

RESEARCH



Computing zeta functions of algebraic curves using Harvey's trace formula

Madeleine Kyng 

*Correspondence:
madeleine.kyng@unsw.edu.au
School of Mathematics and
Statistics, University of New
South Wales, Sydney, NSW 2052,
Australia

Abstract

We present a new method for computing the zeta function of an algebraic curve over a finite field. The algorithm relies on a trace formula of Harvey to count points on a plane model of the curve. The zeta function of the curve is then obtained by making corrections at singular points. We report on an implementation and provide some examples in MAGMA which demonstrate an improvement over Tuitman's algorithm.

1 Introduction

Let \mathbb{F}_q denote the finite field of characteristic p and cardinality $q = p^a$. Let \tilde{X} be a nonsingular projective curve of genus g over \mathbb{F}_q . Recall that the zeta function of \tilde{X} is defined by

$$Z(\tilde{X}, T) = \exp \left(\sum_{r=1}^{\infty} \frac{|\tilde{X}(\mathbb{F}_{q^r})|}{r} T^r \right).$$

It follows from the Weil conjectures [18, Ch. VIII] that $Z(\tilde{X}, T)$ is a rational function of the form $P(T)/(1-T)(1-qT)$ where $P(T) \in \mathbb{Z}[T]$ is a polynomial of degree $2g$ whose roots all have absolute value $q^{-\frac{1}{2}}$. Furthermore $Z(\tilde{X}, T)$ satisfies the functional equation

$$Z(\tilde{X}, q^{-1}T^{-1}) = q^{1-g} T^{2-2g} Z(\tilde{X}, T). \quad (1.1)$$

It follows from these facts that to compute $Z(\tilde{X}, T)$ it suffices to compute $|\tilde{X}(\mathbb{F}_{q^r})|$ for $r = 1, \dots, g$.

The main purpose of this paper is to describe a practical, efficient algorithm for the problem of computing the zeta function of an arbitrary nonsingular curve.

Theorem 1.1 *There exists an explicit deterministic algorithm with the following properties. The input consists of a prime p , a positive integer a , a monic irreducible polynomial $\bar{b} \in \mathbb{F}_p[t]$ of degree a defining the finite field $\mathbb{F}_q \cong \mathbb{F}_p[t]/\bar{b}$, and an absolutely irreducible polynomial $\bar{F} \in \mathbb{F}_q[x, y]$ of degree $d \geq 2$. The output is $Z(\tilde{X}, T)$ where \tilde{X} is the nonsingular projective curve with function field $\mathbb{F}_q(x)[y]/\langle \bar{F} \rangle$. The algorithm has time complexity*

$$a^{O(1)} d^{O(1)} p^{\frac{1}{2} + o(1)}.$$

We do not provide a complete proof of this theorem. We will give an outline of all of the steps involved in the algorithm, from which one could deduce the time complexity of $a^{O(1)}d^{O(1)}p^{\frac{1}{2}+o(1)}$ by referring to the detailed time complexity estimates provided in [14] and [2]. We have not worked out the exponents in the time complexity stated in Theorem 1.1, but intend to present a thorough analysis in a future paper. In Sect. 4 we present concrete examples that demonstrate the utility and generality of our algorithm.

Schoof in [21] was the first to demonstrate a deterministic polynomial-time algorithm for computing the zeta function of an arbitrary genus $g = 1$ curve E/\mathbb{F}_q . The algorithm involves computing the trace of the Frobenius endomorphism modulo a number of small primes ℓ , followed by using the Chinese remainder theorem to determine the exact value of $(q + 1 - \#E(\mathbb{F}_q))$. Schoof's algorithm, and higher genus variants of it such as [19] and [1], are known today as ℓ -adic algorithms for computing the zeta function. These ℓ -adic algorithms have time complexity polynomial in $\log q$ for fixed g , but unfortunately they are badly exponential in g . Extensions of Schoof's algorithm are frequently used for the case of $g = 1$, and specialised ℓ -adic algorithms for the case of $g = 2$ have been useful in practice [11], but as yet ℓ -adic algorithms have been impractical for the case of an arbitrary curve of genus $g \geq 3$.

Kedlaya in [16] demonstrated an efficient p -adic algorithm for the problem of computing the zeta function of an hyperelliptic curve \tilde{X} over a finite field of odd characteristic. Kedlaya showed how one could apply the machinery of Monsky–Washnitzer cohomology to this problem, computing the zeta function by explicitly computing the action of Frobenius on this p -adic cohomology of \tilde{X} . This method proved to be fruitful for extension and generalisation to larger classes of curves, resulting in the development of p -adic point-counting algorithms for the case of superelliptic curves [10], C_{ab} curves [9] and non-degenerate curves [8]. These Kedlaya-style algorithms have time complexity polynomial in g but exponential in $\log p$, and they are used in practice for curves of genus $g \geq 2$.

The most general algorithm among the descendants of [16] is Tuitman's algorithm [23], which can handle *almost all* inputs $\bar{F} \in \mathbb{F}_q[x, y]$. The main drawback of Tuitman's algorithm is that it requires as input a "good" characteristic zero lift of the polynomial $\bar{F} \in \mathbb{F}_q[x, y]$. The properties this lift must have are rather technical; they are described in [23, Ass. 1]. Such a lift always exists (provided that $p > 2$ and allowing for extension of the base field), but the problem of efficiently computing a good lift for arbitrary $\bar{F} \in \mathbb{F}_q[x, y]$ is difficult. Castryck, Tuitman and Vermeulen have shown in [6, 7, 23] how one can compute a suitable lift for inputs \bar{F} that are nondegenerate with respect to their Newton polytope, or that define a curve of geometric genus at most 5, or that define a curve of arithmetic gonality at most 5.

There are more general p -adic algorithms that can be used to compute the zeta function of an *arbitrary* n -dimensional variety over \mathbb{F}_q and that have polynomial time complexity for fixed p . Lauder and Wan were the first to demonstrate an algorithm having these properties; they achieved this by using ideas originally presented by Dwork in his proof of the rationality of the zeta function. For fixed n , in the case of an n -dimensional hypersurface over \mathbb{F}_q , Lauder and Wan achieved a time complexity polynomial in a , p and the degree d of the defining polynomial [17, Thm. 37]. In [14], Harvey developed an algorithm similar in nature to that of [17] but with asymptotically superior time complexity [14, Thm. 1.2].

Our algorithm is based on the trace formula for counting points on a hypersurface from [14, §3]. We present a modified version of this trace formula in Sect. 2. In [14],

Harvey suggested that one could compute the zeta function of a nonsingular curve by using the trace formula to count points on a plane model, which gives the right result except possibly at the singularities of that model, and afterwards making corrections for these singularities. We make explicit in Sect. 3 how one does this.

Let \bar{F} and \tilde{X} be as in Theorem 1.1, and let X_0 be the affine curve cut out by \bar{F} , with projective closure X . By using Harvey’s trace formula we can count the points on X in any extension $\mathbb{F}_{q^r}/\mathbb{F}_q$. If X happens to be nonsingular, then we actually have that X is isomorphic to \tilde{X} over \mathbb{F}_q , and by counting points on X in extensions of \mathbb{F}_q of degree up to g we succeed in computing $Z(\tilde{X}, T)$.

If X is singular, a naive approach to computing $Z(\tilde{X}, T)$ using the trace formula would be to compute $Z(X, T)$ and then remove extraneous factors from the numerator, i.e., remove factors whose roots have absolute value 1 rather than $q^{-\frac{1}{2}}$. One can compute $Z(X, T)$ by bounding the degree of the numerator using Bombieri’s bound [4, Thm. 1A], and then counting points on X in extensions of \mathbb{F}_q of degree up to this bound.

Our algorithm does better than the naive approach — it avoids costly computations of point-counts on X that are performed when using the naive approach in the case that X is singular. In our algorithm we count only the \mathbb{F}_{q^r} -rational points on X for $r = 1, \dots, g$, and we determine using the Montes algorithm [13] precisely how \tilde{X} differs from X by factoring ideals related to the singular points of X in maximal orders \mathcal{O} of $\mathbb{F}_q(\tilde{X})$. With the point-counts on X for $r = 1, \dots, g$ and the extra information about how \tilde{X} and X differ we determine the values $|\tilde{X}(\mathbb{F}_{q^r})|$ for $r = 1, \dots, g$, and hence compute $Z(\tilde{X}, T)$.

The main advantage our algorithm has over Tuitman’s is that we require no assumptions about the lift F of \bar{F} to characteristic zero. To apply Harvey’s trace formula we must lift $\bar{F} \in \mathbb{F}_q[x, y]$ to some $F \in \mathbb{Z}_q[x, y]$, but *any* lift suffices. We have implemented our algorithm in the $q = p$ case in the computer algebra system MAGMA [5] and made the code publicly available. In Sect. 4 we compare the performance of our implementation against the MAGMA implementation of Tuitman’s algorithm.

2 Harvey’s trace formula

In this section, we shall present a generalisation of the trace formula given in Theorem 3.1 of [14]. This version allows us to take into account the shape of the polynomial defining a curve (or more generally, hypersurface), and therefore results in a more efficient computation of point-counts than if one were to use a straightforward implementation of the formula from [14]. Our version works with the actual Newton polytope of the polynomial, whereas [14, Thm. 3.1] works with a dilation of the standard simplex that contains that Newton polytope. For our point-counting purposes we only need the case of curves, but we shall present the general hypersurface case as it is no harder to state or prove.

For any domain R , we denote the Laurent polynomial ring $R[x_1, x_1^{-1}, \dots, x_n, x_n^{-1}]$ by $R[x^\pm]$, and for $F \in R[x^\pm]$ and $u = (u_1, \dots, u_n) \in \mathbb{Z}^n$, we denote by $[F]_u$ the coefficient of $x^u = x_1^{u_1} \dots x_n^{u_n}$ in F . Throughout, we shall use K to denote a convex polytope in \mathbb{R}^n with integral vertices. We will denote by $K_{\mathbb{Z}}$ the set of integral points in K , i.e., $K_{\mathbb{Z}} := K \cap \mathbb{Z}^n$. For $F \in R[x^\pm]$, we denote by $\Delta(F)$ the Newton polytope of F , by which we mean the convex hull in \mathbb{R}^n of the finite set $\{u \in \mathbb{Z}^n : [F]_u \neq 0\} \subseteq \mathbb{Z}^n$. We denote by P_K the free R -module on the set of monomials with exponents in $K \cap \mathbb{Z}^n$:

$$P_K := \bigoplus_{u \in K_{\mathbb{Z}}} R x^u.$$

For $s \in \mathbb{Z}^+$, let sK denote the s -fold dilation of K . For two convex polytopes K_1, K_2 denote by $K_1 \oplus K_2$ the Minkowski sum $\{v_1 + v_2 : v_1 \in K_1, v_2 \in K_2\}$. One can show that for $F, G \in R[x^\pm]$ and $s \in \mathbb{Z}^+$, we have $\Delta(FG) = \Delta(F) \oplus \Delta(G)$ and $\Delta(F^s) = s\Delta(F)$.

For $n \geq 1$, let $\mathbb{A}_{\mathbb{F}_q}^n$ denote affine n -space over \mathbb{F}_q with coordinates x_1, \dots, x_n . Let $\mathbb{T}_{\mathbb{F}_q}^n$ be the affine torus $\{x_1 \cdots x_n \neq 0\} \subseteq \mathbb{A}_{\mathbb{F}_q}^n$. When we refer to a hypersurface V in $\mathbb{T}_{\mathbb{F}_q}^n$ cut out by $\bar{F} \in \mathbb{F}_q[x^\pm]$, we mean the variety with geometric points

$$\{(c_1, \dots, c_n) \in (\overline{\mathbb{F}_q}^*)^n : \bar{F}(c_1, \dots, c_n) = 0\}$$

where $\overline{\mathbb{F}_q}$ is an algebraic closure of \mathbb{F}_q . We denote by $V(\mathbb{F}_{q^r})$ the set of \mathbb{F}_{q^r} -rational points on V , i.e., the geometric points

$$\{(c_1, \dots, c_n) \in (\mathbb{F}_{q^r}^*)^n : \bar{F}(c_1, \dots, c_n) = 0\}.$$

For $q = p^a$, let \mathbb{Q}_q denote the unique unramified extension of degree a of \mathbb{Q}_p , and let \mathbb{Z}_q denote the ring of integers of \mathbb{Q}_q . Let $\phi : \mathbb{F}_q \rightarrow \mathbb{F}_q$ be the absolute Frobenius map $x \mapsto x^p$. We shall also denote by ϕ the unique lift of the Frobenius map to a continuous ring automorphism $\mathbb{Z}_q \rightarrow \mathbb{Z}_q$.

We now define maps ϕ, ψ, T_H, A_H analogous to the maps ϕ, ψ, T_H, A_H in [14, §3.1]. Define $\phi, \psi : \mathbb{Z}_q[x^\pm] \rightarrow \mathbb{Z}_q[x^\pm]$ by

$$\phi(G) = \sum_{u \in \mathbb{Z}^n} \phi([G]_u) x^{pu}, \quad \psi(G) = \sum_{u \in \mathbb{Z}^n} \phi^{-1}([G]_{pu}) x^u.$$

For $H \in \mathbb{Z}_q[x^\pm]$, let $T_H : \mathbb{Z}_q[x^\pm] \rightarrow \mathbb{Z}_q[x^\pm]$ be the multiplication operator $G \mapsto HG$, and let

$$A_H = \psi \circ T_{H^{p-1}}.$$

Note that A_H is a ϕ^{-1} -semilinear map. For $m \geq 1$, we also define

$$H^{(m)} = (H \cdot \phi(H) \cdots \phi^{m-1}(H))^{p-1}.$$

It is straightforward to check that

$$A_H^m = \psi^m \circ T_{H^{(m)}}$$

for any $m \geq 1$, and that A_H maps P_K into P_K for $K \supseteq \Delta(H)$. Note that $A_H^a = \psi^a \circ T_{H^{(a)}}$ is a linear map, it is not just ϕ^{-1} -semilinear.

The following theorem directly generalises [14, Thm. 3.1].

Theorem 2.1 *Let $K \subseteq \mathbb{R}^n$ be a convex polytope with integral vertices, and let $\bar{F} \in \mathbb{F}_q[x^\pm]$ be a non-zero Laurent polynomial with $\Delta(\bar{F}) \subseteq K$. Let V be the hypersurface in $\mathbb{T}_{\mathbb{F}_q}^n$ cut out by \bar{F} . Let $F \in \mathbb{Z}_q[x^\pm]$ be any lift of \bar{F} with $\Delta(F) \subseteq K$. Let r, λ, τ be positive integers satisfying*

$$\tau \geq \frac{\lambda}{(p-1)ar}.$$

Then

$$|V(\mathbb{F}_{q^r})| = (q^r - 1)^n \sum_{s=0}^{\lambda+\tau-1} \alpha_s \text{tr}(A_{F^s}^{ar}) \pmod{p^\lambda},$$

where

$$\alpha_s = (-1)^s \sum_{t=0}^{\tau-1} \binom{-\lambda}{t} \binom{\lambda}{s-t} \in \mathbb{Z},$$

and where $A_{F^s}^a$ is regarded as a linear operator on the \mathbb{Z}_q -module P_{sK} .

Remark 2.2 Harvey mistakenly described the map A_{F^s} as a \mathbb{Z}_q -linear operator in [14, Thm. 3.1]. The map $A_{F^s}^a$ is a \mathbb{Z}_q -linear operator, but in general A_{F^s} is only ϕ^{-1} -semilinear.

We do not give a proof of this theorem. The proof is identical to the one provided by Harvey for [14, Thm. 3.1], except that we remove all references to degree and replace them with the application of the results $\Delta(G^s) = s\Delta(G)$ and $\Delta(G_1G_2) = \Delta(G_1) \oplus \Delta(G_2)$. The key idea of Harvey’s proof is to construct an indicator function J on $(\mathbb{Z}_{q^r}^*)^n$ that takes the value 1 $\in \mathbb{Z}/p^\lambda\mathbb{Z}$ if the input $c \in (\mathbb{Z}_{q^r}^*)^n$ reduces mod p to a zero of \bar{F} , and takes the value 0 otherwise.

As in [14, Lem. 3.2], we now give a more computationally explicit description of $A_{F^s}^{ar}$.

Lemma 2.3 *Let $F \in \mathbb{Z}_q[x^\pm]$ be as in Theorem 2.1. The matrix of $A_{F^s}^a$ on P_{sK} , with respect to the basis $\{x^u : u \in (sK)_{\mathbb{Z}}\}$, is given by*

$$\phi^{a-1}(M_s) \cdots \phi(M_s)M_s$$

where M_s is the square matrix defined by

$$(M_s)_{v,u} = [F^{(p-1)s}]_{pv-u}$$

for $u, v \in (sK)_{\mathbb{Z}}$, and where ϕ acts component-wise on matrices.

The proof of this lemma is identical to the one provided in [14]. The difference between our Theorem 2.1 and Lemma 2.3 and Harvey’s Theorem 3.1 and Lemma 3.2 is that Harvey works with homogeneous \bar{F} in $\mathbb{F}_q[x_0, \dots, x_n]$ rather than Laurent polynomials in $\mathbb{F}_q[x_1, x_1^{-1}, \dots, x_n, x_n^{-1}]$. In Harvey’s proof of Theorem 3.1, homogenous polynomials act as a convenient book-keeping device — the extra variable x_0 serves as a way to keep track of relations between Newton polytopes. However, given that the trace formula counts points on a hypersurface in an affine torus it is in fact more appropriate to work with Laurent polynomials, and thus our Theorem 2.1 is a more natural way to state Harvey’s trace formula.

3 The algorithm

Let \mathbb{F}_q, \bar{F} and \tilde{X} be as in Theorem 1.1. Here we shall present an algorithm based on [14] that takes \bar{F} as input and outputs $Z(\tilde{X}, T)$. Our algorithm consists of two independent sub-routines. The first subroutine counts the number of \mathbb{F}_{q^r} -rational points on the projective plane model X defined by \bar{F} for $r = 1, \dots, g$; the second subroutine determines the errors $|\tilde{X}(\mathbb{F}_{q^r})| - |X(\mathbb{F}_{q^r})|$ for each r . We refer to the former subroutine as **CountPlaneModel**, and the latter as **ComputeCorrections**. We expect the asymptotic time complexity of **CountPlaneModel** to dominate that of **ComputeCorrections**.

3.1 Counting points on a plane model

Here we describe the **CountPlane Model** algorithm. It takes as input a non-zero polynomial $\bar{F} \in \mathbb{F}_q[x, y]$, a positive integer D , and a positive integer λ . The polynomial \bar{F} cuts out an affine plane curve X_0 , whose projective closure we denote by X . The algorithm outputs a list of the point-counts $|X(\mathbb{F}_{q^r})| \pmod{p^\lambda}$ for $r = 1, \dots, D$.

The idea of the algorithm is to first count the points on $X \cap \mathbb{T}^2$ by applying the trace formula from Theorem 2.1 to \bar{F} , followed by locating and counting the points on $X \setminus \mathbb{T}^2$.

Given the input \bar{F}, λ, D , the first step in CountPlaneModel is to choose an $F \in (\mathbb{Z}_q/p^\lambda \mathbb{Z}_q)[x, y]$ satisfying $F \pmod{p} = \bar{F}$ and $\Delta(F) = \Delta(\bar{F})$ to which we can apply the trace formula. Here the ring $\mathbb{Z}_q/p^\lambda \mathbb{Z}_q$ is represented as $(\mathbb{Z}/p^\lambda \mathbb{Z})[t]/b$, where $b \in (\mathbb{Z}/p^\lambda \mathbb{Z})[t]$ is a monic, degree a lift of the irreducible degree a polynomial $\bar{b} \in \mathbb{F}_p[t]$ defining the extension $\mathbb{F}_q/\mathbb{F}_p$. In our application of the trace formula, we will ultimately evaluate

$$|(X \cap \mathbb{T}^2)(\mathbb{F}_{q^r})| \pmod{p^\lambda} = (q^r - 1)^2 \sum_{s=0}^S \alpha_s \text{tr}(A_{F_s}^{ar}) \pmod{p^\lambda} \tag{3.1}$$

for all $r = 1, \dots, D$. In each of these sums we take $S := \lambda + \tau - 1$ with $\tau := \lceil \lambda / (a(p - 1)) \rceil$; by Theorem 2.1 this choice of τ guarantees the correct result modulo p^λ for every $r \in \mathbb{Z}^+$. We can do these evaluations of (3.1) efficiently in the manner described below — this is the same as what is described in [14, §3], with a slight modification to account for the fact that we are working with the Newton polygon $\Delta(F)$ of F rather than the polygon $\text{Conv}\{(0, 0), (0, d), (d, 0)\} \supseteq \Delta(F)$.

As in [14, §3], we iterate s over $1, \dots, S$: for each s we first construct the matrix M_s modulo p^λ where M_s is the matrix described in Lemma 2.3, then we compute the traces $\text{tr}(A_{F_s}^{ar}) \pmod{p^\lambda}$ for $r = 1, \dots, D$. We compute M_s by expanding $F^{(p-1)s}$ and reading off coefficients at monomials $pv - u$ where $u, v \in s\Delta(F) \cap \mathbb{Z}^2$. We compute the traces $\text{tr}(A_{F_s}^{ar}) \pmod{p^\lambda}$ for $r = 1, \dots, D$ exactly as in [14, Lemma 3.4]. That is, we compute the matrix of $A_{F_s}^a$ using a modified binary powering algorithm in accordance with Lemma 2.3, and we naively compute the required traces by computing D successive powers of that matrix. Having computed all of the required traces $\text{tr}(A_{F_s}^{ar}) \pmod{p^\lambda}$ for $1 \leq r \leq D$ and $1 \leq s \leq S$, we iterate r over $1, \dots, D$ and use (3.1) to compute $|(X \cap \mathbb{T}^2)(\mathbb{F}_{q^r})| \pmod{p^\lambda}$ for each r .

Computing $|(X \setminus \mathbb{T}^2)(\mathbb{F}_{q^r})| \pmod{p^\lambda}$ for $r = 1, \dots, D$ can be done by factoring three univariate polynomials over \mathbb{F}_q . We count the points on X in $\mathbb{P}^2 \setminus \mathbb{T}^2$ by factoring the polynomials $\bar{F}_0, \bar{F}_1, \bar{F}_2 \in \mathbb{F}_q[t]$ defined by $\bar{F}_0(t) := \bar{F}_h(t, 1, 0)$, $\bar{F}_1(t) := \bar{F}_h(0, t, 1)$, and $\bar{F}_2(t) := \bar{F}_h(1, 0, t)$, where $\bar{F}_h \in \mathbb{F}_q[x_0, x_1, x_2]$ denotes the homogenisation of \bar{F} .

For simplicity of exposition we have described an unoptimised version of the **CountPlaneModel** algorithm, whose time complexity has a dependence in p of $p^{2+o(1)}$ rather than $p^{\frac{1}{2}+o(1)}$. To achieve the time complexity stated in Theorem 1.1, one can use the deformation recurrence technique from [14, §4] to compute the required matrices M_s rather than using the naive approach of expanding powers of F and reading off the required coefficients. Another optimisation arises from using an individual precision λ_r and upper limit S_r for each r rather than the global precision λ and upper limit $S = \lambda + \lceil \lambda / (a(p - 1)) \rceil - 1$.

The most computationally expensive parts of applying the algorithm described above are: (1) the computation of the coefficients of the powers $F^{(p-1)s}$ that are required to construct the matrices M_s , and (2) the computation of the matrix powers $A_{F_s}^{ar}$ and their traces. On (2), note that the dimension of the matrix M_s is roughly $s^2 \text{vol}(\Delta(F))$ as by construction the matrix M_s has dimension equal to the number of integral points in $s\Delta(F)$, and in \mathbb{R}^2 the polygon $s\Delta(F)$ has s^2 times the volume of $\Delta(F)$.

3.2 Computing corrections for miscounted points

Here we describe the **ComputeCorrections** algorithm. It takes as input an absolutely irreducible polynomial $\bar{F} \in \mathbb{F}_q[x, y]$ of degree $d \geq 2$ and a positive integer D , and it

outputs a list of the differences $|\tilde{X}(\mathbb{F}_{q^r})| - |X(\mathbb{F}_{q^r})|$ for $r = 1, \dots, D$, where X is the projective closure of the affine plane curve cut out by \bar{F} .

The idea of the algorithm is to reduce the problem of computing the differences $|\tilde{X}(\mathbb{F}_{q^r})| - |X(\mathbb{F}_{q^r})|$ to the problem of counting points on \tilde{X} and X above a certain set of points on \mathbb{P}^1 . To do this, we rely on the following theorem.

Theorem 3.1 *Let $\pi : \tilde{X} \rightarrow X$ be the normalisation of X . Let S be the singular subscheme of X . Let Z be any subscheme of X containing S , and let $\tilde{Z} = \pi^{-1}(Z)$. Then we have*

$$|\tilde{X}(\mathbb{F}_{q^r})| - |X(\mathbb{F}_{q^r})| = |\tilde{Z}(\mathbb{F}_{q^r})| - |Z(\mathbb{F}_{q^r})|.$$

Sketch proof Theorem 3.1 follows from: (1) the fact that for any closed point p on X , we have

$$\widetilde{\mathcal{O}}_{X,p} = \bigcap_{q \in \pi^{-1}(p)} \mathcal{O}_{\tilde{X},q}$$

where $\mathcal{O}_{X,p} \subseteq \mathbb{F}_q(\tilde{X})$ is the ring of rational functions on X that are regular at p , $\widetilde{\mathcal{O}}_{X,p}$ is the integral closure in $\mathbb{F}_q(\tilde{X})$ of that ring, and $\mathcal{O}_{\tilde{X},q} \subseteq \mathbb{F}_q(\tilde{X})$ is the ring of rational functions on \tilde{X} that are regular at q [22, Thm. III.2.6], and (2) the fact that $\mathcal{O}_{X,p}$ is integrally closed if p is a nonsingular closed point on X [18, Prop. VII.2.6]. □

We will compute the differences by choosing a subscheme Z of X containing S and counting points on both Z and $\tilde{Z} = \pi^{-1}(Z)$. Our **ComputeCorrections** algorithm will consist of three subroutines: **ComputeY**, **CountPointsOnZ** and **CountPointsAboveZ**, which we describe below.


3.2.1 Counting points on Z

Let $a_0(x), \dots, a_n(x)$ be the coefficients of \bar{F} when regarded as a polynomial in y over $\mathbb{F}_q[x]$, i.e., when we write $\bar{F} = a_n(x)y^n + \dots + a_0(x)$ with $a_n(x) \neq 0$. Let X_0 be the affine plane curve cut out by \bar{F} , and let S_0 be the singular subscheme of X_0 . Let $\varphi : X_0 \rightarrow \mathbb{A}^1$ be the morphism of curves $(x, y) \mapsto x$. Let A be the closed set of \mathbb{A}^1 where $a_n(x)$ vanishes.

We define subschemes Y_0, Y of \mathbb{P}^1 and Z_0, Z of X as follows:

$$\begin{aligned} Y_0 &:= \varphi(S_0) \cup A, \\ Z_0 &:= \varphi^{-1}(Y_0), \\ Y &:= Y_0 \cup \{\infty\}, \\ Z &:= Z_0 \cup (X \setminus X_0). \end{aligned} \tag{3.2}$$

That is, the points of Y_0 are the x -coordinates of singular points on X_0 together with the x -coordinates of “vertical asymptotes” of X_0 . The points on Z_0 are those points on X_0 that share an x -coordinate with a singular point, along with those points at which $a_n(x)$ vanishes. The points on $Z \subseteq X$ are all of the points on $Z_0 \subseteq X_0$ together with all of the points on $X \setminus \mathbb{A}^2 \subseteq \mathbb{P}^2$. Note that by construction, Z contains S . An illustration of the schemes Z_0 and Y_0 is shown below in Fig. 1.



Based on the discussion above we now introduce the subroutine **ComputeY**. It takes as input the polynomial \bar{F} and outputs a list of irreducible polynomials in $\mathbb{F}_q[x]$ representing the closed set $Y_0 \subseteq \mathbb{A}^1$. This subroutine amounts to finding the distinct irreducible factors of the polynomial $a_n(x) \cdot \gcd(\text{res}_y(\bar{F}, \frac{\partial \bar{F}}{\partial y}), \text{res}_y(\bar{F}, \frac{\partial \bar{F}}{\partial x}))$. Note that $\text{res}_y(\cdot, \cdot)$ denotes the resultant where the polynomials are regarded as univariate polynomials in y whose coefficients are polynomials in x , and hence the set of roots of $\gcd(\text{res}_y(\bar{F}, \frac{\partial \bar{F}}{\partial y}), \text{res}_y(\bar{F}, \frac{\partial \bar{F}}{\partial x})) \in \mathbb{F}_q[x]$ includes all of the x -coordinates of common zeroes of \bar{F} , $\frac{\partial \bar{F}}{\partial x}$ and $\frac{\partial \bar{F}}{\partial y}$, i.e., it includes the x -coordinates of singular points on X_0 .

The algorithm **CountPointsOnZ** takes as input the polynomial \bar{F} , the list of irreducible polynomials representing Y_0 , and a positive integer D , and outputs the list of point-counts $|Z(\mathbb{F}_{q^r})|$ for $r = 1, \dots, D$. **CountPointsOnZ** obtains the desired list of point-counts by:

- (1) counting the points on $Z \cap \mathbb{A}^2 = Z_0$ by computing the distinct factors in the factorisation of \bar{F} modulo each of the irreducibles in $\mathbb{F}_q[x]$ that represent Y_0 , and
- (2) counting the points on $Z \setminus \mathbb{A}^2 = X \setminus \mathbb{A}^2$ by computing the distinct factors in the factorisation of the degree d homogeneous part of \bar{F} .

More precisely, we count points on Z_0 by constructing the field $\mathbb{F}_q[x]/\bar{h}(x)$ for each irreducible $\bar{h} \in \mathbb{F}_q[x]$ representing Y_0 , then factoring the univariate in y polynomial $\bar{F} \pmod{\bar{h}}$ over $\mathbb{F}_q[x]/\bar{h}(x)$; each distinct irreducible factor in such a factorisation corresponds to a $\text{Gal}(\overline{\mathbb{F}_q}/\mathbb{F}_q)$ -orbit of points on X_0 whose x -coordinates are roots of $\bar{h}(x)$.

3.2.2 Counting points on \tilde{Z}

Let $\tilde{\varphi} : \tilde{X} \rightarrow \mathbb{P}^1$ be the morphism of curves defined by the rational function $x \in \mathbb{F}_q(x)[y]/\langle \bar{F} \rangle$. Note that when \bar{F} is absolutely irreducible and of degree $d \geq 2$ this morphism is surjective.

The following proposition allows us to count points on $\tilde{Z} = \pi^{-1}(Z)$ by counting the points on \tilde{X} that lie above the scheme $Y \subseteq \mathbb{P}^1$ defined in (3.2).


Proposition 3.2 *The subscheme $\tilde{Z} = \pi^{-1}(Z)$ of \tilde{X} satisfies*

$$\tilde{Z} = \tilde{\varphi}^{-1}(Y).$$

Sketch proof We claim that if $q \in \tilde{X}$ satisfies $\pi(q) \in X_0$ then $\tilde{\varphi}(q) = \varphi(\pi(q))$, otherwise if $\pi(q) \notin X_0$ then $\tilde{\varphi}(q) \in A \cup \{\infty\}$. The equality $\tilde{\varphi}(q) = \varphi(\pi(q))$ for $\pi(q) \in X_0$ follows from the fact that the morphisms φ and $\tilde{\varphi}$ are induced by the same rational function $x \in \mathbb{F}_q(x)[y]/\langle \bar{F} \rangle$; this is illustrated below in Fig. 2. The inclusion $\tilde{\varphi}(\pi^{-1}(X \setminus X_0)) \subseteq \{\infty\} \cup A$ follows from: (1) if $\frac{1}{x}$ vanishes at $p = \pi(q)$ then $\tilde{\varphi}(q) = \infty$, and (2) if $\frac{1}{x}$ does not vanish at $p = \pi(q)$ but $\frac{1}{y}$ does vanish at p , then $\frac{1}{y} \in \mathfrak{m}_{\tilde{X}, q}$, and since $a_n(x)y$ is integral over $\mathbb{F}_q[x]$ it follows that $a_n(x) \in \mathfrak{m}_{\mathbb{P}^1, \tilde{\varphi}(q)}$. \square

To count the points on \tilde{X} lying above Y , we rely on the **Montes** algorithm [13], which we shall refer to as **Montes**. For details on the time complexity of **Montes** see [2]. Montes takes as input a monic irreducible separable polynomial $\bar{F} \in \mathbb{F}_q[x][y]$ defining the function field $\mathbb{F}_q(x)[y]/\langle \bar{F} \rangle$ and an irreducible polynomial $\bar{h}(x) \in \mathbb{F}_q[x]$, and outputs the list of closed points on \tilde{X} that lie above the closed point defined by $\bar{h}(x)$ on $\mathbb{A}^1 = \text{Spec}(\mathbb{F}_q[x])$.

Note that in Theorem 1.1 we do not assume that the input \bar{F} is monic in y , but **Montes** does require an input which is monic in y . Fortunately this is not an issue as it will suffice to apply **Montes** to the polynomial $\bar{F}' = y^n + a_{n-1}(x)y^{n-1} + \dots + a_0(x)a_n(x)^{n-1}$ instead. This works because \bar{F} and \bar{F}' define the same extension of $\mathbb{F}_q(x)$ via the isomorphism $y \mapsto a_n(x)y$.



The algorithm **CountPointsAboveZ** takes as input the polynomial \bar{F} , the list of irreducible polynomials representing Y_0 and a positive integer D , and outputs the list of point-counts $|\tilde{Z}(\mathbb{F}_{q^r})|$ for $r = 1, \dots, D$. In accordance with Proposition 3.2, **CountPointsAboveZ** obtains the desired list of point-counts by:

- (1) counting the points on $\tilde{\varphi}^{-1}(Y_0)$ via applying Montes to \bar{F}' for each of the irreducibles $\tilde{h}(x)$ in $\mathbb{F}_q[x]$ that represent Y_0 , and
- (2) counting the points on $\tilde{\varphi}^{-1}(\infty)$ via applying Montes to a polynomial \bar{F}'' and $\tilde{h}(x) = x$, where \bar{F}'' is the polynomial that we get from \bar{F}' via the change of variables $(x, y) \mapsto (\frac{1}{x}, \frac{y}{x^m})$ with $m := \deg_x(\bar{F}')$.

3.2.3 Computing the corrections

We can now describe the algorithm **ComputeCorrections**. Given the input $\bar{F} \in \mathbb{F}_q[x, y]$ and D , the algorithm first runs the subroutine **ComputeY** to obtain a representation L of the points in $Y \subseteq \mathbb{P}^1$. It then runs the subroutines **CountPointsOnZ** and **CountPointsAboveZ** with arguments \bar{F}, D and L to obtain the point-counts $|Z(\mathbb{F}_{q^r})|$ and $|\tilde{Z}(\mathbb{F}_{q^r})|$ for $r = 1, \dots, D$, from which we obtain the differences $|\tilde{X}(\mathbb{F}_{q^r})| - |X(\mathbb{F}_{q^r})|$ for $r = 1, \dots, D$ by computing $|\tilde{X}(\mathbb{F}_{q^r})| - |X(\mathbb{F}_{q^r})| = |\tilde{Z}(\mathbb{F}_{q^r})| - |Z(\mathbb{F}_{q^r})|$ in accordance with Theorem 3.1.

3.3 The main algorithm

We may now outline an efficient algorithm for computing $Z(\tilde{X}, T)$. Given an absolutely irreducible $\bar{F} \in \mathbb{F}_q[x, y]$ of degree $d \geq 2$, we can compute the genus g of $\mathbb{F}_q(x)[y]/(\bar{F})$ using existing efficient algorithms [15]. From the Weil conjectures, we have the Hasse–Weil bound

$$|q^r + 1 - |\tilde{X}(\mathbb{F}_{q^r})|| \leq 2gq^{r/2}, \tag{3.3}$$

hence we can recover the value $|\tilde{X}(\mathbb{F}_{q^r})|$ from its reduction modulo any integer strictly larger than $4gq^{r/2}$. Based on this, we take $\lambda = \lceil \log_p(4gq^{g/2} + 1) \rceil$, and run **CountPlaneModel** with arguments \bar{F}, λ and g . We get a list of point-counts $|X(\mathbb{F}_{q^r})| \pmod{p^\lambda}$ for $r = 1, \dots, g$ as output. We then run **ComputeCorrections** with arguments \bar{F} and g to obtain a list of the differences $|\tilde{X}(\mathbb{F}_{q^r})| - |X(\mathbb{F}_{q^r})|$ for $r = 1, \dots, g$. Adding these lists together modulo p^λ gives us the point-counts $|\tilde{X}(\mathbb{F}_{q^r})| \pmod{p^\lambda}$ for $r = 1, \dots, g$. We then recover the exact values $|\tilde{X}(\mathbb{F}_{q^r})| \in \mathbb{Z}$ using (3.3). From the power series expansion of

$$(1 - T)(1 - qT) \exp\left(\sum_{r=1}^g \frac{|\tilde{X}(\mathbb{F}_{q^r})|}{r} T^r\right)$$

we obtain the coefficients of the numerator of $Z(\tilde{X}, T)$ up to the monomial T^g . We obtain the remaining coefficients by using the functional equation (1.1) for $Z(\tilde{X}, T)$.

4 Implementation and examples

We have implemented the algorithm described in this paper for prime fields \mathbb{F}_p in the computer algebra system MAGMA [5]. We have implemented a version with time complexity $p^{2+o(1)}d^{O(1)}$. Note that this differs from the time complexity given in Theorem 1.1 — the dependence in p here is p^2 instead of $p^{1/2}$. The reason for this difference is that the current implementation involves computing the polynomials $F^{(p-1)s}$ mentioned in Lemma 2.3 instead of using the deformation-recurrence technique from [14] to obtain

only the coefficients that we need. In the future we hope to make available an updated implementation which incorporates the deformation-recurrence technique and which works for arbitrary finite fields \mathbb{F}_q .

In this section we present and discuss data we have collected about the implementation of our algorithm and the MAGMA implementation of Tuitman’s algorithm. In Sect. 4.1 we provide examples comparing the runtime of our code to that of the MAGMA implementation of Tuitman’s algorithm, and in Sect. 4.2 we present some data on example computations of our code on inputs that cannot be readily dealt with by existing methods.

Our implementation takes as input an absolutely irreducible $\bar{F} \in \mathbb{F}_p[x, y]$. Tuitman’s implementation takes as input a pair (F, p) where $F \in K[x, y]$, F is monic in y , p is a prime, and K is a number field in which p is inert. The polynomial F fed into Tuitman’s code must define a “good” lift of the curve defined by $F \pmod{p}$ over the field $\mathbb{Z}_K/p\mathbb{Z}_K$, i.e., F must define a lift satisfying [23, Ass. 1]. Tuitman’s algorithm has time complexity $\tilde{O}(pd_x^6d_y^4a^3)$, where $a = [\mathbb{Z}_K/p\mathbb{Z}_K : \mathbb{Z}/p\mathbb{Z}]$, d_y is the degree of F in x , and d_x is the degree of F in y .

As mentioned in the introduction, the main advantage of our algorithm over Tuitman’s is that it does not require the computation of a “good” lift of an input $\bar{F} \in \mathbb{F}_q[x, y]$ to characteristic zero. Our code is capable of handling arbitrary absolutely irreducible inputs $\bar{F} \in \mathbb{F}_p[x, y]$, including those for which one cannot readily compute a lift $F \in \mathbb{Q}[x, y]$ of \bar{F} satisfying Assumption 1 of [23]. Furthermore, even when it is feasible to compute a good lift of an input \bar{F} , the lift one obtains may have properties that cause Tuitman’s code to run slowly when compared to the execution of our code on the original \bar{F} .

The inputs \bar{F} for which our code is most likely to outperform Tuitman’s are those where \bar{F} does not meet Baker’s bound on the genus [3, Thm. 2.4], i.e., inputs \bar{F} for which the geometric genus of the curve defined by \bar{F} is strictly less than the number of integral interior points of $\Delta(\bar{F})$. Note that these are precisely those \bar{F} that define a singular curve in the projective toric surface associated with $\Delta(\bar{F})$. When \bar{F} meets Baker’s bound, one can almost always take a naive Newton polygon preserving lift of \bar{F} to obtain a good lift for Tuitman’s algorithm [6, §2.2].

When \bar{F} does not meet Baker’s bound, determining a lift satisfying [23, Ass. 1] is a difficult problem. In some cases this problem can be dealt with by using the methods of [6] or [7]—these methods apply to inputs \bar{F} that define curves of geometric genus at most 5 or whose Newton polygon $\Delta(\bar{F})$ has lattice width at most 5. The *width* of a convex polytope $K \subseteq \mathbb{R}^n$ along a direction $\mathbf{d} \in \mathbb{R}^n$ is defined to be

$$w_{\mathbf{d}}(K) = \max_{\mathbf{x} \in K} \mathbf{d} \cdot \mathbf{x} - \min_{\mathbf{x} \in K} \mathbf{d} \cdot \mathbf{x}.$$

The *lattice width* of a convex polytope $K \subseteq \mathbb{R}^n$ with integral vertices is defined to be

$$lw(K) = \min_{\mathbf{d} \in \mathbb{Z}^n \setminus \{0\}} w_{\mathbf{d}}(K).$$

For $K \subseteq \mathbb{R}^2$ one can equivalently define the lattice width as the minimal height of a horizontal strip inside which K can be mapped by a unimodular transformation.

The timings obtained in this section were obtained using MAGMA V2.25-7 on a computer with an Nvidia RTX 3090 GPU and an Intel Core i9-12900K CPU with 128GB of memory running Ubuntu 20.04. These timings show what is achievable in practice when

running a straightforward MAGMA implementation of our algorithm on high-end consumer hardware. The MAGMA supported linear algebra on an RTX 3090 GPU allows for very fast matrix multiplication, making it feasible to apply our algorithm to a large range of curves.

It should be noted that the MAGMA implementation of Tuitman’s algorithm may be able to be significantly improved. Improvements that appeared in MAGMA V2.25-7 were made by MAGMA developer Allan Steel after being sent test inputs from the present author that were found during the writing of this paper. An advantage of our algorithm (particularly the subroutine **CountPlaneModel**) is that it is comparatively simple and straightforward to implement, and its current implementation has the potential to be sped up substantially through simple low-level optimisations.

The code for the implementation of our algorithm is available at https://github.com/Maddels/zeta_function_. Along with the implementation of our algorithm we provide code for testing the correctness of the output that our code produces. The functions provided for testing correctness do not *guarantee* that the output is correct — failing a test indicates that an error has occurred; passing a test gives partial evidence that the output is consistent with the correct result. We test for correctness by: (1) comparing the point-count that is predicted by the output zeta function for an extension \mathbb{F}_{p^i} where $i > g$ with the point-count that is computed by applying the trace formula (2.1) for the extension \mathbb{F}_{p^i} , and (2) comparing the point-counts that are predicted by the output zeta function for small extensions \mathbb{F}_{p^i} with point-counts that are computed by naive enumeration. These tests were applied for each of the inputs appearing in Sect. 4.2. Correctness was checked for the inputs in Sect. 4.1 by comparing the output obtained by our implementation to the output obtained by the MAGMA implementation of Tuitman’s algorithm.

4.1 Comparison with Tuitman’s algorithm

We now provide examples where we compare the runtime of our implementation with that of the MAGMA implementation of Tuitman’s algorithm. More precisely, we compare the runtime of our algorithm on an input \bar{F} against the combined time of computing a “good” lift F of \bar{F} followed by running Tuitman’s algorithm on input F . For the examples below, we further provide a breakdown of the runtime of both implementations into the runtimes of their key steps. For our implementation, the key steps are:

- (A1) Computing the powers $F^{(p-1)s}$ needed for the construction of the matrices M_s .
- (A2) Computing the matrix powers M_s^r and their traces.
- (A3) All other required computations in the algorithm, including computing the corrections to the point-counts for the plane model.

For the MAGMA implementation of Tuitman’s algorithm combined with lifting to obtain a valid input for Tuitman’s code, the key steps are:

- (B1) Computing a good lift F of \bar{F} to feed into Tuitman’s algorithm.
- (B2) Computing the objects $\Delta, r(x), s(x, y), W^0, W^\infty, G^0, G^\infty$ described in [23, §2].
- (B3) Computing the basis of $H^1(\tilde{X})$.
- (B4) Computing the Frobenius lift.

- (B5) Computing reduction matrices.
- (B6) Computing the Frobenius matrix, from which the zeta function is computed by computing the characteristic polynomial.

Note that steps (B3)–(B6) above correspond to steps I–IV described in [23]. To perform (B1) for a given \bar{F} , we used MAGMA implementations of the methods described in [6, 7], and when those methods were not applicable we instead tried computing a lift whose singular points are lifts of the singular points on the input plane model. That is, for an input \bar{F} to which the methods of [6, 7] do not apply, we attempted to compute an $F \in \mathbb{Z}[x, y]$ that: (1) reduces mod p to \bar{F} , and (2) defines a plane curve over \mathbb{Q} whose singularities reduce mod p to the singularities of X . This method does not guarantee a lift satisfying [23, Ass. 1], but in the case where the input \bar{F} defines a nodal plane curve we often succeed in finding such a lift.

Below we present the runtimes for a selection of 10 inputs \bar{F} that have a variety of features. These 10 inputs along with our code for performing step (B1) can be found on the webpage https://github.com/Maddels/zeta_function_ in the MAGMA files `runtime_comparison_examples.m` and `lift_curve.m`. Inputs 1 and 2 are examples that were made available by Tuitman at the webpage https://github.com/jtuitman/pcc/blob/master/pcc_p/example_p.m, inputs 3 and 4 were the examples explicitly given in [7, §6] for $d = 4$ and $d = 5$, input 5 was generated by specifying its Newton polygon and singular behaviour, and the remaining 5 inputs were generated using the MAGMA function `RandomCurveByGenus`.

Table 1 gives information on the nature of these inputs \bar{F} and the curves they define, as well as which lifting strategy was applied for each input. In Table 1 we denote a naive lift by N, a lift obtained using the methods of [6] by CT, a lift obtained using the methods of [7] by CV, and a singular-point-preserving lift by S. Table 2 gives information on the runtime and memory usage of the implementations of each algorithm on each of these inputs, and Table 3 provides a breakdown of these runtimes into steps A1–A3 and B1–B6. In Table 2 we denote the new algorithm by HK, and Tuitman’s algorithm by T. If the runtime for an example exceeded 12 hours then the computation was terminated, with information provided on the runtimes of completed key steps up until the point of termination.

Table 1 Inputs for runtime comparison

Input #	1	2	3	4	5	6	7	8	9	10
Lifting strategy	N	N	CV	CV	CV	S	S	S	CT	CT
Meets Bakers bound	Yes	Yes	No	No	No	No	No	No	No	No
$lw(\Delta(\bar{F}))$	3	5	4	5	5	6	7	9	5	5
p	1009	11	7	17	101	23	3	19	43	11
g	4	16	10	9	12	6	8	10	4	5
$\#(\text{int}(\Delta\bar{F}) \cap \mathbb{Z}^2) - g$	0	0	47	87	54	4	13	18	2	5
$\text{Vol}(\Delta(\bar{F}))$	7.5	22.5	80	125	87.5	18	30	40.5	12.5	17.5

Table 2 Runtime and memory comparison

curve	HK time	T time	HK memory	T memory
1	37.98s	7.65s	4141MB	122MB
2	27.38s	14.32s	1608MB	469MB
3	13.67s	1.96s	922MB	479MB
4	40.95s	9466.71s	2793MB	735MB
5	199.16s	≥ 12h	8751MB	≥ 4415MB
6	0.32s	437.31s	1835MB	1954MB
7	3.43s	10607.16s	2032MB	6567MB
8	9.55s	≥ 12h	2889MB	≥ 3844MB
9	0.08s	2.01s	2435MB	2435MB
10	0.09s	32.55s	2435MB	2435MB

Bold value indicates which column has the shorter runtime for that row

Table 3 HK and T runtime breakdown (in seconds)

#	A1	A2	A3	B1	B2	B3	B4	B5	B6
1	37.44	0.11	0.43	0.00	0.02	0.04	1.00	1.07	5.52
2	0.46	26.76	0.16	0.00	0.64	2.08	1.42	0.69	9.49
3	0.24	13.39	0.04	0.18	0.10	0.45	0.13	0.16	0.94
4	1.13	39.72	0.10	64.07	73.22	9222.68	1.53	0.62	104.58
5	107.50	90.97	0.69	2236.41	1322.85	-	-	-	-
6	0.15	0.16	0.01	0.10	267.41	125.27	0.69	3.70	40.14
7	0.17	3.22	0.04	0.03	3372.37	7140.78	4.82	19.54	69.59
8	1.13	8.38	0.04	0.58	-	-	-	-	-
9	0.07	0.01	0.00	0.15	0.23	0.37	0.21	0.14	0.91
10	0.02	0.05	0.02	31.82	0.10	0.23	0.05	0.05	0.30

4.2 Examples presently beyond Tuitman’s algorithm

We now provide example computations of our algorithm on inputs \bar{F} that cannot be readily dealt with by existing methods. We selected our examples based on the lifting strategies described in §4.1, in the sense that we selected examples for which these methods cannot be readily applied. One important consideration in the selection of our examples was the lattice width of the Newton polygon $\Delta\bar{F}$. We selected inputs \bar{F} for which it is difficult (or impossible) to find a polynomial $\tilde{G} \in \mathbb{F}_q[x, y]$ that both satisfies $\text{lw}(\Delta(\tilde{G})) \leq 5$ and defines the same nonsingular curve as \bar{F} .

The *arithmetic gonality* of a curve \tilde{X} over \mathbb{F}_q is the minimal degree of the extension $\mathbb{F}_q(\tilde{X})/\mathbb{F}_q(\alpha)$ over all $\alpha \in \mathbb{F}_q(\tilde{X}) \setminus \mathbb{F}_q$. In the case where $g(\tilde{X}) \geq 1$, the arithmetic gonality of \tilde{X} coincides with the minimum lattice width $\text{lw}(\Delta(\tilde{G}))$ among all $\tilde{G} \in \mathbb{F}_q[x, y]$ that define a function field isomorphic to $\mathbb{F}_q(\tilde{X})$.

We chose examples \bar{F} with $\text{lw}(\Delta(\bar{F})) \geq 6$ and $g(\tilde{X}) \geq 9$ as these are likely to have arithmetic gonality exceeding 5, in which case it would be impossible to find a polynomial $\tilde{G} \in \mathbb{F}_q[x, y]$ defining the same function field as \bar{F} to which the methods of [7] apply. The restriction $g(\tilde{X}) \geq 9$ is based on the statements about gonality given in [6, §2.1]. The $\overline{\mathbb{F}_q}$ -gonality (i.e., geometric gonality) of a curve over $\overline{\mathbb{F}_q}$ of genus $g \geq 2$ must lie in the range $2, \dots, \lceil g/2 \rceil + 1$, and inside the moduli space \mathcal{M}_g of curves of genus g the dimension of the locus of curves having fixed $\overline{\mathbb{F}_q}$ -gonality $\gamma \in \{2, \dots, \lceil g/2 \rceil + 1\}$ is $\min\{2g + 2\gamma - 5, 3g - 3\}$ — here we are assuming as in [6] that this result holds in finite characteristic as well as over \mathbb{C} . Thus for $g \geq 9$ the dimension of a locus of curves with an $\overline{\mathbb{F}_q}$ -gonality of $\gamma \leq 5$ in \mathcal{M}_g is strictly smaller than that of a locus of curves with an $\overline{\mathbb{F}_q}$ -gonality of $\gamma \in \{6, \dots, \lceil g/2 \rceil + 1\}$.

Furthermore, even for those \bar{F} that define a curve \tilde{X} of genus $g(\tilde{X}) \geq 9$ with arithmetic gonality at most 5, if the lattice width of $\Delta(\bar{F})$ exceeds 5 then establishing that the curve has smaller gonality is a difficult problem. One can use the algorithm from [20] to compute geometric gonality, but as explained in [12] this becomes impractical beyond genus 7, and even if the geometric gonality is at most 5 it may be the case that the arithmetic gonality is larger.

Another important consideration in the selection of examples \bar{F} was the singular nature of the associated plane curve X . We chose examples with complicated singular behaviour, in particular those with non-ordinary singularities. For these inputs it should be more difficult to compute a lift F of \bar{F} that defines a curve of equal genus.

Ultimately, we selected examples $\bar{F} \in \mathbb{F}_p[x, y]$ with the following properties:

- $lw(\Delta(\bar{F})) \geq 6$,
- $g(\tilde{X}) \geq 9$,
- \bar{F} does not meet Baker’s bound,
- the plane curve cut out by \bar{F} has non-ordinary singularities,
- the naive lift $F \in \mathbb{Z}[x, y]$ of \bar{F} , obtained by lifting each coefficient $\bar{c} \in \mathbb{F}_p$ of \bar{F} to an integer $0 \leq c < p$, defines a curve over \mathbb{Q} of genus larger than $g(\tilde{X})$.

In Table 5 we present the runtimes of our implementation on a selection of 7 inputs having the above properties. All of the inputs below have a Newton polygon of the form $\text{Conv}\{(0, 0), (0, w), (vw, 0)\}$ where $w \geq 6$ and $v \geq 2$. These examples were found by random search among polynomials that: (1) have a Newton polygon that is equal to a specified convex polygon of the form $\text{Conv}\{(0, 0), (0, w), (vw, 0)\}$, and (2) define a curve with specified singular behaviour (e.g. having a non-ordinary singularity that takes several blow-ups to resolve). These 7 inputs can be found on the webpage https://github.com/Maddels/zeta_function in the MAGMA file `examples_beyond_tuitman.m` (https://github.com/Maddels/zeta_function/blob/main/examples_beyond_tuitman.m).

Table 4 Input curves

Input #	1	2	3	4	5	6	7
$lw(\Delta(\bar{F}))$	6	6	6	7	7	8	8
p	19	23	31	19	29	13	17
g	13	20	11	11	12	13	10
$\#(\text{int}(\Delta\bar{F}) \cap \mathbb{Z}^2) - g$	72	50	74	67	66	92	95
$\text{Vol}(\Delta(\bar{F}))$	108	90	108	98	98	128	128

Table 5 Runtime and memory usage

Curve	Overall time	A1	A2	A3	Memory
1	253.28s	4.91s	245.82s	2.55s	15097MB
2	1944.40s	29.07s	1908.90s	6.43s	62568MB
3	119.13s	10.06s	108.22s	0.85s	16985MB
4	59.65s	2.76s	56.76s	0.13s	12200MB
5	176.61s	10.80s	164.75s	1.06s	18723MB
6	632.26s	3.40s	625.03s	3.83s	37145MB
7	90.44s	2.51s	87.69s	0.24s	14111MB

Acknowledgements

Many thanks to David Harvey for his excellent supervision as my doctoral supervisor, for his helpful feedback on earlier versions of this paper, and for suggesting that I pursue the problem of extending his point-counting algorithm. Thanks to Wouter Castryck for helpful discussions on gonality, nondegenerate curves, and the applicability of Tuitman's algorithm. Thanks to Allan Steel and John Cannon of the computational algebra group at the University of Sydney for their assistance with MAGMA and for suggested improvements to the CountPlaneModel algorithm.

Funding Information Open Access funding enabled and organized by CAUL and its Member Institutions

Data availability All data generated or analysed during this study are included in this published article (and its supplementary information files).

Received: 11 September 2022 Accepted: 28 September 2022 Published online: 10 November 2022

References

- Adleman, L.M., Huang, M.-D.: Counting points on curves and abelian varieties over finite fields. In: International Algorithmic Number Theory Symposium (ANTS I), LNCS, vol. 1122, pp. 1–16. Springer, Berlin (1996). https://doi.org/10.1007/3-540-61581-4_36
- Bauch, J.-D., Nart, E., Stainsby, H.D.: Complexity of OM factorizations of polynomials over local fields. *LMS J. Comput. Math.* **16**, 139–171 (2013). <https://doi.org/10.1112/S1461157013000089>
- Beelen, P.: A generalization of Baker's theorem. *Finite Fields Appl.* **15**, 558–568 (2009). <https://doi.org/10.1016/j.ffa.2009.04.003>
- Bombieri, E.: On exponential sums in finite fields. II. *Invent. Math.* **47**, 29–39 (1978). <https://doi.org/10.1007/BF01609477>
- Bosma, W., Cannon, J., Playoust, C.: The Magma algebra system. I. The user language. *J. Symb. Comput.* **24**, 235–265 (1997). <https://doi.org/10.1006/jsc.1996.0125>
- Castryck, W., Tuitman, J.: Point counting on curves using a gonality preserving lift. *Q. J. Math.* **69**, 33–74 (2018). <https://doi.org/10.1093/qmath/hax031>
- Castryck, W., Vermeulen, F.: Lifting low-gonal curves for use in Tuitman's algorithm. In: ANTS XIV-Proceedings of the Fourteenth Algorithmic Number Theory Symposium, Open Book Ser. vol. 4, pp. 109–125. Mathematical Sciences Publishers, Berkeley (2020). <https://doi.org/10.2140/obs.2020.4.109>
- Castryck, W., Deneff, J., Vercauteren, F.: Computing zeta functions of nondegenerate curves. In: International Mathematics Research Papers, vol. 2006, p. 72017 (2006). <https://doi.org/10.48550/arXiv.math/0607308>
- Deneff, J., Vercauteren, F.: Counting points on C_{ab} curves using Monsky-Washnitzer cohomology. *Finite Fields Appl.* **12**, 78–102 (2006). <https://doi.org/10.1016/j.ffa.2005.01.003>
- Gaudry, P., Gürel, N.: An extension of Kedlaya's point-counting algorithm to superelliptic curves. In: Advances in Cryptology—ASIACRYPT: (Gold Coast), LNCS, vol. 2248, pp. 480–494. Springer Berlin (2001). https://doi.org/10.1007/3-540-45682-1_28
- Gaudry, P., Schost, É.: Genus 2 point counting over prime fields. *J. Symb. Comput.* **47**, 368–400 (2012). <https://doi.org/10.1016/j.jsc.2011.09.003>
- Gijswijt, D., Smit, H., van der Wegen, M.: Computing graph gonality is hard. *Discret. Appl. Math.* **287**, 134–149 (2020). <https://doi.org/10.1016/j.dam.2020.08.013>
- Guàrdia, J., Montes, J., Nart, E.: Higher Newton polygons and integral bases. *Finite Fields Appl.* **147**, 549–589 (2015). <https://doi.org/10.1016/j.jnt.2014.07.027>
- Harvey, D.: Computing zeta functions of arithmetic schemes. *Proc. Lond. Math. Soc.* **111**, 1379–1401 (2015). <https://doi.org/10.1112/plms/pdv056>
- Hess, F.: Computing Riemann-Roch spaces in algebraic function fields and related topics. *J. Symb. Comput.* **33**, 425–445 (2002). <https://doi.org/10.1006/jsc.2001.0513>
- Kedlaya, K.S.: Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology. *J. Ramanujan Math. Soc.* **16**, 323–338 (2001). <https://doi.org/10.48550/arXiv.math/0105031>
- Lauder, A.G.B., Wan, D.: Counting points on varieties over finite fields of small characteristic. In: Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography, vol. 44, pp. 579–612. Mathematical Sciences Research Institute Publications (2008). Cambridge University Press, Cambridge. <https://doi.org/10.48550/arXiv.math/0612147>
- Lorenzini, D.: An invitation to arithmetic geometry. In: Graduate Studies in Mathematics, vol. 9. American Mathematical Society (1996). <https://doi.org/10.1090/gsm/009>
- Pila, J.: Frobenius maps of abelian varieties and finding roots of unity in finite fields. *Math. Comp.* **55**, 745–763 (1990). <https://doi.org/10.2307/2008445>
- Schicho, J., Schreyer, F.-O., Weimann, M.: Computational aspects of gonal maps and radical parametrization of curves. *Appl. Algebra Eng. Commun. Comput.* **24**, 313–341 (2013). <https://doi.org/10.1007/s00200-013-0205-0>
- Schoof, R.: Elliptic curves over finite fields and the computation of square roots mod p . *Math. Comput.* **44**, 483–494 (1985). <https://doi.org/10.1090/S0025-5718-1985-0777280-6>
- Stichtenoth, H.: Algebraic function fields and codes. In: Graduate Texts in Mathematics, vol. 254. Springer, Berlin (2009). <https://doi.org/10.1007/978-3-540-76878-4>

23. Tuitman, J.: Counting points on curves using a map to \mathbf{P}^1 , II. *Finite Fields Appl.* **45**, 301–322 (2017). <https://doi.org/10.1016/j.ffa.2016.12.008>

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.