the reduced form of $Z(X)$. The proof is complete.

The above effective algorithm immediately implies the following.

Corollary 2.8 *For fixed n, m, d, q, the number $\#X(\mathbf{F}_{q^k})$ can be computed in time bounded by a polynomial in k.*

We now estimate the output size of any algorithm computing $Z(X)$. Trivially, one has $\#X(\mathbf{F}_{q^k}) \leq \#\mathbf{A}^n(\mathbf{F}_{q^k}) = q^{nk}$. Hence $T \cdot d \log Z(X,T)/dT$ converges as a complex power series for $|T| < q^{-n}$, so its reciprocal poles β_i and γ_j are bounded by q^n in absolute value. Let D denote the total degree of $Z(X)$. Then $R_1(X) = \prod_i (1 - \beta_i T)$ and $R_2(X) = \prod_j (1 - \gamma_j T)$ have altogether $O(D)$ coefficients, each of which is $O(2^D q^{nD})$. Then, regardless of the input size for X, the output size of the algorithm is $O(nD^2 \log q)$. There is no general formula for the total degree D. However, for fixed m, Bombieri's degree bound $(4d+9)^{n+m}$ is reasonably good, and it is comparable to the dense input size $O(m(d+1)^n \log q)$.

We shall be concerned with the case that m is fixed (or small). If m is large, then the problem of computing $\#X(\mathbf{F}_q)$ is of a totally different, more combinatorial nature. The fundamental question that we consider is the following.

Problem 2.9 *Given X with fixed m, compute the zeta function $Z(X)$ in time bounded by a polynomial in $(d+1)^n \log q$.*

Remarks. If X has a sizable automorphism group, then one can often speed up the computation of $Z(X)$ by using a suitable equivariant theory. Examples include diagonal hypersurfaces, certain modular varieties, and certain Calabi-Yau hypersurfaces. In this paper, we do not assume that X is given with any additional structure of this sort.

Currently, the theory of zeta functions over finite fields comprises only two types of methods that are powerful enough to prove the rationality of the zeta function in the general case. These are the ℓ-adic method and the p-adic method, where ℓ denotes a prime number different from the characteristic p of the finite field $\mathbf{F}_q$. It is thus natural to try and exploit these general methods for algorithmic purposes.

In the ℓ-adic method one attempts to compute $Z(X) \bmod \ell^k$ using a suitable ℓ-adic trace formula. The zeta function $Z(X)$ can be recovered from its reduction modulo a single large prime power ℓ^k, or from its reductions modulo many small primes ℓ via the Chinese remainder theorem. Unfortunately, the available ℓ-adic trace formula is in the general case not yet effective. Thus the use of the ℓ-adic method is currently restricted to special varieties, such as curves and abelian varieties. In the cases in which it can be used, the ℓ-adic method usually results in a polynomial time algorithm if d, m, and n are fixed, see Schoof [33], Elkies [13], and Poonen [31] for the first examples. It is an important open problem to make the ℓ-adic method effective in the general case.

In the p-adic method one attempts to compute $Z(X) \bmod p^k$ using a p-adic trace formula, where p^k is chosen so large that one can recover the zeta function $Z(X)$ from its reduction modulo p^k. There are many p-adic trace formulas. All of them can be made effective, although not all of them result in efficient algorithms. The general feeling is that the p-adic method is quite efficient if the characteristic p is suitably small, regardless of the size of the field $\mathbf{F}_q$ of definition and regardless of the degree of X. In the companion paper [22] jointly written with Alan Lauder, we present a p-adic algorithm that proves the following theorem.

Theorem 2.10 *There is an algorithm that, given X, computes the zeta function $Z(X)$ in time bounded by a polynomial in $2^m d^{n^2} p^n (\log q)^n$.*

For fixed m and n, and small p—say, $p = O((d \log q)^c)$ for some positive constant c—the algorithm of Theorem 2.10 runs in polynomial time, and it thus provides a partial solution to Problem 2.9. In Section 7 below, we illustrate some of the basic ideas of the p-adic method by treating an algorithm for the easier problem of computing $Z(X) \bmod p$. For more details the reader is referred to [22].

3 Reduction to hypersurfaces

For the computation of the zeta function, the general case of an affine algebraic set X defined by a system

$$f_1(x) = \ldots = f_m(x) = 0$$

of m polynomial equations in n variables can be reduced to the case of an affine hypersurface defined by one single equation. In the present section we discuss various ways in which this reduction can be accomplished; not all of them are very efficient.

The quickest theoretical reduction depends on the observation that any algebraic set is birational to an affine hypersurface. One then continues by induction on the dimension. As it stands, this method is not very explicit. It may be of interest to make it both explicit and efficient.

A second method, which is explicit, exploits the inclusion-exclusion principle, as follows. For a subset $I \subset \{1, 2, \ldots, m\}$, let $H(I)$ be the affine hypersurface defined by

$$\prod_{i \in I} f_i = 0,$$

and let $H(I)^c$ be the complement of $H(I)$ in $\mathbf{A}^n$. Thus,

$$H(I) = \bigcup_{i \in I} \{f_i = 0\},$$
$$H(I)^c = \bigcap_{i \in I} \{f_i \neq 0\}.$$

In particular, we have $H(\emptyset) = \emptyset$ and $H(\emptyset)^c = \mathbf{A}^n$. The inclusion-exclusion principle implies that

$$Z(X) = \prod_{I \subset \{1,2,\ldots,m\}} Z(H(I)^c, T)^{(-1)^{\#I}}.$$

Since

$$Z(H(I)^c, T) = \frac{1}{(1 - q^n T) Z(H(I), T)},$$

we conclude that for $m > 0$ we have

$$Z(X) = \prod_{\substack{I \subset \{1,2,\ldots,m\} \\ I \neq \emptyset}} Z(H(I), T)^{(-1)^{\#I-1}},$$

where the factor 1 corresponding to $I = \emptyset$ has been dropped. Each factor of the above product is now the zeta function of an affine hypersurface. Note that this reduction uses $2^m - 1$ hypersurfaces, but they are all in the original affine space $\mathbf{A}^n$.

If we apply Theorem 2.6 to each factor in the above identity, then we find that the total degree of $Z(X)$ is bounded by

$$\sum_{k=1}^{m} \binom{m}{k} (4kd+9)^{n+1} < 2^m (4md+9)^{n+1}.$$

In some cases where m is large this is better than what one obtains by applying Theorem 2.6 directly.

If one is willing to work with the slightly more general situation of L-functions of exponential sums, then one can use the single polynomial

$$g(x,y) = y_1 f_1(x) + \ldots + y_m f_m(x)$$

in $n+m$ variables. Let ζ_p denote a fixed primitive pth root of unity in an extension field of $\mathbf{Q}$. For each positive integer k, define the exponential sum

$$S_k(g) = \sum_{x_i, y_j \in \mathbf{F}_{q^k}} \zeta_p^{\mathrm{Tr}_k(g(x,y))},$$

where Tr_k denotes the absolute trace from $\mathbf{F}_{q^k}$ to the prime field $\mathbf{F}_p$. The L-function associated to g is defined to be

$$L(g,T) = \exp\Big(\sum_{k=1}^{\infty} \frac{S_k(g)}{k} T^k\Big).$$

It is straightforward to check that

$$\#X(\mathbf{F}_{q^k}) = \frac{1}{q^{mk}} S_k(g).$$

This gives the desired reduction

$$Z(X) = L(g, \frac{1}{q^m} T).$$

Replacing f_i by af_i, one also deduces that

$$Z(X) = L(ag, \frac{1}{q^m}T)$$

for each non-zero $a \in \mathbf{F}_q$.

One can avoid the L-function in the above reduction by using the zeta function of the following Artin-Schreier hypersurface in $\mathbf{A}^{m+n+1}$:

$$Y\colon\ z^p - z = y_1 f_1(x) + \ldots + y_m f_m(x).$$

In fact, a direct calculation gives that

$$\#Y(\mathbf{F}_{q^k}) = \sum_{a\in\mathbf{F}_p} S_k(ag) = q^{(m+n)k} + \sum_{a\in\mathbf{F}_p^*} S_k(ag).$$

It follows that

$$Z(Y, q^{-m}T) = \frac{1}{1-q^nT} \cdot \prod_{a\in\mathbf{F}_p^*} L(ag, q^{-m}T) = \frac{1}{1-q^nT} Z(X)^{p-1}.$$

We obtain the formula

$$Z(X) = ((1-q^nT)\cdot Z(Y, q^{-m}T))^{1/(p-1)}.$$

For large p this reduction is not likely to be very efficient.

One can also use the affine hypersurface $H \subset \mathbf{A}^{m+n}$ defined by

$$H\colon\ g(x,y) = y_1 f_1(x) + \ldots + y_m f_m(x) = 0.$$

As S. Gao observed, one has

$$\#H(\mathbf{F}_q) = q^{m+n-1} + \#X(\mathbf{F}_q)(q^m - q^{m-1}).$$

This shows that for the purpose of counting rational points, we can work with a single hypersurface in the affine space $\mathbf{A}^{m+n}$. In terms of zeta functions, Gao's formula says that

$$Z(H,T) = \frac{Z(X, q^mT)}{(1-q^{m+n-1}T)Z(X, q^{m-1}T)}.$$

One can inductively solve for $Z(X)$ in terms of $Z(H,T)$. Doing this from the complex point of view, one gets the infinite complex product

$$Z(X) = \prod_{k=0}^{\infty}(1-q^{n-1-k}T)\cdot \prod_{k=0}^{\infty} Z(H, q^{-m-k}T).$$

Doing it from the p-adic point of view, one gets the infinite p-adic product

$$\frac{1}{Z(X)} = \prod_{k=0}^{\infty}(1 - q^{n+k}T) \cdot \prod_{k=0}^{\infty} Z(H, q^{1-m+k}T).$$

There is a standard manner, as in Hilbert's tenth problem [25], to define X by means of a system of equations of degree two. For example, if one of the equations defining X has a term $ax_1x_3^3x_4^2$, with $a \in \mathbf{F}_q$, then one can introduce new variables $x_{1,3}$, $x_{3,3}$, $x_{4,4}$, $x_{1,3,3,3}$, as well as new equations

$$x_{1,3} = x_1x_3, \qquad x_{3,3} = x_3^2, \qquad x_{4,4} = x_4^2, \qquad x_{1,3,3,3} = x_{1,3}x_{3,3},$$

and replace the term $ax_1x_3^3x_4^2$ by $ax_{1,3,3,3}x_{4,4}$; and one can proceed similarly with other terms. If one next applies Gao's reduction, then one obtains a hypersurface H that is defined by a cubic polynomial.

An amusing application of Gao's formula is the reduction of the Hasse-Weil meromorphy conjecture to the case of a cubic hypersurface. Let the polynomials f_i have integer coefficients, and let the affine algebraic sets X and H be defined as above, but now over $\mathbf{Z}$. Let

$$\zeta(X, z) = \prod_{p \text{ prime}} Z(X \otimes \mathbf{F}_p, p^{-z})$$

be the global complex Hasse-Weil zeta function of X; it is defined for complex numbers z whose real part is sufficiently large. The Hasse-Weil conjecture asserts that $\zeta(X)$ can be extended to a meromorphic function on all of $\mathbf{C}$. Gao's formula implies that

$$\zeta(H, z) = \frac{\zeta(X, z - m)\zeta(z + 1 - m - n)}{\zeta(X, z - m + 1)},$$

where $\zeta(z)$ is the Riemann zeta function. From this relation, one deduces by induction or iteration that if the Hasse-Weil conjecture is valid for all cubic hypersurfaces H, then it is valid for all algebraic sets X. (By contrast, for Hilbert's tenth problem on diophantine equations over the integers, one only obtains a reduction to quartic hypersurfaces.) If we write

$$\zeta(X, z) = \sum_{k=1}^{\infty} \frac{a_k(X)}{k^z}$$

as a Dirichlet series, then Corollary 2.7 shows that each coefficient $a_k(X)$ is effectively computable.

4 Hypersurface examples

In this section, we focus on the crucial case of hypersurfaces. We discuss them by increasing dimension.

Example 1. Let X be the zero-dimensional hypersurface defined by $f(x) = 0$, where $f(x)$ is a non-constant monic polynomial over $\mathbf{F}_q$ in one variable. Write

$$f(x) = P_1(x)^{k_1} \cdots P_e(x)^{k_e},$$

where the $P_i(x)$ are pairwise distinct monic irreducible polynomials in $\mathbf{F}_q[x]$ and the k_i are positive integers. Then the Euler product reads

$$Z(X) = \prod_{j=1}^{e} \frac{1}{1 - T^{\deg P_j(x)}}.$$

It is not hard to show that one can compute $Z(X)$ in polynomial time using the Frobenius map. Note that if one factors the polynomial $f(x)$ first, one does not get a polynomial time algorithm for computing $Z(X)$. This is because there is currently no known (deterministic) polynomial time algorithm for factoring univariate polynomials over $\mathbf{F}_q$ if p is large, see Wan [35] for a further discussion and for the close relation of $Z(X)$ to various algorithms for factoring univariate polynomials over finite fields. This example also occurred as an exercise ascribed to Lenstra in Cohen's book ([6], Chap. 6, Exerc. 8).

Example 2. Let $f(x_1, x_2) \in \mathbf{F}_q[x_1, x_2]$ be of degree d, and suppose that the homogenization of $f(x_1, x_2)$ defines a smooth projective plane curve C_d over $\mathbf{F}_q$. The genus g of the curve C_d is well known to be

$$g = \frac{(d-1)(d-2)}{2}.$$

From the Riemann-Roch theorem one can deduce (see Monsky [27])

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

where $P(C_d, T)$ is a polynomial of degree $2g$. Weil proved that we further have

$$P(C_d, T) = \prod_{j=1}^{2g} (1 - \alpha_j T), \quad |\alpha_j| = \sqrt{q}, \quad \alpha_j \alpha_{2g+1-j} = q.$$

The Riemann-Roch theorem also shows that the special value $P(C_d, 1)$ has the following arithmetic meaning:

$$P(C_d, 1) = \#J(C_d)(\mathbf{F}_q) \in \mathbf{Z}_{>0},$$

where $J(C_d)$ is the Jacobian variety of C_d, which is a g-dimensional abelian variety over $\mathbf{F}_q$. All these results hold for any smooth projective geometrically irreducible curve over $\mathbf{F}_q$, not just for plane curves.

The ℓ-adic method can be made effective in the case of curves and abelian varieties. One can then use the Chinese remainder theorem as mentioned before. In this way, one obtains an algorithm for computing $Z(C_d, T)$ with running time $O((\log q)^{\Delta_d})$, where Δ_d is in general an exponential function of d, see Schoof [33] and Pila [29]. Thus, for fixed d, the ℓ-adic method computes the zeta function $Z(C_d, T)$ in polynomial time, although the algorithm is still doubly exponential in d. For hyperelliptic curves, the exponent Δ_d has been improved to a polynomial in d, see Adleman-Huang [1]. On the other hand, using the p-adic algorithm in [22], one can compute the zeta function $Z(C_d, T)$ in time $(dp \log q)^{O(1)}$, which is polynomial in d but exponential in $\log p$. In particular, the zeta function $Z(C_d, T)$ can be computed in polynomial time if $p = O((d \log q)^{O(1)})$. This example is important because of its many applications in number theory and cryptography, see Koblitz [19] and Blake et al. [4]. For special types of curves in small characteristic, more practical versions of various p-adic algorithms have been designed by a number of authors, see Satoh [32], Kedlaya [17], Lauder-Wan [23], Denef-Vercauteren [10], and the references listed in those papers. Restricting Problem 2.9 to plane curves, we obtain the following.

Problem 4.1 *Given a smooth projective plane curve C_d over $\mathbf{F}_q$, compute $Z(C_d, T)$ in time $O((d \log q)^c)$, where c is an explicit absolute positive constant.*

Example 3. Let $f(x_1, \ldots, x_n) \in \mathbf{F}_q[x_1, \ldots, x_n]$ be of degree d, and suppose that the homogenization of f defines a smooth projective hypersurface H_d of dimension $n-1$ over $\mathbf{F}_q$. Then by the Weil conjectures [7] we can write

$$Z(H_d, T) = \frac{P(H_d, T)^{(-1)^n}}{\prod_{j=0}^{n-1}(1 - q^j T)},$$

where $P(H_d, T) \in 1 + T\mathbf{Z}[T]$ is a polynomial of degree (see Monsky [27])

$$D = \frac{1}{d}\{(d-1)^{n+1} + (-1)^{n+1}(d-1)\},$$

and

$$P(H_d, T) = \prod_{j=1}^{D}(1 - \alpha_j T), \quad |\alpha_j| = \sqrt{q}^{n-1}, \quad \alpha_j \alpha_{D+1-j} = q^{n-1}.$$

This higher dimensional example is undoubtedly more difficult than the previous example of curves, and it has attracted a smaller amount of attention. Both theoretically and computationally, and from an applied point of view, our understanding of this case leaves a great deal to be desired. In particular, the ℓ-adic method has not been made effective for higher dimensional smooth projective hypersurfaces. We do know that by the p-adic algorithm in [22], the zeta function $Z(H_d, T)$ can be computed in time $(d^n p \log q)^{O(n)}$ for any $(n-1)$-dimensional hypersurface H_d, not necessarily smooth or projective. This gives a polynomial time algorithm for small p and fixed n. In the smooth projective case, it should be possible to improve the exponent $O(n)$ by finer p-adic cohomological methods. In fact, Lauder [20][21] has gone further and used the deformation method on the cohomology space to show that the zeta function can be computed in time $(d^n p \log q)^{O(1)}$ for suitable smooth projective $(n-1)$-dimensional hypersurface H_d. It would be interesting to explore possible applications of higher dimensional hypersurfaces. As a special case of Problem 2.9, we have the following.

Problem 4.2 *Given a smooth projective $(n-1)$-dimensional hypersurface H_d defined over $\mathbf{F}_q$, compute $Z(H_d, T)$ in time bounded by a polynomial in $(d+1)^n \log q$.*

Lauder's recent results solve this problem if p is small.

5 Pure weight decomposition

In this section, we consider the problem of computing the numbers of zeros and poles of $Z(X)$ with a given complex absolute value, for any algebraic set X defined over $\mathbf{F}_q$. Let $R(X, T)$ be the numerator or the denominator of the zeta function $Z(X)$.

Over the complex numbers $\mathbf{C}$, we can write

$$R(X, T) = \prod_i (1 - \alpha_i T),$$

where each α_i is a non-zero algebraic integer. Define the *weight* of a non-zero complex number α by

$$w(\alpha) = \log_q(\alpha\bar{\alpha}),$$

where $\bar{\alpha}$ denotes the complex conjugate of α. An elementary archimedean estimate shows that

$$\#X(\mathbf{F}_{q^k}) = O(q^{k \dim(X)}),$$

where the implied constant depends on the degree of X. As in Section 2, one deduces the elementary estimate

$$|\alpha_i| \le q^{\dim(X)}, \quad w(\alpha_i) \le 2\dim(X).$$

The following much deeper result is a consequence of Deligne's main theorem on the Weil conjectures [8].

Theorem 5.1 *The weights of the reciprocal roots α_i of $R(X,T)$ are integers in the interval $[0, 2\dim(X)]$. That is, for each α_i, there is an integer w_i with $0 \le w_i \le 2\dim(X)$ such that*

$$\alpha_i \bar{\alpha}_i = q^{w_i}.$$

In particular, each α_i is an ℓ-adic unit for all prime numbers $\ell \neq p$. Furthermore, each α_i and its Galois conjugates have the same weight.

The last part of the theorem can be deduced from the first part in an elementary manner, as shown in the following result of Lenstra.

Lemma 5.2 *Let $f \in \mathbf{Q}[T]$ be an irreducible polynomial. Suppose α and β are two complex roots of f with $\alpha\bar{\alpha} \in \mathbf{Q}$ and $\beta\bar{\beta} \in \mathbf{Q}$. Then one has $\alpha\bar{\alpha} = \beta\bar{\beta}$.*

Proof. Let $\alpha\bar{\alpha} = a$, $\beta\bar{\beta} = b$, and $\deg(f) = n > 0$. We need to show that $a = b$. Since there is a field isomorphism $\mathbf{Q}(\alpha) \cong \mathbf{Q}(\beta)$ that sends α to β, we have

$$N_{\mathbf{Q}(\alpha)/\mathbf{Q}}(\alpha) = N_{\mathbf{Q}(\beta)/\mathbf{Q}}(\beta),$$

where N denotes the norm map. From $a \in \mathbf{Q}$ and $\mathbf{Q}(\alpha) = \mathbf{Q}(\bar{\alpha})$ one deduces

$$a^n = N_{\mathbf{Q}(\alpha)/\mathbf{Q}}(a) = N_{\mathbf{Q}(\alpha)/\mathbf{Q}}(\alpha) N_{\mathbf{Q}(\bar{\alpha})/\mathbf{Q}}(\bar{\alpha}) = N_{\mathbf{Q}(\alpha)/\mathbf{Q}}(\alpha)^2.$$

Similarly,

$$b^n = N_{\mathbf{Q}(\beta)/\mathbf{Q}}(\beta)^2.$$

Putting the above together, one deduces that $a^n = b^n$. Since $a \ge 0$ and $b \ge 0$, we conclude that $a = b$. The proof is complete.

For an integer w with $0 \le w \le 2\dim(X)$, let

$$R(w, X, T) = \prod_{w(\alpha_i)=w} (1 - \alpha_i T).$$

This is called the *pure weight w part* of $R(X,T)$. By the above theorem of Deligne, each $R(w,X,T)$ is a polynomial in $1+T\mathbf{Z}[T]$. The *pure weight decomposition* is

$$R(X,T)=\prod_{w=0}^{2\dim(X)} R(w,X,T), \quad R(w,X,T)\in 1+T\mathbf{Z}[T].$$

Theorem 5.3 *Given the zeta function $Z(X)$, one can compute all pure parts $R(w,X,T)$ in polynomial time.*

Proof. By the LLL (Lenstra-Lenstra-Lovász [24]) factorization algorithm, the polynomial $R(X,T)$ can be factored as a product of irreducible polynomials in polynomial time:

$$R(X,T)=\prod_i g_i(T),$$

where each factor

$$g_i\in 1+T\mathbf{Z}[T]$$

is irreducible over $\mathbf{Q}$. Write

$$g_i(T)=1+a_{i1}T+\ldots+a_{ie_i}T^{e_i}=\prod_{j=1}^{e_i}(1-\beta_{ij}T), \quad a_{ie_i}\neq 0.$$

By Deligne's theorem 5.1, each β_{ij} is pure of some integer weight w_i. Hence $a_{ie_i}^2=\prod_j \beta_{ij}\bar{\beta}_{ij}=q^{e_iw_i}$, and therefore

$$a_{ie_i}=\pm\sqrt{q}^{w_ie_i}\in\mathbf{Z}.$$

One can recover the integer weight w_i from

$$w_i=2\frac{\log_q|a_{ie_i}|}{e_i}.$$

The pure weight w part of $Z(X)$ is then

$$R(w,X,T)=\prod_{w_i=w} g_i(T).$$

The theorem is proved.

Clearly, for each irreducible factor $g_i(T)$ of $R(X,T)$, the map $\beta_{ij}\mapsto\bar{\beta}_{ij}=q^{w_i}/\beta_{ij}$ permutes the reciprocal roots of $g_i(T)$. This gives the following functional equation.

Corollary 5.4 *For each w, the pure weight part $R(w,X,T)$ satisfies*

$$R(w,X,1/(q^wT)) = \pm q^{-wd(w,R)/2}T^{-d(w,R)}R(w,X,T),$$

where $d(w,R)$ denotes the degree of $R(w,X,T)$.

The following is also immediate from Theorem 5.3.

Corollary 5.5 *Given the zeta function, the degrees $d(w,R)$ of all pure weight parts $R(w,X,T)$ can be computed in polynomial time.*

The pure degrees $d(w,R)$ contain important geometric information about the variety X. For instance, if R_2 denotes the denominator of $Z(X)$, then $d(2\dim(X),R_2)$ is the number of top-dimensional components of $X \otimes \bar{\mathbf{F}}_q$. If X is a geometrically irreducible curve and R_1 denotes the numerator of $Z(X)$, then $d(1,R_1)$ is twice the genus of the non-singular model of X. If X is smooth and projective, the pure degrees $d(w,R)$ are geometric invariants, that is, they depend only on $X\otimes\bar{\mathbf{F}}_q$. In fact, in this smooth projective case, they are given by the ℓ-adic Betti numbers, as shown in the proof of the following result.

Corollary 5.6 *For any smooth projective variety X defined over $\mathbf{F}_q$, the ℓ-adic Betti numbers of X can be effectively computed, where $\ell \neq p$.*

Proof. For a smooth projective variety X over $\mathbf{F}_q$, the ℓ-adic trace formula states

$$Z(X) = \prod_{i=0}^{2\dim(X)} \det(I - (\mathrm{Frob}|H^i(X\otimes\bar{\mathbf{F}}_q,\mathbf{Q}_\ell))T)^{(-1)^{i-1}},$$

where Frob denotes the geometric Frobenius map. By the Weil conjectures as proved in [7], each (complex) eigenvalue of Frob acting on H^i has weight equal to i. Thus, the ith Betti number

$$B_i(X,\ell) = \dim_{\mathbf{Q}_\ell} H^i(X\otimes\bar{\mathbf{F}}_q,\mathbf{Q}_\ell)$$

is given by the formula

$$B_i(X,\ell) = \begin{cases} d(i/2,R_1), & \text{if } i \text{ is odd,} \\ d(i/2,R_2), & \text{if } i \text{ is even,} \end{cases}$$

where R_1 and R_2 are the numerator and denominator of $Z(X)$, respectively. The corollary follows.

Let Y be a smooth projective scheme over $\mathbf{Z}$ and let $X = Y \otimes \mathbf{F}_p$ be the reduction modulo p of Y. The cohomological comparison theorem shows that for all large primes p, the pure degrees $d(w, R(X,T))$ depend only on the geometry of $Y \otimes \bar{\mathbf{Q}}$, not on the chosen large prime p.

If X is singular or open, the pure degrees $d(w, R)$ may not be geometric invariants, as it may happen that a root and a pole of $Z(X)$ have a quotient that is a root of unity. On the other hand, it is clear that the pure difference degrees $d(w, R_1) - d(w, R_2)$ are geometric invariants, where we recall that R_1 and R_2 denote the numerator and the denominator of $Z(X)$, respectively.

It would be interesting to know if the pure degrees $d(w, R)$ can be computed in polynomial time without the zeta function being given. Even for a singular plane curve this seems to be unknown. In the case of a smooth projective complete intersection there is a well-known formula showing that the pure degrees can be computed from the total degree of $Z(X)$ and the dimension of X; see Example 3 of Section 4 for the case of a smooth projective hypersurface and [9] (page 39–61) for the general smooth projective complete intersection case.

For arbitrary X over $\mathbf{F}_q$, not necessarily smooth or projective, Grothendieck [14] has shown that for $\ell \neq p$ there is a similar ℓ-adic formula in terms of ℓ-adic cohomology with compact support:

$$Z(X) = \prod_{i=0}^{2\dim(X)} \det(I - (\mathrm{Frob}|H_c^i(X \otimes \bar{\mathbf{F}}_q, \mathbf{Q}_\ell))T)^{(-1)^{i-1}}.$$

But in this generality, even the conjectured independence of the ℓ-adic Betti numbers on ℓ is unknown. The following result of Katz [16] provides weak evidence in this direction. It gives an explicit upper bound for the ℓ-adic Betti numbers with compact support that is independent of ℓ. It is obtained by means of an inductive reduction to the Bombieri-Adolphson-Sperber degree bound for $Z(X)$, which in turn is p-adic in nature.

Theorem 5.7 *Let the polynomials $f_1, \ldots, f_m$ form a system of defining equations for the affine algebraic set X, and put $d = \max_i \deg(f_i)$. Then, for every prime number $\ell \neq p$, we have*

$$\sum_{i \geq 0} \dim_{\mathbf{Q}_\ell} H_c^i(X \otimes \bar{\mathbf{F}}_q, \mathbf{Q}_\ell) \leq 2^{m+2}(md+3)^{n+1}.$$

Turning to global zeta functions, it may be of interest to point out that the real parts of the zeros of the classical Riemann zeta function $\zeta(z)$ are not known to be effectively computable in a finite region. To make this precise, denote, for real numbers $w \in (0,1)$ and $t > 0$, by $d(w, \zeta; t)$ the number of zeros of $\zeta(z)$ that lie on the line segment

$$\mathrm{Re}(z) = w, \quad 0 \leq \mathrm{Im}(z) \leq t.$$

For any given t, there are only finitely many w for which $d(w, \zeta; t) > 0$. The following computational problem is now analogous to finding the pure weight decomposition: given t, determine the finitely many positive integers among the numbers $d(w, \zeta; t)$, for $0 \leq w \leq 1$, as well as approximations to the corresponding numbers w. No effective algorithm for doing this is currently available. The main difficulty is caused by the possibility of a multiple zero or zeros that are very close to each other, see Odlyzko [28]. Of course, the Riemann hypothesis says that $d(w, \zeta; t) = 0$ for all t and all $w \neq 1/2$.

6 Pure slope decomposition

In the previous section, we considered the problem of purity decomposition of $Z(X)$ from the complex point of view. One can also consider the purity decomposition from a non-archimedean point of view. If ℓ is a prime number different from p, then Deligne's Theorem 5.1 shows that $Z(X)$ is already pure from an ℓ-adic point of view. Thus, we will consider the remaining p-adic case in this section. Let $R(X, T)$ again denote either the numerator or the denominator of $Z(X)$.

Let $\mathbf{C}_p$ be the completion of a fixed algebraic closure of $\mathbf{Q}_p$. Over $\mathbf{C}_p$, we can write

$$R(X, T) = \prod_i (1 - \alpha_i T),$$

where each α_i is a non-zero algebraic integer in $\mathbf{C}_p$. Define the slope of a non-zero element $\alpha \in \mathbf{C}_p$ by

$$s(\alpha) = \mathrm{ord}_q(\alpha) = -\log_q |\alpha|_p,$$

where $|\cdot|_p$ denotes the p-adic absolute value, normalized such that $|p| = 1/p$. The slopes $s(\alpha_i)$ are non-negative rational numbers since the α_i are algebraic integers. One immediately derives the bound

$$0 \leq s(\alpha_i) \leq s(\alpha_i \bar{\alpha}_i) = w(\alpha_i) \leq 2\dim(X).$$

This bound can be improved somewhat. Deligne's integrality theorem [9] (page 384–400) states that $q^{\dim(X)}/\alpha_i$ is an algebraic integer. We deduce the following result.

Theorem 6.1 *The slopes $s(\alpha_i)$ are rational numbers in* $[0, \dim(X)]$.

Other than in the complex absolute value case, α_i and its Galois conjugates over $\mathbf{Q}$ may have different slopes. Of course, each α_i and its Galois conjugates over $\mathbf{Q}_p$ do have the same slope. The slopes $s(\alpha_i)$ are not integers (or half integers) in general. They are merely rational numbers. It would be interesting to get good bounds for the denominators of the slopes $s(\alpha_i)$.

For a rational number s with $0 \leq s \leq \dim(X)$, let

$$R(s, X, T) = \prod_{s(\alpha_i)=s} (1 - \alpha_i T).$$

This is called the pure slope s part of $R(X, T)$. We have the p-adic purity decomposition

$$R(X, T) = \prod_{s \in \mathbf{Q}} R(s, X, T).$$

The question is then to understand each pure slope part $R(s, X, T)$.

Note that in general, the pure slope parts $R(s, X, T)$ do not have coefficients in $\mathbf{Z}$ any more. The theory of Newton polygons implies that $R(s, X, T)$ is a polynomial in $1 + T(\bar{\mathbf{Q}} \cap \mathbf{Z}_p)[T]$, as we shall see now.

Write

$$R(X, T) = 1 + a_1 T + \ldots + a_e T^e = \prod_{i=1}^{e} (1 - \alpha_i T), \quad a_e \neq 0.$$

Definition 6.2 *The Newton polygon* $\mathrm{NP}(R)$ *of* $R(X, T)$ *is the lower convex hull in the plane* $\mathbf{R}^2$ *of the points*

$$(k, \mathrm{ord}_q(a_k)), \quad k = 0, 1, \ldots, e.$$

A basic property of the Newton polygon is the following result, see Koblitz [18].

Theorem 6.3 *The polynomial* $R(X, T) \in 1 + T\mathbf{Z}[T]$ *has exactly h reciprocal zeros α_i with slope* $\mathrm{ord}_q(\alpha_i) = s$ *(counting multiplicities) if and only if* $\mathrm{NP}(R)$ *has a side of slope s and horizontal length h. Furthermore, the coefficients of $R(s, X, T)$ are in $\mathbf{Z}_p$ for each s.*

The following is an immediate consequence.

Corollary 6.4 *Let $d(s, R)$ denote the degree of $R(s, X, T)$. Then, $d(s, R)$ is the horizontal length of the slope s side of* $\mathrm{NP}(R)$. *In particular, when $Z(X)$ is given, the p-adic pure degrees $d(s, R)$ can be computed in polynomial time.*

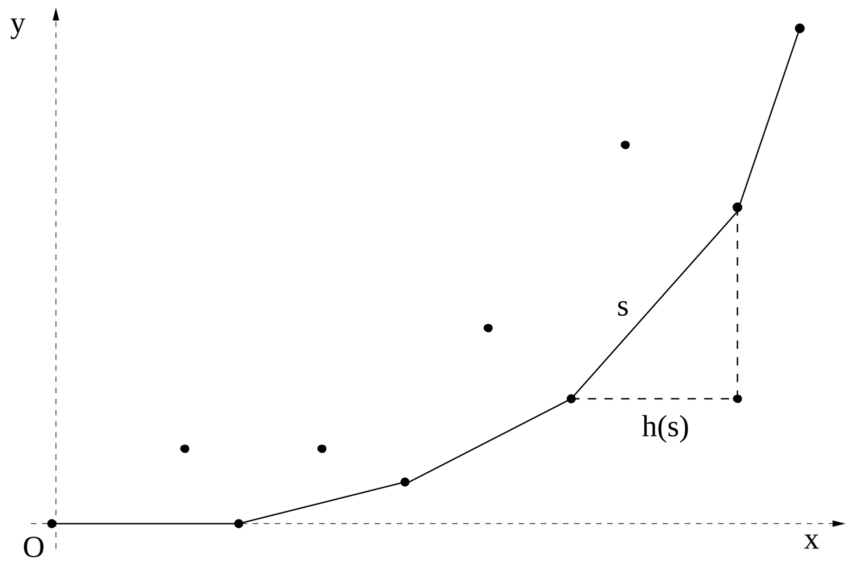

Figure 1: Newton polygon

It would be interesting to know if the p-adic pure degrees $d(s, R)$ can be computed in polynomial time without the zeta function being given. This amounts to computing the Newton polygons of the numerator and the denominator of the zeta function. Even in a very well-behaved situation such as the smooth projective hypersurface case, we do not have a complete answer in general; see Wan [36] for a theoretical introduction to Newton polygons for zeta functions and L-functions. In the smooth projective case, the p-adic pure degrees $d(s, R)$ are geometric invariants.

If X is the good reduction modulo p of some smooth projective scheme over $\mathbf{Z}$, the p-adic pure degrees $d(s, R)$ depend not just on the geometry of the generic fibre $Y \otimes \bar{\mathbf{Q}}$, but also on the chosen prime p. Thus, the p-adic pure degrees $d(s, R)$ contain arithmetic information on Y. They are related to but much deeper than the topological Hodge numbers of $Y \otimes \mathbf{Q}$ defined in terms of the De Rham cohomology of $Y \otimes \mathbf{Q}$, see Mazur [26] for an introductory account. Describing the variation of $d(s, R)$ as p varies is a very subtle arithmetic problem, already in the special case that $Y \otimes \mathbf{Q}$ is an elliptic curve.

As a polynomial with coefficients in $\mathbf{Z}_p$, the pure slope part $R(s, X, T)$ cannot be written down in a finite amount of time. However, given $R(X, T)$, one can compute the pure slope parts $R(s, X, T)$ modulo any given power of p in polynomial time using the Newton polygon and Hensel lifting. Note that here we do not factor $R(X, T)$ into a product of irreducible factors over $\mathbf{Q}_p$, which is a harder problem.

7 Zeta functions modulo p

In [22] a p-adic algorithm for computing the zeta function $Z(X)$ is given, where p is the characteristic of $\mathbf{F}_q$. In this final section, we describe some of the basic ideas behind that algorithm by showing how $Z(X)$ may be computed modulo p. The method expands the outline given in Wan [35].

Without loss of generality we may restrict to hypersurfaces. Thus, let X be the affine hypersurface defined by a polynomial $f(x_1, \ldots, x_n)$ over $\mathbf{F}_q$ of total degree d in n variables. Let A_d be the $\mathbf{F}_q$-vector space of polynomials in $x_1, \ldots, x_n$ of total degree at most d that are divisible by the product $x_1 \cdots x_n$:

$$A_d = (x_1 \cdots x_n \mathbf{F}_q[x_1, \ldots, x_n])_{\leq d}.$$

It has a row basis $\vec{e}$ consisting of monomials

$$\vec{e} = \{x^u | u = (u_1, \ldots, u_n),\ u_i \geq 1,\ |u| \leq d\},$$

where

$$x^u = x_1^{u_1} \cdots x_n^{u_n}, \quad |u| = u_1 + \ldots + u_n.$$

One computes that

$$\dim_{\mathbf{F}_q} A_d = \binom{d}{n}.$$

Definition 7.1 *Let σ be the pth power Frobenius map acting on $\mathbf{F}_q$:*

$$\sigma(a) = a^p, \quad a \in \mathbf{F}_q.$$

Let ψ_p be the σ^{-1}-linear operator on the $\mathbf{F}_q$-vector space $\mathbf{F}_q[x_1, \ldots, x_n]$ defined by

$$\psi_p\Big(\sum_u a_u x^u\Big) = \sum_u \sigma^{-1}(a_u)\psi_p(x^u),$$

where

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

Define

$$\psi_q = \psi_p^r = \psi_p \circ \cdots \circ \psi_p$$

to be the rth iterate of ψ_p, where $q = p^r$.

Since σ^r is the identity on $\mathbf{F}_q$, the operator ψ_q is actually $\mathbf{F}_q$-linear, although ψ_p is only σ^{-1}-linear. The operator ψ_q is a left inverse of the qth power Frobenius map on $\mathbf{F}_q[x] = \mathbf{F}_q[x_1, \ldots, x_n]$. For any polynomial $g \in \mathbf{F}_q[x]$, multiplication by g is also an $\mathbf{F}_q$-linear map from $\mathbf{F}_q[x]$ to itself.

Lemma 7.2 *The $\mathbf{F}_q$-linear composed operator $\psi_q \circ f^{q-1}$ maps the finite dimensional $\mathbf{F}_q$-subspace A_d of $\mathbf{F}_q[x]$ to itself; here f^{q-1} denotes multiplication by f^{q-1}. Similarly, the σ^{-1}-linear composed operator $\psi_p \circ f^{p-1}$ maps A_d to itself, and we have the relation*

$$\psi_q \circ f^{q-1} = (\psi_p \circ f^{p-1})^r.$$

Proof. Let $h \in A_d$. Then h has degree at most d, so $f^{q-1}h$ has degree at most $d(q-1) + d = dq$. Thus, the degree of $\psi_q(f^{q-1}h)$ is at most d. Furthermore, if h is divisible by $x_1 \cdots x_n$, then $\psi_q(f^{q-1}h)$ is also divisible by $x_1 \cdots x_n$. This proves that $\psi_q(f^{q-1}h) \in A_d$. The proof of the second part of the lemma is the same. To prove the last part, we write

$$q - 1 = (p-1) + p(p-1) + \ldots + p^{r-1}(p-1).$$

This gives

$$f(x)^{q-1} = \prod_{i=0}^{r-1} f^{\sigma^i}(x^{p^i})^{p-1},$$

where $f^{\sigma^i}(x)$ denotes the polynomial obtained by applying σ^i to the coefficients of f. Using the easily checked relation

$$\psi_p \circ g^{\sigma}(x^p) = g(x) \circ \psi_p,$$

one deduces

$$\psi_q \circ f^{q-1} = \psi_p^r \circ \prod_{i=0}^{r-1} f^{\sigma^i}(x^{p^i})^{p-1} = (\psi_p \circ f^{p-1})^r.$$

The lemma is proved.

Theorem 7.3 *Let X be the affine hypersurface defined by a polynomial $f(x_1, \ldots, x_n)$ over $\mathbf{F}_q$ of total degree d in n variables. Then we have the congruence formula*

$$(Z(X)^{(-1)^n} \bmod p) = \det(I - (\psi_q \circ f^{q-1} | A_d)T).$$

Before we prove this result, we first recall some facts on L-functions in characteristic p. Let $g \in \mathbf{F}_q[x] = \mathbf{F}_q[x_1, \ldots, x_n]$. For a geometric point $x \in \mathbf{A}^n(\bar{\mathbf{F}}_q)$, the product $g(x)g(x^q)\cdots g(x^{q^{\deg(x)-1}})$ is an element of $\mathbf{F}_q$ that clearly depends only on the orbit (the closed point) of x under the qth power Frobenius map. We define the L-function of g by

$$L(g,T) = \prod_x \frac{1}{1 - g(x)g(x^q)\cdots g(x^{q^{\deg(x)-1}})T^{\deg(x)}} \in 1 + T\mathbf{F}_q[[T]],$$

where x runs over the set of closed points of the affine space $\mathbf{A}^n$ over $\mathbf{F}_q$.

For a real number c, write A_c for the finite dimensional $\mathbf{F}_q$-subspace of $\mathbf{F}_q[x]$ generated by the monomials of total degree at most c that are divisible by the product $x_1 \cdots x_n$. If the total degree of g is at most e, one checks as in the above lemma that the operator $\psi_q \circ g$ maps the subspace $A_{e/(q-1)}$ to itself. Furthermore, the (matrix of the) induced map $\psi_q \circ g$ on the quotient vector space $x_1 \cdots x_n \mathbf{F}_q[x]/A_{e/(q-1)}$ is strictly triangular with respect to the monomial basis $\{x^u | u_i \geq 1, |u| > e/(q-1)\}$. Thus, the composed operator $\psi_q \circ g$ acting on the infinite dimensional $\mathbf{F}_q$-vector space $x_1 \cdots x_n \mathbf{F}_q[x]$ has a well defined characteristic power series $\det(I - (\psi_q \circ g | x_1 \cdots x_n \mathbf{F}_q[x])T)$, which is given by the polynomial

$$\det(I - (\psi_q \circ g | x_1 \cdots x_n \mathbf{F}_q[x])T) = \det(I - (\psi_q \circ g | A_{e/(q-1)})T).$$

The characteristic p version of Dwork's trace formula for the affine space $\mathbf{A}^n$, as given in [34], implies that

$$L(g,T)^{(-1)^{n-1}} = \det(I - (\psi_q \circ g | x_1 \cdots x_n \mathbf{F}_q[x])T)$$

and thus

$$L(g,T)^{(-1)^{n-1}} = \det(I - (\psi_q \circ g | A_{e/(q-1)})T).$$

The reader is referred to [22] to see the Dwork trace formula over the n-torus, which is cleaner than the Dwork trace formula over the affine n-space $\mathbf{A}^n$.

We now return to the proof of the theorem. Taking $g = f^{q-1}$ and $e = d(q-1)$, we deduce that $e/(q-1) = d$ and

$$L(f^{q-1}, T)^{(-1)^{n-1}} = \det(I - (\psi_q \circ f^{q-1} | A_d)T).$$

For a geometric point $x \in \mathbf{A}^n(\bar{\mathbf{F}}_q)$, one has

$$g(x)g(x^q)\cdots g(x^{q^{\deg(x)-1}}) = f(x)^{q^{\deg(x)}-1},$$

which is 0 or 1 according as $x \in X(\bar{\mathbf{F}}_q)$ or not. It follows that $L(f^{q-1}, T)$ is the reduction modulo p of the zeta function of the complement of X in $\mathbf{A}^n$. Hence for $n > 0$ we obtain

$$L(f^{q-1}, T) = (1/Z(X) \bmod p).$$

Substituting this into the above formula for $L(f^{q-1}, T)^{(-1)^{n-1}}$, we obtain the theorem.

Corollary 7.4 *The zeta function $Z(X)$ modulo p can be computed in time bounded by a polynomial in $p\binom{d}{n}\log q$.*

Proof. Recall that $q = p^r$. Let M_1 be the matrix of the σ^{-1}-linear map $\psi_p \circ f^{p-1}$ with respect to the row monomial basis $\vec{e}$ of A_d. That is,

$$(\psi_p \circ f^{p-1})(\vec{e}) = \vec{e}M_1.$$

Then, by the σ^{-1}-linearity of $\psi_p \circ f^{p-1}$, we deduce

$$(\psi_p \circ f^{p-1})^2(\vec{e}) = (\psi_p \circ f^{p-1})(\vec{e}M_1) = \vec{e}M_1 M_1^{\sigma^{-1}},$$

where $M_1^{\sigma^{-1}}$ is obtained from M_1 by applying σ^{-1} to each entry (and similarly with $M_1^{\sigma^{-i}}$ below). By iteration, one finds that the matrix of the $\mathbf{F}_q$-linear map

$$\psi_q \circ f^{q-1} = (\psi_p \circ f^{p-1})^r$$

with respect to the row basis $\vec{e}$ is given by

$$M_r = M_1 M_1^{\sigma^{-1}} \cdots M_1^{\sigma^{-(r-1)}}.$$

The matrix M_1 can be written down in time $(p\binom{d}{n}\log q)^{O(1)}$. It follows that the matrix M_r can also be computed in time $(p\binom{d}{n}\log q)^{O(1)}$. The zeta function modulo p is essentially just the characteristic polynomial of the matrix M_r:

$$(Z(X)^{(-1)^n} \bmod p) = \det(I - M_r T).$$

The corollary is proved. Alternatively, applying σ^{r-1} to the above congruence formula, we have

$$(Z(X)^{(-1)^n} \bmod p) = \det(I - M_1^{\sigma^{r-1}} \cdots M_1^{\sigma} M_1 T).$$

This formula may be slightly more efficient from a computational point of view.

If p is small, that is, $p = (\binom{d}{n}\log q)^{O(1)}$, the above corollary gives a polynomial time algorithm for computing $Z(X)$ modulo p. If p is large, we do not get a polynomial time algorithm.

References

[1] L.M. Adleman and M.D. Huang, Counting rational points on curves and abelian varieties over finite fields, Lecture Notes in Computer Science, 1122 (1996), 1–16.

[2] A. Adolphson and S. Sperber, On the degree of the L-function associated with an exponential sum, Compositio Math., 68 (1988), no. 2, 125–159.

[3] R.E. Blahut, Decoding of cyclic codes and codes on curves, in Handbook of Coding Theory, edited by V.S. Pless and W.C. Huffman, 1998, 1569–1633.

[4] I.F. Blake, G. Seroussi, N.P. Smart, Elliptic Curves in Cryptography, London Mathematical Society Lecture Note Series, 265, Cambridge University Press, Cambridge, 2000.

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

[6] H. Cohen, A Course in Computational Number Theory, Springer-Verlag, Berlin, 1993.

[7] P. Deligne, La Conjecture de Weil, I, Publ. Math., IHES, 43 (1974), 273–307.

[8] P. Deligne, La Conjecture de Weil, II, Publ. Math., IHES, 52 (1980), 137–252.

[9] P. Deligne and N. Katz, Groupes de Monodromie en Géométrie Algébrique, SGA 7 II, Lecture Notes in Math., 340 (1973).

[10] J. Denef and F. Vercauteren, An extension of Kedlaya's algorithm to Artin-Schreier curves in characteristic 2, preprint, 2002.

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

[12] D. Eisenbud, Commutative Algebra, GTM, Springer-Verlag, New York, 1995.

[13] N.D. Elkies, Elliptic and modular curves over finite fields and related computational issues, in Computational perspectives on number theory (Chicago, IL, 1995), 21–76, AMS/IP Stud. Adv. Math., 7, Amer. Math. Soc., Providence, RI, 1998.

[14] A. Grothendieck, Formule de Lefschetz et rationalité des fonctions L, Séminaire Bourbaki, exposé 279, 1964/65.

[15] N. Katz, On a theorem of Ax, Amer. J. Math., 93 (1971), 485–499.

[16] N. Katz, Sums of Betti numbers in arbitrary characteristic, Finite Fields Appl., 7 (2001), 29–44.

[17] K.S. Kedlaya, Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology, J. Ramanujan Math. Soc., 16 (2001), 323–338.

[18] N. Koblitz, p-adic Numbers, p-adic Analysis, and Zeta Functions, GTM, Springer, 1996.

[19] N. Koblitz, Hyperelliptic cryptosystems, J. Cryptology, 1 (1989), 139–150.

[20] A. Lauder, Deformation theory and the computation of zeta functions, Proc. London Math. Soc., to appear.

[21] A. Lauder, Counting solutions to equations in many variables over fintie fields, prperpint.

[22] A. Lauder and D. Wan, Counting points on varieties over finite fields of small characteristic, these proceedings.

[23] A. Lauder and D. Wan, Computing zeta functions of Artin-Schreier curves over finite fields, LMS J. Comput. Math., 5 (2002), 34–55.

[24] A.K. Lenstra, H.W. Lenstra, Jr., L. Lovász, Factoring polynomials with rational coefficients. Math. Ann., 261 (1982), no. 4, 515–534.

[25] Y.V. Matiyasevich, Hilbert's Tenth Problem, The MIT Press, 1996.

[26] B. Mazur, Frobenius and the Hodge filtration, Bull. Amer. Math. Soc., 78 (1972), 653–667.

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

[28] A.M. Odlyzko, Analytic computations in number theory, Mathematics of Computation 1943–1993: a half-century of computational mathematics (Vancouver, BC, 1993), 451–463, Proc. Sympos. Appl. Math., 48, Amer. Math. Soc., Providence, RI, 1994.

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

[30] C. Pomerance, Primality testing: variations on a theme of Lucas, these proceedings.

[31] B. Poonen, Computational aspects of curves of genus at least 2, Lecture Notes in Computer Science, 1122 (1996), 283–306.

[32] 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.

[33] R. Schoof, Elliptic curves over finite fields and the computation of square root modulo p, Math. Comp., 44 (1985), 483–494.

[34] D. Wan, Meromorphic continuation of L-functions of p-adic representations, Ann. Math., 143 (1996), 469–498.

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

[36] D. Wan, Variation of p-adic Newton polygons for L-functions of exponential sums, preprint.

[37] A. Weil, Numbers of solutions of equations in finite fields, Bull. Amer. Math. Soc., 55 (1949), 497–508.