Point counting after Kedlaya, EIDMA-Stieltjes Graduate course, Leiden, September 22–26, 2003

Bas Edixhoven

October 23, 2007

1 Motivation, plans

This series of three lectures of one hour each was preceded by two introductory lectures by Henk van Tilborg about applications of discrete logarithms (in multiplicative groups as well as elliptic curves) to cryptography. Those introductory lectures should now serve as motivation for the coming three lectures. Large cyclic subgroups of prime order in elliptic curves or in Jacobians of higher genus curves are useful in the cryptographic applications because at present there seems to be no sub-exponential algorithm known for the discrete log problem in this context (except in some very special cases), in contrast to the index calculus algorithm for the discrete log problem in the multiplicative group of a finite field. This state of affairs is then used as justification for a smaller block size, preserving the security.

Let us consider the case of elliptic curves. As one needs subgroups of prime order, it is crucial to know the exact order of groups such as $E(\mathbb{F}_q)$, where E is an elliptic curve over a finite field $\mathbb{F}_q$, as well as the factorization into prime factors of them. In practice one uses finite fields $\mathbb{F}_q$ with q of size at least about 200 binary digits.

At the start of elliptic curve cryptography around 1986 there was only one algorithm to compute $\#E(\mathbb{F}_q)$, namely, Schoof's algorithm, running in time $O((\log q)^4)$ (after improvements by Atkin, Elkies and Couveignes; Schoof's original algorithm had running time $O((\log q)^5)$). Schoof's algorithm is often said to use l-adic methods, because it uses the torsion points of all prime orders l up to a suitable bound.

But, in 2000[1], Satoh came up with a so-called p-adic method. Here, p is the characteristic

[1]Note added during the workshop: Vercauteren showed me the article [5] by Goro Kato and Saul Lubkin,

of $\mathbb{F}_q$, i.e., one has $q = p^n$, with p prime and $n \geq 1$. His method has been improved (Satoh, Vercauteren, Preneel, Vandewalle, Harley, Gaudry, Lercier and Lubicz), and now, for p fixed, the method has running time $O(n^{2+\varepsilon})$. For details, see Satoh's overview article [10] in ANTS-V.

A common feature of these p-adic methods is that they are not polynomial in $\log p$, but that for fixed p (and varying n) they are faster than Schoof's algorithm. Instead of using torsion points of small and varying order l, they use the derivative of the q-Frobenius endomorphism of the canonical lift of E modulo p^m for large enough m. In other words, Schoof uses étale cohomology with $\mathbb{F}_l$-coefficients (in the form of l-torsion points), and Satoh uses crystalline cohomology (in the form of de Rham cohomology of the canonical lift).

In 2001, Kedlaya came up with an algorithm (see [6]) for computing the zeta function of hyper-elliptic curves $C/\mathbb{F}_q$ of arbitrary genus g with running time $O(g^{4+\varepsilon}n^{3+\varepsilon})$, when $p \neq 2$ is fixed (the condition $p \neq 2$ was there for convenience, it was removed by Vercauteren and Denef). Kedlaya's algorithm uses so-called Monsky-Washnitzer cohomology (a variant of crystalline cohomology). We should also mention that at the same time Lauder and Wan came up with algorithms with a similar running time, using p-adic methods that in some sense are the Fourier transform of Kedlaya's. We want to stress that it is quite remarkable that the running time is polynomial in the genus. In particular, the big open problem in this field is to find an algorithm for general q (with p *not* fixed) and g, that is polynomial in $\log q$ and g

The aim of these three lectures is to present Kedlaya's algorithm in as much detail as possible. Our reasons for doing that are the following: (1) it is actually theoretically much simpler even in the case of elliptic curves than Satoh's method using canonical lifts (in particular, there is almost no geometry necessary), and (2) it is a much more general method (it works for arbitrary genus, and also in higher dimension). In other words, our aim is to make publicity for p-adic cohomology, because we think that it has an undeserved reputation of being complicated and not so useful; the truth is precisely the opposite.

In the exposition here we will try to keep everything as simple as possible. We will emphasize the things that actually are to be computed, and pay a lot less attention to the theory behind. But we hope that this course will give the audience a good appetite for the theory. Some experience with the theory of complex analytic functions in one variable is very useful, plus some basic facts about polynomials in one variable over a field. For the rest, we suppose barely anything else (probably, knowing what finite fields are should be sufficient). We will use some facts about zeta functions of non-singular projective curves, but those will be explained (without proof).

---

from 1982

2 A class of point counting problems

Now that we are done with the introduction, let us first describe exactly the problem that we want to solve. For $n \geq 1$ an integer, we define $\mathbb{Z}/n\mathbb{Z} := \{0, 1, \ldots, n-1\}$, with the usual operations $+$ (addition) and $\cdot$ (multiplication) performed "modulo n", i.e., one performs these operations in the ring of integers $\mathbb{Z}$, but then reduces again mod n by composing with the function taking the remainder upon division by n. The sets $\mathbb{Z}/n\mathbb{Z}$, equipped with these operations, are *rings*, i.e., the usual properties of addition and multiplication are true.

Let us recall that a ring is called a *field* if $1 \neq 0$ and every non-zero element has a multiplicative inverse. Then $\mathbb{Z}/n\mathbb{Z}$ is a field if and only if n is prime, and, in that case, we also denote it by $\mathbb{F}_n$. We have the following basic result about finite fields.

2.1 Theorem. *Let F be a finite field. Then $q := \#F$ is of the form p^n with p a prime number and $n \geq 1$. For every q of this form there does exist a finite field F with $\#F = q$; it can be constructed as $\mathbb{F}_p[x]/(f)$, with f an irreducible polynomial of degree n over $\mathbb{F}_p$. Two finite fields with the same number of elements are isomorphic (they all are splitting fields over $\mathbb{F}_p$ of $x^q - x$). For $q = p^n$ a prime power, we denote by $\mathbb{F}_q$ a field with q elements.*

We note that for practical purposes one implements $\mathbb{F}_q$ as $\mathbb{F}_p[X]/(f)$, with

$$f = x^n + f_{n-1}x^{n-1} \cdots + f_1 x + f_0$$

monic and irreducible over $\mathbb{F}_p$. Then $(1, x, \ldots, x^{n-1})$ is an $\mathbb{F}_p$-basis of $\mathbb{F}_q$. Addition and multiplication are performed as in $\mathbb{F}_p[x]$, followed by reduction mod f. Inversion can be performed by an extended gcd algorithm (or by powering, or by inverting a linear map).

We can now describe the class of point counting problems considered by Kedlaya. Let p be a prime different from 2, and let $q = p^n$ for some $n \geq 1$. Let $f \in \mathbb{F}_q[x]$ be monic, of some odd degree $d = 2g + 1$, and without multiple roots (in some algebraic closure of $\mathbb{F}_q$), i.e., $\gcd(f, f') = 1$. Then we let C_f denote the algebraic curve in the affine plane over $\mathbb{F}_q$ given by the equation:

$$y^2 = f.$$

The condition $\gcd(f, f') = 1$ ensures that C_f is non-singular (indeed, the equation $y^2 - f$ and its partial derivatives have no common zero). For $d = 3$ (i.e., $g = 1$), the curves C_f are elliptic curves, minus the point at infinity in the projective plane. For each finite extension $\mathbb{F}_{q^m}$ of $\mathbb{F}_q$ we can then consider the set of points of C_f with coordinates in $\mathbb{F}_{q^m}$:

$$C_f(\mathbb{F}_{q^m}) = \{(a, b) \in \mathbb{F}_{q^m}^2 \mid b^2 = f(a)\}.$$

The general theory of zeta functions of curves then implies the following (see [2], or [12]).

2.2 Theorem. *(Weil) For q and f as above, there exists a monic polynomial P_f with integer coefficients, of degree $2g$, with the following properties. Its complex roots α_i, $1 \leq i \leq 2g$, satisfy $|\alpha_i| = q^{1/2}$, and can be numbered such that $\alpha_i \alpha_{g+i} = q$. For every $m \geq 1$, one has:*

$$\#C_f(\mathbb{F}_{q^m}) = q^m - \sum_i \alpha_i^m. \tag{2.3}$$

For J_f the Jacobian variety of C_f completed with a non-singular point one has:

$$\#J_f(\mathbb{F}_q) = P_f(1).$$

We note that cyclic subgroups of large prime order of the $J_f(\mathbb{F}_q)$ can be used for cryptography, hence it is important to be able to compute $\#J_f(\mathbb{F}_q)$ efficiently, or, even better, P_f itself. This is precisely what Kedlaya's theorem in [6] is about.

2.4 Theorem. *(Kedlaya) Let $p > 2$ be a prime number. For $n \geq 1$ and f in $\mathbb{F}_q[x]$ of degree $2g+1$ as above, P_f can be computed in time $O(g^{4+\varepsilon} n^{3+\varepsilon})$.*

Note that for elliptic curves (i.e., $g = 1$ and p fixed) Kedlaya's algorithm has better complexity than Schoof's, but worse than Satoh's. In the next sections we will describe Kedlaya's algorithm. We remark that there seems to be no simple way to deduce P_f from q and f.

3 ***Some cohomological background (curves only)***

This section is included only to describe the general framework used to do point counting (and also to educate those Dutch mathematicians who always seem to work in the projective non-singular case only!); it is not necessary in order to understand the computational aspects of Kedlaya's algorithm. The title of the section is surrounded by stars in order to discourage the reader from reading the section. *Advice: do not let this section discourage you from reading the rest of this text.* After this section everything will be very explicit, and easier than you may think.

The notation remains as above. In particular, $q = p^n$. The first observation is to view $C_f(\mathbb{F}_{q^m})$ as the set of fixed points of the q^m-power Frobenius map. Let $\mathbb{F}$ be an algebraic closure of $\mathbb{F}_q$, and let:

$$F \colon \mathbb{F} \to \mathbb{F}, \quad a \mapsto a^p$$

be the p-power Frobenius map. It is an automorphism of the field $\mathbb{F}$. We let $F_q := F^n$ be the q-power Frobenius. Then, for any $m \geq 1$ we have:

$$\mathbb{F}_{q^m} = \{a \in \mathbb{F} \mid F_q^m(a) = a\}.$$

As f has coefficients in $\mathbb{F}_q$, F_q also gives a map, still denoted F_q and called the q-power Frobenius of C_f:

$$F_q \colon C_f(\mathbb{F}) \to C_f(\mathbb{F}), \quad (a,b) \mapsto (F_q(a), F_q(b)),$$

of which $C_f(\mathbb{F}_{q^m})$ is precisely the set of fixed points. Now, in algebraic topology, one knows that counting fixed points can be done using cohomology. In the setting of algebraic varieties that we are in now, the required theory does exist, and in fact in many variants (l-adic étale, p-adic (Monsky-Washnitzer, rigid)). For our curves C_f, each of these theories associates to C_f two vector spaces $\mathrm{H}^1_{\mathrm{c}}(C_f)$ and $\mathrm{H}^2_{\mathrm{c}}(C_f)$ over a field of characteristic zero, such that F_q induces endomorphisms:

$$F_q^* \colon \mathrm{H}^i_{\mathrm{c}}(C_f) \to \mathrm{H}^i_{\mathrm{c}}(C_f),$$

with the property that:

$$\#C_f(\mathbb{F}_{q^m}) = \mathrm{trace}((F_q^m)^* \mid \mathrm{H}^2_{\mathrm{c}}(C_f)) - \mathrm{trace}((F_q^m)^* \mid \mathrm{H}^1_{\mathrm{c}}(C_f)).$$

This last formula is the Lefschetz trace formula in the case of affine curves. It is valid in great generality (for all separated varieties, and also for "constructible sheaves" replacing the constant coefficients). In particular, the variety does not need to be complete or projective, and it may have singularities. The subscript "c" stands for "compact support", hence the cohomology in the formula is "cohomology with compact supports". In formula 2.3 above, q^m is the trace on $\mathrm{H}^2_{\mathrm{c}}$ and the sum of the α_i^m is the trace on $\mathrm{H}^1_{\mathrm{c}}$. In particular, the important polynomial P_f is the characteristic polynomial of F_q^* acting on $\mathrm{H}^1_{\mathrm{c}}(C_f)$. Now cohomology with compact supports is not hard to define (one takes the ordinary cohomology of the sheaf extended by zero on some compactification, and in the case of de Rham cohomology of differentiable varieties one can also just replace the de Rham complex of differential forms by those with compact support), but it can be avoided for non-singular varieties. Indeed, Poincaré duality says that $\mathrm{H}^2_{\mathrm{c}}(C_f)$ is one-dimensional with maps acting as their degree, and that the usual product:

$$\mathrm{H}^1_{\mathrm{c}}(C_f) \times \mathrm{H}^1(C_f) \to \mathrm{H}^2_{\mathrm{c}}(C_f)$$

is a perfect pairing. This implies that P_f is the characteristic polynomial of $q(F_q^*)^{-1}$ on $\mathrm{H}^1(C_f)$. But this is still not how we are going to compute P_f. We put:

$$C'_f := C_f - \{\text{zeros of } y\}.$$

The curves C_f and C'_f have an automorphism ι of order two:

$$\iota \colon (a,b) \mapsto (a,-b).$$

Then ι induces automorphisms ι^* of $\mathrm{H}^1(C'_f)$ and $\mathrm{H}^1(C_f)$, decomposing them into two eigenspaces on which ι^* acts as 1 and -1, respectively:

$$\mathrm{H}^1(C'_f) = \mathrm{H}^1(C'_f)^+ \oplus \mathrm{H}^1(C'_f)^-, \qquad \mathrm{H}^1(C_f) = \mathrm{H}^1(C_f)^+ \oplus \mathrm{H}^1(C_f)^-.$$

Then we have the following description of P_f, that will actually be used for computing it.

3.1 Proposition. *P_f is the characteristic polynomial of F_q^* on $\mathrm{H}^1(C'_f)^-$.*

Let us try to understand this. Let $\overline{C_f}$ be the compactification of C_f, and let R denote the set of zeros of f plus the point ∞ at infinity (it is the ramification locus of the map from $\overline{C_f}$ to $\mathbb{P}^1$ extending $(a, b) \mapsto a$). Then

$$\mathrm{H}^1(\overline{C_f})^+ = \mathrm{H}^1(\mathbb{P}^1) = \{0\}, \quad \mathrm{H}^1(\overline{C_f})^- = \mathrm{H}^1(\overline{C_f}) = \mathrm{H}^1_{\mathrm{c}}(\overline{C_f}).$$

And we have an exact sequence:

$$\mathrm{H}^0_{\mathrm{c}}(R) \to \mathrm{H}^1_{\mathrm{c}}(C'_f) \to \mathrm{H}^1_{\mathrm{c}}(\overline{C_f}) \to \{0\}.$$

As ι acts trivially on R, it follows that the map $\mathrm{H}^1_{\mathrm{c}}(C'_f)^- \to \mathrm{H}^1_{\mathrm{c}}(\overline{C_f})^-$ is an isomorphism. Moreover, $\mathrm{H}^1_{\mathrm{c}}(\overline{C_f})^- = \mathrm{H}^1(\overline{C_f})$.

Let us end this section with the remark that the reason for working with C'_f instead of C_f is to get a Frobenius lift that is given by a simple formula. This will become clear later.

4 Monsky-Washnitzer cohomology

In this section we give an explicit description of the first Monsky-Washnitzer cohomology group of the curves of the form C'_f as above ($p \neq 2$). In order to construct these groups it does not suffice to work over $\mathbb{F}_p$, because that would only lead to cohomology groups that are $\mathbb{F}_p$-vector spaces (and hence would only give information about P_f modulo p). We want cohomology groups that are vector spaces over a field of characteristic zero. In this case, that field will be $\mathbb{Q}_p$, the field of p-adic numbers, and we will now first explain what that field is.

4.1 The p-adic numbers

Let p be a prime number. Writing integers in base p means that each integer can be written in a unique way as $a_r \cdots a_1 a_0 = \sum a_i p^i$, with a_i in $\{0, 1, \ldots, p-1\}$ (of course, one may also choose another set of representatives, such that $|a_i| < p/2$ if $p > 2$). The operations of addition and multiplication are done in the usual way. Now nothing prevents us from enlarging the ring $\mathbb{Z}$ by extending these sequences arbitrarily to the *left*.

4.1.1 Definition. The set of p-adic integers is the set $\{\cdots a_2a_1a_0\}$ of all sequences with $0 \leq a_i < p$.

4.1.2 Proposition. *Addition and multiplication done in the usual way make the set of p-adic integers into a ring, that we will denote by $\mathbb{Z}_p$. The morphism $\mathbb{Z} \to \mathbb{Z}_p$ is injective. Multiplication by p is a shift to the left. We have $\mathbb{Z}_p/p^n\mathbb{Z}_p = \mathbb{Z}/p^n\mathbb{Z}$. An element a of $\mathbb{Z}_p$ is invertible if and only if $a_0 \neq 0$ (one can construct the inverse digit by digit).*

In more fancy language, $\mathbb{Z}_p$ is the inverse limit of the $\mathbb{Z}/p^n\mathbb{Z}$, i.e., the p-adic *completion* of $\mathbb{Z}$.

4.1.3 Definition. The set of p-adic numbers is the set $\{\cdots a_2a_1a_0, a_{-1}\cdots a_{-r}\}$ of all sequences infinite to the left, and finite to the right, with $0 \leq a_i < p$.

4.1.4 Proposition. *Addition and multiplication done in the usual way make the set of p-adic numbers into a field, that we will denote by $\mathbb{Q}_p$. The map $v_p\colon \mathbb{Q}_p^* \to \mathbb{Z}$, $a \mapsto \min\{i \mid a_i \neq 0\}$ is a* valuation. *The map $|\cdot|\colon \mathbb{Q}_p \to \mathbb{R}$ sending 0 to 0 and non-zero a to $p^{-v_p(a)}$ is an* absolute value *on $\mathbb{Q}_p$, for which $\mathbb{Q}_p$ is complete and for which $\mathbb{Q}$ is dense. And $\mathbb{Z}_p = \{a \mid |a| \leq 1\}$ is compact.*

With respect to this absolute value, we can speak of converge of sequences etc. Then:

$$\cdots a_2a_1a_0, a_{-1}\cdots a_{-r} = \sum_i a_i p^i$$

(note that $|p| = 1/p$, hence the series converges). It is clear that, just as with the real numbers $\mathbb{R}$, the ring operations are easy to implement with any desired precision. Later in this article, we will see that all computations in Kedlaya's algorithm can be done in some *fixed* $\mathbb{Z}/p^m\mathbb{Z}$, i.e., within $\mathbb{Z}_p$ and with a precision of m digits.

Let $q = p^n$ for some $n \geq 1$, and let $\overline{f}$ in $\mathbb{F}_p[x]$ be monic and irreducible of degree n. Then we write $\mathbb{F}_q = \mathbb{F}_p[x]/(f)$, as before, and we know that another choice of $\overline{f}$ gives an isomorphic field. Now we lift this over $\mathbb{Z}_p$.

4.1.5 Definition. Let p, q and $\overline{f}$ as above, and let f in $\mathbb{Z}_p[x]$ be a monic lift of $\overline{f}$. Then we define:

$$\mathbb{Z}_q := \mathbb{Z}_p[x]/(f), \qquad \mathbb{Q}_q := \mathbb{Q}_p[x]/(f).$$

Concretely, $\mathbb{Z}_q$ has $\mathbb{Z}_p$-basis $(1, x, \ldots, x^{n-1})$, and this is also a $\mathbb{Q}_p$-basis of $\mathbb{Q}_q$. Operations are done modulo f.

4.1.6 Proposition. *The ring $\mathbb{Q}_q$ is a field (and hence $\mathbb{Z}_q$ is an integral domain). The constructions are independent of the choice of the lift f in the sense that any other lift leads to rings that are canonically isomorphic to the ones gives by f.*

4.2 A dagger ring

Let $p > 2$ be a prime number, let $n \geq 1$ and let $q := p^n$. Let $\overline{Q}$ be monic in $\mathbb{F}_q[x]$, of odd degree $d = 2g + 1$, and such that $\gcd(\overline{Q}, \overline{Q}') = 1$. We put $\overline{A} := \mathbb{F}_q[x, y, y^{-1}]/(y^2 - \overline{Q})$. This is the ring of functions on the curve $C'_{\overline{Q}}$ (notation as before) that we want to consider. Note that as y is invertible, $\overline{Q}$ is invertible too. Another way to write this ring is:

$$\begin{aligned}\overline{A} = (\mathbb{F}_q[x, \overline{Q}^{-1}])[y]/(y^2 - \overline{Q}) &= \bigoplus_{i,j,k} \mathbb{F}_q \; x^i \, \overline{Q}^j \, y^k, \quad (0 \leq i < d, j \in \mathbb{Z}, 0 \leq k < 2) \\ &= \bigoplus_{i,j} \mathbb{F}_q \; x^i y^j, \quad (0 \leq i < d, j \in \mathbb{Z}).\end{aligned}$$

Just as we have lifted $\mathbb{F}_q$ to $\mathbb{Z}_q$, we will now lift $\overline{A}$ to a ring A. In contrast to Satoh's method for elliptic curves, where one has to choose the canonical lift that is not so easy to compute, we can just *choose* a monic lift Q in $\mathbb{Z}_q[x]$ of $\overline{Q}$, and the whole theory will work. Then we put:

$$A := \mathbb{Q}_q[x, y, y^{-1}]/(y^2 - Q) = \bigoplus_{i,j} \mathbb{Q}_q \; x^i \, y^j, \quad (0 \leq i < d, j \in \mathbb{Z}).$$

This ring A is the ring of functions of the algebraic curve C'_Q over $\mathbb{Q}_q$. It definitely depends on the choice of our lift Q if $g > 0$ (just think of the case of elliptic curves). In order to get rings that do not depend on the choice of Q we have to get out of the world of algebraic curves and move to-wards more analytic objects. We put:

$$A^\infty := \left\{ \sum_{i,j} a_{i,j} x^i y^j \mid a_{i,j} \in \mathbb{Q}_q, |a_{i,j}| \to 0 \text{ as } |j| \to \infty \right\},$$

with $0 \leq i < d$ and $j \in \mathbb{Z}$. The elements of A^∞ are the arbitrary series $f = \sum_{i,j} a_{i,j} x^i y^j$ with the property that for every integer k almost all $a_{i,j}$ are in $p^k \mathbb{Z}_q$. A more fancy way to say this is that $A^\infty = \mathbb{Q}_q \otimes_{\mathbb{Z}_q} A^\infty_+$, with A^∞_+ the p-adic completion of $\mathbb{Z}_q[x, y, y^{-1}]/(y^2 - Q)$. One can show that A^∞_+, and hence also A^∞, do not depend on the choice of our lift Q: any lift gives an isomorphic result, but the isomorphism is not unique. (Just for information: this comes from the fact that the set of lifts over $\mathbb{Z}_q/p^2\mathbb{Z}_q$ form a torsor under $\mathrm{H}^1(C'_{\overline{Q}}, \mathrm{T})$, which is zero because $C'_{\overline{Q}}$ is affine and the tangent sheaf T is coherent, and so on.)

The ring A^∞ is the ring of series $f = \sum_{i,j} a_{i,j} x^i y^j$ ($0 \leq i < d$ and $j \in \mathbb{Z}$) that converge on the part of $C_Q(\overline{\mathbb{Q}}_q)$ given by the condition $|y| = 1$ (implying $|x| \leq 1$). In complex analysis, the power series that converge on the closed unit disk give functions that are continuous but not necessarily differentiable (example: $(1 - z)^{1/2}$). In the p-adic case, one gets functions that one cannot integrate (example: $\sum_{n \geq 0} p^n z^{p^n - 1}$). So, in both cases, de Rham cohomology will

not give the results that we want for such rings. To get the right cohomology, one uses *overconvergent* functions, i.e., elements of A^∞ that converge on a neighborhood of the closed part that one considers (the neighborhood depends on the element). This gives the ring that we will work with:

$$A^\dagger := \left\{ \sum_{i,j} a_{i,j} x^i y^j \mid a_{i,j} \in \mathbb{Q}_q, \liminf_{|j|\to\infty} v_p(a_{i,j})/|j| > 0 \right\}.$$

One can show that $A^\dagger$ does not depend on the choice of our lift Q (see for example Corollary 7.5.10 in [1]; in this explicit case, one can also use an explicit method, just as for the construction of the Frobenius lift in section 5).

4.3 Differentials and de Rham cohomology

We start with the differential forms on C'_Q. The functions on C'_Q that we consider are the elements of A. Each element f of A has to have a differential df, such that the Leibniz rule $d(fg) = fdg + gdf$ holds and such that $da = 0$ for $a \in \mathbb{Q}_q$. The differentials should form an A-module. In other words, d should be a $\mathbb{Q}_q$-derivation from A to an A-module. Now there exists a *universal* such derivation $d\colon A \to \Omega$, such that for any $\mathbb{Q}_q$-derivation $D\colon A \to M$ there is a unique A-linear map $l\colon \Omega \to M$ such that $D = ld$. The universal derivation is easy to describe. First of all, Ω is generated by the df for f in A, hence by dx and dy. But we have:

$$0 = y^2 - Q, \qquad \text{hence} \quad 0 = 2ydy - Q'dx, \quad dy = \frac{Q'dx}{2y}.$$

(Note that $2y$ is invertible in A.) It follows that Ω is a free A-module and that $(Q'dx)/2y$ is an A-basis:

$$\Omega = A \cdot \frac{dx}{2y}.$$

The de Rham complex of C'_Q is then:

$$A \xrightarrow{d} A \cdot \frac{dx}{2y}, \quad x^i y^j \mapsto \left(2ix^{i-1}y^{j+1} + jx^i Q' y^{j-1}\right) \frac{dx}{2y}$$

with A in degree zero and $A(dx)/2y$ in degree one. The algebraic de Rham cohomology is simply the homology of this complex:

$$\mathrm{H}^0_{\mathrm{dR}}(C'_Q) = \ker(d) = \{f \in A \mid df = 0\}, \qquad \mathrm{H}^1_{\mathrm{dR}}(C'_Q) = \mathrm{coker}(d) = \left(A \cdot \frac{dx}{2y}\right)/dA.$$

4.3.1 Proposition. *We have* $\mathrm{H}^0_{\mathrm{dR}}(C'_Q) = \mathbb{Q}_q$. *The classes* $[x^i y^{-1}(dx)/y]$ *with* $0 \le i \le 2g$ *form a basis for* $\mathrm{H}^1_{\mathrm{dR}}(C'_Q)^+$, *and the classes* $[x^i(dx)/y]$ *with* $0 \le i < 2g$ *form a basis for* $\mathrm{H}^1_{\mathrm{dR}}(C'_Q)^-$.

Proof. We first split the complex into its two eigen-spaces for ι. Recall that the ring A has $\mathbb{Q}_q$-basis $x^i y^j$, $0 \leq i < d$, $j \in \mathbb{Z}$, that $\iota x = x$ and that $\iota y = -y$. It follows that the decompositions in ι-eigen-spaces are:

$$\begin{aligned}
A^+ &= \bigoplus_{\substack{0\leq i<d\\ j\equiv 0(2)}} \mathbb{Q}_q x^i y^{2j} = \bigoplus_{\substack{0\leq i<d\\ j\in\mathbb{Z}}} \mathbb{Q}_q x^i Q^j = \mathbb{Q}_q[x, Q^{-1}],\\
\Omega^+ &= \bigoplus_{\substack{0\leq i<d\\ j\equiv 1(2)}} \mathbb{Q}_q x^i y^j \frac{dx}{2y} = \bigoplus_{\substack{0\leq i<d\\ j\equiv 1(2)}} \mathbb{Q}_q x^i Q^{(j-1/2)} dx = \mathbb{Q}_q[x, Q^{-1}]dx,\\
A^- &= \bigoplus_{\substack{0\leq i<d\\ j\equiv 1(2)}} \mathbb{Q}_q x^i y^j,\\
\Omega^- &= \bigoplus_{\substack{0\leq i<d\\ j\equiv 0(2)}} \mathbb{Q}_q x^i y^j \frac{dx}{2y}.
\end{aligned}$$

Let us deal with the $+$-part first. In fact, this part *is* the de Rham complex for the ring $\mathbb{Q}_q[x, Q^{-1}]$, i.e., for the curve $\mathbb{A}^1$ minus the zeros of Q. In the exercises it is proved that $\mathrm{H}^0_{\mathrm{dR}}(C_Q)^+ = \mathbb{Q}_q$ and that the $x^i Q^{-1} dx$, $0 \leq i < d$, form a basis for $\mathrm{H}^1_{\mathrm{dR}}(C'_Q)^+$.

Let us now deal with the $-$-part. We order the monomials $x^i y^j$, $0 \leq i < d$ and $j \in \mathbb{Z}$ lexicographically, with "j first and then i". (On the $x^i y^j$ in A^- this ordering is given by the order at the unique point ∞ at infinity of $\overline{C_Q}$: $x^i y^j$ has order $-2i - jd$. We also note that if α is a root in $\overline{\mathbb{Q}}_q$ of Q, then the order of $x^i y^j$ at $(\alpha, 0)$ is j if $\alpha \neq 0$ and $2i + j$ if $\alpha = 0$.) It is a good idea to draw a picture of these points (i, j). For $0 \leq i < d$, we consider the division of $x^i Q'$ by Q in $\mathbb{Q}_q[x]$:

$$x^i Q' = a_i Q + b_i, \quad \deg(b_i) < d.$$

We have $a_0 = 0$ and $b_0 = Q'$. For $1 \leq i < d$ we have $a_i = dx^{i-1} + \cdots$, hence $\deg(a_i) = i - 1$. Then we have:

$$d\colon x^i y^j \mapsto \left(2ix^{i-1}y^{j+1} + jx^i Q' y^{j-1}\right)\frac{dx}{2y} = \left(2ix^{i-1}y^{j+1} + ja_i y^{j+1} + jb_i y^{j-1}\right)\frac{dx}{2y}.$$

For $0 \leq i < d$ and $j \equiv 1(2)$ it follows that the *highest* monomial of $d(x^i y^j)$ is $x^{i-1}y^{j+1}$ if $1 \leq i < d$, and $x^{d-1}y^{j-1}$ if $i = 0$ (note that $2i + jd \neq 0$ and $jd \neq 0$). We also note that the multiplication by Q' on $\mathbb{Q}_q[x]/(Q)$ is an isomorphism, because $\gcd(Q, Q') = 1$. Hence $b_0, \ldots, b_{d-1}$ form a $\mathbb{Q}_q$-basis of $\mathbb{Q}_q[x]_{<d}$. It follows that the *lowest* monomial of $d(x^i y^j)$ is of the form $x^k y^{j-1}$ with $0 \leq k < d$. Graphically speaking, $d(x^i y^j)$ is concentrated on the $j + 1$th and $j - 1$th rows of the monomial space, and on the $j - 1$th row only if $i = 0$.

Let us now prove the claim that $\mathrm{H}^0_{\mathrm{dR}}(C'_Q)^- = \ker(d\colon A^- \to \Omega^-) = 0$. This is immediate from the fact that all $d(x^iy^j)$ with $0 \leq i < d$ and $j \equiv 1(2)$ have different highest monomials, so that the $d(x^iy^j)$ are linearly independent.

It remains to prove that the classes $[x^i(dx)/y]$ with $0 \leq i < d-1$ form a basis for $\mathrm{H}^1_{\mathrm{dR}}(C'_Q)^- = \Omega^-/dA^-$. Let us first show that these elements are linearly independent. So suppose that we have:

$$\sum_{0\leq k<d-1} \lambda_k x^k(dx)/y = \sum_{\substack{0\leq i<d \\ j\equiv 1(2)}} \mu_{i,j} d(x^iy^j) \neq 0,$$

with almost all $\mu_{i,j}$ equal to zero. Then the highest monomials occurring on the right hand side are of the form x^iy^{-1} or y and $i = 0$. However, y does not occur because on the left hand side the coefficient of $x^{d-1}(dy)/x$ is zero. The linear independence of the b_i implies that the lowest monomial occurring on the right hand side is of the form x^iy, $0 \leq i < d$. So we have proved that the right hand side is zero, hence also the left hand side.

We prove that the $[x^i(dx)/y]$ generate by giving a ***reduction algorithm*** to write an arbitrary element $f(dx)/y$ of Ω^- as the sum of an element of dA^- and a linear combination of the $x^i(dx)/y$. So, let $f(dx)/y$ be given. As long as f has a monomial with $j < 0$, do the following: let x^iy^j be the lowest monomial of f; use the unique linear combination dg of the $d(x^ky^{j+1})$ such that $f(dx)/y - dg$ has no monomials x^my^n with $n = j$. Now f has no monomials x^iy^j with $j < 0$. As long as f has monomials x^iy^j with $j > 0$, do the following: let x^iy^j be the highest monomial of f; let x^ky^l be the monomial such that $d(x^ky^l)$ has highest monomial x^iy^j and replace $f(dx)/y$ by $f(dx)/y$ minus the appropriate multiple of $d(x^ky^l)$. Now f has no monomials x^iy^j with $j \neq 0$. Now subtract from $f(dx)/y$ the appropriate multiple of dy so that the monomial $x^{d-1}y$ does not occur in the difference. Now $f(dx)/y$ is a linear combination of the $x^i(dx)/y$ with $0 \leq i < d-1$. □

At this point, we have the vector space $\mathrm{H}^1_{\mathrm{dR}}(C'_Q)^- = \oplus_{0\leq i<d-1}\mathbb{Q}_q x^i(dx)/y$ that will give us the desired polynomial $P_{\overline{Q}}$. But, unfortunately, we do not yet have the Frobenius operator of which $P_{\overline{Q}}$ is the characteristic polynomial. Namely, the Frobenius endomorphism F_q of $C'_{\overline{Q}}$ cannot be lifted to C'_Q if $g > 1$, and only for a very special choice of Q if $g = 1$ (the canonical lift). (The reason that for $g > 1$ there is no lift is that curves of genus > 1 over a field of characteristic zero have no endomorphisms of degree > 1 by Hurwitz's formula.) In order to get a Frobenius operator, we will replace the ring A by $A^\dagger$.

By definition, the de Rham complex of $A^\dagger$ is given by:

$$d\colon A^\dagger \longrightarrow A^\dagger\cdot\frac{dx}{2y},$$
$$\sum_{i,j} a_{i,j}x^iy^j \mapsto \sum_{i,j} a_{i,j}d(x^iy^j) = \sum_{i,j} a_{i,j}\left(2ix^{i-1}y^{j+1} + jx^iQ'y^{j-1}\right)\frac{dx}{2y}.$$

Indeed, one checks that if $\liminf_{|j|\to\infty} v_p(a_{i,j})/|j| > 0$, then the same holds for the coefficients of $\sum_{i,j} a_{i,j}(2ix^{i-1}y^{j+1} + jx^iQ'y^{j-1})$ (the "liminf" does not even get smaller). We will denote the cohomology groups of this complex by $\mathrm{H}^i(C'_{\overline{Q}})$; note that they are $\mathbb{Q}_q$-vector spaces. Just as before, the automorphism ι splits them into two parts.

4.3.2 Proposition. *The elements* $[x^i(dx)/y]$, $0 \leq i < d-1$, *form a basis for* $\mathrm{H}^1(C'_{\overline{Q}})^-$.

Proof. There is a general theory that guarantees that passing from A to $A^\dagger$ does not change the cohomology (the technical condition is that we are dealing with an open sub-scheme of a proper and smooth $\mathbb{Z}_q$-scheme with complement a relative divisor with normal crossings). But in this case we give a proof by computation that will also give us a very precise control of the denominators that arise in the reduction algorithm. In any case, we need the reduction algorithm in order to compute the matrix of the Frobenius operator with respect to the basis that we have.

Let us show that the $[x^i(dx)/y]$, $0 \leq i < d-1$, generate $\mathrm{H}^1(C'_{\overline{Q}})^-$. So let $\sum_m a_m y^m(dx)/y$ be an element of $A^\dagger\cdot(dx)/y$, with a_m in $\mathbb{Q}_q[x]$ and $\deg(a_m) < d$. After multiplication by a suitable power of p we have that a_m is in $\mathbb{Z}_q[x]$ for all m. For each $m \neq 0$, the following two lemmas give us a b_m in $\mathbb{Q}_q[x]$ and f_m in A^- such that $a_m y^m(dx)/y = b_m(dx)/y + df_m$ and $\deg(b_m) < d-1$. By definition, there exists an $\varepsilon > 0$ and an integer m_0 such that a_m is divisible by $p^{\lfloor \varepsilon|m|\rfloor}$ when $|m| > m_0$. It follows that we have, in $A^\dagger\cdot(dx)/y$:

$$\sum_m a_m y^m(dx)/y = a_0(dx)/y + \left(\sum_{m\neq 0} b_m\right)(dx)/y + d\left(\sum_m f_m\right).$$

In order to get rid of $a_{0,d-1}x^{d-1}y$ one uses the identity:

$$dy = Q'\frac{dx}{2y} = (dx^{d-1} + \cdots)\frac{dx}{2y}.$$

This finishes the proof that the $[x^i(dx)/y]$, $0 \leq i < d-1$, generate. In order to see that they are linearly independent, one multiplies a hypothetical relation by a suitable power of p to make it integral, then reduces modulo an arbitrarily high power of p, and proceeds as in the proof for independence in $\mathrm{H}^1_{\mathrm{dR}}(C'_Q)^-$. □

4.3.3 Remark. The proofs of the two following lemmas are quite technical. If we just admit the fact that the proposition above is true, then we could use a more direct but worse bound on the denominators that does not change the exponents of g and $\log_p q$ in Kedlaya's algorithm (the computations would then be done with a much higher precision than necessary).

4.3.4 Lemma. *Consider $\omega = ay^{-m}(dx)/y$, with $m > 0$, $m \equiv 0(2)$, and a in $\mathbb{Z}_q[x]$ with $\deg(a) < d$. Then there is a unique b in $\mathbb{Q}_q[x]$ with $\deg(b) < d$, and a unique f in A^- such that:*

$$\omega = ay^{-m}(dx)/y = b(dx)/y + df,$$

with $f = \sum_{j=-m+1}^{-1} f_j y^j$, $f_j \in \mathbb{Q}_q[x]$, $\deg(f_j) < d$. Then one has:

$$p^{\lfloor \log_p(m-1) \rfloor} b \in \mathbb{Z}_q[x] \quad \text{and} \quad p^{\lfloor \log_p(m-1) \rfloor} f_j \in \mathbb{Z}_q[x] \quad \text{for all } j.$$

Proof. The existence of b and f follows directly from the reduction algorithm above. It remains to bound their denominators.

Let us first remark that one can bound the denominators of b and the f_j by inspecting the divisions that have to be done in the reduction algorithm above. If one just bounds the denominators by putting all the divisions together, one finds a bound of the form $p^{f(m)}$ with $f(m)$ linear in m. Hence we need to do better. (Also, somewhere the fact that in our situation we have a smooth compactification with a normal crossings divisor as boundary should show up somewhere).

As we are dealing with negative powers of y, so that it is a natural idea to look at what happens at the zeros of y, i.e., at the points $(\alpha, 0)$ with α a root of Q. So let α be a root of Q in some $\mathbb{Z}_{q^r}$, and let $P = (\alpha, 0)$. At P, we have y as a local coordinate. Hence the completion of $\mathbb{Z}_{q^r}[x, y]/(Q)$ at P is $\mathbb{Z}_{q^r}[[y]]$. We can then write df and f as series:

$$df = \sum_{j \geq -m} c_j y^j dy, \qquad f = \sum_{j \geq -m} \frac{c_j}{j+1} y^{j+1},$$

with c_j in $\mathbb{Z}_{q^r}$, and with $c_j = 0$ if $j \equiv 1(2)$ (because f is in A^-). The series expansion of $b(dx)/y$ in y at P has no pole, and the series of $ay^{-m}(dx)/y$ has coefficients in $\mathbb{Z}_{q^r}$. It follows that c_j is in $\mathbb{Z}_{q^r}$ if $j < 0$. We put $n = p^{\lfloor \log_p(m-1) \rfloor}$. Then it follows that $nc_j/(j+1)$ is in $\mathbb{Z}_{q^r}$ for $j < 0$. Evaluating the coefficient of y^{-m+1} in the identity $f = \sum_{j=-m+1}^{-1} f_j y^j$ gives that $f_{-m+1}(P) = c_{-m}/(-m+1)$, hence that $nf_{-m+1}(P)$ is in $\mathbb{Z}_{q^r}$. Such an integrality statement holds at each zero of y. It follows that nf_{-m+1} is in $\mathbb{Z}_q[x]$ (use that the d roots of $\overline{Q}$ are distinct and that the reduction modulo p of $p^e f_{-m+1}$ does not have more than $d-1$ roots if it is non-zero).

Now the series of $nf - nf_{-m+1}y^{-m+1}$ has integral coefficients at y^j with $j < 0$. The same argument then gives that nf_{-m+2} is integral, and so forth. □

4.3.5 Lemma. *Consider $\omega = ay^m(dx)/y$, with $m > 0$, $m \equiv 0(2)$, and a in $\mathbb{Z}_q[x]$ with $\deg(a) < d$. Then there is a unique b in $\mathbb{Q}_q[x]$ with $\deg(b) < d$, and a unique f in A^- such that:*

$$\omega = ay^m(dx)/y = b(dx)/y + df,$$

with $f = \sum_{j=1}^{m-1} f_j y^j$, $f_j \in \mathbb{Q}_q[x]$, $\deg(f_j) < d$. Then one has:

$$p^{\lfloor \log_p(dm+d-2) \rfloor} b \in \mathbb{Z}_q[x] \quad \text{and} \quad p^{\lfloor \log_p(dm+d-2) \rfloor} f_j \in \mathbb{Z}_q[x] \quad \text{for all } j.$$

Proof. The existence and uniqueness of b and f is guaranteed by the reduction algorithm above. In this case we are dealing with positive powers of y, hence we study what happens at the unique pole ∞ of y. Let v_∞ denote the valuation at this point. Then we have $v_\infty(x) = -2$, $v_\infty(y) = -2$ (because $y^2 = Q$ and $\deg(Q) = d$), $v_\infty(dx) = -3$, $v_\infty((dx)/y) = d-3$, $v_\infty(f) \geq -md-d+2$, and $v_\infty(b(dx)/y) \geq -d+1$. A suitable coordinate at ∞ is given by $z := x^{(d-1)/2}/y$; note that $\iota(z) = -z$. Then we have expansion:

$$df = \sum_{j \geq -dm-d+1} c_j z^j, \qquad f = \sum_{j \geq -dm-d+1} \frac{c_j}{j+1} z^{j+1},$$

with c_j in $\mathbb{Q}_q$, and $c_j = 0$ if $j \equiv 1(2)$. As the expansion of $ay^m(dx)/y$ has coefficients in $\mathbb{Z}_q$, and $v_\infty(b(dx)/y) \geq -d+1$, it follows that c_j is in $\mathbb{Z}_q$ for $j \leq -d$. We put $n = p^{\lfloor \log_p(dm+d-2) \rfloor}$. Then $nc_j/(j+1)$ is in $\mathbb{Z}_q$ if $j \leq -d$. As all $v_\infty(x^i y^j)$ for $0 \leq i < d$ and $j > 0$ are distinct and $\leq -d$, it follows that nf_j is in $\mathbb{Z}_q[x]$ for all j. □

5 Frobenius lifts

In this section, we will show that the Frobenius endomorphism $F_q \colon C'_{\overline{Q}} \to C'_{\overline{Q}}$ can be lifted to $A^\dagger$, and how its action on the cohomology space $\mathrm{H}^1_{\mathrm{dR}}(C'_{\overline{Q}})^-$ gives us the polynomial $P_{\overline{Q}}$ that we are after. As the cohomology space is given with the explicit basis $[x^i(dx)/2y]$, $0 \leq i < d-1$, computing the action of F_q means computing its matrix with respect to this basis. Then $P_{\overline{Q}}$ is the characteristic polynomial.

5.1 Frobenius on points and rings

Let us now first how we interpret the morphism $F_q \colon C'_{\overline{Q}} \to C'_{\overline{Q}}$ that we have seen in section 3 as a map from $C_{\overline{Q}}(\mathbb{F})$ to itself, with $\mathbb{F}$ an algebraic closure of $\mathbb{F}_q$, as a morphism from the $\mathbb{F}_q$-algebra $\overline{A} := \mathbb{F}_q[x, y, y^{-1}]/(y^2 - \overline{Q})$ to itself. The morphism F_q sends an element (a, b) of $C_{\overline{Q}}(\mathbb{F})$ to (a^q, b^q). By the definition of $\overline{A}$ we have:

$$C_{\overline{Q}}(\mathbb{F}) = \mathrm{Hom}_{\mathbb{F}_q}(\overline{A}, \mathbb{F}),$$

where (a, b) corresponds to the morphism that sends x to a, and y to b. It is then immediate that the map $(a, b) \mapsto (a^q, b^q)$ is induced by the $\mathbb{F}_q$-algebra morphism $F_q \colon \overline{A} \to \overline{A}$ that sends x to x^q and y to y^q. Recall now that $q = p^n$. It follows then that $F_q = F_p^n$, with F_p the $\mathbb{F}_p$-algebra morphism from $\overline{A}$ to itself that sends any element a to its pth power a^p, i.e., F_p is the absolute p-power Frobenius endomorphism. Computationally, it is then a good idea to work with F_p. Note that F_p is not an $\mathbb{F}_q$-algebra morphism if $n > 1$; indeed, we have a commutative diagram:

$$\begin{array}{ccc} \overline{A} & \xrightarrow{F_p} & \overline{A} \\ \uparrow & & \uparrow \\ \mathbb{F}_q & \xrightarrow{\sigma} & \mathbb{F}_q \end{array}, \qquad \sigma \colon \mathbb{F}_q \to \mathbb{F}_q, \quad a \mapsto a^p,$$

where, as usual, we let σ denote the absolute Frobenius on $\mathbb{F}_q$.

5.2 A lifting of Frobenius on $A^\dagger$

We start with lifting σ to an automorphism, still denoted σ, from $\mathbb{Z}_q$ to itself. That this can be done, in a unique way, results from Proposition 4.1.6. As $\mathbb{Q}_q = \mathbb{Z}_q[1/p]$, σ extends uniquely to an automorphism σ of $\mathbb{Q}_q$. We extend σ to $F_p \colon \mathbb{Z}_q[x] \to \mathbb{Z}_q[x]$ by sending x to x^p. We will show that F_p can be extended uniquely to an endomorphism F_p of A^∞ that is p-adically continuous and compatible with the $\mathbb{Q}_q[x]$-algebra structure on A^∞. These last conditions mean that F_p maps the sub-ring $A_+^\infty = (\mathbb{Z}_q[x, y, z]/(y^2 - Q, yz - 1))^\wedge$ to itself, and that for each $m \geq 0$ the restriction of F_p on $\mathbb{Z}_q[x, y, z]/(y^2 - Q, yz - 1, p^m)$ to $\mathbb{Z}_q[x]/(p^m)$ is the F_p that we already had.

Let us write $Q = x^d + Q_{d-1}x^{d-1} + \cdots + Q_0$. Then we have:

$$F_pQ = x^{pd} + \sigma(Q_{d-1})x^{p(d-1)} + \cdots + \sigma(Q_0) =: Q^\sigma(x^p).$$

As F_p lifts the p-power Frobenius endomorphism of $\mathbb{F}_q[x]$, F_pQ and Q^p have the same image in $\mathbb{F}_q[x]$, hence their difference is divisible by p in $\mathbb{Z}_q[x]$. We put:

$$E := \frac{F_pQ - Q^p}{p} \qquad (E \in \mathbb{Z}_q[x]).$$

By construction, extending F_p as we want to A_+^∞ means the following:

1. $F_pa = \sigma a$ for all a in $\mathbb{Z}_q$,
2. $F_px = x^p$,
3. F_py is an element of A_+^∞ that satisfies $(F_py)^2 = F_pQ$ and has image y^p in $\overline{A}$,
4. F_pz is an element of A_+^∞ that satisfies $(F_py)(F_pz) = 1$.

In A_+^∞ we have:

$$F_pQ = Q^p + pE = y^{2p} + pE = y^{2p}(1 + pEz^{2p}),$$

hence we can put:

$$F_py = y^p(1 + pEz^{2p})^{1/2} = y^p \sum_{k\geq 0} \binom{1/2}{k} p^k E^k z^{2pk}.$$

Here $\binom{1/2}{k}$ is the polynomial $\binom{t}{k} = t(t-1)\cdots(t-k+1)/k!$ evaluated at $t = 1/2$. Note that $v(\binom{1/2}{k}) \geq 0$, i.e, $\binom{1/2}{k}$ is in $\mathbb{Z}_q$ (argument: approximate $1/2$ p-adically with elements of $\mathbb{Z}$). We also put:

$$F_pz = z^p(1 + pEz^{2p})^{-1/2} = z^p \sum_{k\geq 0} \binom{-1/2}{k} p^k E^k z^{2pk}.$$

So this extends F_p to A_+^∞, and hence to A^∞. Let now $a = \sum_{i,j} a_{i,j}x^iy^j$ be an element of $A^\dagger$ (the sum ranging over the (i,j) with $0 \leq i < d$, $j \in \mathbb{Z}$). Then we have:

$$F_pa = \sum_{i,j} \sigma(a_{i,j})x^{pi}(F_py)^j \in A^\infty,$$

where $(F_py)^j = (F_pz)^{-j}$ if $j < 0$. We leave it as an exercise to show that F_pa is again an element of $A^\dagger$. So now we have our lift of Frobenius F_p on $A^\dagger$. We remark that the reason to invert y was precisely to be able to impose $F_px = x^p$ on our Frobenius lift. One can also lift Frobenius without inverting y, and it might be interesting to see if this improves the constants (not the exponents) in the running time of Kedlaya's algorithm. We also remark that wanting a Frobenius lift forces us to work with $A^\dagger$, and not with the ring of functions on a fixed neighborhood of the region given by $|y| = 1$, as the Frobenius lift sends each neighborhood to a *larger* neighborhood. Indeed, consider power series $\sum_{i\geq 0} a_ix^i$, $a_i \in \mathbb{Z}_p$ and the Frobenius lift that sends x to x^p, then one has:

$$\sum_i a_ix^i \mapsto \sum_i a_ix^{pi}, \quad \liminf v(a_i)/ip = p^{-1} \liminf v(a_i)/i.$$

5.3 Frobenius action on H^1 and consequences

Our Frobenius lift F_p on $A^\dagger$ induces a σ-linear endomorphism of the de Rham complex over $A^\dagger$, and hence a σ-linear endomorphism of the $\mathbb{Q}_q$-vector space $\mathrm{H}^1(C_{\overline{Q}})^-$. Recall that the $[x^i(dx)/y]$, $0 \leq i < d-1$, form a basis. We have:

$$F_p\colon x^i(dx)/y \mapsto px^{pi+p-1}y(F_py)^{-1}(dx)/y = px^{p(i+1)-1}y(F_pz)(dx)/y,$$

with

$$y(F_pz) = y^{-p+1} \sum_{k\geq 0} \tbinom{-1/2}{k} p^k E^k y^{-2pk}.$$

Hence:

$$F_p(x^i(dx)/y) = \sum_{k\geq 0} \binom{-1/2}{k} p^{k+1} E^k x^{p(i+1)-1} y^{-(2k+1)p+1} \cdot (dx)/y.$$

The degree of E is at most $pd-1$, hence that of $E^k x^{p(i+1)-1}$ is at most $p(d-1)-1+k(pd-1)$. We note that:

$$\frac{k(pd-1)+p(d-1)-1}{d} < (k+1)p$$

So, for $k \geq 0$ we can write:

$$\binom{-1/2}{k} E^k x^{p(i+1)-1} y^{-(2k+1)p+1} = \sum_{-(2k+1)p<j<p} c_{i,k,j} y^j, \quad \text{with } c_{i,k,j} \text{ in } \mathbb{Z}_q[x] \text{ and } \deg(c_{i,k,j}) < d.$$

Then we have, for $0 \leq i < d-1$:

$$F_p(x^i(dx)/y) = \sum_{\substack{k\geq 0 \\ -(2k+1)p<j<p}} p^{k+1} c_{i,k,j} y^j \cdot (dx)/y.$$

Now we need apply Lemmas 4.3.4 and 4.3.5 in order to rewrite $[F_p(x^i(dx)/y)]$ in terms of our basis. We find:

$$[F_p(x^i(dx)/y)] = \left[\sum_{k\geq 0} p^{k+1} c'_{i,k}(dx)/y\right],$$
$$\text{with } c'_{i,k} \text{ in } \mathbb{Q}_q[x],\ \deg(c'_{i,k}) < d,\ p^{m_k} c'_{i,k} \in \mathbb{Z}_q[x],$$
$$m_k = \max(\lfloor \log_p((2k+1)p) \rfloor, \lfloor \log_p(pd-2) \rfloor).$$

But we have not reduced $F_p(x^i(dx)/y)$ completely, because of the term $x^{d-1}(dx)/y$ that can still occur in the last formula. The last reduction step, that uses dy, needs a division by d. If d is not divisible by p, then this division gives an integral result, and the action of Frobenius on $\mathrm{H}^1(C'_{\overline{Q}})$ is given by a matrix with coefficients in $\mathbb{Z}_q$. But in general we do have to take this division into account, and we get the following. Let $c''_{i,k}$ be in $\mathbb{Q}_q[x]$ with $\deg(c''_{i,k}) < d-1$ such that $[c'_{i,k}(dx)/y] = [c''_{i,k}(dx)/y]$. Then:

$$[F_p(x^i(dx)/y)] = \left[\sum_{k\geq 0} p^{k+1} c''_{i,k}(dx)/y\right],$$
$$\text{with } c''_{i,k} \text{ in } \mathbb{Q}_q[x],\ \deg(c''_{i,k}) < d-1,\ p^{m_k+v_p(d)} c''_{i,k} \in \mathbb{Z}_q[x].$$

As we are going to consider the nth power of F_p acting on $\mathrm{H}^1(C'_{\overline{Q}})^-$, it is quite useful (but not essential for the complexity) to know whether or not the $\mathbb{Z}_q$-module $\oplus_{0\leq i<d-1} \mathbb{Z}_q \cdot x^i(dx)/y$ is

stable under F_p, i.e., whether or not the matrix of F_p with respect to this basis has coefficients in $\mathbb{Z}_q$. In the notation of above $k+1-m_k-v_p(d)$ can take negative values, in fact:

$$\min\{k+1-m_k-v_p(d) \mid k \geq 0\} = -v_p(d) - \lfloor \log_p(d-2/p) \rfloor.$$

So our estimates do not show that the coefficients in question are in $\mathbb{Z}_q$. In fact Vercauteren tells me that his computations show that very often the matrix in question does not have all coefficients in $\mathbb{Z}_q$. In order to not have to worry about the growth of the denominators when computing the action of F_p^n, the following proposition is quite useful, and seems new.

5.3.1 Proposition. *Let t be a parameter at the point ∞ at infinity such that $\iota t = -t$ (e.g., $t := x^g/y$). Let L be the sub-$\mathbb{Z}_q$-module of $\oplus_{0\leq i<d-1}\mathbb{Z}_q{\cdot}x^i(dx)/y$ consisting of those ω whose image in $(t^{-2g}\mathbb{Z}_q[[t]]dt)/(t^{-1}\mathbb{Z}_q[[t]]dt)$ can be integrated, i.e., are in the image of:*

$$d\colon \frac{t^{-2g+1}\mathbb{Z}_q[[t]]}{\mathbb{Z}_q[[t]]} \longrightarrow \frac{t^{-2g}\mathbb{Z}_q[[t]]dt}{t^{-1}\mathbb{Z}_q[[t]]dt}.$$

Then the action of F_p on $\mathrm{H}^1(C'_{\overline{Q}})^-$ respects the $\mathbb{Z}_q$-module L. Moreover, there is an isomorphism:

$$\frac{\oplus_{0\leq i<d-1}\mathbb{Z}_q{\cdot}x^i(dx)/y}{L} \longrightarrow \bigoplus_{\substack{-(2g-1)\leq i<0\\ i\equiv 1(2)}} \mathbb{Z}_q/i\mathbb{Z}_q,$$

hence the quotient $(\oplus_{0\leq i<d-1}\mathbb{Z}_q{\cdot}x^i(dx)/y)/L$ is annihilated by $p^{\lfloor \log_p(2g-1)\rfloor}$.

Proof. We just sketch the proof, as it is much too technical anyway. Let $\overline{C_Q}$ be the smooth projective curve over $\mathbb{Z}_q$ given by Q. Then there is a canonical isomorphism:

$$\mathrm{H}^1_{\mathrm{dR}}(\overline{C_Q}/\mathbb{Z}_q) = \mathrm{H}^1_{\mathrm{crys}}(\overline{C_{\overline{Q}}}/\mathbb{Z}_q),$$

between crystalline cohomology and de Rham cohomology (see [3]). Hence $\mathrm{H}^1_{\mathrm{dR}}(\overline{C_Q}/\mathbb{Z}_q)$ is canonically equipped with an action of Frobenius. Then we consider:

$$\mathrm{H}^1_{\mathrm{dR}}(\overline{C_Q}/\mathbb{Z}_q)\hookrightarrow \mathrm{H}^1_{\mathrm{dR}}(\overline{C_Q}/\mathbb{Q}_q) = \mathrm{H}^1_{\mathrm{dR}}(C'_Q/\mathbb{Q}_q)^- = \mathrm{H}^1(C'_{\overline{Q}})^-.$$

Then one shows that L is the image of $\mathrm{H}^1_{\mathrm{dR}}(\overline{C_Q}/\mathbb{Z}_q)$, by considering the complexes:

$$C(m) := \big(\mathcal{O}(m\infty) \to \Omega^1((m+1)\infty)\big)$$

on $\overline{C_Q}$, with $m \geq 0$. One finds that $\mathrm{H}^1_{\mathrm{dR}}(\overline{C_Q}/\mathbb{Z}_q) = \mathrm{H}^1(C(0))$. The quotient of $C(2g-1)$ by its subcomplex $C(0)$ is the complex given by the map d in the proposition that we are proving. One

finds the following diagram, whose first two rows and first two columns are exact:

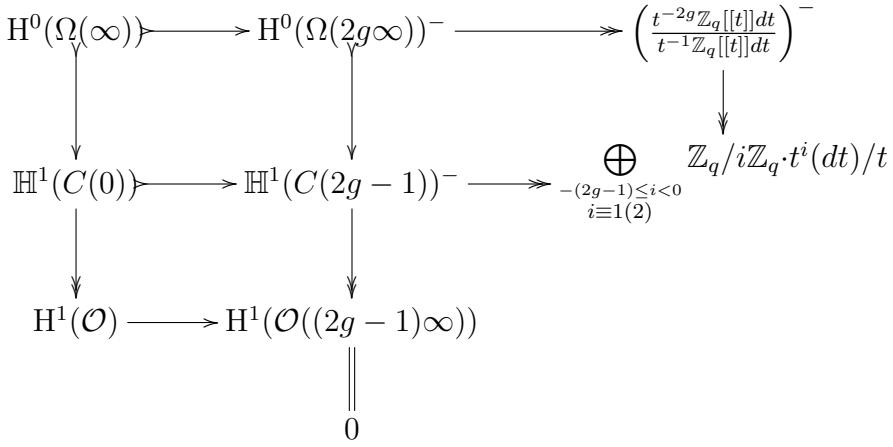

□

Let $e = (e_1, \ldots, e_{2g})$ be a $\mathbb{Z}_q$-basis of L, and let m be the matrix that gives the action of F_p on L with respect to e. This means that:

$$F_p e_j = \sum_i m_{i,j} e_i.$$

We note that Proposition 5.3.1 shows that m is in $\mathrm{M}_{2g}(\mathbb{Z}_q)$. As F_p is σ-linear, we have, for λ_j in $\mathbb{Z}_q$:

$$F_p \sum_j \lambda_j e_j = \sum_{i,j} \sigma(\lambda_j) m_{i,j} e_i.$$

Also, for a and b σ-linear endomorphisms of a $\mathbb{Q}_q$-vector space V with a basis e, we have:

$$\mathrm{mat}(ab)_e = (\mathrm{mat}\, a)_e (\sigma(\mathrm{mat}\, b))_e.$$

(Exercise left to the reader.) It follows that the matrix with respect to e of the linear endomorphism $F_q = F_p^n$ of $\mathrm{H}^1(C'_{\overline{Q}})^-$ is given by:

$$(\mathrm{mat}\, F_q)_e = m \cdot (\sigma m) \cdots (\sigma^{n-1} m).$$

5.3.2 Theorem. *The characteristic polynomial of F_p^n on the $\mathbb{Q}_q$-vector space $\mathrm{H}^1(C'_{\overline{Q}})^-$ is the polynomial:*

$$P_{\overline{Q}} = \prod_{i=1}^{2g} (t - \alpha_i) = t^{2g} - a_1 t^{2g-1} + \cdots - a_{2g-1} t + a_{2g}.$$

in $\mathbb{Z}[t]$ of Theorem 2.2. We have:

$$a_{2g-i} = q^{g-i} a_i, \qquad |a_i| \leq 2^{2g} q^{i/2}.$$

Proof. The Lefschetz fixed point formula for Monsky-Washnitzer cohomology gives that P_f is the characteristic polynomial of F_p (for a proof see the new book [1] by Fresnel and van der Put). The identity $\alpha_i \alpha_{g+i} = q$ is a consequence of Poincaré duality. This implies $a_{2g-i} = q^{g-i} a_i$, see [2] or [12]. That $|\alpha_i| = q^{1/2}$ is due to Weil (because this is for a curve, Deligne has proved such a thing for general non-singular projective varieties) (see the appendix of Hartshorne's book). We use it as follows:

$$|a_i| = |\sum_{j_1 < \cdots < j_{2g}} \alpha_{j_1} \cdots \alpha_{j_i}| \leq \sum_{j_1 < \cdots < j_{2g}} |\alpha_{j_1} \cdots \alpha_{j_i}| = \binom{2g}{i} q^{i/2} \leq 2^{2g} q^{i/2}.$$

□

It follows that it suffices to compute the a_i, for $1 \leq i \leq g$, as elements of $\mathbb{Z}_q / p^N \mathbb{Z}_q$, where N is chosen such that:

$$p^N > 2 \cdot 2^{2g} \cdot q^{g/2}.$$

The computations earlier in this subsection show that the power series in y and y^{-1} can be truncated at an appropriate place, so that the computation becomes finite. This will be done in detail in the next section.

6 The algorithm

The *input data* are the following: a finite field $\mathbb{F}_q$, given by means of a prime number $p > 2$, an integer $n \geq 1$ and a monic irreducible polynomial $\overline{f}$ in $\mathbb{F}_p[z]$ of degree n, an odd integer $d = 2g+1 \geq 1$ and an element $\overline{Q}$ of $\mathbb{F}_q[x]$, monic and of degree d, such that $\gcd(\overline{Q}, \overline{Q}') = 1$. The *output* is: the polynomial $P_{\overline{Q}}$ in $\mathbb{Z}[t]$ described in Theorem 2.2. Now we describe the algorithm.

6.1 Initialisation

First lift $\overline{f}$ to f in $\mathbb{Z}[z]$ by lifting the coefficients in $\{0, 1, \ldots, p-1\}$. Let $\mathbb{Z}_q := \mathbb{Z}_p[z]/(f)$. Lift $\overline{Q}$ to Q in $\mathbb{Z}_q[x]$ in the obvious way: $Q = x^d + Q_{d-1} x^{d-1} + \cdots + Q_0$, with $Q_i = \sum_{0 \leq j < n} Q_{i,j} z^j$ in $\mathbb{Z}_q = \mathbb{Z}_p 1 \oplus \cdots \oplus \mathbb{Z}_p x^{n-1}$ and the $Q_{i,j}$ in $\{0, 1, \ldots, p-1\}$.

Put:

$$N := \lceil \log_p(2^{2g+1} q^{g/2}) \rceil, \quad \text{so that} \quad p^N > 2^{2g+1} q^{g/2}$$
$$N_1 := N + v_p(d) + \lfloor \log_p(d - 2/p) \rfloor + \lfloor \log_p(2g-1) \rfloor.$$

Let M be the smallest integer such that:

$$M - \lfloor \log_p(2M+1) \rfloor \geq N_1.$$

6.2 Compute the action of Frobenius on the differentials

For $0 \leq i < d-1$, compute:

$$\sum_{0 \leq k < M} \binom{-1/2}{k} p^{k+1} E^k x^{p(i+1)-1} y^{-(2k+1)p+1} = \sum_{\substack{0 \leq k < M \\ -(2k+1)p < j < p}} p^{k+1} c_{i,k,j} y^j$$

as an element of:

$$\mathbb{Z}_q[x, y, y^{-1}]/(y^2 - Q) = \bigoplus_{\substack{0 \leq k < d \\ l \in \mathbb{Z}}} \mathbb{Z}_q x^k y^l,$$

with a precision of N_1 "digits", i.e., modulo p^{N_1} if you wish. Recall that $E = (FQ - Q^p)/p$. The $c_{i,k,j}$ are in $\mathbb{Z}_q[x]$ of degree $< d$.

6.3 Applying the reduction algorithm

Apply the reduction algorithm of page 11:

$$\sum_{\substack{0 \leq k < M \\ -(2k+1)p < j < p}} p^{k+1} c_{i,k,j} y^j \cdot (dx)/y \equiv \widetilde{m}_i \cdot (dx)/y,$$

with $\widetilde{m}_i$ in $\mathbb{Z}_q[x]$, $\deg(\widetilde{m}_i) < d-1$. Here we work with N_1 digits left of the comma, if at some point a division by p has to be done, one just shifts the number one place to the right, inserting whatever one wants at the left. Let $\widetilde{m}$ be the matrix whose columns are the $\widetilde{m}_i$; it is square, of size $2g$. Now compute a $\mathbb{Z}_q$-basis e for the $\mathbb{Z}_q$-module L as in Proposition 5.3.1, with a precision of N digits, and compute the matrix m of F_p with respect to e. Note that we have seen that m is in $\mathrm{M}_{2g}(\mathbb{Z}_q)$, and that we have its N most significant digits.

6.4 Computing the characteristic polynomial

Compute:

$$m' = m \cdot (\sigma m)(\sigma^2 m) \cdots (\sigma^{n-1} m).$$

Then m' is in $\mathrm{M}_{2g}(\mathbb{Z}_q)$. Compute:

$$\det(t \cdot \mathrm{id} - m) = P = t^{2g} + P_{2g-1} t^{2g-1} + \cdots + P_0 \quad \text{in } \mathbb{Z}_q[t].$$

For $1 \leq i \leq g$ let a_i be the unique integer with $|a_i| \leq 2^{2g} q^{i/2}$ such that $a_i = (-1)^i P_i$ in $\mathbb{Z}/p^N\mathbb{Z}$ (indeed, we know that the images of these P_i in $\mathbb{Z}_q/p^N\mathbb{Z}_q$ are in $\mathbb{Z}/p^N\mathbb{Z}$). For $g < i \leq 2g$, put $a_i := q^{i-g} a_{2g-i}$. Then one has:

$$P_{\overline{Q}} = t^{2g} - a_1 t^{2g-1} + \cdots - a_{2g-1} t + a_{2g}.$$

7 Running time and space requirement of the algorithm

As minimisation of the running time of an algorithm is not my specialty, I just give the facts that are stated in Kedlaya's article.

Note that $N_1 = O(gn)$. An element of $\mathbb{Z}_q/p^{N_1}\mathbb{Z}_q$ can be stored in $O(gn^2)$ bits. All ring operations in $\mathbb{Z}_q/p^{N_1}\mathbb{Z}_q$ can be done in time $O(g^{1+\varepsilon}n^{2+\varepsilon})$, using fast integer multiplication.

For $0 \leq k < n$ the matrix of the automorphism σ^k of $\mathbb{Z}_q/p^{N_1}\mathbb{Z}_q = (\mathbb{Z}/p^{N_1}\mathbb{Z})[z]/(f)$ can be computed in time $O(g^{1+\varepsilon}n^{3+\varepsilon})$: first compute in $\mathbb{F}_q$, then lift using Newton's method.

Step 6.2 can be done in time $O(g^{3+\varepsilon}n^{3+\varepsilon})$, and the storage space is $O(g^3n^3)$ bits. Note that $M = O(gn)$, $\deg(E) = O(g)$, $d = \deg(Q) = O(g)$. Basically, this step means computing a polynomial of degree $O(g^2n)$ in x and writing it as a polynomial in Q with coefficients of degree $< d$.

Step 6.3, the reduction step, can be done in time $O(g^{4+\varepsilon}n^{3+\varepsilon})$, no additional space is required. We have g forms to reduce, each takes $M = O(gn)$ multiplications by $(Q')^{-1}$ in $\mathbb{Z}_q[x]/(Q, p^N)$. So: $g{\cdot}gn{\cdot}g^{2+\varepsilon}n^{2+\varepsilon}$.

Step 6.4 can be done in time $O(g^{4+\varepsilon}n^{2+\varepsilon}) + O(g^{3+\varepsilon}n^{3+\varepsilon})$, and no additional space is required. One computes $m' = m{\cdot}(\sigma m)(\sigma^2 m)\cdots(\sigma^{n-1}m)$ by repeated squaring ($m_1 := m{\cdot}(\sigma m)$, $m_2 := m_1{\cdot}(\sigma^2 m_1)$, etc.). So one computes $O(\log n)$ multiplications of square matrices of size $2g$, with coefficients of size gn^2, and $O(g^2 \log n)$ applications of powers of σ. One computes the characteristic polynomial of m' by choosing an element v and computing v, $m'v$, $(m')^2v, \ldots$ until these are linearly dependent. That takes $O(g^3)$ ring operations.

One concludes that the dominant step is the reduction step, and that the total running time of the algorithm is $O(g^{4+\varepsilon}n^{3+\varepsilon})$, and that it requires $O(g^3n^3)$ bits of storage.

It seems (see Lauder's articles) that if one does not fix p, one has a running time of $(pg^{4+\varepsilon}n^{3+\varepsilon})$. As the algorithm manipulates polynomials that are not sparse and of degree at least p, this factor p is not easy to get rid of.

8 Acknowledgements

I take this opportunity to thank Frederik Vercauteren for informing me that the $\mathbb{Z}_q$-module generated by the cohomology classes of the $x^iy^{-1}dx$, $0 \leq i < 2g$, is *almost never* stable under the q-power Frobenius map, contrary to what I asserted in an older version of these notes. I also thank Robert Carls for discussions that we had preparing the graduate course and the exercises, and Marius van der Put for giving me a preliminary version of his book [1].

References

[1] J. Fresnel and M. van der Put. *Rigid geometry and applications.* To appear in Progress in Mathematics. Birkhäuser, Boston, Mass. (Goes to press on October 10, 2003.)

[2] R. Hartshorne. *Algebraic geometry.* Graduate Texts in Mathematics, No. 52. Springer-Verlag, New York-Heidelberg, 1977. xvi+496 pp. ISBN: 0-387-90244-9

[3] L. Illusie. *Crystalline cohomology.* Motives (Seattle, WA, 1991), 43–70, Proc. Sympos. Pure Math., 55, Part 1, Amer. Math. Soc., Providence, RI, 1994.

[4] B. Iversen. *Cohomology of sheaves.* Universitext. Springer-Verlag, Berlin, 1986. xii+464 pp. ISBN: 3-540-16389-1

[5] G. Kato and S. Lubkin. *Zeta matrices of elliptic curves.* Journal of Number Theory **15**, 318–330 (1982).

[6] K. Kedlaya. *Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology.* Journal of the Ramanujan Mathematical Society 16 (2001), 323-338. Also available as arXiv:math.AG/0105031.

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

[8] A. Lauder. *Deformation theory and the computation of zeta functions.* To appear in the Proceedings of the London Mathematical Society. Available on: `http://web.comlab.ox.ac.uk/oucl/work/alan.lauder/`

[9] T. Satoh. *The canonical lift of an ordinary elliptic curve over a finite field and its point counting.* J. Ramanujan Math. Soc. 15 (2000), no. 4, 247–270.

[10] T. Satoh. *On p-adic point counting algorithms for elliptic curves over finite fields.* In "Algorithmic Number Theory, 5th International Symposium, ANTS-V", Springer Lecture Notes in Computer Science 2369 (2002), 43–66.

[11] R. Schoof. *Counting points on elliptic curves over finite fields.* Les Dix-huitièmes Journées Arithmétiques (Bordeaux, 1993). J. Théor. Nombres Bordeaux 7 (1995), no. 1, 219–254.

[12] H. Stichtenoth. *Algebraic function fields and codes.* Universitext. Springer-Verlag, Berlin, 1993. x+260 pp. ISBN: 3-540-56489-6