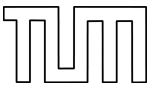

TECHNISCHE UNIVERSITÄT MÜNCHEN
FAKULTÄT FÜR MATHEMATIK

Master's Thesis in Mathematics

Efficient point counting and Monsky-Washnitzer cohomology

Lorenz Panny

Advisor: Prof. Dr. Christian Liedtke
Date: March 15, 2017

I assure the single-handed composition of this thesis only supported by declared resources.

Garching, March 15, 2017

Contents

1 Point counting

Ranging all the way from theoretical investigations which culminated in the famous *Weil conjectures* to highly practical applications in cryptography and coding theory, *point counting* is a central problem of algebraic geometry over finite fields. This thesis covers both extremes: Section 2 gives a theoretical account of the Weil conjectures and presents a partial solution by means of an abstract *Weil cohomology theory*. Section 3 introduces *Monsky-Washnitzer cohomology*, which satisfies some of the properties of Weil cohomologies and has the additional benefit of being very explicit and computable. In Section 4, the theory of Monsky-Washnitzer cohomology will be applied to algorithmically counting the points on a concrete family of varieties, leading to *Kedlaya's algorithm* for hyperelliptic curves.

1.1 Applications in cryptography

Many modern cryptographic primitives are formulated in terms of operations over a finite abelian group, together with certain hardness assumptions. The most prominent example is the *Diffie-Hellman key agreement* [17] invented in 1976, which is still essential to secure communications over the internet.[1] Traditionally (ever since the dawn of public-key cryptography), this algorithm was performed in multiplicative groups of finite fields, but over the years, cryptanalytic progress suggested to migrate to something else. Hence, abelian varieties became relevant to cryptography: In 1986, the existing primitives were reformulated to use elliptic curves over finite fields in place of multiplicative groups [38]. Only three years later, it was proposed to use the Jacobians of hyperelliptic curves (that is, Kummer surfaces) for cryptography, but this has not gained much traction for the following reason: The resulting cryptosystem (for Diffie-Hellman and others) does not stand any chance to be secure unless the group's order has a large prime divisor,[2] so it is important to be able to compute the number of points on the underlying abelian variety. Although well-chosen hyperelliptic curves do actually offer speed advantages over elliptic curves, point counting is still highly expensive for the required cardinalities: In fact, back in 2006, plans to search for a secure hyperelliptic curve over $\mathbb{F}_{2^{127}-1}$ seem to have been abandoned due to the high cost, i. e. practical infeasibility, of point counting [3]. Only as late as 2012, Gaudry and Schost's efforts to find a secure curve of the desired size were fruitful, using a highly optimized algorithm and a considerable amount of computational power [24]. This shows hyperelliptic point counting for cryptographic sizes is still an active and practically relevant research area.

1.2 Algorithms

Over the years, numerous approaches to point counting on algebraic varieties have been developed. Covering each of them in detail would lead us too far astray, hence we only mention the core concepts involved and forward the reader to the referenced publications for details.

For elliptic curves, René Schoof in 1985 presented a point counting algorithm based on computing the *trace of Frobenius* τ, which is intimately related to the number of points,[3] modulo many small primes $\ell \neq p$ until it is uniquely defined by the Chinese remainder theorem. In order to compute $\tau \bmod \ell$, one makes use of a functional equation $\varphi^2 - [\tau]\varphi + q = 0$ of the Frobenius endomorphism φ modulo *division polynomials* ψ_ℓ whose zeroes are just the coordinates of ℓ-torsion points of the curve [43]. Schoof's original algorithm, exhibiting a complexity of $\mathcal{O}(\log^8 q)$ operations, has subsequently been optimized by Noam Elkies and A. O. L. Atkin; the result has a heuristic running time of $\tilde{\mathcal{O}}(\log^4 q)$ [44]. Note this algorithm is the only one which is well-suited for the large-characteristic case.

Another approach suggested by Satoh in 1999 for small primes p involves computing the *canonical lift* of an elliptic curve over a certain extension field of $\mathbb{Q}_p$. This lift is characterized by its endomorphism ring, which is isomorphic to the original curve's — in particular, the canonical lift admits a lift of the Frobenius endomorphism. The core observation then is that the action of the dual of the Frobenius

[1] This is a protocol allowing two communicating parties to securely establish a shared secret over a monitored channel.

[2] Unfortunately, going into the details of this would carry us away too far from the topic of this thesis. However, the interested reader may want to search for the terms 'discrete logarithm problem' and 'Pohlig-Hellman algorithm'.

[3] The motif that linear-algebraic properties of the Frobenius endomorphism are related to point counting is a recurring phenomenon — ultimately, the reason for this lies in the Weil conjectures presented in Section 2.

endomorphism on a holomorphic differential is related to the zeta function. For fixed p, the resulting algorithm has a running time of $\tilde{\mathcal{O}}(n^{2.5})$ operations, where n is the extension degree $\log_p q$. [42]

When the genus of the curve grows, the algorithm zoo quickly becomes smaller. Schoof's algorithm should *in theory* generalize to curves of arbitrary genus, but this has only been made practical for hyperelliptic curves of genus 2 by Gaudry and Schost. [24]

For fields of small characteristic, there is *Kedlaya's algorithm* for hyperelliptic curves, based on a Frobenius action on Monsky-Washnitzer cohomology. The theory behind Kedlaya's algorithm works quite generally, namely for arbitrary smooth affine varieties, but in practice the necessary prerequisites and estimates are very tedious to obtain for generic examples.

Tuitman in 2014 generalized Kedlaya's algorithm to 'almost all' curves, i. e., those for which a 'good' lift to characteristic zero can be found. The time complexity for a plane curve given by a polynomial of degrees d_x, d_y is stated as $\tilde{\mathcal{O}}(pd_x^4 d_y^3 n^3)$. [47]

Despite considerable research effort, there is currently no known algorithm for point counting on curves whose time complexity is simultaneously polynomial in the bit length of p, the extension degree n, and the genus g.

Finally, for varieties of higher dimension, Abbott, Kedlaya, and Roe in 2006 generalized Kedlaya's algorithm to smooth hypersurfaces $X \subseteq \mathbb{P}^{m+1}\mathbb{F}_q$. In that case, the complement $\mathbb{P}^{m+1}\mathbb{F}_q \setminus X$ is smooth and affine, thus Monsky-Washnitzer cohomology can be employed to count its points. However, the running time of this algorithm grows very quickly, hence it is only practical for very small examples. [1]

A more performant algorithm was presented in 2004 by Lauder, who embeds the target variety into a one-parameter family and considers the Frobenius action on all of them at once. The action on the target variety is then computed from the action on a simple variety in the family as the solution of a differential equation, the *Gauß-Manin connection*. [37]

For an overview on p-adic point counting, see Tuitman [46].

Notation. Throughout this document, we fix the following notation:

- Unless otherwise noted, the term *ring* shall refer to a commutative unitary ring.
- The natural numbers $\mathbb{N}$ include zero.
- Usually, p will denote a prime number and $q = p^e$ one of its powers. ℓ is a prime distinct from p.
- The notation $\mathcal{O}$ shall always refer to *bit* complexity if not stated otherwise. The notation $\tilde{\mathcal{O}}$ ignores logarithmic terms; see von zur Gathen and Gerhard [51, Definition 25.8].
- The symbols $\square/\triangle/\circ$ mark the end of a proof/example/remark.

2 Weil cohomology and the Weil conjectures

One of the most important cornerstones of 20th-century algebraic geometry are the *Weil conjectures*, which were put forward in 1949 by André Weil and are nowadays proven. These are statements about the *zeta functions* of algebraic varieties over a finite field — a generating function involving the number of points of the variety under field extensions. An intuitive interpretation of the Weil conjectures is that the growth of the number of points of a variety over an extension field is 'well-behaved', in a sense made precise below.

Until its resolution, the proof of the Weil conjectures was one of the major problems in algebraic geometry. Weil himself proved them for curves and abelian varieties in 1948 [52], but the general case took over twenty years to be solved. Dwork proved the rationality and the functional equation of the zeta function in 1960 using p-adic methods [18]. The general case for the other two conjectures remained open until Grothendieck, building on previous results of Serre, introduced the ℓ-adic cohomology (for a short summary, see Section 2.2.2) to algebraic geometry in order to tackle these problems. This shows the Weil conjectures were a strong catalyst to the development of modern algebraic geometry: In fact, the abstract proof of the Lefschetz trace formula given in Section 2.2 was known before a suitable cohomology theory had even been constructed! Building on those new ideas, Deligne finally proved the analogue of the Riemann hypothesis in 1974 [13], thus completing the solution of the Weil conjectures.

Besides being of theoretical interest, the cohomological proof of the Weil conjectures has clearly served as an inspiration for current point-counting algorithms, such as Kedlaya's algorithm described in Section 4. Moreover, such algorithms typically make use of the guarantees obtained from the Weil conjectures as part of their correctness proof (for instances of this, see Section 4.6 and Lemma 50).

The historical developments revolving around various approaches to, and proofs of, the Weil conjectures are well summarized in Hartshorne [26, Appendix C.2]. Readers with a deeper interest in the Weil cohomology assumptions and their implications may consult the axiomatic treatment of Kleiman [35].

2.1 The Weil conjectures

We shall now define the zeta function and state the Weil conjectures. Note that we content ourselves with the smooth projective case: Some of the definitions or results generalize to singular or non-projective schemes, but we omit those to simplify the presentation. The material in this section can be found in Hartshorne [26, Appendix C].

Definition 1. The *zeta function* of an algebraic variety X over a finite field $\mathbb{F}_q$ is the formal power series

$$Z(X;t) = \exp\left(\sum_{r=1}^{\infty} \frac{\#X(\mathbb{F}_{q^r})}{r} \cdot t^r\right) \quad \in \mathbb{Q}[\![t]\!].$$

Theorem 2 (Weil conjectures [26, Appendix C.1])**.** Let X be a smooth projective variety of dimension n defined over a finite field $\mathbb{F}_q$. Then

(i) **Rationality:** $Z(X;t)$ is a rational function, that is, $Z(X;t) \in \mathbb{Q}(t)$.

(ii) **Functional equation:** If E is the self-intersection number of the diagonal $\Delta \subseteq X \times X$, then

$$Z(X;q^{-n}t^{-1}) = \pm q^{nE/2} t^E Z(X;t).$$

(iii) **Analogue of the Riemann hypothesis:** The zeta function can be written as

$$Z(X;t) = \frac{P_1(t) \cdot P_3(t) \cdots P_{2n-1}(t)}{P_0(t) \cdot P_2(t) \cdots P_{2n}(t)}$$

with $P_0(t) = 1 - t$, $P_{2n}(t) = 1 - q^n t$, and such that each $P_i(t)$ has integer coefficients and is of the form $P_i(t) = \prod_j (1 - \alpha_{ij} t)$ for algebraic integers α_{ij} of absolute value $q^{i/2}$.

(iv) **Betti numbers:** Let $b_i(X) := \deg P_i(t)$. Then $E = \sum_i (-1)^i b_i$. Moreover, if X is a reduction of a variety Y over an algebraic number ring R modulo a prime ideal, then b_i equals the ith Betti number of the analytification of $Y \times_R \mathbb{C}$.

2.2 Weil cohomology theories

In this section, we introduce the formalism of a *Weil cohomology theory*, which is a general set of assumptions that leads to a proof of the Weil conjectures. We shall then derive the Lefschetz trace formula (a major intermediate step in proving the Weil conjectures) from such a cohomology, prove the rationality of the zeta function, and close the section by giving an overview of the most prominent examples of algebraic Weil cohomologies.

Definition 3 [26, Appendix A.1]**.** Let X be a smooth quasi-projective variety. An element of the free abelian group generated by the closed irreducible subvarieties of codimension r of X is called *cycle of codimension* r on X. Note that Weil divisors on a subvariety of X form cycles on X: Define $A^r(X)$ to be the group of codimension-r cycles on X modulo the subgroup generated by principal divisors on subvarieties of X. Finally, the *Chow group* of X is the graded group $A^*(X) := \bigoplus_{r\geq 0} A^r(X)$.

Theorem 4 [26, Section A.1]**.** There exists an *intersection product* denoted '.' on $A^*(X)$ that makes it a graded ring, the *Chow ring*. It satisfies the following properties:

(i) Each map $f\colon X \to Y$ between smooth quasi-projective varieties induces ring homomorphisms

$$f_*\colon A^*(X) \to A^*(Y) \qquad\qquad f^*\colon A^*(Y) \to A^*(X)$$

such that both of these assignments are functorial.

(ii) If f is proper, then, for any $x \in A^*(X)$ and $y \in A^*(Y)$,

$$f_*(x.f^*(y)) = f_*(x).y.$$

Intuitively, the product of two cycles is given by the subvarieties in the intersection, with the appropriate (signed — reflecting orientation) multiplicities.

Definition 5 [11, Chapter 4, §1, Section 1.7]**.** Let k and K denote fields with $\operatorname{char} K = 0$. By a *Weil cohomology* we shall refer to a family of contravariant functors

$$H^i\colon \{\text{ smooth projective varieties over } k\} \longrightarrow \{\text{ finite-dimensional vector spaces over } K\}$$

that satisfy the following list of properties. Let $H^*(X)$ denote the graded K-vector space $\bigoplus_{i\geq 0} H^i(X)$. We occasionally write f^* for a map $H^*(f)$ induced on cohomology. Throughout the list, let X and Y be smooth projective varieties over k of dimensions n and m.

(i) $H^i(X) = 0$ for $i < 0$ or $i > 2n$.

(ii) **Cup product:** There are maps $\smile\colon H^i(X) \times H^j(X) \to H^{i+j}(X)$ which render $H^*(X)$ a graded-commutative ring.[1]

(iii) **Cycle map:** There is a grade-doubling morphism of rings $\mathrm{cl}\colon A^*(X) \to H^*(X)$ that commutes with f^* for f an endomorphism of X.[2]

(iv) **Orientation map:** There exists an isomorphism $\eta\colon H^{2n}(X) \xrightarrow{\sim} K$ such that $\eta(\mathrm{cl}(P)) = 1$ for any closed point $P \in X$.

(v) **Poincaré duality:** The cup-product pairing $H^i(X) \times H^{2n-i}(X) \to H^{2n}(X) \cong K$ is non-degenerate.

(vi) **Künneth formula:** The map $H^*(X) \otimes_K H^*(Y) \to H^*(X \times Y)$, $\alpha \otimes \beta \mapsto p^*(\alpha) \smile q^*(\beta)$, where p and q denote the projections from $X \times Y$, is an isomorphism of rings.

(vii) **Gysin map:** Let $r = m - n$. For any morphism $f\colon X \to Y$, Poincaré duality induces maps $f_*\colon H^i(X) \to H^{i-2r}(Y)$ going the 'wrong' direction. If f is an endomorphism, those morphisms commute with the cycle map cl.

[1] That is: If $x \in H^i(X)$ and $y \in H^j(X)$, then $yx = (-1)^{ij}xy$. Note that in this context, the term 'ring' does *not* refer to a commutative ring (contrary to the convention we established for the rest of the document).

[2] Recall that outside of cl, the notations f_* and f^* refer to pushforwards and pullbacks on the Chow ring A^*.

Our first goal now is to prove from these assumptions a *Lefschetz trace formula*, i. e., a representation of the number of fixed points of an endomorphism in terms of traces of the induced morphisms on cohomology. First, we state an additional lemma (without proof):

Lemma 6 [39, Lemma 25.4]**.** For any morphism $\varphi\colon X \to Y$ of smooth projective varieties and $y \in H^*(y)$,

$$p_*(\mathrm{cl}(\Gamma_\varphi) \smile q^*(y)) = \varphi^*(y).$$

This auxiliary lemma enables us to prove the Lefschetz trace formula for a general Weil cohomology theory. The following proof, modulo minor clarifications, can be found in sources such as Milne [39, Section 25], de Jong [12], or Mustaţă [41, Theorem 4.7].

Theorem 7 (Lefschetz formula)**.** Let H^* denote any Weil cohomology. Then, for an endomorphism φ of a smooth projective variety X of dimension n, the intersection number[1] of Γ_φ and the diagonal Δ is

$$(\Gamma_\varphi \cdot \Delta) = \sum_{i=0}^{2n} (-1)^i \operatorname{tr} H^i(\varphi).$$

Proof. Note that Γ_φ and Δ have codimension n in $X \times X$, hence $\mathrm{cl}(\Gamma_\varphi)$ and $\mathrm{cl}(\Delta)$ lie in $H^{2n}(X \times X)$, which is (by the Künneth formula) isomorphic to $H^{2n}(X) \otimes_K H^{2n}(X)$. For each $i \in \{0 \ldots 2n\}$, fix a basis $\{e_t^{(i)}\}$ of $H^i(X)$ and let $\{f_t^{(2n-i)}\}$ denote the corresponding dual basis[2] of $H^{2n-i}(X)$, such that $f_u^{(2n-i)} \smile e_t^{(i)} = \delta_{tu}$. Then $\mathrm{cl}(\Gamma_\varphi)$ can be written as $\sum_{i=0}^{2n} \sum_t a_t^{(i)} \otimes f_t^{(2n-i)}$ with $a_t^{(i)} \in H^i(X)$. Using Lemma 6 and the Künneth formula, we get

$$\varphi^*(e_t^{(i)}) = p_*(\mathrm{cl}(\Gamma_\varphi) \smile q^*(e_t^{(i)})) = \sum_{j=0}^{2n} \sum_u p_*(p^*(a_u^{(j)})) \smile p_*(q^*(f_u^{(2n-j)} \smile e_t^{(i)})).$$

Clearly, the element $\alpha := f_u^{(2n-j)} \smile e_t^{(i)}$ is of degree $2n + i - j$, hence $p_*(q^*(\alpha)) \in H^{i-j}(X)$. The first observation shows $\alpha = 0$ for $i > j$, while the second implies $p_*(q^*(\alpha)) = 0$ for $i < j$. In other words, only the summands with $i = j$ contribute to the sum. By construction, we have $f_u^{(2n-i)} \smile e_t^{(i)} = \delta_{tu}$, thus $\varphi^*(e_t^{(i)}) = a_t^{(i)}$, and therefore

$$\mathrm{cl}(\Gamma_\varphi) = \sum_{i=0}^{2n} \sum_t \varphi^*(e_t^{(i)}) \otimes f_t^{(2n-i)}.$$

On the other hand, since $\Delta = \Gamma_{\mathrm{id}}$ and $\smile$ is graded-commutative,

$$\mathrm{cl}(\Delta) = \sum_{i=0}^{2n} \sum_t e_t^{(i)} \otimes f_t^{(2n-i)} = \sum_{i=0}^{2n} (-1)^i \sum_t f_t^{(2n-i)} \otimes e_t^{(i)}.$$

Combining those two, we get (using the Künneth formula)

$$\begin{aligned}
\mathrm{cl}(\Gamma_\varphi.\Delta) &= \mathrm{cl}(\Gamma_\varphi) \smile \mathrm{cl}(\Delta) \\
&= \left(\sum_{i=0}^{2n} \sum_t \varphi^*(e_t^{(i)}) \otimes f_t^{(2n-i)} \right) \smile \left(\sum_{j=0}^{2n} (-1)^i \sum_u f_u^{(2n-j)} \otimes e_u^{(j)} \right) \\
&= \sum_{i=0}^{2n} \sum_{j=0}^{2n} (-1)^i \sum_t \sum_u p^*(\varphi^*(e_t^{(i)})) \smile q^*(f_t^{(2n-i)}) \smile p^*(f_u^{(2n-j)}) \smile q^*(e_u^{(j)}) \\
&= \sum_{i=0}^{2n} \sum_{j=0}^{2n} (-1)^i \sum_t \sum_u (-1)^{ij} p^*(\varphi^*(e_t^{(i)}) \smile f_u^{(2n-j)}) \smile q^*((-1)^{ij} e_u^{(j)} \smile f_t^{(2n-i)}) \\
&= \sum_{i=0}^{2n} \sum_{j=0}^{2n} (-1)^i \sum_t \sum_u p^*(\varphi^*(e_t^{(i)}) \smile f_u^{(2n-j)}) \smile q^*(e_u^{(j)} \smile f_t^{(2n-i)}).
\end{aligned}$$

[1] If Γ_φ and Δ intersect *transversely* (which roughly means their tangent spaces at the intersection point only meet at that point), this equals the number of intersection points.

[2] Poincaré duality and finite-dimensionality yield isomorphisms $H^i(X) \cong \mathrm{Hom}(H^{2n-i}(X), K) \cong H^{2n-i}(X)$. If $\{e_t\}$ is a basis of $H^i(X)$, and $\{f_t\}$ the corresponding basis of $H^{2n-i}(X)$ under that isomorphism, then $e_t \smile f_u = \delta_{tu}$ by construction.

As above, $e_u^{(j)} \smile f_t^{(2n-i)} \in H^{2n+j-i}(X)$ vanishes if $i < j$, and similarly, $\varphi^*(e_t^{(i)}) \smile f_u^{(2n-j)} \in H^{2n+i-j}(X)$ vanishes for $i > j$; thus, we are left with

$$\begin{aligned}\mathrm{cl}(\Gamma_\varphi.\Delta) &= \sum_{i=0}^{2n}(-1)^i \sum_t \sum_u p^*(\varphi^*(e_t^{(i)}) \smile f_u^{(2n-i)}) \smile q^*(\overbrace{e_u^{(i)} \smile f_t^{(2n-i)}}^{=(-1)^i\delta_{tu}}) \\ &= \sum_{i=0}^{2n}\sum_t (\varphi^*(e_t^{(i)}) \smile f_t^{(2n-i)}) \otimes 1.\end{aligned}$$

Therefore, if $H^i(\varphi)(e_t^{(i)}) = \varphi^*(e_t^{(i)}) = \sum_u a_{tu}^{(i)} e_u^{(i)}$ for some set of coefficients $a_{tu}^{(i)} \in K$, then

$$\begin{aligned}\mathrm{cl}(\Gamma_\varphi.\Delta) &= \sum_{i=0}^{2n}\sum_t\sum_u a_{tu}^{(i)} e_u^{(i)} \smile f_t^{(2n-i)} \otimes 1 = \sum_{i=0}^{2n}\sum_t\sum_u (-1)^i a_{tu}^{(i)}\delta_{tu} \otimes 1 \\ &= \sum_{i=0}^{2n}(-1)^i \sum_t a_{tt}^{(i)} \otimes 1 = \sum_{i=0}^{2n}(-1)^i \operatorname{tr} H^i(\varphi) \otimes 1.\end{aligned}$$

Applying the orientation map η to both sides yields the claim. □

Using the following linear-algebraic lemma, the Lefschetz trace formula almost immediately implies the first part of the Weil conjectures, the rationality of the zeta function:

Lemma 8 [26, Appendix C, Lemma 4.1]**.** If V is a finite-dimensional vector space over a field k and φ an endomorphism of V, then

$$\det(1 - t\varphi)^{-1} = \exp\left(\sum_{r=1}^{\infty} \frac{\operatorname{tr}\varphi^r}{r} \cdot t^r\right).$$

Theorem 9 [41, Theorem 4.11]**.** Let X be a smooth projective variety of dimension n over $\mathbb{F}_q$ with q-power Frobenius endomorphism $\sigma\colon X \to X$. Then $Z(X;t)$ is a rational function; in particular,

$$Z(X;t) = \frac{P_1(t) \cdot P_3(t) \cdots P_{2n-1}(t)}{P_0(t) \cdot P_2(t) \cdots P_{2n}(t)} \in \mathbb{Q}(t),$$

and each $P_i(t)$ equals $\det\left(1 - tH^i(\sigma)\right)$, where $H^i(\sigma)$ is σ's induced endomorphism on $H^i(X)$.

Proof. It is clear that for all $r \in \mathbb{N}$, the $\mathbb{F}_{q^r}$-rational points of X are precisely those fixed by φ^r, hence

$$\#X(\mathbb{F}_{q^r}) = |\{\, x \in X \mid \varphi^r(x) = x \,\}|.$$

The graph of φ and the diagonal intersect transversely [41, Prop. 2.4], therefore Theorem 7 implies

$$\#X(\mathbb{F}_{q^r}) = \sum_{i=0}^{2n}(-1)^i \operatorname{tr} H^i(\varphi^r).$$

Substituting this into the definition of the zeta function yields

$$\begin{aligned}Z(X;t) &= \exp\left(\sum_{r=1}^{\infty} \frac{\#X(\mathbb{F}_{q^r})}{r} \cdot t^r\right) = \exp\left(\sum_{r=1}^{\infty}\sum_{i=0}^{2n}(-1)^i \cdot \frac{\operatorname{tr} H^i(\sigma^r)}{r} \cdot t^r\right) \\ &= \prod_{i=0}^{2n}\exp\left((-1)^i \sum_{r=1}^{\infty} \frac{\operatorname{tr} H^i(\sigma^r)}{r} \cdot t^r\right) = \prod_{i=0}^{2n}\exp\left(\sum_{r=1}^{\infty} \frac{\operatorname{tr}\left(H^i(\sigma)^r\right)}{r} \cdot t^r\right)^{(-1)^i},\end{aligned}$$

which by Lemma 8 equals

$$Z(X;t) = \prod_{i=0}^{2n}\left(\det\left(1 - tH^i(\sigma)\right)\right)^{(-1)^{i+1}}.$$

Thus, setting $P_i(t) := \det\left(1 - tH^i(\sigma)\right) \in K[t]$ proves the claim $Z(X;t) = \prod_{i=0}^{2n} P_i(t)^{(-1)^{i+1}}$, which lies in $K(t)$. On the other hand, $Z(X;t) \in \mathbb{Q}[\![t]\!]$ by construction, hence the claim $Z(X;t) \in \mathbb{Q}(t)$ follows. □

In the preceding section, we have used the general axioms of a Weil cohomology to prove the rationality of the zeta function of a smooth projective variety over a finite field. The proof of the functional equation easily follows from Poincaré duality along with a little linear algebra. The properties of the Betti numbers are not too hard, but require a few more preparations. Both are not immediately relevant to the topic of this thesis, so we shall omit them. Finally, the analogue of the Riemann hypothesis is unfortunately a lot harder than the other three conjectures: It was solved by Deligne more than ten years after the rationality of the zeta function had been shown, and the methods used in the proof are quite a bit more advanced than those demonstrated above [13].

2.2.1 Algebraic de Rham cohomology

The most obvious construction of a cohomology for algebraic varieties is the *algebraic de Rham cohomology*, which mimicks the classical de Rham cohomology of differentiable manifolds. One replaces the sheaf of smooth differential forms by a sheaf of formally constructed modules of *Kähler differentials*, and takes the cohomology of the complex given by a formal analogue of the exterior derivative.

For smooth and proper varieties over a field of characteristic zero, this cohomology coincides with the classical de Rham cohomology. However, for varieties over a finite field, the resulting functor is not well-behaved: The exterior derivative annihilates pth powers, hence algebraic de Rham cohomology is typically infinite-dimensional [31, Section 2.1]. Moreover, even if the cohomology was 'good', the traces in the Lefschetz formula would still be elements of $\mathbb{F}_q$, i. e., reduced modulo p, hence unsuitable for point counting. Therefore, algebraic de Rham cohomology is not a Weil cohomology in positive characteristic.

2.2.2 ℓ-adic cohomology

The *ℓ-adic cohomology* outlined in this section is the historically first example of a Weil cohomology theory over a finite field. The construction proceeds as follows, taking the *étale cohomology* $H^*_{\text{ét}}$ introduced by Grothendieck [25] as a starting point. Let X be a smooth projective variety over $\mathbb{F}_q$ and $\ell \neq p$ a prime. For all $n \in \mathbb{N}$, the cohomology group $H^i_{\text{ét}}(X, \mathbb{Z}/(\ell^n))$ is a module over the ring $\mathbb{Z}/(\ell^n)$, and they form an inverse system

$$H^i_{\text{ét}}(X, \mathbb{Z}/(\ell)) \leftarrow H^i_{\text{ét}}(X, \mathbb{Z}/(\ell^2)) \leftarrow H^i_{\text{ét}}(X, \mathbb{Z}/(\ell^3)) \leftarrow H^i_{\text{ét}}(X, \mathbb{Z}/(\ell^4)) \leftarrow H^i_{\text{ét}}(X, \mathbb{Z}/(\ell^5)) \leftarrow \cdots.$$

The *ℓ-adic cohomology* of X is the inverse limit of that system modulo torsion, i. e.,

$$H^i(X, \mathbb{Q}_\ell) := \left(\varprojlim_n H^i_{\text{ét}}(X, \mathbb{Z}/(\ell^n))\right) \otimes_{\mathbb{Z}_\ell} \mathbb{Q}_\ell.$$

Note that taking cohomology does *not* commute with inverse limits, hence this is different from the étale cohomology $H^i_{\text{ét}}(X, \mathbb{Q}_\ell)$ of X with respect to the constant sheaf $\mathbb{Q}_\ell$.

Theorem 10 [14]**.** ℓ-adic cohomology is a Weil cohomology with $k = \mathbb{F}_q$ and $K = \mathbb{Q}_\ell$.

2.2.3 Crystalline cohomology

In 1974, Berthelot presented *crystalline cohomology*, a p-adic Weil cohomology for smooth proper varieties over finite fields. [4] This theory is strongly related to the cohomology of Monsky and Washnitzer presented in Section 3, in the sense that they share a common generalization: *Rigid cohomology*, introduced in 1986 by Berthelot, applies to arbitrary smooth varieties and reduces to crystalline cohomology in the proper case and to Monsky-Washnitzer cohomology in the affine case. [5] For a thorough summary of the historical development of p-adic cohomology, we refer the reader to Section 1.2 of Kedlaya [32].

Theorem 11 [4]**.** Crystalline cohomology is a Weil cohomology with $k = \mathbb{F}_q$ and $K = \mathbb{Q}_q$.

3 Monsky-Washnitzer cohomology

In the previous sections, we have demonstrated that a sufficiently well-behaved cohomology theory for algebraic varieties over finite fields admits a Lefschetz trace formula, that is, an expression for the number of rational points on the variety in terms of traces of a Frobenius action on the cohomology spaces. However, the classical ℓ-adic cohomology à la Grothendieck is by no means practically computable — while it lends itself nicely for theoretical problems (such as the proof of the Weil conjectures), it is unhelpful for concrete calculations.

Therefore, *Monsky-Washnitzer cohomology* enters the picture: This is a cohomology theory for smooth *affine* varieties over finite fields $\mathbb{F}_q$, whose cohomology groups are finite-dimensional vector spaces over a certain characteristic-zero field $\mathbb{Q}_q$ (which we'll soon learn more about). As vector spaces, the cohomology spaces are isomorphic to the much simpler algebraic de Rham cohomology of a lift to $\mathbb{Q}_q$, but the latter generally does not admit a Frobenius lift.[1] In contrast, Monsky-Washnitzer cohomology does come with a Frobenius lift, and this action satisfies a Lefschetz trace formula. Even though computing in Monsky-Washnitzer cohomology involves manipulating infinite objects, it can be approximated well enough in finite time to give a provably correct result for the number of points.

First, we need to introduce some new objects in order to define Monsky-Washnitzer cohomology in Section 3.3. To illustrate the theory, we shall afterwards compute cohomology groups and work out the Lefschetz trace formula for a few examples.

3.1 The Witt ring

The main idea of Monsky-Washnitzer cohomology (and its generalizations) is to lift a smooth variety from characteristic p to a field of characteristic zero in order to obtain a useful cohomology theory. However, we cannot make much use of arbitrary characteristic-zero fields: Simply lifting the coefficients of, for instance, an elliptic curve $y^2 = x^3 + ax + b$ from $\mathbb{F}_p$ to $\mathbb{Q}$ by taking any representant will give a very different object than the original curve over a finite field. Intuitively, $\mathbb{Q}$ is unsuitable for our purposes since p does not play a special rôle in $\mathbb{Z}$ compared to any other prime number, hence lifting to $\mathbb{Q}$ fails at exposing much of the information that is specific to the variety over $\mathbb{F}_p$. Instead, it seems promising to lift to the p-adic numbers $\mathbb{Q}_p$: This is a field of characteristic zero, and its ring of integers $\mathbb{Z}_p$ is a discrete valuation ring with unique maximal ideal (p) and residue field $\mathbb{Z}_p/(p) \cong \mathbb{F}_p$. In that sense, the p-adics 'interpolate' between fields of positive characteristic and characteristic zero. [2, Section 3.4.1]

However, the p-adic numbers only make sense for *prime* p, and depending on the particular construction used to obtain them, there might not necessarily be an obvious analogue for composite fields. The *ring of Witt vectors* $\mathbb{Z}_q$ of $\mathbb{F}_q$ presented in this section provides a very natural generalization of the p-adic integers to a prime power $q = p^e$; in particular, the Witt vectors form a discrete valuation ring of characteristic zero with maximal ideal (p) and residue field $\mathbb{F}_q$.

Fix a prime number p. We shall abbreviate a countable sequence $\{X_0, X_1, \dots\}$ of variables or polynomials by $\underline{X}$. The general structure and exposition of this section roughly follows Serre [45, §II.6].

Definition 12. Define the family of *Witt polynomials* $(w_n)_{n\in\mathbb{N}}$ by

$$w_n := \sum_{j=0}^{n} p^j X_j^{p^{n-j}} \in \mathbb{Z}[\underline{X}].$$

Lemma 13. For any $f \in \mathbb{Z}[X, Y]$, there exists a unique sequence of polynomials $\underline{F} \in \mathbb{Z}[\underline{X}, \underline{Y}]$ such that

$$w_n(\underline{F}) = f(w_n(\underline{X}), w_n(\underline{Y})).$$

Proof. See Serre [45, §II.6, Theorem 6], or for a more elementary proof Hazewinkel [28, Theorem 5.2] □

[1] For instance, an endomorphism φ of a smooth projective curve of genus $g > 1$ (such as a hyperelliptic curve) must satisfy the Hurwitz formula $(2g-2)(\deg\varphi - 1) + \deg R = 0$, where R is the *ramification divisor* of φ. By construction, that divisor is non-negative; thus necessarily $\deg\varphi \leq 1$. However, the q-power Frobenius endomorphism is a morphism of degree q. [19]

Definition 14. The *(p-typical) Witt ring* or *ring of Witt vectors* $W_p(R)$ of a ring R is the set $R^{\mathbb{N}}$ of countable sequences of elements of R with with addition and multiplication given by the unique polynomial sequences $\underline{S}, \underline{P} \in \mathbb{Z}[\underline{X}, \underline{Y}]$ defined (via Lemma 13) by

$$w_n(\underline{S}) = w_n(\underline{X}) + w_n(\underline{Y}); \qquad w_n(\underline{P}) = w_n(\underline{X}) \cdot w_n(\underline{Y}).$$

Note this implies that each w_n is a ring homomorphism from $W_p(R)$ to R. Unsurprisingly, the zero and one elements of this ring are given by the vectors $(0, 0, 0, \dots)$ and $(1, 0, 0, \dots)$.

Explicit addition and multiplication laws for the Witt vectors can be obtained simply by rearranging the defining equations, as shown in the following lemmas.

Lemma 15. The polynomials $S_n \in \mathbb{Z}[\underline{X}, \underline{Y}]$ may be recursively computed as

$$S_0 = X_0 + Y_0$$
$$S_n = X_n + Y_n + \sum_{i=0}^{n-1} p^{i-n}\big(X_i^{p^{n-i}} + Y_i^{p^{n-i}} - S_i^{p^{n-i}}\big).$$

Proof. Definition 14 requires that

$$w_n(\underline{S}) = \sum_{i=0}^{n} p^i S_i^{p^{n-i}} = \sum_{i=0}^{n} p^i X_i^{p^{n-i}} + \sum_{i=0}^{n} p^i Y_i^{p^{n-j}} = w_n(\underline{X}) + w_n(\underline{Y}).$$

Solving this equation for $p^n S_n$ is easy:

$$p^n S_n = p^n X_n + p^n Y_n - \sum_{i=0}^{n-1} p^i\big(X_i^{p^{n-i}} + Y_i^{p^{n-i}} - S_i^{p^{n-i}}\big),$$

and Lemma 13 asserts the right-hand side is divisible by p^n. □

Lemma 16. The polynomials $P_n \in \mathbb{Z}[\underline{X}, \underline{Y}]$ may be recursively computed as

$$P_0 = X_0 \cdot Y_0$$
$$P_n = p^{-n} \cdot w_n(\underline{X}) \cdot w_n(\underline{Y}) - \sum_{i=0}^{n-1} p^{i-n} P_i^{p^{n-i}}.$$

Proof. According to Definition 14,

$$w_n(\underline{P}) = \sum_{i=0}^{n} p^i P_i^{p^{n-i}} = w_n(\underline{X}) \cdot w_n(\underline{Y}).$$

Hence

$$p^n P_n = w_n(\underline{X}) \cdot w_n(\underline{Y}) - \sum_{i=0}^{n-1} p^i P_i^{p^{n-i}},$$

and as before, Lemma 13 guarantuees the right-hand side is a multiple of p^n. □

The previous lemmas hint the intuition behind the Witt ring's rather technical definition: Witt vectors formalize the idea of taking the information that would be lost (due to positive characteristic) when adding or multiplying one component of the vectors, and *carrying* that 'overflow' into the next component. That this resembles the p-adic integers is no accident; in fact, the Witt ring is a proper generalization of that construction:

Lemma 17. The p-typical Witt ring of a prime field $\mathbb{F}_p$ is the ring of p-adic integers $\mathbb{Z}_p$.

Proof. Compare the carrying rules for $\mathbb{Z}_p$ to those of $W_p(\mathbb{F}_p)$. □

Definition 18. As justified by the previous lemma, let $\mathbb{Z}_q$ denote the p-typical Witt ring of $\mathbb{F}_q$, and let $\mathbb{Q}_q = \operatorname{Quot} \mathbb{Z}_q$ denote its field of fractions.

Lemma/Definition 19. The Witt ring construction W_p is a covariant endofunctor of the category of rings; in particular, any map $\varphi\colon R \to S$ of rings lifts to a morphism

$$W_p(R) \to W_p(S),\ (x_0, x_1, x_2, \dots) \mapsto (\varphi(x_0), \varphi(x_1), \varphi(x_2), \dots).$$

If R is a ring of characteristic p, this allows us to lift p-power Frobenius endomorphism from R to its ring of Witt vectors:

$$\sigma\colon\ W_p(R) \to W_p(R),\ (x_0, x_1, x_2, \dots) \mapsto (x_0^p, x_1^p, \dots).$$

Note that by construction, the Witt vector Frobenius endomorphism of $\mathbb{Z}_q$ is an automorphism of order $\log_p q$ of $\mathbb{Z}_q$. In particular, as a consequence to Fermat's little theorem, the Frobenius endomorphism of $W_p(\mathbb{F}_p) = \mathbb{Z}_p$ is the identity.

Theorem 20 [45]**.** The Witt ring $\mathbb{Z}_q$ is a complete discrete valuation ring of characteristic zero. Its unique maximal ideal is generated by the uniformizing element p, and the resulting residue field $\mathbb{Z}_q/(p)$ is the finite field $\mathbb{F}_q$. The field of fractions $\mathbb{Q}_q$ is an unramified extension of the p-adic numbers $\mathbb{Q}_p$ of degree $\log_p q$, and its ring of integers is $\mathbb{Z}_q$.

Note that implementing the Witt vectors on a computer using the formulas from Lemmas 15 and 16 is not the best way to proceed: Hubrechts [29] provides efficient representations and algorithms for the rings $\mathbb{Z}_q$ and $\mathbb{Q}_q$.

3.2 Weak completion

Now that the Witt vectors provide a good ring of characteristic zero to lift to, we can tackle the next problem: After a variety has been lifted from $\mathbb{F}_q$ to $\mathbb{Z}_q$, naïvely raising each variable to the pth power and applying the Witt vector Frobenius to the coefficients, so as to obtain a Frobenius lift, generally does not result in a well-defined ring endomorphism. For example, the smooth affine curve $y = x + 1$ over $\mathbb{F}_q$ lifts to the spectrum of $\mathbb{Z}_q[x, y]/(y - x - 1)$, but in the latter, we have

$$y^p = (x+1)^p = \sum_{i=0}^{p} \binom{p}{i} x^i \neq x^p + 1$$

since $\operatorname{char} \mathbb{Z}_q = 0$. In a sense, one needs to mitigate the fact that $p \neq 0$ in $\mathbb{Z}_q$ by adding 'correction terms' to the Frobenius lift. Since finitely many such error terms are not enough in general, we shall pass into a power series ring. However, simply taking the p-adic completion is not desirable for our purposes because it would distort the de-Rham-like cohomology groups we will define later[1] — it turns out that a suitable construction is the *weak completion*, which will be introduced in this section. This material can be found in the original publication by Monsky and Washnitzer [40], as well as (in part) in van der Put's survey article [49].

Definition 21 [40, Section 1]**.** Let R be a Noetherian ring, $I \subseteq R$ an ideal, and A an R-algebra with I-adic completion $A^\infty = \varprojlim_j A/I^jA$. We define the ($I$-adic) *weak completion* of A to be the subalgebra $A^\dagger$ of A^∞ consisting of those elements which can be written as $\sum_{j=0}^{\infty} f_j(a_1 \dots a_n)$ for some $a_1 \dots a_n \in A$, with each $f_j \in I^j \cdot R[x_1 \dots x_n]$, and such that there exists a bound C with all $\deg f_j \leq C(j+1)$.

Lemma 22. The p-adic weak completion of the polynomial ring $\mathbb{Z}_q[x_1 \dots x_n]$ is the $\mathbb{Z}_q$-algebra of *overconvergent power series*

$$\mathbb{Z}_q\langle x_1 \dots x_n\rangle^\dagger := \left\{ \sum_{\alpha \in \mathbb{N}^n} c_\alpha x^\alpha \in \mathbb{Z}_q[\![x_1 \dots x_n]\!] \;\middle|\; \exists \gamma \in \mathbb{R},\ \varrho > 0.\ \forall \alpha \in \mathbb{N}^n.\ \nu_p(c_\alpha) \geq \gamma + \varrho|\alpha| \right\}.$$

The condition on the p-adic valuations roughly means that $\nu_p(c_\alpha)$ grows at least linearly with $|\alpha|$. In terms of the p-adic absolute value $|\cdot|_p$ on $\mathbb{Z}_q$, it can be rephrased as follows: There exist $C > 0$ and $R \in (0, 1)$ such that $|c_\alpha|_p \leq CR^{|\alpha|}$ for all $\alpha \in \mathbb{N}^n$.

[1] This issue is best exemplified by a p-adically convergent series such as $\sum_{i=0}^{\infty} p^i x^{p^i - 1}$, whose integral $\sum_{i=0}^{\infty} x^{p^i}$ diverges.

Proof. Consider a series $Z = \sum_{\alpha\in\mathbb{N}^n} c_\alpha x^\alpha$ in the weak completion of $\mathbb{Z}_q[x_1...x_n]$. By definition, Z admits a representation as $\sum_{j=0}^\infty f_j(a_1...a_m)$ with $a_i \in \mathbb{Z}_q[x_1...x_n]$ and each $f_j \in p^j \cdot \mathbb{Z}_q[y_1...y_m]$ of degree $\leq C(j+1)$. Let $d := \max\{\deg a_1 ... \deg a_m\}$. Hence, the degree of $f_j(a_1...a_m)$ is bounded by $Cd(j+1)$, and this implies that Z's higher coefficients c_α with $|\alpha| > Cd(j+1)$ must come only from the terms $f_i(a_1...a_m)$ with $i > j$, thus $c_\alpha \in p^j \cdot \mathbb{Z}_q[x_1...x_n]$, that is, $\nu_p(c_\alpha) \geq j$. For $\alpha \in \mathbb{N}^n$, let $j^{(\alpha)}$ denote the largest j such that $|\alpha| > Cd(j+1)$: thus, $|\alpha| \leq Cd(j^{(\alpha)}+2)$. Therefore, $\nu_p(c_\alpha) \geq j^{(\alpha)} \geq |\alpha|/Cd - 2$, and we have shown that Z is in $\mathbb{Z}_q\langle x_1...x_n\rangle^\dagger$.

As to the other inclusion: Let $Z = \sum_{\alpha\in\mathbb{N}^n} c_\alpha x^\alpha \in \mathbb{Z}_q\langle x_1...x_n\rangle^\dagger$, i. e., there exist $\gamma \in \mathbb{R}$ and $\varrho > 0$ such that all $\nu_p(c_\alpha) \geq \gamma + \varrho|\alpha|$. This implies that for each $j \in \mathbb{N}$, there exist only finitely many $\alpha \in \mathbb{N}^n$ with $\nu_p(c_\alpha) = j$, hence we may set

$$f_j := \sum_{\substack{\alpha\in\mathbb{N}^n \\ \nu_p(c_\alpha)=j}} c_\alpha y^\alpha.$$

Clearly, $Z = \sum_{j=0}^\infty f_j(x_1...x_n)$, and f_j lies in $p^j\mathbb{Z}_q[y_1...y_n]$ by construction. Moreover, the degree of f_j is the maximal $|\alpha|$ with $\nu_p(c_\alpha) = j$, and by assumption, we have $|\alpha| \leq \frac{1}{\varrho}(j-\gamma)$ for such α. Therefore, $\deg f_j \leq \frac{1}{\varrho}(j-\gamma) \leq \frac{1}{\varrho}(j+1)$. □

Lemma 23. Polynomials in $\mathbb{Z}_q[x_1...x_n]$ are overconvergent.

Proof. Let $f = \sum_{\alpha\in\mathbb{N}^n} c_\alpha x^\alpha \in \mathbb{Z}_q[x_1...x_n]$. Since only finitely many $c_\alpha \neq 0$, we can take $\varrho := 1$ and

$$\gamma := \min\{\, \nu_p(c_\alpha) - |\alpha| \;\mid\; \alpha \in \mathbb{N}^n \text{ with } c_\alpha \neq 0 \,\}.$$ □

Lemma 24. The integral of a univariate overconvergent series $Z \in \mathbb{Z}_q\langle x\rangle^\dagger$ lies in $\mathbb{Z}_q\langle x\rangle^\dagger \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$.

Proof. Let $Z = \sum_{i=0}^\infty c_i x^i$ be overconvergent, hence there exist $\gamma \in \mathbb{R}$ and $\varrho > 0$ with $\nu_p(c_i) \geq \gamma + \varrho i$ for all $i \in \mathbb{N}$. Note

$$\int Z \,\mathrm{d}x = \sum_{i=0}^\infty \frac{1}{i+1} c_i x^{i+1} = \sum_{i=1}^\infty \frac{c_{i-1}}{i} x^i,$$

in $\mathbb{Q}_q[\![x]\!]$, so we are interested in a lower bound for $\nu_p(c_{i-1}/i)$:

$$\nu_p(c_{i-1}/i) = \nu_p(c_{i-1}) - \nu_p(i) \geq \gamma + \varrho(i-1) - \log_p i.$$

Eventually (say for $i > K$), we will have $\log_p i \leq \varrho i/2$. Clearly, $m := \max_{i\in\{1...K\}} \nu_p(i)$ exists, thus

$$\nu_p(p^m c_{i-1}/i) \geq m + \gamma + \varrho(i-1) - \frac{\varrho}{2} i \geq (\gamma - \varrho) + \frac{\varrho}{2} i \qquad \text{for } i > K; \text{ and}$$
$$\nu_p(p^m c_{i-1}/i) \geq m + \nu_p(c_{i-1}) - \nu_p(i) \geq \nu_p(c_{i-1}) \geq (\gamma - \varrho) + \varrho i \geq (\gamma - \varrho) + \frac{\varrho}{2} i \qquad \text{for } i \leq K.$$

Therefore,

$$\int Z \,\mathrm{d}x = \sum_{i=1}^\infty \frac{c_{i-1}}{i} x^i = p^{-m} \cdot \overbrace{\sum_{i=1}^\infty \frac{p^m c_{i-1}}{i} x^i}^{\in \mathbb{Z}_q\langle x\rangle^\dagger},$$

which may be identified with an element of $\mathbb{Z}_q\langle x\rangle^\dagger \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$. □

Lemma 25. Each homomorphism of R-algebras $\varphi\colon A \to B$ induces a unique R-algebra homomorphism $\varphi^\dagger\colon A^\dagger \to B^\dagger$ compatible with φ, that is, such that the square

$$\begin{array}{ccc} A & \xrightarrow{\varphi} & B \\ \downarrow{\scriptstyle \iota_A} & & \downarrow{\scriptstyle \iota_B} \\ A^\dagger & \xrightarrow[\varphi^\dagger]{} & B^\dagger \end{array}$$

commutes.

Proof. By component-wise application to the inverse limit, φ extends to a compatible homomorphism of R-algebras $\varphi^\infty\colon A^\infty \to B^\infty$ with the property that $\varphi^\infty(A^\dagger) \subseteq B^\dagger$. Thus we can choose $\varphi^\dagger := \varphi^\infty|_{A^\dagger}$. Suppose $\psi^\dagger$ is another such morphism. Then $(\varphi^\dagger - \psi^\dagger) \circ \iota_A = \varphi^\dagger \circ \iota_A - \psi^\dagger \circ \iota_A = \iota_B \circ \varphi - \iota_B \circ \varphi = 0$, thus $A \subseteq \ker(\varphi^\dagger - \psi^\dagger)$ and therefore $\varphi^\dagger - \psi^\dagger = 0$. □

3.3 The Monsky-Washnitzer complex

To be able to define Monsky-Washnitzer cohomology for all smooth affine varieties over finite fields, it needs to be guaranteed that there exists a suitable lift to characteristic zero:

Lemma 26. For any smooth finitely generated $\mathbb{F}_q$-algebra $\bar{A}$ there exists a smooth finitely generated $\mathbb{Z}_q$-algebra A such that $A/(p)$ is canonically isomorphic to $\bar{A}$.

Proof. A generalized statement is proved in Elkik [20, Théorème 6]. □

Finally, we arrive at the definition of Monsky-Washnitzer cohomology:

Definition 27 [49, §2]**.** Let $X = \operatorname{Spec} \bar{A}$ be a smooth affine variety over $\mathbb{F}_q$, and let $\mathbb{Z}_q$ denote the p-typical ring of Witt vectors of $\mathbb{F}_q$. According to Lemma 26, the ring $\bar{A}$ lifts to a smooth finitely generated $\mathbb{Z}_q$-algebra $A = \mathbb{Z}_q[x_1 \ldots x_n]/(f_1 \ldots f_m)$ such that $A/(p) = \bar{A}$. Instead of A, we consider its weak completion

$$A^\dagger = \mathbb{Z}_q\langle x_1 \ldots x_n\rangle^\dagger \,/\, (f_1 \ldots f_m)$$

and construct the *module of finite differentials* [36, §11]

$$\widetilde{\Omega}_{A^\dagger/\mathbb{Z}_q} := \bigoplus_{i=1}^{n} A^\dagger \cdot \mathrm{d}x_i \Big/ \left\langle \sum_{j=1}^{n} \frac{\partial f_k}{\partial x_j} \cdot \mathrm{d}x_j \;\middle|\; k \in \{1 \ldots m\} \right\rangle$$

of $A^\dagger$ over $\mathbb{Z}_q$. The *Monsky-Washnitzer cohomology* $H^*(X, \mathbb{Q}_q)$ of X is finally obtained by taking de Rham cohomology with respect to this differential module and tensoring over $\mathbb{Z}_q$ with $\mathbb{Q}_q$.

In more detail, this means we construct the following cohomology spaces. For shorthand, let

$$\Omega^i := \bigwedge^i \widetilde{\Omega}_{A^\dagger/\mathbb{Z}_q}$$

denote the ith exterior power[1] of $\widetilde{\Omega}_{A^\dagger/\mathbb{Z}_q}$ — This is an $A^\dagger$-module. To aid readability, let $\mathrm{d}\underline{x}$ abbreviate some $\mathrm{d}x_{j_1} \wedge \ldots \wedge \mathrm{d}x_{j_i}$. For each $i \in \mathbb{N}$, we define the *exterior derivative*

$$\mathrm{d}^i\colon\ \Omega^i \longrightarrow \Omega^{i+1},\ \ f \cdot \mathrm{d}\underline{x} \longmapsto \left(\sum_{j=1}^{n} \frac{\partial f}{\partial x_j} \cdot \mathrm{d}x_j \right) \wedge \mathrm{d}\underline{x}.$$

Note that $\mathrm{d}^0\colon A^\dagger \to \Omega^1$ is taking the gradient and that all d^i are $\mathbb{Z}_q$-linear (but of course usually not $A^\dagger$-linear). By abuse of notation, we often write just d for any d^i. The elements of Ω^i are referred to as *i-forms*; an i-form is said to be *closed* if it lies in the kernel of d^i, and *exact* if it lies in the image of d^{i-1}. Now it is easy to see that $\mathrm{d}^i \circ \mathrm{d}^{i-1} = 0$, hence there is the complex of $\mathbb{Z}_q$-modules

$$0 \longrightarrow A^\dagger \xrightarrow{\mathrm{d}^0} \Omega^1 \xrightarrow{\mathrm{d}^1} \Omega^2 \xrightarrow{\mathrm{d}^2} \Omega^3 \xrightarrow{\mathrm{d}^3} \Omega^4 \xrightarrow{\mathrm{d}^4} \ldots.$$

Finally, the Monsky-Washnitzer cohomology of X is defined as the family of $\mathbb{Q}_q$-vector spaces

$$H^i(X, \mathbb{Q}_q) := \left(\ker \mathrm{d}^i \,/\, \operatorname{im} \mathrm{d}^{i-1}\right) \underset{\mathbb{Z}_q}{\otimes} \mathbb{Q}_q.$$

---

[1] The nth *exterior power* $\bigwedge^n M$ of an R-module M is the n-fold tensor product of M with itself modulo the submodule generated by all simple tensors $m_1 \otimes \cdots \otimes m_n$ in which two factors m_i, m_j are equal.

Note if $M = \bigoplus_{i=1}^{r} R \cdot m_i$ is free of rank r, then $\{\, m_{j_1} \wedge \ldots \wedge m_{j_n} \mid 1 \le j_1 < \cdots < j_n \le r \,\}$ forms a $\binom{r}{n}$-element basis of $\bigwedge^n M$. Moreover, there exists an associative, bilinear, graded-commutative *wedge product*

$$\wedge\colon \bigwedge^i M \times \bigwedge^j M \to \bigwedge^{i+j} M,\ (m_1 \otimes \cdots \otimes m_i, \mu_1 \otimes \cdots \otimes \mu_j) \mapsto m_1 \otimes \cdots \otimes m_i \otimes \mu_1 \otimes \cdots \otimes \mu_j.$$

Two elements of $\ker \mathrm{d}^i$ which differ by an element of $\operatorname{im} \mathrm{d}^{i-1}$ are said to be *cohomologous*.

As one would expect from a cohomology theory, this construction is functorial: If $\varphi\colon X \to Y$ is a map of smooth affine varieties over $\mathbb{F}_q$, and if $A^\dagger$ and $B^\dagger$ are as above for X and Y, then φ induces a morphism of rings $\varphi^\dagger\colon B^\dagger \to A^\dagger$, and therefore an endomorphism

$$H^i(\varphi, \mathbb{Q}_q)\colon\ H^i(Y, \mathbb{Q}_q) \to H^i(X, \mathbb{Q}_q),\ f\,\mathrm{d}\underline{x} \mapsto \varphi^\dagger(f)\mathrm{d}\varphi^\dagger(\mathrm{d}\underline{x})$$

on Monsky-Washnitzer cohomology.

It is implicit in this definition that the resulting cohomology is well-defined — in particular, it technically demands proof that everything is independent of the choice of the lift A to characteristic zero, and that the induced maps are independent of the chosen lift to dagger rings. This is not at all obvious; a proof can be found in Monsky and Washnitzer [40].

Furthermore, note that some authors prefer to tensor the above complex with $\mathbb{Q}_q$ *before* taking cohomology: As tensoring with the field of fractions $\mathbb{Q}_q = \operatorname{Quot} \mathbb{Z}_q$ is the same thing as localizing away from the zero ideal, and since localization is an exact functor, the sequences

$$0 \longrightarrow \operatorname{im} \mathrm{d}^{i-1} \underset{\mathbb{Z}_q}{\otimes} \mathbb{Q}_q \longrightarrow \ker \mathrm{d}^i \underset{\mathbb{Z}_q}{\otimes} \mathbb{Q}_q \longrightarrow H^i(X, \mathbb{Q}_q) \longrightarrow 0$$

are exact, thus both definitions are equivalent. In computations, we shall freely choose either point of view depending on the context.

Remark. It ought to be noted that both the Witt ring $\mathbb{Z}_q$ and the weak completion $A^\dagger$ consist of *infinite* sequences. That is, the differential forms representing elements of the Monsky-Washnitzer cohomology are infinite in not only one, but two 'directions': The elements of $A^\dagger$ are infinite power series over $\mathbb{Z}_q$ due to the weak completion, and each of its coefficients has an infinite p-adic expansion coming from the construction of the Witt vectors.

To be completely explicit: We have seen in Lemma 22 that elements of the weak completion $A^\dagger$ are overconvergent series $\sum_{\alpha\in\mathbb{N}^n} c_\alpha x^\alpha$ with each $c_\alpha \in \mathbb{Z}_q$ an infinite Witt vector. Note that a $A^\dagger$-basis S_i of Ω^i is given by the forms $\mathrm{d}x_{j_1}\wedge \ldots \wedge \mathrm{d}x_{j_i}$ with $1 \le j_1 < j_2 < \cdots < j_i \le n$; thus, each element of the ith Monsky-Washnitzer cohomology space $H^i(X, \mathbb{Q}_q)$ may be represented by a sum $\sum_{\omega\in S_i} f_\omega \cdot \omega$ such that each $f_\omega \in A^\dagger$ is an overconvergent series as described above. (See also Section 2.2 of Harvey [27].)

Of course, when performing computations involving Monsky-Washnitzer cohomology on a real — necessarily *finite* — machine or sheet of paper, it will be necessary to approximate those objects. Determining the precision required to obtain correct results contributes a significant portion of the theoretical effort in developing such an algorithm, as the reader shall see in Section 4.7. ∘

Luckily, we may (for now) blissfully ignore those peculiarities and rejoice in the beauties of theory:

Example 28. We compute the Monsky-Washnitzer cohomology of n-dimensional affine space $\mathbb{A}^n\mathbb{F}_q$ over a finite field $\mathbb{F}_q$. The coordinate ring $\bar{A} = \mathbb{F}_q[x_1 \ldots x_n]$ admits the lift $A = \mathbb{Z}_q[x_1 \ldots x_n]$ over the Witt vectors $\mathbb{Z}_q$: Evidently, A is finitely generated and smooth over $\mathbb{Z}_q$, and the short exact sequence

$$0 \longrightarrow \mathbb{Z}_q[x_1 \ldots x_n] \xrightarrow{\cdot p} \mathbb{Z}_q[x_1 \ldots x_n] \xrightarrow{\bmod p} \mathbb{F}_q[x_1 \ldots x_n] \longrightarrow 0$$

shows $A/(p) \cong \bar{A}$. Thus, $A^\dagger = \mathbb{Z}_q\langle x_1 \ldots x_n\rangle^\dagger$ is the $\mathbb{Z}_q$-algebra of overconvergent power series. The module of finite differentials of $A^\dagger$ over $\mathbb{Z}_q$ is $\Omega = \bigoplus_{i=1}^n A^\dagger \cdot \mathrm{d}x_i$, and in this case, all closed forms are exact.[1] Thus, the Monsky-Washnitzer complex of $\mathbb{A}^n\mathbb{F}_q$ is an exact sequence, and we see

$$H^i(\mathbb{A}^n\mathbb{F}_q, \mathbb{Q}_q) = \begin{cases} \mathbb{Q}_q & \text{if } i = 0 \\ 0 & \text{for } i \ge 1. \end{cases}$$ △

[1] Recall that a $A^\dagger$-basis of Ω^k is given by the forms $S_k := \{\, \mathrm{d}x_{j_1}\wedge \ldots \wedge \mathrm{d}x_{j_k} \mid 1 \le j_1 < j_2 < \cdots < j_k \le n \,\}$, and let S_k' denote the set of all forms in S_k which do not contain $\mathrm{d}x_1$.

We perform induction on n: For $n = 0$, there is nothing to show. Suppose the claim has been proven for $n-1$. Clearly, a closed i-form ω can be written as $\omega = \mathrm{d}x_1 \wedge \chi + \psi$ such that neither χ nor ψ contain $\mathrm{d}x_1$. Write $\chi = \sum_{\lambda\in S'_{i-1}} f_\lambda \cdot \lambda$ with coefficients $f_\lambda \in A^\dagger$ and define

$$\varphi := \sum_{\lambda\in S'_{i-1}} \left(\int f_\lambda \,\mathrm{d}x_1 \right) \cdot \lambda.$$

After this little 'sanity check', let us continue with a more interesting case: an elliptic curve. This example illustrates the difficulties introduced by taking the weak completion — suddenly, one has to invoke topological arguments to argue about infinite series, even though the problem was originally about polynomials over a finite field. Luckily, the following result generalizes nicely to other varieties, so we can do away with these (tedious) considerations after having seen the method once.

Example 29 [49, Prop. 7.1]**.** Consider an elliptic curve

$$E\colon\ y^2 = x^3 + \bar{a}x + \bar{b}$$

over $\mathbb{F}_q$. We treat E as an affine curve by removing the point at infinity. The coordinate ring $\bar{A} = \mathbb{F}_p[x,y]/(y^2 - x^3 - \bar{a}x - \bar{b})$ may be lifted to

$$A := \mathbb{Z}_q[x,y] \,/\, (y^2 - x^3 - ax - b)$$

by choosing arbitrary preimages a, b of $\bar{a}, \bar{b}$ under the projection $\mathbb{Z}_q \twoheadrightarrow \mathbb{F}_q$. Clearly, A is a finitely generated $\mathbb{Z}_q$-algebra, and it is smooth over $\mathbb{Z}_q$ since E is. Its weak completion is

$$A^\dagger = \mathbb{Z}_q\langle x,y\rangle^\dagger \,/\, (y^2 - x^3 - ax - b),$$

thus we have the $A^\dagger$-modules[1]

$$\Omega^1 = (A^\dagger \cdot \mathrm{d}x \oplus A^\dagger \cdot \mathrm{d}y) \,/\, \langle 2y\mathrm{d}y - (3x^2 + a)\mathrm{d}x\rangle;$$
$$\Omega^2 = A^\dagger \cdot \mathrm{d}x \wedge \mathrm{d}y \,/\, \langle 2y \cdot \mathrm{d}x \wedge \mathrm{d}y, (3x^2 + a)\mathrm{d}x \wedge \mathrm{d}y\rangle.$$

At this point, note that it is *not* possible to just repeatedly reduce terms of high degree via the given relations to compute a basis of the cohomology spaces: We are dealing with infinite series of differentials, hence such a reduction algorithm will not terminate in general. Of course, when employing Monsky-Washnitzer cohomology for practical computations, those series are approximated by finite sums — however, for our application, it is necessary to write a certain morphism with respect to a basis of the cohomology (up to some error), thus we first need to establish a basis of the (full) cohomology to get provably correct results.

To begin with, note that $\Omega^1 \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$ is a free $(A^\dagger \otimes_{\mathbb{Z}_q} \mathbb{Q}_q)$-module of rank 1: Since A is smooth, Hilbert's Nullstellensatz implies there exist polynomials $\alpha, \beta \in \mathbb{Q}_q[x]$ such that

$$\alpha \cdot (x^3 + ax + b) + \beta \cdot (3x^2 + a) = 1.$$

Define $\omega := \alpha y \mathrm{d}x + 2\beta \mathrm{d}y \in \Omega^1 \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$. It is easily verified that[2]

$$\mathrm{d}y = \frac{3x^2 + a}{2}\omega \quad \text{and} \quad \mathrm{d}x = y\omega,$$

(Note this is a well-defined element of $\Omega^{i-1} \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$ according to Lemma 24.) Then,

$$\mathrm{d}\varphi = \sum_{\lambda \in S'_{i-1}} (f_\lambda \cdot \mathrm{d}x_1 + [\dots]) \wedge \lambda = \sum_{\lambda \in S'_{i-1}} f_\lambda \cdot \mathrm{d}x_1 \wedge \lambda + [\dots] = \mathrm{d}x_1 \wedge \chi + [\dots],$$

where the $[\dots]$ represent some terms not containing $\mathrm{d}x_1$; therefore $\tau := \omega - \mathrm{d}\varphi$ does not contain $\mathrm{d}x_1$.

Since ω is closed, τ is as well, and we use this to show that τ cannot contain x_1: Write $\tau = \sum_{\lambda \in S'_{i-1}} g_\lambda \cdot \lambda$ with $g_\lambda \in A^\dagger$. Thus,

$$0 = \mathrm{d}\tau = \sum_{j=1}^{n} \sum_{\lambda \in S'_{i-1}} \frac{\partial g_\lambda}{\partial x_j} \cdot \mathrm{d}x_j \wedge \lambda,$$

so all terms of the double sum containing $\mathrm{d}x_1$ must come from the $j = 1$ terms, implying

$$0 = \sum_{\lambda \in S'_{i-1}} \frac{\partial g_\lambda}{\partial x_1} \cdot \mathrm{d}x_1 \wedge \lambda.$$

Now since $\{\, \mathrm{d}x_1 \wedge \lambda \mid \lambda \in S'_{i-1} \,\}$ is $A^\dagger$-linearly independent, this in turn yields $\partial g_\lambda / \partial x_1 = 0$ for all $\lambda \in S'_{i-1}$. In other words, all g_λ are in $\mathbb{Z}_q\langle x_2 \dots x_n\rangle^\dagger$, and we may conclude by induction that τ — and thus ω — is exact. □

[1] In Ω^2, observe $(3x^2 + a)\mathrm{d}x \wedge \mathrm{d}y = 2y \cdot \mathrm{d}y \wedge \mathrm{d}y = 0$ and $2y \cdot \mathrm{d}x \wedge \mathrm{d}y = \mathrm{d}x \wedge (3x^2 + a)\mathrm{d}x = 0$.

[2] Due to this, ω is often referred to as $\frac{\mathrm{d}x}{y}$ in the literature, but we deem this rather confusing as y need not be invertible in $A^\dagger$.

thus

$$\Omega^1 \otimes_{\mathbb{Z}_q} \mathbb{Q}_q = A^\dagger\omega \otimes_{\mathbb{Z}_q} \mathbb{Q}_q.$$

We first examine the restricted map

$$\delta = \mathrm{d}|_{A\otimes_{\mathbb{Z}_q}\mathbb{Q}_q} \colon A \otimes_{\mathbb{Z}_q} \mathbb{Q}_q \to A\omega \otimes_{\mathbb{Z}_q} \mathbb{Q}_q.$$

Its kernel and cokernel are nothing but E's zeroth and first *algebraic de Rham cohomology* over $\mathbb{Q}_q$, which will turn out to be strongly related to the Monsky-Washnitzer cohomology of E. The kernel is easy: It consists of the elements of A represented by constant functions, i. e., $\ker\delta = \mathbb{Q}_q$. To compute the cokernel, consider an arbitrary form $(f+gy)\cdot\omega \in A\omega \otimes_{\mathbb{Z}_q}\mathbb{Q}_q$, where $f, g \in \mathbb{Q}_q[x]$. The $gy\omega$ part is always identified with zero, as $gy\omega = g\,\mathrm{d}x$ is integrable. Note that for $n \in \mathbb{N}$,

$$\delta(x^n y) = nx^{n-1}y\,\mathrm{d}x + x^n\,\mathrm{d}y = nx^{n-1}y^2\omega + \tfrac{3x^2+a}{2}x^n\omega = \Big(\big(n+\tfrac{3}{2}\big)x^{n+2} + r\Big)\cdot\omega,$$

where $r \in \mathbb{Q}_q[x]$ is of degree $\leq n+1$. Since $n+\frac{3}{2}$ is invertible in $\mathbb{Q}_q$, one can reduce $f\omega$ to an equivalent form by subtracting an appropriate multiple of $\delta(x^n y)$, until f is at most linear.[1] Therefore,

$$\operatorname{coker}\delta = \omega\mathbb{Q}_q \oplus x\omega\mathbb{Q}_q.$$

The remainder of the example will consist of proving that a $\mathbb{Q}_q$-basis of the Monsky-Washnitzer cohomology is given by the same forms ω and $x\omega$. To this end, we require infinitesimal methods, hence we start off by constructing a topology on Ω^1. For any $\varrho > 1$, let R_ϱ denote the subring of $\mathbb{Z}_q\langle x, y\rangle^\dagger$ consisting of series $\sum_{i,j\geq 0} a_{ij}x^i y^j$ such that $\lim_{i,j\to\infty}|a_{ij}|_p\varrho^{i+j} = 0$. On R_ϱ, we define the norm[2]

$$\Big\|\sum_{i,j\geq 0} a_{ij}x^i y^j\Big\|_\varrho := \max_{i,j\geq 0}|a_{ij}|_p\varrho^{i+j}.$$

Using the properties of the p-adic absolute value, one may check this is indeed a non-Archimedean norm. We let $A^\dagger_\varrho := R_\varrho/(y^2-x^3-ax-b)$ and use the induced seminorm[3] on $A^\dagger_\varrho$, also denoted by $\|\cdot,\cdot\|_\varrho$. Note that R_ϱ, and therefore $A^\dagger_\varrho$, are complete with respect to their (semi)norm [7, Section 1.1.7]. For any $\varrho' > \varrho$, we have an inclusion $A^\dagger_{\varrho'} \hookrightarrow A^\dagger_\varrho$, and the direct limit of that system is just $\bigcup_{\varrho>1} A^\dagger_\varrho = A^\dagger$. Hence, we may give $A^\dagger$ the final topology, that is: the finest topology such that all the inclusions $A^\dagger_\varrho \hookrightarrow A^\dagger$ are continuous. This topology carries over to Ω^1 via the identification $\Omega^1 = A^\dagger\omega$.

Armed with these definitions, we can now show that the image of $\mathrm{d}\colon A^\dagger \to A^\dagger\omega$ is closed. Since $A^\dagger\omega$ is equipped with the direct limit topology, it suffices to prove the following: If $(\mathrm{d}f_k)_{k\in\mathbb{N}}$ is a sequence of exact differentials in $A^\dagger_\varrho\omega$ converging to $\varphi \in A^\dagger_\varrho\omega$, then φ is exact as well. Using the definition of d, it is not hard (but a bit tedious) to show that there exists a constant $C > 0$ such that $\|f_k\|_\varrho \leq C\|\mathrm{d}f_k\|_\varrho$ for all $k \in \mathbb{N}$.[4] Since $(\mathrm{d}f_k)_{k\in\mathbb{N}}$ is a convergent sequence, this implies that $(f_k)_{k\in\mathbb{N}}$ is a Cauchy sequence, hence (by completeness) converges to some $f \in A^\dagger_\varrho$. Due to the bound above, f must satisfy $\mathrm{d}f = \varphi$.

---

[1] Note this argument would not make sense for infinite series.

[2] With this definition, we make R_ϱ a *normed ring* in the sense of non-Archimedean analysis [7, Chapter 1]. That is a ring R together with a map $|\cdot|\colon R \to \mathbb{R}_{\geq 0}$ such that the following axioms hold for all $a, b \in R$:

- $|a| = 0$ if and only if $a = 0$;
- $|a+b| \leq \max\{|a|, |b|\}$ (note this implies the triangle inequality $|a+b| \leq |a| + |b|$);
- $|1| = 1$;
- $|ab| \leq |a||b|$.

If the first condition is relaxed to $|0| = 0$, then $|\cdot|$ is called a *seminorm*.
A *(semi)normed group* is an abelian group G with a map $|\cdot|\colon G \to \mathbb{R}_{\geq 0}$ such that the first two axioms are satisfied.

[3] For an ideal $\mathcal{I}$ of a seminormed ring R (or, more generally, a subgroup of a seminormed group), this is the seminorm on the quotient $R/\mathcal{I}$ given by [7, Section 1.1.6]

$$|a+\mathcal{I}| := \inf_{b\in\mathcal{I}}|a+b|.$$

[4] This boils down to computing an explicit expression for $\mathrm{d}f_k$, rewriting it as a multiple of ω, and obtaining a bound on the p-adic absolute values $|(i+j)a_{ij}|_p$ for the coefficients a_{ij} of the resulting power series (using the fact that f_k is overconvergent). This is similar to the argument made in the proof of Lemma 24.

Now it is fairly easy to prove from the axioms of seminormed groups that the induced seminorm on a quotient is a norm if and only if the quotient is taken modulo a closed subset. This implies the topology on $\operatorname{coker} \mathrm{d} = A^\dagger\omega \,/\, \operatorname{im} \mathrm{d} = H^1(E, \mathbb{Q}_q)$ is Hausdorff. Clearly $\operatorname{im}\delta \subseteq \operatorname{im} \mathrm{d}$, hence $\operatorname{coker}\delta$ embeds into $H^1(E, \mathbb{Q}_q)$ as a dense subset, which — endowed with the subspace topology — is Hausdorff as well. As we have previously shown, $\operatorname{coker}\delta$ is a finite-dimensional vector space, hence its topology must be the product topology induced from the underlying topological field $\mathbb{Q}_q$ [22, Theorem 5.0.3]. However, for that topology, $\operatorname{coker}\delta$ is complete, hence it is closed in $H^1(E, \mathbb{Q}_q)$. Since dense closed subsets must be the whole space, we may conclude $H^1(E, \mathbb{Q}_q) = \operatorname{coker}\delta = \omega\mathbb{Q}_q \oplus x\omega\mathbb{Q}_q$. △

In the preceding example, we observed (via a tedious argument) that the Monsky-Washnitzer cohomology has the same forms as a basis as the algebraic de Rham cohomology of a lift over the Witt vectors. This is not by pure chance:

Theorem 30 [31, Theorem 1]**.** Let X be a smooth affine variety over $\mathbb{F}_q$ with coordinate ring $\bar{A}$ and let A denote a lift of $\bar{A}$ over $\mathbb{Z}_q$. Then, the Monsky-Washnitzer cohomology of X is canonically isomorphic to the algebraic de Rham cohomology of $\operatorname{Spec}(A \otimes_{\mathbb{Z}_q} \mathbb{Q}_q)$.

Remark. Given that correspondence, a reader might be tempted to wonder why we burdened ourselves with the laborious construction of the Monsky-Washnitzer cohomology in the first place. The reason is that this theorem only provides us with an *isomorphism of $\mathbb{Q}_q$-vector spaces*, but it does not relate functorial properties of the two cohomologies: In particular, the de Rham cohomology (in characteristic zero!) usually does not admit a lift of the Frobenius endomorphism of X, but Monsky-Washnitzer cohomology does — this is what the detour to the weak completion rewards us with.

Nevertheless, we gain a lot from Theorem 30: It shows that considering only finite differentials is sufficient when computing the basis of the cohomology, thus we can refrain from reiterating the cumbersome topological arguments given above in future examples. ○

Example 31 [50, Example 2.1.2]**.** As another rather simple example, let us consider a punctured affine line: Fix a squarefree monic polynomial $\bar{r} \in \mathbb{F}_q[x]$ of degree d over $\mathbb{F}_q$, and let

$$X := \operatorname{Spec}\ \mathbb{F}_q[x,y]/(\bar{r}y - 1) = \mathbb{A}^1\mathbb{F}_q \setminus \{\bar{r} = 0\}.$$

A lift of its coordinate ring is given by $A^\dagger = \mathbb{Z}_q\langle x, y\rangle^\dagger/(ry - 1)$, where r is any monic lift of $\bar{r}$, and we get

$$\Omega^1 = (A^\dagger \cdot \mathrm{d}x \oplus A^\dagger \cdot \mathrm{d}y)/\langle r'y\mathrm{d}x + r\mathrm{d}y\rangle.$$

Since $\mathrm{d}y = y \cdot r\mathrm{d}y = y \cdot (-r'y\mathrm{d}x) = -r'y^2\mathrm{d}x$ in Ω^1, we recognize that Ω^1 is a free $A^\dagger$-module of rank 1 generated by $\mathrm{d}x$. According to Theorem 30, a basis of the algebraic de Rham cohomology of A over $\mathbb{Q}_q$ is also a basis of the Monsky-Washnitzer cohomology. In other words: To get a basis for $H^1(X, \mathbb{Q}_q)$, it is sufficient to reduce the $\mathbb{Z}_q$-generators $x^iy^j\mathrm{d}x$ of $A \cdot \mathrm{d}x$ instead of $\Omega^1 = A^\dagger \cdot \mathrm{d}x$, which is not spanned by these forms as a $\mathbb{Z}_q$-module. To begin with, note that a form $f\mathrm{d}x$ with $f \in \mathbb{Q}_q[x]$ is always exact. For any $i \in \mathbb{N}$ and $j \geq 2$, we have

$$\mathrm{d}(x^iy^{j-1}) = ix^{i-1}y^{j-1}\mathrm{d}x + (j-1)x^iy^{j-2}\mathrm{d}y = ix^{i-1}y^{j-1}\mathrm{d}x - (j-1)x^iy^jr'\mathrm{d}x,$$

hence in $H^1(X, \mathbb{Q}_q)$ the relation

$$x^iy^jr'\mathrm{d}x = \tfrac{i}{j-1}x^{i-1}y^{j-1}\mathrm{d}x = \tfrac{i}{j-1}x^{i-1}y^jr\mathrm{d}x. \qquad (*)$$

Since r is assumed squarefree, there exist polynomials $\alpha, \beta \in \mathbb{Q}_q[x]$ such that $\alpha r + \beta r' = 1$. For $i \in \mathbb{N}$ and $j \geq 2$, we therefore get the relation

$$\begin{aligned} x^iy^j\mathrm{d}x &= x^iy^j(\alpha r + \beta r')\mathrm{d}x = \alpha x^iy^jr\mathrm{d}x + \beta x^iy^jr'\mathrm{d}x \\ &\underset{(*)}{=} \alpha x^iy^jr\mathrm{d}x + \beta\tfrac{i}{j-1}x^{i-1}y^jr\mathrm{d}x = \big(\alpha x + \beta\tfrac{i}{j-1}\big)x^{i-1}y^{j-1}\mathrm{d}x = \overbrace{\big(\alpha x + \beta\tfrac{i}{j-1}\big)x^{i-1}}^{\in\mathbb{Q}_q[x]}y^{j-1}\mathrm{d}x \end{aligned}$$

in $H^1(X, \mathbb{Q}_q)$. This can be used to reduce a form $x^iy^j\mathrm{d}x$ with $j \geq 2$ to an equivalent form $fy\mathrm{d}x$, where $f \in \mathbb{Q}_q[x]$. Note we have so far reduced our generating set of $H^1(X, \mathbb{Q}_q)$ to the forms $x^iy\mathrm{d}x$ with $i \in \mathbb{N}$. Moreover, for any $i \geq d$, Euclidean division of x^i by r yields $u, v \in \mathbb{Q}_q[x]$ such that $x^i = ur + v$ with

$\deg v < d$, hence $x^i y \mathrm{d}x = u\mathrm{d}x + vy\mathrm{d}x$. Since $u\mathrm{d}x$ is exact, this implies that the $x^i y\mathrm{d}x$ with $i \in \{0 \ldots d-1\}$ generate $H^1(X, \mathbb{Q}_q)$. Thus, we have shown

$$H^i(X, \mathbb{Q}_q) = \begin{cases} \mathbb{Q}_q & \text{if } i = 0 \\ \bigoplus_{i=0}^{d-1} \mathbb{Q}_q \cdot x^i y \mathrm{d}x & \text{if } i = 1 \\ 0 & \text{for } i > 1. \end{cases}$$ △

3.4 The Lefschetz trace formula

The main reason for our interest in the cohomology theory of Monsky and Washnitzer is the following theorem, parallel to Theorem 7 for Weil cohomologies (Section 2.2). It enables us to express the number of q-rational points on a smooth affine variety in terms of traces of linear maps induced on the Monsky-Washnitzer cohomology by the Frobenius endomorphism.

Note that, as we shall see in the subsequent examples, as well as in Section 4, these traces are perfectly computable, making the Monsky-Washnitzer cohomology suitable for practical point counting. Moreover, observe that the apparent limitation to affine varieties need not concern us: The method easily generalizes to projective varieties by partitioning into an affine patch and its complement 'at infinity', which is projective of smaller dimension, until the problem becomes trivial.

Theorem 32 (Lefschetz trace formula)**.** Let $X = \operatorname{Spec} \bar{A}$ be a smooth affine variety of dimension n over a finite field $\mathbb{F}_q$, let $A^\dagger$ denote the p-adic weak completion of a lift of $\bar{A}$ over $\mathbb{Z}_q$, and suppose the q-power Frobenius endomorphism of $\bar{A}$ lifts to an endomorphism σ of $A^\dagger$.[1] Then

$$\#X(\mathbb{F}_q) = \sum_{i=0}^{n} (-1)^i \operatorname{tr}\left(q^n \sigma_*^{-1} \mid H^i(X, \mathbb{Q}_q)\right).$$

Proof. See van der Put [49, Theorem 4.1] or Fresnel and van der Put [21, Theorem 7.6.4]. □

This theorem comes with an easy corollary which yields an expression for the full zeta function at once, rather than only the number of points over a single field. The proof proceeds as that of Theorem 9.

Corollary 33. With the notation and prerequisites of Theorem 32, the zeta function of X is

$$Z(X; t) = \prod_{i=0}^{n} \left(\det\left(1 - tq^n \sigma_*^{-1} \mid H^i(X, \mathbb{Q}_q)\right)\right)^{(-1)^{r+1}}.$$

Remark. Observe that the above formulation of the Lefschetz trace formula differs from Theorem 7 — it follows from the usual statement via Poincaré duality.

In particular, note it is not necessary to invert the matrix of σ_* to compute the zeta function; see Section 4.6 for a concrete instance of this argument. ○

To illustrate these results, let us work through some examples. Write a^σ as shorthand for $\sigma(a)$.

Example 34. Continuing our toy Example 28, we illustrate the usage of the Lefschetz trace formula to count the $\mathbb{F}_q$-rational points of n-dimensional affine space $\mathbb{A}^n\mathbb{F}_q$ over a finite field $\mathbb{F}_q$. By Lemma 19, the q-power Frobenius on $\mathbb{Z}_q$ is just the identity, thus the induced map σ_* on $H^0(\mathbb{A}^n\mathbb{F}_q, \mathbb{Q}_q)$ is as well.[2] Moreover, $H^i(\mathbb{A}^n\mathbb{F}_q) = 0$ for all $i \geq 1$; therefore the Lefschetz trace formula simplifies to

$$\#\mathbb{A}^n\mathbb{F}_q(\mathbb{F}_q) = \operatorname{tr}\left(q^n \sigma_*^{-1} \mid H^0(\mathbb{A}^n\mathbb{F}_q \mid \mathbb{Q}_q)\right) = \operatorname{tr}(q^n \mid \mathbb{Q}_q) = q^n.$$

The zeta function is

$$Z(\mathbb{A}^n\mathbb{F}_q; t) = \left(\det(1 - tq^n \mid \mathbb{Q}_q)\right)^{-1} = \frac{1}{1 - q^n t}.$$ △

[1] This means: The endomorphism induced by σ on the residue ring $A^\dagger/(p) = \bar{A}$ is just the regular Frobenius endomorphism.

[2] In fact, we have $H^0(X, \mathbb{Q}_q) = \mathbb{Q}_q$ for *any* smooth affine variety X, hence the action of σ on $H^0(X, \mathbb{Q}_q)$ is *always* trivial.

Example 35. As another simple example, let us consider a punctured affine line $X := \mathbb{A}^1\mathbb{F}_q \setminus \{0\}$ with one point removed over $\mathbb{F}_q$. Recall from Example 31 that in this case

$$A^\dagger = \mathbb{Z}_q\langle x, y\rangle^\dagger/(xy - 1)$$

and

$$H^1(X, \mathbb{Q}_q) = \mathbb{Q}_q \cdot y\mathrm{d}x.$$

The map $\sigma\colon A^\dagger \to A^\dagger$ given by the substitutions $x \mapsto x^q$ and $y \mapsto y^q$ is a lift of the q-power Frobenius. Its action on $H^0(X, \mathbb{Q}_q)$ is (as always) the identity, and on $H^1(X, \mathbb{Q}_q)$ we get

$$\sigma_*(y\mathrm{d}x) = \sigma(y) \cdot \mathrm{d}\sigma(x) = y^q \mathrm{d}x^q = y^q q x^{q-1}\mathrm{d}x = q \cdot y\mathrm{d}x,$$

hence σ_* acts as multiplication by q on $H^1(X, \mathbb{Q}_q)$. Therefore, the Lefschetz trace formula gives

$$\#X(\mathbb{F}_q) = \mathrm{tr}\left(q^1 \mid H^0(X, \mathbb{Q}_q)\right) - \mathrm{tr}\left(q^1\sigma_*^{-1} \mid H^1(X, \mathbb{Q}_q)\right) = \mathrm{tr}(q \mid \mathbb{Q}_q) - \mathrm{tr}(1 \mid \mathbb{Q}_q) = q - 1,$$

as expected. For the zeta function of X, we get

$$Z(X; t) = \left(\det(1 - tq^1 \mid \mathbb{Q}_q)\right)^{-1} \cdot \det(1 - tq^1q^{-1} \mid \mathbb{Q}_q) = \frac{1 - t}{1 - qt}$$

which makes sense as $\frac{1}{1-t}$ and $\frac{1}{1-qt}$ are the zeta functions of a point and an affine line, respectively. △

The next two examples serve to demonstrate that obtaining a lift of Frobenius is not always as simple as in the previous cases; in fact, this is typically quite involved for generic instances (which were not specifically crafted to behave nicely).

Example 36. Now consider a more general punctured line, i. e., continue in the setting of Example 31. For notational convenience, we assume $q = p$ and point out that the general case works exactly the same.

Recall that $r \in \mathbb{Z}_p[x]$ is a lift of a monic squarefree polynomial $\bar{r} \in \mathbb{F}_p[x]$ and our variety X is the spectrum of $\mathbb{F}_p[x, y]/(\bar{r}y - 1)$, hence

$$A^\dagger = \mathbb{Z}_p\langle x, y\rangle^\dagger/(ry - 1); \qquad H^1(X, \mathbb{Q}_p) = \bigoplus_{i=0}^{d-1} \mathbb{Q}_p \cdot x^i y\mathrm{d}x.$$

We construct a p-power Frobenius lift $\sigma\colon A^\dagger \to A^\dagger$ by setting $x^\sigma := x^p$. However, also sending y to y^p would in general not give a well-defined ring endomorphism, as pointed out in Section 3.2. Instead, we must find find an inverse y^σ of r^σ in $A^\dagger$, such that $r^\sigma y^\sigma = 1$. To this end, is convenient to set

$$E := \frac{r^p - r^\sigma}{p},$$

which is well-defined in $\mathbb{Z}_q[x]$ because r^p and r^σ have the same image modulo p, and rewrite r^σ as

$$r^\sigma = r^p - pE = r^p(1 - py^pE).$$

Clearly, the inverse of r^p is y^p. To invert the parenthesized term $1 - py^pE$, note that py^pE has positive p-adic valuation, hence the geometric series

$$\frac{1}{1 - py^pE} = \sum_{k=0}^{\infty} p^k y^{pk} E^k$$

converges — in fact, it is overconvergent. Therefore, setting

$$y^\sigma := \sum_{k=0}^{\infty} p^k y^{pk+p} E^k$$

completes our lift of the Frobenius endomorphism. Its action on our basis $\{ x^i y\mathrm{d}x \mid i \in \{0...d-1\} \}$ of the cohomology $H^1(X, \mathbb{Q}_p)$ of X is

$$\sigma_*(x^i y\mathrm{d}x) = (x^i y)^\sigma \mathrm{d}x^\sigma = x^{pi} y^\sigma \mathrm{d}x^p = y^\sigma x^{pi} p x^{p-1} \mathrm{d}x = \sum_{k=0}^{\infty} p^{k+1} x^{pi+p-1} y^{pk+p} E^k \mathrm{d}x.$$

Unfortunately, rewriting the image of $x^i y\mathrm{d}x$ as a $\mathbb{Q}_p$-linear combination of our basis depends on the concrete value of r, and presenting the computation is not particularly instructive, thus we refrain from showing this — one would normally implement the calculation on a computer and be done with it. See Section 4.5 for such a reduction algorithm. △

Example 37. Let us attempt to apply Theorem 32 to the (affine) elliptic curve $E\colon y^2 = x^3 + \bar{a}x + \bar{b}$ over $\mathbb{F}_p$. Lift the defining equation of E to $f = y^2 - x^3 - ax - b \in \mathbb{Z}_q[x]$.

The main challenge is lifting the p-power Frobenius to

$$A^\dagger = \mathbb{Z}_p\langle x, y\rangle^\dagger/(f).$$

It would be nice to just set $x^\sigma := x^p$ as before and compute y^σ using a Newton iteration; however, this would come down to evaluating the recursion

$$\eta_{k+1} := \eta_k - \frac{f(x^p, \eta_k)}{2\eta_k},$$

and the first approximation $\eta_0 = y^p$ is not invertible in $A^\dagger$. Thus, it is necessary to construct x^σ and y^σ at the same time.

The standard method to solve this problem is a two-variable Newton lift, see for instance Cohen et al. [9, Algorithm 12.23].

Alternatively, Besser, Escriva, and Jeu [6, Section 3, cf. Example 4.2] propose the following procedure to find a solution: Let f_x, f_y denote the partial derivatives of f with respect to x, y. Choose polynomials $\delta_x, \delta_y, \Delta \in \mathbb{Z}_q[x, y]$ such that $\delta_x f_x + \delta_y f_y = 1 + \Delta f$ and assume

$$x^\sigma = x^p + \delta_x Z; \qquad y^\sigma = y^p + \delta_y Z$$

for an overconvergent series $Z\langle x, y\rangle^\dagger$. Then apply a Newton iteration, such as described for instance in Hubrechts [29, Prop. 2], to the univariate polynomial

$$g(z) := f(x^p + \delta_x(x^p, y^p)z,\ y^p + \delta_y(x^p, y^p)z) - f^p - f^p\Delta(x^p, y^p)z$$

to compute a suitable Z; by Hensel's lemma this process converges. One can prove that the resulting Z is overconvergent.

The remaining steps are as above: We get a representation of the images ω^σ and $x^\sigma\omega^\sigma$ in terms of the basis $\{\omega, x\omega\}$, and applying the Lefschetz formula yields the number of points or the zeta function. In practice, one has to terminate the infinite series at some point, but if the precision is large enough, the result will be correct. △

The foregoing example outlined a procedure for lifting the Frobenius endomorphism to the whole of an affine elliptic curve — a rather laborious task, as one needs to solve for the images of x and y under σ simultaneously. This lifting process can be greatly simplified by removing a few more points from the curve (namely those with y-coordinate zero), which corresponds to rendering y invertible in the coordinate ring. Kedlaya [30] makes use of this technique; see Section 4.2 for details.

4 Kedlaya's algorithm for hyperelliptic curves

In Section 3, we have demonstrated how to apply the Lefschetz trace formula to a smooth affine variety over a finite field in order to count its rational points. The general strategy is:

1. Lift the coordinate ring to the Witt vectors, yielding a smooth finitely generated $\mathbb{Z}_q$-algebra A.
2. Determine a basis of the non-vanishing algebraic de Rham cohomology spaces of $A \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$, along with an explicit algorithm to represent any given element with respect to that basis. According to Theorem 30, this is also a basis of the Monsky-Washnitzer cohomology.
3. Obtain a bound on the required p-adic precision, taking into account the loss of precision in the reduction algorithm.
4. Construct an approximate lift of the Frobenius endomorphism to the p-adic weak completion $A^\dagger$. Calculate the images of the basis elements under the resulting Frobenius action on cohomology.
5. Using the reduction algorithm from step 2, obtain a matrix for the Frobenius action on each cohomology space, and compute the traces in the Lefschetz formula to get the number of points.

Now since electronic computing devices currently exist (as of 2017), automating this procedure as much as possible appears to be a rewarding endeavour. These ideas are what *Kedlaya's algorithm* for computing the zeta function of a hyperelliptic curve is based on: After dealing with steps 1 through 3 theoretically for the whole class of varieties at once, the remaining steps 4 and 5 are executed algorithmically for each concrete instance of a hyperelliptic curve. The runtime is polynomial in the genus of the curve as well as in $\log_p q$, but not in the bit length of p. However, for small p, the algorithm — and its enhanced variants even more so — are highly practical.

Hyperelliptic curves lie in the sweet spot of intersection between simplicity, practical relevance, and performance of the resulting algorithm, which seems to be the reason why they were chosen to demonstrate the applicability of Monsky-Washnitzer cohomology to point counting. In fact, Kedlaya does not only count points, but instead computes the full zeta function of the curve: This also yields the order of the Jacobian (Lemma 39), which is desirable for hyperelliptic curve cryptography (cf. Section 1.1). Meanwhile, Kedlaya's algorithm has been generalized to other classes of curves, such as hyperelliptic curves in characteristic 2 [15], superelliptic curves [23], and C_{ab} curves [16], underlining the success of the approach.

Among the available literature, Kedlaya's original paper [30] from 2001 is rather succinct and technical. The later survey article Kedlaya [31] contains more context, but only summarizes the major steps of the algorithm. The expositions in Edixhoven [19] and Cohen et al. [9, Section 17.3.3] are detailed and precise. For a short and practically oriented description, see Chapter 2 of Harvey's PhD thesis [27].

Our exposition of the algorithm shall first introduce the general setting we're working in. After showing a few neat tricks to simplify the algorithm and improve the performance at the same time, we formulate and analyze the resulting algorithm. Note that in the literature, these 'tricks' to simplify or optimize the algorithm are typically presented as a necessary step of the algorithm and quickly glanced over, but putting a little more emphasis on those to clarify their motivation and implications seems worthwhile to the author.

Throughout the following, assume p is an odd prime and fix an integer $g \geq 1$ and a monic squarefree polynomial $\bar{h} \in \mathbb{F}_q[x]$ of degree $2g+1$,[1] such that the smooth projective model $\mathcal{C}$ of the affine curve

$$y^2 = \bar{h}(x)$$

is a hyperelliptic curve of genus g.[2] However, since Monsky-Washnitzer cohomology only deals with affine varieties, we will mainly work with the *affine* curve

$$C := \operatorname{Spec} \mathbb{F}_q[x,y]/(y^2 - \bar{h}(x)).$$

[1] Degree $2g+2$ also gives rise to a hyperelliptic curve, but for simplicity, we restrict ourselves to the case of $2g+1$.

[2] Note the projective closure of $y^2 = h(x)$ is *not* always smooth: For $g > 1$, it has a singularity at the point at infinity.

Consistent with the notation we've used throughout Section 3, let $h(x) \in \mathbb{Z}_q[x]$ denote a monic lift of $\bar{h}(x)$ to the Witt vectors $\mathbb{Z}_q$ and set

$$A := \mathbb{Z}_q[x,y] \,/\, (y^2 - h(x));$$
$$A^\dagger := \mathbb{Z}_q\langle x,y\rangle^\dagger \,/\, (y^2 - h(x)).$$

The Monsky-Washnitzer cohomology of C has the following structure. Clearly, $H^0(C,\mathbb{Q}_q) = \mathbb{Q}_q$ and $H^i(C,\mathbb{Q}_q) = 0$ for $i \geq 2$. For the only interesting cohomology space, we have:

Lemma 38. Let $\alpha, \beta \in \mathbb{Q}_q[x]$ such that $\alpha h(x) + \beta h'(x) = 1$, and set $\omega := \alpha y \mathrm{d}x + 2\beta \mathrm{d}y \in \Omega^1 \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$. Then

$$H^1(C,\mathbb{Q}_q) = \bigoplus_{i=0}^{2g-1} \mathbb{Q}_q \cdot x^i \omega.$$

Proof. Invoking Theorem 30, this works exactly like the computation of $\operatorname{coker}\delta$ in Example 29. □

We mention the following property of the zeta function in passing, which permits obtaining the number of points on the Jacobian of a hyperelliptic curve via Kedlaya's algorithm:

Lemma 39 [9, Corollary 5.70]**.** Evaluating either

- the numerator of the zeta function $Z(\mathcal{C};t)$; or equivalently
- the characteristic polynomial of the Frobenius action $\sigma_* \colon H^1(C,\mathbb{Q}_q) \to H^1(C,\mathbb{Q}_q)$

at the value $t = 1$ yields the number of $\mathbb{F}_q$-rational points on the Jacobian of $\mathcal{C}$.
(The equivalence of both variants will follow from Corollary 48.)

Let us now work out the details of Kedlaya's algorithm — the general structure of the algorithm is as described above.

4.1 Making use of the p-power Frobenius

This section describes an optimization for the case $q \neq p$, which we shall assume throughout. While Kedlaya's algorithm is about hyperelliptic curves, it should be emphasized that there are no obstacles in applying this method to more general varieties. We abbreviate 'r-power Frobenius' by 'r-Frobenius'.

The core idea of this section is the observation that lifting the p-Frobenius to $A^\dagger$ is computationally cheaper than lifting the q-Frobenius, since it converges faster and one thus needs fewer terms of the expansion. As the latter is given by repeated applications of the former, it appears to be possible to derive a matrix for the q-Frobenius action on cohomology from a matrix of the p-Frobenius, circumventing the need for a lift to $A^\dagger$ of the full q-Frobenius. However, there is a problem: A lift of the p-Frobenius is *not* a $\mathbb{Z}_q$-algebra endomorphism of $A^\dagger$, as the Witt vector Frobenius of $\mathbb{Z}_q$ is a nontrivial automorphism (see Lemma 19). The closest to linearity we can get is a $\mathbb{Z}_q$-*semilinear* map σ, where the semilinearity[1] is with respect to the Frobenius automorphism of $\mathbb{Z}_q$. This situation may be expressed by means of the following commutative diagram: [19, Section 5.1]

$$\begin{array}{ccc} \mathbb{Z}_q & \xrightarrow{\text{Frob.}} & \mathbb{Z}_q \\ \downarrow & & \downarrow \\ A^\dagger & \xrightarrow{\sigma} & A^\dagger \end{array}$$

Although σ does not constitute a homomorphism of $\mathbb{Z}_q$-algebras, it still (by the usual pushforward construction on differentials) induces a map σ_* on cohomology, which is $\mathbb{Q}_q$-semilinear with respect to the p-Frobenius automorphism of $\mathbb{Q}_q$. To get a matrix representation of the q-Frobenius action on cohomology, which *is* a $\mathbb{Q}_q$-linear map, we can use the following lemma:

[1] Let M and N be R-modules. A group homomorphism $\varphi \colon M \to N$ is called *R-semilinear* if there exists an automorphism τ of R such that $\varphi(\alpha m) = \tau(\alpha) f(m)$ for all $\alpha \in R$ and $m \in M$. (Roughly speaking, φ is R-linear 'up to an automorphism'.)

Lemma 40. Consider a free R-module F of rank r and let $\varphi\colon F \to F$ be R-semilinear with respect to an automorphism $\tau \in \mathrm{Aut}(R)$ of order k. (Note this implies that φ^k is R-linear, i. e., an R-module homomorphism.) Fix a basis $B = \{b_1 ... b_r\}$ of F and let $M \in R^{r\times r}$ be the matrix such that $Mb_i = \varphi(b_i)$ for all $i \in \{1...r\}$. Then, a matrix representing the endomorphism φ^k with respect to B is given by the product $M \cdot M^\tau \cdots M^{\tau^{k-1}}$, where M^{τ^i} is the matrix obtained from M by applying the automorphism τ^i to each coefficient.

Proof. Let $\alpha_{ij} \in R$ denote the entry in the ith row and jth column of M, such that $\varphi(b_i) = \sum_{j=1}^r \alpha_{ij} b_j$. We prove by induction on t that $\varphi^t(b_i) = M \cdot M^\tau \cdots M^{\tau^{t-1}} \cdot b_i$ holds for all $t \in \mathbb{N}_{\geq 1}$. The base case $t = 1$ is obvious. Suppose the claim has been shown for t; then for all $i \in \{1...r\}$,

$$\varphi^{t+1}(b_i) = \varphi^t(\varphi(b_i)) = \varphi^t\Big(\sum_{j=1}^r \alpha_{ij} b_j\Big) = \sum_{j=1}^r \tau^t(\alpha_{ij})\varphi^t(b_j) = \sum_{j=1}^r M \cdot M^\tau \cdots M^{\tau^{t-1}} \cdot \tau^t(\alpha_{ij}) b_j.$$

By definition of matrix multiplication, we have $\sum_{j=1}^r \tau^t(\alpha_{ij}) b_j = M^{\tau^t} b_j$, therefore the sum above equals $M \cdot M^\tau \cdots M^{\tau^{t-1}} \cdot M^{\tau^t} \cdot b_i$ as claimed. □

According to the preceding discussion, it is sufficient to lift the p-Frobenius to $A^\dagger$ as a $\mathbb{Z}_q$-semilinear map and use a matrix of that to obtain the matrix required for the Lefschetz formula. Note that it is this shortcut in particular that makes the running time of Kedlaya's algorithm depend mostly on the characteristic p instead of the field size q: Taking larger extension degrees $\log_p q$ does not impede the performance much, since we only need to lift the p-Frobenius either way [33, Remark 2.3.3]. The details of how to compute such a lift will be worked out in Section 4.4, since this depends on a few more implementation tricks we need to discuss first.

4.2 Inverting y

Recall that one of the major complications in Example 37 arose from the denominator appearing in the Newton iteration we had to perform when solving for the image y^σ under a Frobenius lift. This pitfall can be avoided by cutting out all points with y-coordinate zero from the curve C, which amounts to adjoining an inverse of y to the coordinate ring. That is: Let

$$\bar{A} := \mathbb{F}_q[x, y, z]/(y^2 - h(x), yz - 1) \cong \mathbb{F}_q[x, y, y^{-1}]/(y^2 - h(x))$$

and consider the variety $C' := \operatorname{Spec} \bar{A}$ instead of C.[1] Clearly, C' may (and usually will) have fewer $\mathbb{F}_q$-rational points than the original curve C. In theory, this defect could be computed by extracting and counting linear factors over $\mathbb{F}_q$ from h, for instance using the algorithm of Cantor and Zassenhaus [8]; however, Section 4.3 will show that this is not even necessary, since the introduced error cancels out owing to another optimization.

To operate on the cohomology of C', we need a basis:

Lemma 41 [9, Lemma 17.76]**.** A basis of $H^1(C', \mathbb{Q}_q)$ is given by the forms

$$\Big\{ x^i \frac{\mathrm{d}x}{y} \;\Big|\; i \in \{0...2g-1\} \Big\} \cup \Big\{ x^i \frac{\mathrm{d}x}{y^2} \;\Big|\; i \in \{0...2g\} \Big\}.$$

Note that contrary to Example 29, the element y is now invertible in $A^\dagger$, hence the notation $\frac{\mathrm{d}x}{y}$ and $\frac{\mathrm{d}x}{y^2}$ does make sense.

4.3 Decomposing the cohomology into eigenspaces

By definition, a hyperelliptic curve C has a degree-2-map onto the projective line $\mathbb{P}^1$, hence there exists an automorphism of order 2 of C: the *hyperelliptic involution* $\iota \in \mathrm{Aut}(C)$. On the usual model $y^2 = h(x)$ of C, the involution ι is given simply by flipping the sign of y, which geometrically corresponds to

[1] Any point (x, y) on C with $y \neq 0$ corresponds to the point (x, y, y^{-1}) on C'.

mirroring each point of the curve at the x-axis. Obviously, ι restricts to an automorphism of the smaller curve C'. The action of ι on our basis (cf. Lemma 41) of the first Monsky-Washnitzer cohomology space of C' is given by

$$\iota_*\colon H^1(C',\mathbb{Q}_q) \longrightarrow H^1(C',\mathbb{Q}_q),\ \begin{cases} x^i\frac{\mathrm{d}x}{y} & \longmapsto\ -x^i\frac{\mathrm{d}x}{y} \\ x^i\frac{\mathrm{d}x}{y^2} & \longmapsto\ x^i\frac{\mathrm{d}x}{y^2}, \end{cases}$$

hence $H^1(C',\mathbb{Q}_q)$ splits as the direct sum of the two eigenspaces

$$H^1_-(C',\mathbb{Q}_q) = \left\langle x^i\frac{\mathrm{d}x}{y} \;\middle|\; i\in\{0,\dots,2g-1\}\right\rangle; \qquad H^1_+(C',\mathbb{Q}_q) = \left\langle x^i\frac{\mathrm{d}x}{y^2} \;\middle|\; i\in\{0,\dots,2g\}\right\rangle$$

under ι_*, with eigenvalues -1 and 1 respectively. Clearly, the lifted q-power Frobenius endomorphism σ respects the action of ι, hence its action on cohomology restricts to endomorphisms of each eigenspace.

The point of these observations is that the cohomology $H^1(C,\mathbb{Q}_q)$ of the *full* affine curve C is precisely the negative eigenspace $H^1_-(C',\mathbb{Q}_q)$, and that this isomorphism clearly respects the Frobenius action. This is a consequence to Lemma 38 above; see also Fresnel and van der Put [21, Prop. 7.6.10(vi)].[1] Therefore, according to the Lefschetz trace formula,

$$\begin{aligned}\#C(\mathbb{F}_q) &= \operatorname{tr}\left(q\sigma_*^{-1} \mid H^0(C,\mathbb{Q}_q)\right) - \operatorname{tr}\left(q\sigma_*^{-1} \mid H^1(C,\mathbb{Q}_q)\right)\\ &= \operatorname{tr}\left(q\sigma_*^{-1} \mid \mathbb{Q}_q\right) - \operatorname{tr}\left(q\sigma_*^{-1} \mid H^1_-(C',\mathbb{Q}_q)\right),\end{aligned}$$

and since σ is the identity on $\mathbb{Q}_q$, this equals

$$\#C(\mathbb{F}_q) = q - \operatorname{tr}\left(q\sigma_*^{-1} \mid H^1_-(C',\mathbb{Q}_q)\right).$$

Notice that — as advertised back in Section 4.2 — the defect introduced while passing from C to C', i. e. the number of $\mathbb{F}_q$-rational points of $C\setminus C'$, does not appear in this formula. Moreover, keep in mind that this is the number of points on the *affine* version of a hyperelliptic curve; for the projective model, add 1 to account for the point at infinity.

We shall see in Section 4.6 how to recover the zeta function of C, hence of the completed curve $\mathcal{C}$, from a matrix of the Frobenius action on $H^1_-(C',\mathbb{Q}_q)$.

4.4 Lifting Frobenius

Another benefit of inverting y in the coordinate ring is that we may now perform the following computation to get a Frobenius lift, instead of the Newton iteration described in Example 37. Due to Section 4.1, it is sufficient to lift only the p-power Frobenius. This section is based on Edixhoven [19, Section 5.2].

We start off with the p-power Witt vector Frobenius σ on $\mathbb{Z}_q$, and extend this to $\mathbb{Z}_q\langle x\rangle^\dagger$ by setting

$$x^\sigma := x^p$$

The images of x and y must satisfy the equation of the curve, hence we need to solve the equation

$$(y^\sigma)^2 = h(x)^\sigma$$

in $A^\dagger = \mathbb{Z}_q\langle x,y,y^{-1}\rangle^\dagger/(y^2-h(x))$ for y^σ. Again, one way to do this is by a Newton iteration, but in this case, there exists a simpler formula: Modulo p, the polynomials $h(x)^\sigma$ and $h(x)^p$ have the same image, thus

$$E := \frac{h(x)^\sigma - h(x)^p}{p}$$

[1] Furthermore, the positive eigenspace $H^1_+(C',\mathbb{Q}_q)$ is just the cohomology of the affine line $\{y=0\}$ minus its points of intersection with C — cf. Example 31, in which we computed the Monsky-Washnitzer cohomology of a punctured affine line. In general, these and similar statements arise from an excision exact sequence for rigid cohomology which puts the cohomology spaces of a variety, a closed subset, and its complement into relation [2, Section 3.3.1], but in the case at hand they are evident. In sum, we have isomorphisms

$$H^1_-(C',\mathbb{Q}_q) = H^1(C,\mathbb{Q}_q); \qquad H^1_+(C',\mathbb{Q}_q) = H^1(\mathbb{A}^1\mathbb{F}_q\setminus(C\setminus C'),\mathbb{Q}_q)$$

which respect the action of the Frobenius endomorphism on each of the spaces.

is a well-defined element of $\mathbb{Z}_q[x]$. Hence, in $A = \mathbb{Z}_q[x, y, y^{-1}]/(y^2 - h(x))$, we see

$$h(x)^\sigma = h(x)^p + pE = (y^2)^p + pE = y^{2p}\big(1 + py^{-2p}E\big).$$

The point of these transformations is that the right-hand side yields a square root in A^∞ of $h(x)^\sigma$ by the binomial series:[1]

$$y^\sigma := \sqrt{h(x)^\sigma} = \Big(y^{2p}\big(1 + py^{-2p}E\big)\Big)^{1/2} = y^p \cdot \sum_{k=0}^{\infty} \binom{1/2}{k} p^k E^k y^{-2pk}.$$

Since $\deg E^k = k \cdot \deg E \leq kp \cdot \deg h(x)$, the total degree (in x, y, y^{-1}) of the kth term is bounded by $(2 + \deg h(x)) \cdot pk$. Moreover, $\binom{1/2}{k}$ is always integral,[2] hence the series is overconvergent, that is, defines an element of $A^\dagger$. The image of the inverse of y is obtained the same way:

$$y^{-\sigma} := \Big(y^{2p}\big(1 + py^{-2p}E\big)\Big)^{-1/2} = y^{-p} \cdot \sum_{k=0}^{\infty} \binom{-1/2}{k} p^k E^k y^{-2pk}.$$

Again, $\binom{-1/2}{k} \in \mathbb{Z}_q$, hence $y^{-\sigma} \in A^\dagger$.

Remark. Notice that p affects the number of terms of an approximation (modulo some power of p) of these series linearly — this is the fundamental reason for the running time of Kedlaya's algorithm to depend on the characteristic p rather than the field size q, and for the linear (that is, exponential in the bit length) growth with p. ∘

The next step is to compute the action of this map on our basis $\left\{ x^i \frac{\mathrm{d}x}{y} \;\middle|\; i \in \{0 \ldots 2g-1\} \right\}$ of $H^1_-(C', \mathbb{Q}_q)$:

$$\begin{aligned}
\sigma_*\Big(x^i \frac{\mathrm{d}x}{y}\Big) &= \sigma(x^i)\frac{\mathrm{d}x^\sigma}{y^\sigma} = x^{pi} y^{-\sigma}\,\mathrm{d}x^p = x^{pi} y^{-\sigma} p x^{p-1}\,\mathrm{d}x = p x^{pi+p-1} y^{1-\sigma}\frac{\mathrm{d}x}{y} \\
&= \Big(p x^{pi+p-1} y^{1-p} \cdot \sum_{k=0}^{\infty} \binom{-1/2}{k} p^k E^k y^{-2pk}\Big)\frac{\mathrm{d}x}{y} \\
&= \sum_{k=0}^{\infty} \binom{-1/2}{k} p^{k+1} E^k x^{pi+p-1} y^{-2pk-p+1}\frac{\mathrm{d}x}{y}.
\end{aligned}$$

Algorithmically, it makes sense to precompute and cache a sufficiently good approximation of $y^{-\sigma}$ once in the beginning, instead of starting anew with the expansion for each basis element.

4.5 Reduction of differentials

To represent the action of Frobenius on $H^1_-(C', \mathbb{Q}_q)$ as a matrix, one needs to write the image of each basis element as a linear combination of the basis. While everything presented so far could be done by hand in advance, this is the first step that depends on the concrete polynomial $h(x)$, hence is reasonable to perform programmatically on a computer. Of course, it is impossible to reduce the full infinite series for $\sigma_*(x^i \frac{\mathrm{d}x}{y})$ obtained above — in practice, one works with approximations at this point. The precision required to obtain correct results from the Lefschetz formula will be determined in Section 4.7.

By applying the results of the previous section, we get (an approximation of) the image of each basis element of $H^1_-(C', \mathbb{Q}_q)$ under the Frobenius action, and need to rewrite it as a $\mathbb{Q}_q$-linear combination of the basis elements to construct a matrix of the Frobenius. To be explicit, the input to the reduction algorithm consists of a finite sum

$$\sum_{j=-M}^{M} a_j(x) y^{2j} \frac{\mathrm{d}x}{y}$$

[1] For $\alpha \in \mathbb{Q} \cap \mathbb{Z}_p$ and $z \in \mathbb{Z}_q$ of positive valuation, $(1+z)^\alpha = \sum_{k=0}^{\infty} \binom{\alpha}{k} z^k$, where $\binom{\alpha}{k} = \frac{1}{k!}\prod_{j=0}^{k-1}(\alpha - j)$. [10, Section 9]

[2] It suffices to prove this for the subring $\mathbb{Z}_p \subseteq \mathbb{Z}_q$: Let be $(\alpha_i)_{i\in\mathbb{N}}$ any sequence in $\mathbb{N}$ which converges p-adically to some $\alpha \in \mathbb{Z}_p$, e. g. this could be the partial sums of α's series expansion. The map $\mathbb{Z}_p \to \mathbb{Q}_p$, $\gamma \mapsto \binom{\gamma}{k}$ is given by a polynomial, hence continuous, thus $\binom{\alpha}{k} = \lim_{i\to\infty}\binom{\alpha_i}{k}$. However, all $\binom{\alpha_i}{k}$ are in $\mathbb{N}$ by basic combinatorics, hence their p-adic limit must lie in $\mathbb{Z}_p$.

with $a_j(x) \in \mathbb{Q}_q[x]$, and the output is a $\mathbb{Q}_q$-linear combination of the basis elements $x^0\frac{\mathrm{d}x}{y}$ ... $x^{2g-1}\frac{\mathrm{d}x}{y}$ of $H^1_-(C', \mathbb{Q}_q)$ which is cohomologous to the input. Note we only need to consider even powers of y since the image of each basis element of $H^1_-(C', \mathbb{Q}_q)$ under the Frobenius action is guaranteed to lie in the same eigenspace. Another way to see this is to directly examine the power series expansion we obtained for σ_* on $H^1_-(C', \mathbb{Q}_q)$ in Section 4.4.

For this purpose, we use the following relations as ingredients to a reduction algorithm:

Lemma 42. Let $f(x) \in \mathbb{Q}_q[x]$ and let $\alpha, \beta \in \mathbb{Q}_q[x]$ such that $\alpha h(x) + \beta h'(x) = f(x)$. Then, for all $s \geq 2$,

$$f(x)y^{-s}\frac{\mathrm{d}x}{y} = \big(\alpha + \tfrac{2}{s-1}\beta'\big)y^{-s+2}\frac{\mathrm{d}x}{y}$$

holds in $H^1_-(C', \mathbb{Q}_q)$.

Proof. Recalling that $\mathrm{d}y = \frac{1}{2}h'(x)\frac{\mathrm{d}x}{y}$ in Ω^1, we have

$$\mathrm{d}\big(\beta y^{-s+1}\big) = \beta' y^{-s+1}\mathrm{d}x + (-s+1)\beta y^{-s}\mathrm{d}y = \beta' y^{-s+2}\frac{\mathrm{d}x}{y} - \tfrac{s-1}{2}\beta h'(x)y^{-s}\frac{\mathrm{d}x}{y}, \qquad (*)$$

hence

$$f(x)y^{-s}\frac{\mathrm{d}x}{y} = \alpha y^{-s+2}\frac{\mathrm{d}x}{y} + \beta h'(x)y^{-s}\frac{\mathrm{d}x}{y} \underset{(*)}{=} \alpha y^{-s+2}\frac{\mathrm{d}x}{y} + \tfrac{2}{s-1}\beta' y^{-s+2}\frac{\mathrm{d}x}{y}. \qquad \square$$

Lemma 43. Let $f(x) \in \mathbb{Q}_q[x]$ of degree $i \geq 2g$. Then, in $H^1_-(C', \mathbb{Q}_q)$,

$$f(x)\frac{\mathrm{d}x}{y} = \Big(f(x) - \tfrac{2}{2i-2g+1}\big((i-2g)x^{i-2g-1}h(x) + \tfrac{1}{2}x^{i-2g}h'(x)\big)\Big)\frac{\mathrm{d}x}{y},$$

which removes the highest power of x from $f(x)\frac{\mathrm{d}x}{y}$. Thus, by repeatedly applying this relation, one may reduce $f(x)$ to a polynomial of degree $< 2g$ without changing the represented differential.

Proof. Observe

$$\mathrm{d}(x^{i-2g}y) = (i-2g)x^{i-2g-1}y\mathrm{d}x + x^{i-2g}\mathrm{d}y = \big((i-2g)x^{i-2g-1}h(x) + \tfrac{1}{2}x^{i-2g}h'(x)\big)\frac{\mathrm{d}x}{y}.$$

The polynomial of the rightmost differential has degree i and lead coefficient $i-2g+\frac{2g+1}{2} = i-g+\frac{1}{2}$. $\square$

Lemma 44. If $f(x) \in \mathbb{Q}_q[x]$ and $j \geq 0$, then

$$f(x)y^{2j}\frac{\mathrm{d}x}{y} = f(x)h(x)^j\frac{\mathrm{d}x}{y}.$$

Clearly, these calculations permit reducing any finite sum of differentials in $H^1_-(C', \mathbb{Q}_q)$ as above to a $\mathbb{Q}_q$-linear combination of our basis. In summary, we have the following algorithm:

Algorithm 45 (Reduction of differentials)**.**
Input. A finite sum $\sum_{j=-M}^{M} a_j(x)y^{2j}\frac{\mathrm{d}x}{y} \in H^1_-(C', \mathbb{Q}_q)$ of differentials, where $a_j \in \mathbb{Q}_q[x]$.
Output. The unique $\mathbb{Q}_q$-linear combination of $x^0\frac{\mathrm{d}x}{y}$ ... $x^{2g-1}\frac{\mathrm{d}x}{y}$ that is cohomologous to the input.

1. Initialize $\omega := 0$.
2. **For each** $j \in \{-M \ldots M\}$:
   1) **If** $j \geq 0$, **then** use Lemma 44 to rewrite $a_j(x)y^j\frac{\mathrm{d}x}{y}$ as $b_j(x)\frac{\mathrm{d}x}{y}$ with $b_j \in \mathbb{Q}_q[x]$.
   2) **If** $j \leq -1$, **then** repeatedly apply Lemma 42 to rewrite $a_j(x)y^j\frac{\mathrm{d}x}{y}$ as $b_j(x)\frac{\mathrm{d}x}{y}$ with $b_j \in \mathbb{Q}_q[x]$.
   3) Repeatedly apply Lemma 43 to reduce $b_j(x)\frac{\mathrm{d}x}{y}$ to $c_j\frac{\mathrm{d}x}{y}$ with $\deg c_j < 2g$.
   4) Add $c_j\frac{\mathrm{d}x}{y}$ to ω.
3. **Return** ω.

We postpone the analysis of runtime and space complexity to Section 4.9.

For now, it should only be pointed out that, while the above procedure is certainly correct and comprehensible, it is not optimal: In 2008, Harvey devised a method known as *controlled reduction*, which causes the reduction algorithm to deal with sparse instead of dense polynomials. This decreases the dependency on p of the runtime to $\mathcal{O}(\sqrt{p})$, rather than $\mathcal{O}(p)$ as in Kedlaya's original algorithm. [27]

4.6 Recovering the zeta function

This section is concerned with calculating the zeta function of the completed curve $\mathcal{C}$ from a matrix of the Frobenius action on $H^1_-(C', \mathbb{Q}_q)$, as obtained by the reduction Algorithm 45. To this end, we make excessive use of the Weil conjectures (Theorem 2) and the isomorphism $H^1_-(C', \mathbb{Q}_q) = H^1(C, \mathbb{Q}_q)$.

In this section, let σ be a lift of the q-power Frobenius to $A^\dagger$ and let $\sigma_*\colon H^1(C, \mathbb{Q}_q) \to H^1(C, \mathbb{Q}_q)$ denote the induced map on cohomology.

First of all, there is the following general expression for the zeta function:

Lemma 46. The zeta function of the *projective* model $\mathcal{C}$ of C is of the form

$$Z(\mathcal{C}; t) = \frac{L(t)}{(1-t)(1-qt)}$$

where the numerator $L(t) \in \mathbb{Z}[t]$ is a polynomial of degree $2g$; it may be computed as

$$L(t) = \det\left(1 - tq\sigma_*^{-1} \mid H^1(C, \mathbb{Q}_q)\right).$$

Proof. Corollary 33 yields

$$Z(C; t) = \frac{\det\left(1 - tq\sigma_*^{-1} \mid H^1(C, \mathbb{Q}_q)\right)}{\det(1 - tq)}.$$

Since $\mathcal{C} \setminus C$ consists of exactly one point, this implies

$$Z(\mathcal{C}; t) = Z(C; t) \cdot Z(\mathcal{C} \setminus C; t) = \frac{L(t)}{1-qt} \cdot \frac{1}{1-t}.$$

By Lemma 38, the dimension of $H^1(C, \mathbb{Q}_q)$ is $2g$, hence $\deg L(t) = \dim H^1(C, \mathbb{Q}_q) = 2g$. □

Using the analogue of the Riemann hypothesis and the functional equation, this may be related to the eigenvalues of σ_*:

Lemma 47. The algebraic integers $\alpha_1 \ldots \alpha_{2g}$ in the numerator $L(t) = \prod_{i=1}^{2g}(1 - \alpha_i t)$ of the zeta function $Z(\mathcal{C}; t)$ may be rearranged such that

$$\alpha_i \alpha_{i+g} = q$$

for any $i \in \{1 \ldots g\}$. Furthermore, the α_i are the eigenvalues of $\sigma_*\colon H^1(C, \mathbb{Q}_q) \to H^1(C, \mathbb{Q}_q)$.

Proof. For a smooth projective curve of genus g, the self-intersection number of the diagonal is the Euler characteristic $E = 2 - 2g$, hence the functional equation of the zeta function (Theorem 2(ii)) states

$$Z(\mathcal{C}; q^{-1}t^{-1}) = \pm q^{1-g} t^{2-2g} Z(\mathcal{C}; t).$$

Since on the other hand

$$Z(\mathcal{C}; q^{-1}t^{-1}) = \frac{L(q^{-1}t^{-1})}{(1-t^{-1})(1-q^{-1}t^{-1})} = \frac{qt^2 L(q^{-1}t^{-1})}{(1-t)(1-qt)},$$

this implies

$$L(t) = \pm q^g t^{2g} L(q^{-1}t^{-1}).$$

Clearly, these polynomials must have the same set of roots: For each root $1/\alpha_i$ of $L(t)$, the number α_i/q must therefore also be a root of $L(t)$. Grouping those dual pairs together shows the first claim.

For the second claim, we use the formula for $L(t)$ from Lemma 46:

$$L(t) = \det\left(1 - tq\sigma_*^{-1} \mid H^1(C, \mathbb{Q}_q)\right) = t^{-2g} \cdot \det\left(t - q\sigma_*^{-1} \mid H^1(C, \mathbb{Q}_q)\right).$$

According to this, the roots $1/\alpha_i$ of $L(t)$ are eigenvalues of $q\sigma_*^{-1}$; hence, each $\alpha_{i+g} = \alpha_i/q$ is an eigenvalue of σ_* with the same eigenspace (and vice-versa). □

From the preceding results, we finally obtain an easy way to compute $Z(\mathcal{C}; t)$ from a matrix of σ_*:

Corollary 48. Let

$$\chi(t) := \det\left(t - \sigma_* \mid H^1(C, \mathbb{Q}_q)\right)$$

denote the characteristic polynomial of σ_*. Then the zeta function of $\mathcal{C}$ is

$$Z(\mathcal{C}; t) = \frac{t^{2g}\chi(1/t)}{(1-t)(1-qt)}.$$

Notice that $t^{2g}\chi(1/t)$ is the 'reverse' of $\chi(t)$: the coefficient of $t^{2g}\chi(1/t)$ corresponding to t^j is just the coefficient of $\chi(t)$ corresponding to t^{2g-j}.

Using the analogue of the Riemann hypothesis, we can even go one step further and exploit the symmetry of the zeta function to rely only on half of the coefficients of $\chi(t)$:

Lemma 49. The coefficients of the numerator $L(t) = \sum_{j=0}^{2g} a_j t^j$ of the zeta function $Z(\mathcal{C}; t)$ satisfy

$$a_{2g-j} = q^{g-j} a_j$$

for all $j \in \{0 \ldots g\}$.

Proof. Via Lemma 47, let $L(t) = \prod_{i=1}^{2g}(1 - \alpha_i t)$ such that $\alpha_{i+g} = q/\alpha_i$ for $i \in \{1 \ldots g\}$. Note this implies the product of all α_i is q^g. Let $\mathcal{S}_r$ denote the set of all subsets of cardinality r of $\{1 \ldots 2g\}$. Then

$$(-1)^j a_{2g-j} = \sum_{S \in \mathcal{S}_{2g-j}} \prod_{i \in S} \alpha_i = \sum_{S \in \mathcal{S}_j} q^g \prod_{i \in S} \frac{1}{\alpha_i} = \sum_{S \in \mathcal{S}_j} q^g \prod_{i \in S} \frac{\alpha_{i+g}}{q} = q^{g-j} \sum_{S \in \mathcal{S}_j} \prod_{i \in S} \alpha_i = q^{g-j}(-1)^j a_j,$$

where the indices of the α_i are taken modulo $2g$ and we use that S ranges over all subsets — hence shifting the indices by any amount does not change the sum. □

Summarily, the results in this section imply that $Z(\mathcal{C}; t)$ is uniquely defined by the $g + 1$ high-order coefficients of the characteristic polynomial $\chi(t)$ of the Frobenius action $\sigma_* \colon H^1(C, \mathbb{Q}_q) \to H^1(C, \mathbb{Q}_q)$.

4.7 Estimating precision

Up until now, all treatment of the algorithm has been ignorant of the difficulties introduced by truncating the Witt vectors and power series at some point, as enforced by the finiteness of a computing machine — this shall be remedied in this section. The analysis is twofold: We first derive a lower bound on the required precision for the Frobenius matrix, and then investigate its implications on where to cut off the power series expansion of $y^{-\sigma}$.

4.7.1 The Frobenius matrix

To bound the required precision such that the approximated Frobenius matrix yields the correct zeta function, the Weil conjectures (Theorem 2) turn out immensely useful. Recall that $Z(\mathcal{C}; t) = \frac{L(t)}{(1-t)(1-qt)}$.

Lemma 50. If $L(t) = \sum_{j=0}^{2g} a_j t^j$, then each $|a_j| \leq \binom{2g}{j} q^{j/2}$.

Proof. From the analogue of the Riemann hypothesis, we know that $L(t) = \prod_{i=1}^{2g}(1 - \alpha_i t)$ for complex α_i of absolute value $\sqrt{q}$. The coefficient a_j of t^j in $L(t)$ is (modulo sign) a sum of all possible products of j such coefficients. Any product of j complex numbers of absolute value $\sqrt{q}$ clearly has absolute value $q^{j/2}$, and since there are $\binom{2g}{j}$ such products, the claim follows using the triangle inequality. □

Corollary 51. Let $n = \log_p q$. To uniquely determine the zeta function of $\mathcal{C}$, it suffices to compute the matrix of Frobenius modulo p^B, where

$$B := \left\lceil (g/2) \cdot n + \log_p \left(2 \cdot \tbinom{2g}{g}\right) \right\rceil.$$

Proof. According to Lemma 49, computing only the coefficients $a_0 ... a_g$ of $L(t) = \sum_{j=0}^{2g} a_j t^j$ is enough. By Lemma 50, these are uniquely determined by their image in an interval of size at least $2 \cdot \binom{2g}{g} q^{g/2}$. □

Note these estimates are by no means optimal: In fact, Kedlaya [34] has shown that

$$p^B > \max\{ \tfrac{4g}{i} q^{i/2} \mid i \in \{1...g\} \}$$

suffices to uniquely determine the zeta function.

Algorithm 52 (Recovering the zeta function)**.**
Input. An approximation M modulo p^B of the matrix of the q-power Frobenius action on $H^1_-(C', \mathbb{Q}_q)$.
Output. The zeta function of the completed hyperelliptic curve $\mathcal{C}$.

1. Compute the characteristic polynomial $\chi(t) := \det(t - M) \in (\mathbb{Z}/(p^B))[t]$.
2. Let $\gamma_j \in \mathbb{Z}$ denote symmetric representants of the coefficients of $\chi(t)$, such that $\chi(t) = \sum_{j=0}^{2g} \gamma_j t^j$.
3. **For each** $j \in \{0...g\}$
   1) Set $a_j := \gamma_{2g-j}$.
   2) Set $a_{2g-j} := q^{g-j} a_j$.
4. Set $L(t) := \sum_{j=0}^{2g} a_j t^j$ and **return** $\frac{L(t)}{(1-t)(1-qt)}$.

The correctness is evident; the complexity will be treated in Section 4.9.

4.7.2 Denominators introduced by the reduction

The second part of the precision analysis is concerned with the loss of p-adic precision in the reduction Algorithm 45 due to multiples of p appearing in denominators. To illustrate this issue, consider the differential $5^2 x^3 \frac{\mathrm{d}x}{y}$ in the setting $h(x) = x^3 + 1 \in \mathbb{Z}_5[x]$: This superficially looks as if it reduced to zero modulo $p^B = 5^2$, so one might be tempted to dismiss it right away in the algorithm. However, applying the reduction step from Lemma 43 brings the presumed dead differential back into the realms of what we're interested in:

$$5^2 x^3 \frac{\mathrm{d}x}{y} = 5^2 \big(x^3 - \tfrac{2}{5}(h(x) + \tfrac{1}{2} x h'(x))\big) \frac{\mathrm{d}x}{y} = 5^2 \big(x^3 - \tfrac{2}{5}(\tfrac{5}{2} x^3 + 1)\big) \frac{\mathrm{d}x}{y} = -5 \cdot 2 \frac{\mathrm{d}x}{y}$$

hence one needs to expand the series for y^σ further than it might seem judging only from the powers of p in the formula derived in Section 4.4. The question now is: How many terms do we need?

To answer this, Kedlaya proves the following bounds:

Lemma 53 [50, Prop. 3.2.1]**.** Let $f(x) \in \mathbb{Z}_q[x]$ of degree $\leq 2g$ and $j \in \mathbb{Z}$. Then the reduction of $f(x) y^{2j} \frac{\mathrm{d}x}{y}$ becomes integral upon multiplication by p^e, where

$$e = \begin{cases} \lfloor \log_p(1 - 2j) \rfloor & \text{if } j < 0; \\ \lfloor \log_p((2g+1)(2j+1)) \rfloor & \text{if } j \geq 0. \end{cases}$$

The proofs are quite technical, hence omitted.

Remark. Note these bounds provide an alternate direct proof for Theorem 30 in this instance: Roughly speaking, the reduction of an overconvergent series of differentials must converge, since terms with high powers of p can not change the coefficients of lower powers of p in the reduced differential. Hence $H^1(C, \mathbb{Q}_q)$ is generated by the same forms as the algebraic de Rham cohomology. [19, Prop. 4.3.2] ∘

4.7.3 Denominators in the Frobenius matrix

Note the matrix M of the p-power Frobenius action on $H^1_-(C', \mathbb{Q}_q)$ generally does not have integral coefficients. Based on a proof sketch by Edixhoven [19, Prop. 5.3.1], van den Bogaart proves that the denominators in this matrix are small:

Lemma 54 [48, Prop. 5.4]**.** The matrix of the p-power Frobenius action on $H^1_-(C', \mathbb{Q}_q)$ becomes integral upon multiplication by $p^{\lfloor \log_p(2g-1) \rfloor}$.

4.7.4 Approximating the power series

Using the results from the preceding sections, a tedious calculation [19, Section 5.3] shows that it suffices to approximate

$$\sigma_*\Big(x^i \frac{\mathrm{d}x}{y}\Big) = \sum_{k=0}^{\infty} \binom{-1/2}{k} p^{k+1} E^k x^{pi+p-1} y^{-2pk-p+1} \frac{\mathrm{d}x}{y} \qquad \text{(cf. Section 4.4)}$$

to p-adic precision

$$N := B + \nu_p(2g+1) + \big\lfloor \log_p(2g+1-2/p) \big\rfloor + \big\lfloor \log_p(2g-1) \big\rfloor.$$

An only slightly less tedious computation using Lemma 53 then shows that the reduction of

$$\sigma_*\Big(x^i \frac{\mathrm{d}x}{y}\Big) \approx \sum_{k=0}^{K} \binom{-1/2}{k} p^{k+1} E^k x^{pi+p-1} y^{-2pk-p+1} \frac{\mathrm{d}x}{y}$$

is correct modulo p^N if K is an (e. g. the smallest) integer such that

$$K - \lfloor \log_p(2K+1) \rfloor > N.$$

4.8 The full algorithm

We summarize Kedlaya's algorithm, as resulting from the preceding discussion.

Algorithm 55 (Kedlaya)**.**
Input. An odd prime power $q = p^n$ and a monic squarefree polynomial $\bar{h}(x) \in \mathbb{F}_q[x]$ of degree $2g+1$.
Output. The zeta function $Z(\mathcal{C}; t)$ of the hyperelliptic curve $\mathcal{C}$ given by $y^2 = \bar{h}(x)$ over $\mathbb{F}_q$.

1. Arbitrarily lift the coefficients of $\bar{h}(x)$ to $\mathbb{Z}_q$ to obtain $h(x) \in \mathbb{Z}_q[x]$.
2. Set $B := \lceil (g/2) \cdot n + \log_p \left(2 \cdot \binom{2g}{g}\right) \rceil$.
3. Set $N := B + \nu_p(2g+1) + \lfloor \log_p(2g+1-2/p) \rfloor + \lfloor \log_p(2g-1) \rfloor$.
4. Determine K as the smallest integer such that $K - \lfloor \log_p(2K+1) \rfloor > N$.
5. Compute $\eta := \sum_{k=0}^{K} \binom{-1/2}{k} p^k E^k y^{-2pk}$ modulo p^N. (This is an approximation of $y^{-\sigma}$.)
6. Initialize the matrix $M \in \mathbb{Q}_q^{2g \times 2g}$ with zeroes. (This will be a matrix of the p-power Frobenius action.)
7. **For each** $i \in \{0 \ldots 2g-1\}$:
   1) Set $\omega_i := p x^{pi+p-1} \eta \mathrm{d}x \in H^1_-(C', \mathbb{Q}_q)$. (This is an approximation of $\sigma_*(x^i \frac{\mathrm{d}x}{y})$.
   2) Using Algorithm 45, compute a form $\sum_{j=0}^{2g-1} m_j x^j \frac{\mathrm{d}x}{y}$ that is cohomologous to ω_i.
   3) For all $j \in \{0 \ldots 2g-1\}$, set the entry of M at column $i+1$ and row $j+1$ to m_j.
8. Compute $M' := M \cdot M^{\sigma} \cdots M^{\sigma^{n-1}}$. (This is a matrix of the q-power Frobenius action.)
9. Use Algorithm 52 to compute the zeta function $Z(\mathcal{C}; t)$ from M' and **return** it.

4.9 Complexity analysis

Let $n := \log_p q$. We keep the characteristic p fixed — it is clear that the running time grows essentially linearly with p, which is why the algorithm performs a lot better for small characteristic. To begin with, note that B, N, K are all in $\mathcal{O}(gn)$. Moreover, [29]

- Arithmetic in $\mathbb{Z}_q$ modulo p^N requires $\mathcal{O}(n^{1+\varepsilon}N^{1+\varepsilon}) \subseteq \mathcal{O}(g^{1+\varepsilon}n^{2+\varepsilon})$ elementary operations, and storing an element of $\mathbb{Z}_q/(p^N)$ uses $\mathcal{O}(nN) \subseteq \mathcal{O}(gn^2)$ bits of memory.
- Applying any power of the Witt vector Frobenius σ to an element of $\mathbb{Z}_q$ takes time $\mathcal{O}(g^{1+\varepsilon}n^{3+\varepsilon})$.

This allows us to deduce the asymptotic complexity of Kedlaya's algorithm:

- The computation of $y^{-\sigma}$ up to p-adic precision N in step 5 can be efficiently implemented using a Newton iteration [9, Equation 17.34]. We compute $\mathcal{O}(gn)$ terms of $y^{-\sigma}$, each of which contains a polynomial of degree $< 2g$; hence the computation of η takes time $\mathcal{O}(g^{3+\varepsilon}n^{3+\varepsilon})$ and η requires space $\mathcal{O}(g^3n^3)$.
- In each reduction step performed by Algorithm 45, the most expensive operation is the extended $\gcd$ computation to obtain $\alpha, \beta \in \mathbb{Q}_q[x]$ with $\alpha h(x) + \beta h'(x) = f(x)$. This can be optimized by precomputing $\alpha, \beta \in \mathbb{Q}_q[x]$ for the case $f(x) = 1$, from which one easily obtains coefficients for $f(x)$ in each iteration. The extended $\gcd$ can be computed in time $\mathcal{O}(g^{2+\varepsilon}n^{2+\varepsilon})$, and since we reduce $2g$ forms using $\mathcal{O}(gn)$ reduction steps each, the total time for step 2) is $\mathcal{O}(g^{4+\varepsilon}n^{3+\varepsilon})$.
- The computation of M' from M in step 8 can be optimized using a variant of the well-known exponentiation-by-squaring algorithm which interleaves twists of the matrix by powers of σ; this takes time $\mathcal{O}(g^{3+\varepsilon}n^{3+\varepsilon})$.
- The characteristic polynomial of M' in Algorithm 52 can be computed in time $\mathcal{O}(g^{4+\varepsilon}n^{2+\varepsilon})$.

In summary, the running times of all individual steps are dominated by the cost $\mathcal{O}(g^{4+\varepsilon}n^{3+\varepsilon})$ of the reduction Algorithm 45. The space requirement is essentially that of η. Therefore:

Theorem 56 [30]**.** Algorithm 55 is correct and requires time $\mathcal{O}(g^{4+\varepsilon}n^{3+\varepsilon})$ and space $\mathcal{O}(g^{3+\varepsilon}n^{3+\varepsilon})$.

References

[1] Timothy G. Abbott, Kiran S. Kedlaya, and David Roe. 'Bounding Picard numbers of surfaces using p-adic cohomology'. Version 2 (Jan. 2007). URL: https://arxiv.org/abs/math/0601508.

[2] Zubair Ashraf, Ali Juma, and Pramathanath Sastry. 'Report on the Denef-Vercauteren/Kedlaya Algorithm'. *Algebraic Curves and Cryptography*. Ed. by Kumar Murty. Vol. 58. Fields Institute Communications. American Mathematical Society, Nov. 2010, pp. 65–82. ISBN: 978-0-8218-4311-6.

[3] Daniel J. Bernstein. *Surface1271: high-speed genus-2-hyperelliptic-curve cryptography*. Sept. 2006. URL: https://cr.yp.to/hecdh.html.

[4] Pierre Berthelot. *Cohomologie Cristalline des Schémas de Caractéristique $p > 0$*. French. Vol. 407. Lecture Notes in Mathematics. Springer, 1974. ISBN: 978-3-540-37807-5.

[5] Pierre Berthelot. 'Géométrie rigide et cohomologie des variétés algébriques de caractéristique p'. French. *Mémoires de la Société Mathématique de France* 23 (1986), pp. 7–32.

[6] Amnon Besser, François-Renaud Escriva, and Rob de Jeu. 'Frobenius lifts and point counting for smooth curves'. Version 1 (June 2013). URL: https://arxiv.org/abs/1306.5102.

[7] Siegfried Bosch, Ulrich Güntzer, and Reinhold Remmert. *Non-Archimedean Analysis. A Systematic Approach to Rigid Analytic Geometry*. Vol. 261. Grundlehren der Mathematischen Wissenschaften. Springer, 1984. ISBN: 978-3-540-12546-4.

[8] David G. Cantor and Hans Zassenhaus. 'A new algorithm for factoring polynomials over finite fields'. *Mathematics of Computation* 36 (1981), pp. 587–592.

[9] Henri Cohen et al. *Handbook of Elliptic and Hyperelliptic Curve Cryptography*. 1st ed. Discrete Mathematics and its Applications. Chapman & Hall/CRC, 2006. ISBN: 978-1-58488-518-4.

[10] Keith Conrad. *Infinite series in p-adic fields*. Expository paper. Version 2016-12-23. URL: http://www.math.uconn.edu/~kconrad/blurbs/gradnumthy/infseriespadic.pdf.

[11] Vladimir I. Danilov. 'Cohomology of Algebraic Varieties'. *Algebraic Geometry II*. Ed. by Igor R. Shafarevich. Vol. 35. Encyclopaedia of Mathematical Sciences. Springer, 1992. ISBN: 978-3-642-60925-1.

[12] Aise Johan de Jong. *Weil cohomology theories*. Seminar notes. Columbia University, Oct. 2007. URL: http://math.columbia.edu/~dejong/seminar/note_on_weil_cohomology.pdf.

[13] Pierre Deligne. 'La conjecture de Weil: I'. French. *Publications Mathématiques de l'IHÉS* 43 (1974), pp. 273–307.

[14] Pierre Deligne et al. *Cohomologie étale. Séminaire de Géométrie Algébrique du Bois-Marie (SGA $4\frac{1}{2}$)*. French. Vol. 569. Lecture Notes in Mathematics. Springer, 1977.

[15] Jan Denef and Frederik Vercauteren. 'An Extension of Kedlaya's Algorithm to Hyperelliptic Curves in Characteristic 2'. *Journal of Cryptology* 19.1 (2006), pp. 1–25.

[16] Jan Denef and Frederik Vercauteren. 'Counting Points on C_{ab} Curves Using Monsky-Washnitzer Cohomology'. *Finite Fields and Their Applications* 12.1 (Jan. 2006), pp. 78–102.

[17] Whitfield Diffie and Martin E. Hellman. 'New directions in cryptography'. *IEEE Transactions on Information Theory* 22.6 (Nov. 1976), pp. 644–654.

[18] Bernard Dwork. 'On the Rationality of the Zeta Function of an Algebraic Variety'. *American Journal of Mathematics* 82.3 (July 1960), pp. 631–648.

[19] Bas Edixhoven. *Point counting after Kedlaya*. Lecture notes. Version 2006-10-25. Leiden, Sept. 2003. URL: http://pub.math.leidenuniv.nl/~edixhovensj/oww/mathofcrypt/carls_edixhoven/kedlaya.pdf.

[20] Renée Elkik. 'Solutions d'équations à coefficients dans un anneau hensélien'. French. *Annales scientifiques de l'École Normale Supérieure* 6.4 (1973), pp. 553–603.

[21] Jean Fresnel and Marius van der Put. *Rigid Analytic Geometry and Its Applications*. Vol. 218. Progress in Mathematics. Birkhäuser Boston, 2004. ISBN: 978-1-4612-0041-3.

[22] Paul Garrett. *Topological vectorspaces*. Lecture notes. University of Minnesota, July 2011. URL: http://www-users.math.umn.edu/~Garrett/m/fun/Notes/05_tvs.pdf.

[23] Pierrick Gaudry and Nicolas Gürel. 'An Extension of Kedlaya's Point-Counting Algorithm to Superelliptic Curves'. *Proceedings of the 7th International Conference on the Theory and Application of Cryptology and Information Security: Advances in Cryptology*. ASIACRYPT '01. Springer, 2001, pp. 480–494. ISBN: 3-540-42987-5.

[24] Pierrick Gaudry and Éric Schost. 'Genus 2 point counting over prime fields'. *Journal of Symbolic Computation* 47 (2012), pp. 368–400.

[25] Alexander Grothendieck et al. *Théorie des Topos et Cohomologie Étale des Schémas. Séminaire de Géométrie Algébrique du Bois-Marie (SGA 4)*. French. Vol. 269, 270, 305. Lecture Notes in Mathematics. Springer, 1972–1973.

[26] Robin Hartshorne. *Algebraic Geometry*. 1st ed. Graduate Texts in Mathematics 52. Springer, 1977. ISBN: 978-1-4419-2807-8.

[27] David Michael Harvey. 'Algorithms for p-adic cohomology and p-adic heights'. PhD thesis. Harvard University, Apr. 2008.

[28] Michiel Hazewinkel. 'Witt Vectors. Part 1'. Version 1 (Apr. 2008). URL: https://arxiv.org/abs/0804.3888.

[29] Hendrik Hubrechts. 'Fast arithmetic in unramified p-adic fields'. *Finite Fields and Their Applications* 16 (3 May 2010), pp. 155–162.

[30] Kiran S. Kedlaya. 'Counting Points on Hyperelliptic Curves using Monsky-Washnitzer Cohomology'. Version 2 (Nov. 2001). URL: https://arxiv.org/abs/math/0105031.

[31] Kiran S. Kedlaya. 'Computing Zeta Functions via p-Adic Cohomology'. *Algorithmic Number Theory: 6th International Symposium*. Ed. by Duncan Buell. Vol. 3076. Lecture Notes in Computer Science. Springer, June 2004, pp. 1–17. ISBN: 978-3-540-24847-7.

[32] Kiran S. Kedlaya. 'Fourier transforms and p-adic "Weil II"'. Version 3 (July 2005). URL: https://arxiv.org/abs/math/0210149.

[33] Kiran S. Kedlaya. 'p-adic cohomology: from theory to practice'. Arizona Winter School 2007 Lecture Notes (2007). URL: http://swc.math.arizona.edu/aws/2007/notes.html.

[34] Kiran S. Kedlaya. *Getting precise about precision*. Presentation slides. Oxford, Mar. 2010.

[35] Steven L. Kleiman. 'Algebraic cycles and the Weil conjectures'. *Dix Exposés sur la Cohomologie des Schémas*. Ed. by A. Grothendieck and N. H. Kuiper. Vol. 3. Advanced Studies in Pure Mathematics. North-Holland Publishing Company, 1968, pp. 359–386.

[36] Ernst Kunz. *Kähler Differentials*. Advanced Lectures in Mathematics. Springer, 1986. ISBN: 978-3-663-14074-0.

[37] Alan G. B. Lauder. 'Deformation Theory and The Computation of Zeta Functions'. *Proceedings of the London Mathematical Society* 88 (3 May 2004), pp. 565–602.

[38] Victor S. Miller. 'Use of Elliptic Curves in Cryptography'. *Advances in Cryptography — CRYPTO '85 Proceedings*. Ed. by Hugh C. Williams. Springer, 1986, pp. 417–426. ISBN: 978-3-540-39799-1.

[39] James S. Milne. *Lectures on Étale Cohomology*. Lecture notes. Version 2.21. University of Michigan, Mar. 2013. URL: http://jmilne.org/math/CourseNotes/lec.html.

[40] Paul Monsky and Gerard Washnitzer. 'Formal Cohomology: I'. *Annals of Mathematics*. 2nd ser. 88.2 (Sept. 1968), pp. 181–217.

[41] Mircea Mustaţă. *Zeta functions in algebraic geometry*. Lecture notes. University of Michigan, 2011. URL: http://math.lsa.umich.edu/~mmustata/zeta_book.pdf.

[42] Takakazu Satoh, Berit Skjernaa, and Yuichiro Taguchi. 'Fast computation of canonical lifts of elliptic curves and its application to point counting'. *Finite Fields and Their Applications* 9 (1 Jan. 2003), pp. 89–101.

[43] René Schoof. 'Elliptic curves over finite fields and the computation of square roots $\bmod p$'. *Mathematics of Computation* 44.170 (Apr. 1985), pp. 483–494.

[44] René Schoof. 'Counting points on elliptic curves over finite fields'. *Journal de Théorie des Nombres de Bordeaux* 7 (1995), pp. 219–254.

[45] Jean-Pierre Serre. *Local Fields*. 1st ed. Graduate Texts in Mathematics 67. Springer, 1979. ISBN: 978-1-4757-5673-9.

[46] Jan Tuitman. *A survey of p-adic point counting*. Presentation slides. Katholieke Universiteit Leuven, Mar. 2016.

[47] Jan Tuitman. 'Counting points on curves using a map to $\mathbb{P}^1$, II'. Version 2 (Sept. 2016). URL: https://arxiv.org/abs/1412.7217.

[48] Theo van den Bogaart. 'About the choice of a basis in Kedlaya's algorithm'. Version 1 (Sept. 2008). URL: https://arxiv.org/abs/math/0809.1243.

[49] Marius van der Put. 'The cohomology of Monsky and Washnitzer'. *Mémoires de la Société Mathématique de France* 23 (1986), pp. 33–59.

[50] Masha Vlasenko. *p-adic cohomology and counting points on varieties over finite fields*. Lecture notes. ICTP Trieste, Sept. 2014. URL: http://imath.kiev.ua/~mariyka/publ/trieste_notes.pdf.

[51] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. 3rd Edition. Cambridge University Press, 2013. ISBN: 978-1-107-03903-2.

[52] André Weil. *Sur les courbes algébriques et les variétés qui s'en déduisent*. French. Vol. 1041. Actualités scientifiques et industrielles. Hermann & Cie, 1948.