

AN EXTENSION OF KEDLAYA'S ALGORITHM FOR HYPERELLIPTIC CURVES.

MICHAEL HARRISON

ABSTRACT. In this paper we describe a generalisation and adaptation of Kedlaya's algorithm for computing the zeta function of a hyperelliptic curve over a finite field of odd characteristic that the author used for the implementation of the algorithm in the MAGMA library. We generalise the algorithm to the case of an even degree model. We also analyse the adaptation of working with the $x^i dx/y^3$ rather than the $x^i dx/y$ differential basis. This basis has the computational advantage of always leading to an integral transformation matrix whereas the latter fails to in small genus cases. There are some theoretical subtleties that arise in the even degree case where the two differential bases actually lead to different redundant eigenvalues that must be discarded.

1. INTRODUCTION

Kedlaya's algorithm for hyperelliptic curves in odd characteristic was one of the first practical computational algorithms for computing the zeta function of a curve of genus greater than 1 over a large finite field of small characteristic [Ked01], [Ked04]. It was generalised by Denef and Vercauteren to characteristic two [DV06b] and has also been extended to more general curves like C_{ab} curves [DV06a]. Kedlaya's algorithm is based on the calculation of the Frobenius action on an appropriate p -adic cohomology group that can be described in sufficiently concrete terms for explicit computer computations to be made. In the hyperelliptic case, Kedlaya used Monsky-Washnitzer cohomology on the open affine subset of the curve defined by the removal of all Weierstrass points.

In 2003, the author wrote the implementation of Kedlaya's algorithm in the standard user library of the MAGMA computer algebra system [BCP97]. In practical terms, there appeared to be two issues with the algorithm as it stood.

Firstly, it only covered the odd degree case of a hyperelliptic model with a single point at infinity. Following Kedlaya's analysis, we extended the algorithm in a natural way to also cover the even degree case. The extension is fairly straightforward and the algorithm runs as before except that a degree one term has to be removed from the final characteristic polynomial corresponding to an extra eigenvalue q (the field size) arising from the extra point at infinity removed from the complete curve.

More seriously, if p , the characteristic of the base finite field \mathbf{F}_q , is small compared to the genus g of the hyperelliptic curve C - specifically if $p \leq 2g - 1$ in the odd degree case and $p \leq g$ in the even degree case - then the matrix M representing the σ -linear transformation of p -Frobenius

Key words and phrases. Kedlaya's Algorithm, Monsky-Washnitzer Cohomology, Magma.

on Kedlaya's chosen differential basis of cohomology is non-integral. That is, the integral lattice generated by the basis is not stable under Frobenius. Because M has to be σ -powered to a large degree to get to the final result, this presents obvious p -adic precision problems. If M represented a *linear* transformation, it could be easily replaced by an integral conjugate before powering (though then the characteristic polynomial of the power could be computed without matrix powering, anyway!), but the semi-linear situation is not so easy to work with. This issue is remarked upon in [Ked03] and can be dealt with in a number of ways. One approach is to try to analyse the situation using high-powered techniques like crystalline cohomology (or general F -module theory) to find an integral lattice to work with that is invariant under Frobenius. For example, Edixhoven gives a general criterion for stability under Frobenius of a sub \mathbf{Z}_p -module L of the \mathbf{Z}_p -module of differentials generated by Kedlaya's differential basis in Prop. 5.3.1 of [Edi06]. A full proof of the criterion can be found in [vdB08]. See also [CDV06], which is described further below, for more general plane curves.

In this very concrete situation, however, we computed that a slightly different differential basis for the minus part of the H^1 cohomology always works: namely differentials of the form $x^i dx/y^3$ rather than $x^i dx/y$. The computation is again straightforward but, as far as we are aware, it has not appeared in detail before in the literature so, for completeness, we will show that Kedlaya's reduction process applied to this space of differentials always leads to an integral matrix M .

The interesting technical point is that the y^3 differentials only form a basis for the minus part of the cohomology in the odd degree case. In the even degree case, the map from this space of differentials into H_-^1 actually has a 1-dimensional kernel and cokernel. It turns out that the kernel has eigenvalue 1 and cokernel eigenvalue q under q -Frobenius, so in this case we have to remove a factor of $t - 1$ rather than $t - q$ from the final characteristic polynomial. This is demonstrated in the final section of the paper.

In summary, in the even degree case, one additional eigenvalue of Frobenius occurs on the affine Monsky-Washnitzer cohomology because of the additional point removed at infinity. This merely has to be removed at the end in order to get the numerator of the zeta function. Our alternative set of differentials generate a \mathbf{Z}_p -module V with Frobenius action. This space genuinely gives an F -stable lattice in H_-^1 for odd degree and the algorithm goes through as before, except with guaranteed p -integral matrices. For even degree, V also gives p -integrality but $V \otimes \mathbf{Q}$ doesn't quite coincide with H_-^1 as a Frobenius module. However, an explicit analysis in this case shows that the difference between V and H_-^1 results in just having to remove a different additional eigenvalue at the end.

A very general Kedlaya-style algorithm applying to non-degenerate plane curves is presented by Castryck, Denef and Vercauteren in [CDV06]. There, a deterministic algorithm is given where a basis for cohomology is determined and an integrality analysis is performed involving Edixhoven's criterion and consideration of the Newton polygon of the curve. The hyperelliptic case, however, with its particular choices of differential bases, is still an important special case

amenable to the specific original analysis of Kedlaya and that presented here, and I have had requests from a number of people to publish details of the MAGMA implementation.

We should also mention some of the other point-counting methods which have been developed over the last decade for curves of genus greater than 1 and that use different techniques to that of Kedlaya.

Generalising the elliptic curve case, Mestre devised an algorithm for ordinary hyperelliptic curves in characteristic 2 based on the theory of the *canonical lift*. This computes a 2-adic approximation of a particular function of the eigenvalues of Frobenius from which a finite number of possibilities for the characteristic polynomial of Frobenius can be obtained by rational reconstruction in many cases (e.g. when the Jacobian is irreducible). Again generalising their algorithm for the genus 1 case, Lercier and Lubicz found a way to efficiently effect the lifting stage to obtain a quasi-quadratic algorithm [LL06]. The author implemented this algorithm for the standard MAGMA user library. Following the work of Robert Carls on theta structures of canonical lifts [Car07], Carls and Lubicz have generalised the algorithm to odd characteristic [CL09].

Another important p -adic method is the *deformation* method of Lauder and Wan [LW08]. This generalises from the curve case to higher-dimensional hypersurfaces and provides the basis for the computation of zeta-functions of fairly general varieties over finite fields. The ideas go back to Dwork and use his approach to p -adic cohomology theory, working with parametrised families of hypersurfaces and continuously deforming to ones of special form (diagonal in Dwork's original work). R. Gerkmann has further studied the method, considering relations to rigid cohomology and practical p -adic precision analysis [Ger07]. He has written an implementation in MAGMA. Fuller details for the deformation method in the particular case of hyperelliptic curves have been worked out by H. Hubrechts [Hub08] who provided the implementation that appears in the standard MAGMA user library.

A brief outline of the paper is as follows. In the next section, we introduce basic notation, summarise Kedlaya's original algorithm and describe our extension of it. We also give a brief overview of Monsky-Washnitzer cohomology and explain Kedlaya's reduction procedure on differentials which remains formally the same in the extended version.

In Section 3, we consider our alternative (pseudo)-basis and give a proof of the integrality of the reduction of the image of p -Frobenius on its elements alongside an analysis of Kedlaya's original basis. We also give the short proof of the generalisation of the point-counting formula to even degree hyperelliptic models.

Finally, in the last section we give proofs of the slightly more technical result relating the space spanned by our pseudo-basis to its image in Monsky-Washnitzer cohomology and giving the difference between the eigenvalues of Frobenius on these two spaces.

Acknowledgements I would like to thank the referees for many useful suggestions.

2. REVIEW OF KEDLAYA'S ALGORITHM

In this section we give a summary of Kedlaya's algorithm as it appears in [Ked01] as well as describing our extension of it. The basic notation introduced below will be used throughout the paper.

Basic notation.

Throughout, $q := p^n$ will denote a positive power of an odd prime p . k will denote the finite field \mathbf{F}_q , unless otherwise indicated. R will denote $W(k)$, the ring of integers of K , the unique unramified degree n extension of the local field \mathbf{Q}_p . σ_p will denote the p -Frobenius automorphism of R or K that reduces to $a \mapsto a^p$ on k .

C will denote the hyperelliptic curve which is the projective normalisation of the smooth plane affine curve C_1 with defining equation

$$y^2 = Q(x)$$

where $Q(x) = a_d x^d + \dots + a_0$ is a separable polynomial of degree d in $k[x]$. To simplify notation, we also use $Q(x)$ to denote some arbitrary lift of $Q(x)$ to R (i.e. a degree d polynomial over R such that reduction mod p of the coefficients gives $Q(x)$). It will always be clear from the context which polynomial is being referred to.

We let g denote the genus of C , so that $d = 2g + 1$ or $d = 2g + 2$. We refer to the $d = 2g + 1$ case as the odd case and the $d = 2g + 2$ case as the even case. In the odd case $C \setminus C_1$ consists of a single k -rational point, which is a Weierstrass point of C and will sometimes be referred to as ∞ . In the even degree case, $C \setminus C_1$ consists of a pair of non-Weierstrass points, ∞_1 and ∞_2 , which are either k -rational or conjugate points over \mathbf{F}_{q^2} . Computationally, it is easiest to transform the initial Q over $k[x]$ so that $a_d = 1$ and the lift to R of a_d is also 1. This may involve working with the quadratic twist of C in the even case, but there is no problem converting back the final result (by substituting $t \mapsto -t$ in the numerator of the zeta function). So from now on, we assume that a_d is 1 and C has two k -rational points at infinity in the even case.

Following Kedlaya, we define C^a as the open affine subset of C_1 given by inverting y ; i.e. $C^a = \text{Spec}(A_k)$ where

$$A_k := k[x, y, y^{-1}]/(y^2 - Q(x))$$

and we will let $A_R := R[x, y, y^{-1}]/(y^2 - Q(x))$ which is a finitely-generated, R -smooth R -algebra with $A_R \otimes_R k \simeq A_k$. C^a is just C with all Weierstrass points and points at infinity removed.

Basic outline of the algorithm.

Given an odd degree model of a hyperelliptic curve C over \mathbf{F}_q as above, Kedlaya's algorithm computes the degree $2g$ monic polynomial $L(X)$ that gives the numerator of the zeta-function of C [$\zeta_C(s) = L(q^{-s})/(1 - q^{-s})(1 - q^{1-s})$]. The number of points on C , $\#C(\mathbf{F}_{q^r})$, or the order of its Jacobian, $\#Jac(C)(\mathbf{F}_{q^r})$, over any finite extension \mathbf{F}_{q^r} of the base field can be simply computed from $L(X)$ in the usual way (e.g. see Appendix C [Har77]).

The main stages are given in Algorithm 1.

Algorithm 1 Kedlaya's original algorithm

Step 0: Input $Q(x)$.

Step 1: Working in $W(k)[x][[1/y]]$, compute $(1/y^\sigma)$ to sufficiently large p -adic and $(1/y)$ -adic precision by formally expanding

$$y^{-p} \left(1 + \frac{Q(x)^\sigma - Q(x)^p}{y^{2p}} \right)^{-1/2}$$

This gives a finite approximation of the image of the differential basis of cohomology $x^i(dx/y)$, $0 \leq i \leq d-2$ under p -Frobenius.

Step 2: Reexpress the images of the differentials as $W(k) \otimes \mathbf{Q}$ -linear combinations of the differential basis using the **RednA** and **RednB** reduction processes described below. This results in a $(2g)$ -by- $(2g)$ matrix M for the action of p -Frobenius to finite p -adic approximation.

Step 3: By binary-powering or similar, compute the product $N = MM^\sigma \dots M^{\sigma^{n-1}}$ and the characteristic polynomial $F_p(X)$ of N . This is actually equal to $L(X) \in \mathbf{Z}[X]$ but will have been determined here in $\mathbf{Z}_p[X]$ to a large, finite p -adic precision.

Step 4: Recover and return $L(X)$ from the p -adic approximation in step 3, using the Weil bound to guarantee correct integer coefficients.

Extension of the original algorithm.

We adapt/extend the original algorithm in two ways.

- Even degree models (d even) are allowed.
- When $p < 2g$, d odd, or $p \leq g$, d even, the differential pseudo-basis $x^i(dx/y^3)$, $0 \leq i \leq d-2$ is used rather than $x^i(dx/y)$.

The first change extends the algorithm to an arbitrary hyperelliptic curve (possibly after applying a quadratic twist as described earlier).

The second change guarantees that we always work with a p -Frobenius matrix M with p -integral coefficients. In the cases where we use the alternative differential basis, Kedlaya's original basis generally leads to a p -adically non-integral M . Strictly speaking, the alternative set of differential forms only form a basis for cohomology when d is odd. This is why we refer to it as a pseudo-basis. It still leads to correct results in the new algorithm. All of this, along with the justification for the new Step 4 below in the even degree case, is demonstrated in Sections 3 and 4.

The new algorithm is formally very similar to the original, so we will just state the changes that need to be made.

Steps 1 and 2. These are unaffected except that the expression to be formally expanded in step 1 has exponent $-3/2$ rather than $-1/2$ in the cases where the alternative differential

pseudo-basis is used. The matrix M in step 2 will be of size $2g + 1$ rather than $2g$ when d is even.

Step 4 Compute $L_1(X)$ from the p -adic approximation to $F_p(X)$ coming from Step 3. If d is even then let $L(X) = L_1(X)/(X - q)$ if using the dx/y basis or $L(X) = L_1(X)/(X - 1)$ if using the dx/y^3 pseudo-basis. If d is odd, just let $L(X) = L_1(X)$. Return $L(X)$.

The linear factor that has to be removed in the even case comes from an extra eigenvalue of the action of Frobenius on cohomology (see Section 3.1). That the factor is different for the pseudo-basis comes from the relation between it and an actual cohomology basis. The extra q eigenvalue is lost but a new eigenvalue 1 appears (see Section 4).

In practice, only half of the coefficients of $F_p(X)$ (those of the higher powers of X) need to be computed (because of the $\alpha \leftrightarrow q/\alpha$ symmetry of the algebraic roots of $L(X)$) and we can effectively remove the extra $X - q$ or $X - 1$ factor from $F_p(X)$ (rather than from $L_1(X)$ at the end) in the even degree case during these computations. The coefficients can be computed from the traces of the first g powers of N as a matrix over \mathbf{Z}_p . Removing the extra factors at this stage means that there is no necessity to increase the p -adic precision to which we need to know N beyond the same lower bound used in the odd-degree case. This is determined from the upper bound for the size of the (top) coefficients of $L(X)$ that comes from all of its roots (over \mathbf{C}) having absolute value \sqrt{q} . Expressions in g and n for the p -adic precision needed in the initial series expansion computed in Step 1 are given near the end of Section 3.

In the remainder of this section - which relates to Steps 1 and 2 - where we describe the Monsky-Washnitzer cohomology groups and the reduction procedures for Step 2, no distinction need be made between the even and odd degree cases except where indicated.

That the differential reductions of Step 2 take p -Frobenius transforms of elements of the pseudo-basis back into linear combinations of such elements will be demonstrated in Lemma 3.4.

Monsky-Washnitzer Cohomology. [MW68], [Mon68], [Mon71].

Let X be a non-singular affine scheme over k . Monsky and Washnitzer defined a p -adic cohomology theory for such X with appropriate fixed-point theorems for proving zeta-function results. Kedlaya used this (originally at least) to provide the technical basis for his algorithm. Monsky-Washnitzer cohomology agrees with Berthelot's more general rigid cohomology in the affine case and is pleasantly explicit in its definition. We will need some of its properties for later proofs and so we give a brief description of the theory here.

Let A_k temporarily represent the affine coordinate ring of our general X and A_R/R be a lift to an R -smooth R -algebra as above and $A_K = A_R \otimes_R K$.

Definition 2.1. Let F_q be the k -linear Frobenius endomorphism $A_k \rightarrow A_k$ given by $a \mapsto a^q$. Similarly, let F_p be the k -semilinear endomorphism of A_k , $a \mapsto a^p$.

The goal is to define a good cohomology group which simulates de Rham cohomology of A_K and to which F_q lifts. F_q lifts to the p -adic completion, $\hat{A}_R = \varprojlim_n A_R/p^n A_R$, but the de Rham cohomology of $\hat{A}_K = \hat{A}_R \otimes_R K$ is usually bigger than that of A_K . Monsky-Washnitzer define a subalgebra A^\dagger of \hat{A}_R , referred to as the weak completion, as follows. If x_1, \dots, x_r are R -algebra generators of A_R then

$$A^\dagger := \left\{ \sum_{n=0}^{\infty} a_n p_n(x_1, \dots, x_r) : a_n \in p^n R, p_n \text{ of total degree} \leq C(n+1) \text{ for some } C > 1 \right\}$$

and $A_K^\dagger = A^\dagger \otimes_R K$. Up to isomorphism, A^\dagger is shown to be independant of the lift A_R and the generators x_i .

$\tilde{\Omega}_{A_K^\dagger/K}^i$ is the separated i th differential module, the plain differential module $\Omega_{A^\dagger}^i$ of A^\dagger divided out by the intersection $\cap_n p^n \Omega_{A^\dagger}^i$ and tensored with K . There is the usual differential complex

$$0 \rightarrow A_K^\dagger \xrightarrow{d} \tilde{\Omega}_{A_K^\dagger/K}^1 \xrightarrow{d} \tilde{\Omega}_{A_K^\dagger/K}^2 \xrightarrow{d} \dots$$

the homology groups of which give the MW cohomology groups $H^i(A_k; K)$.

If A_k is of Krull dimension 1, as in our case, then $\tilde{\Omega}_{A_K^\dagger/K}^i = 0$ for all $i \geq 2$ and so $H^1(A_k; K) = \tilde{\Omega}_{A_K^\dagger/K}^1/d(A_K^\dagger)$ and all higher cohomology is trivial.

If F_q lifts to F on A^\dagger then F functorially induces a K -linear automorphism F_* of the H^i , which is independent of the choice of lift, and there is a cohomological trace formula for $\#X(\mathbf{F}_{q^m})$ for all $m \geq 1$ (see next section). Furthermore, if F_p lifts to a σ -semilinear map $F_p^\dagger : A^\dagger \rightarrow A^\dagger$, then F_p^\dagger induces a σ -semilinear automorphism F_{p*} of the H^i with $F_* = F_{p*}^n$.

Now let A_k, A_R refer to the hyperelliptic algebras again. The inversion of y allows Kedlaya to define a lift of F_p to A^\dagger by

$$x \mapsto x^p \quad y \mapsto y^p \left(1 + \frac{Q^\sigma(x^p) - [Q(x)]^p}{y^{2p}} \right)^{1/2} \quad y^{-1} \mapsto y^{-p}(1 + \dots)^{-1/2}$$

The congruence $Q^\sigma(x^p) \equiv Q(x)^p \pmod{pR[x]}$ means that the standard power series expansions of $(1 + \dots)^{1/2}$ and $(1 + \dots)^{-1/2}$ converge to elements in A_K^\dagger .

In fact, Kedlaya gives the following explicit description of A^\dagger :

$$A^\dagger = \left\{ \sum_{n=-\infty}^{\infty} S_n(x) y^n : \deg(S_n) \leq d-1 \quad \liminf_{n \rightarrow \infty} \frac{v_p(S_n)}{n} > 0 \quad \liminf_{n \rightarrow \infty} \frac{v_p(S_{-n})}{n} > 0 \right\}$$

where $v_p(f)$, $f \in R[x]$ is the smallest m such that $f \in p^m R[x]$.

The hyperelliptic involution $\omega : x \mapsto x, y^{\pm 1} \mapsto -y^{\pm 1}$ extends to A^\dagger (and $A^\dagger \otimes_R K$) giving the direct sum decomposition

$$A^\dagger = A_+^\dagger \oplus A_-^\dagger \quad \text{with} \quad A_+^\dagger = \left\{ \sum S_{2n} y^{2n} \right\}, \quad A_-^\dagger = \left\{ \sum S_{2n+1} y^{2n+1} \right\}$$

and a corresponding decomposition of $H^1(A_k; K)$ into $+$ and $-$ components. Kedlaya shows that the Monsky-Washnitzer trace formula leads to the result that the numerator of the zeta-function of C is just the characteristic polynomial of F_* on H_-^1 when d is odd. We will show in Section 3.1 that the same analysis gives only a minor difference when d is even.

Reduction steps in the computation of F_*

Kedlaya shows that a K -basis for the finite-dimensional $H^1(A_k; K)$ is given by the A_K differentials

$$\{x^i dx/y : 0 \leq i \leq d-2\} \cup \{x^i dx/y^2 : 0 \leq i \leq d-1\}$$

the first set giving a basis for H_-^1 and the second for H_+^1 . We come back to this in the next section where we note that it also holds for d even.

The first stage of the algorithm consists of expanding the series for $F_{p*}(1/y)$ to sufficient p -adic precision. We will give a precise value for the precision required at the end of Section 3.

The second stage consists of applying two types of reduction to reexpress these images as K -linear combinations of basis elements. The two basic relations are

$$y^2 = Q(x) \quad \text{and} \quad dy = (Q'(x)/2y)dx$$

where the prime denotes the standard derivative.

As Q and Q' are relatively prime in $k[x]$, there exist $U, V \in R[x]$ such that $UQ + VQ'$ is 1. Therefore, for any $S \in R[x]$, there exist $A, B \in R[x]$ with $S = AQ + BQ'$. Then, for $m \neq 2$,

$$S \frac{dx}{y^m} = A \frac{dx}{y^{m-2}} + 2B \frac{dy}{y^{m-1}} = A \frac{dx}{y^{m-2}} + \left(\frac{2}{m-2} \right) B' \frac{dx}{y^{m-2}} - \left(\frac{1}{m-2} \right) d \left(\frac{2B}{y^{m-2}} \right)$$

This gives the first reduction type:

$$\boxed{\text{RednA}} \quad S \frac{dx}{y^m} \equiv \left(A + \left(\frac{2}{m-2} \right) B' \right) \frac{dx}{y^{m-2}} \quad \text{if } S = AQ + BQ'$$

to reduce m by 2 when $m > 2$. Note that in practice, we only apply this for $\deg(S) < d$ because we begin by recursively dividing S by Q (which is monic) to express S as $S_0 + S_1Q + S_2Q^2 + \dots$ with $S_i \in R[x]$, $\deg(S_i) < d$ and then replace Q^i by y^{2i} . In fact, we only divide out by Q and replace by y^2 while this leads to negative powers of y in the expression. Note also that if $\deg(S) < d$ (in fact, if $\deg(S) < 2d-1$), then A and B can be chosen as $SU \bmod Q'$ and $SV \bmod Q$, so with $\deg(A) < d-1$ and $\deg(B) < d$.

In this way, **RednA** applied recursively reduces $S(dx/y^m)$ to a $T(dx/y)$ or $T(dx/y^2)$ depending on the parity of m . Note also, that if the initial m was ≤ 0 , then we could shift up instead by replacing a positive power y^{2i} by Q^i , but this case doesn't occur.

The second reduction uses the differential equalities (for $r \geq 0$)

$$\begin{aligned} d(x^r) &= rx^{r-1}dx = rx^{r-1}Q(x)(dx/y^2) \quad \text{leading term } rx^{r+d-1} \\ d(2x^r y) &= [2rx^{r-1}Q(x) + x^r Q'(x)](dx/y) \quad \text{leading term } (2r+d)x^{r+d-1} \end{aligned}$$

Subtracting multiples of the right hand sides of these from $T(dx/y^2)$ or $T(dx/y)$, reduces the degree of T until we are back to linear combinations of basis elements. This will be referred to as **RednB**.

Applying a number of **RednA** followed by a number of **RednB** steps thus reduces any $S(dx/y^m)$ to a linear combination of basis elements. The reductions can clearly introduce a power of p into the denominator of the final expression. This should be accurately estimated in order to gauge *a priori* what the loss of p -adic precision may be and if there will be non-integral terms at the end. We give the analysis in Section 3.

Stages 1 and 2 of the algorithm give an explicit $(d-1)$ -by- $(d-1)$ matrix M over R which represents the σ -linear transformation F_{p*} on H_-^1 with respect to the chosen $x^i(dx/y)$ basis. Computationally, the entries of M will be finite approximations of the exact values which are correct mod p^N for some N depending on the p -adic precision that we carried out the stage 1 expansion to and on the loss of precision in stage 2. The final stage is to compute the n th twisted power of M : $M^{\sigma^{n-1}} M^{\sigma^{n-2}} \dots M$. This gives the matrix of F_* on H_-^1 and we just need its characteristic polynomial, $P_F(t)$.

If M is p -integral, $P_F(t)$ will be correct mod p^N and the Weil bound tells us how large N should be taken for this to determine the numerator of the zeta function of C . If M is non-integral, it is hard to give good *small* estimates of the p -adic precision lost in the twisted powering. Therefore, it is highly desirable to have a p -integral M . As we show in Section 3, for small p , the $x^i(dx/y)$ basis will usually lead to M with denominators whereas the $x^i(dx/y^3)$ pseudo-basis never does.

3. ADAPTATION OF THE BASIC ALGORITHM

In this section we describe in detail the adaptations to Kedlaya's algorithm outlined in the introduction and previous section, and provide proofs of correctness.

3.1. Zeta function formula: even or odd case.

Let $P_C(t)$ be the numerator of the zeta-function of C/k (see, eg, App. C, [Har77]). The polynomial $P_C(t) = t^{2g} + c_{2g-1}t^{2g-1} + \dots + c_0$, a monic polynomial over \mathbf{Z} . Its roots over \mathbf{C} , $\{\alpha_i\}$, all have absolute value $q^{1/2}$ and this set is invariant under $\alpha \mapsto q/\alpha$. Furthermore, if $S_r(\alpha) = \alpha_1^r + \dots + \alpha_{2g}^r$ then

$$\#C(\mathbf{F}_{q^r}) = q^r + 1 - S_r(\alpha) \quad \forall r \geq 1$$

Lemma 3.1. *The characteristic polynomial of F_* acting K -linearly on $H^1(A_k; K)_-$ is $P_C(t)$ when d is odd, and is $(t-q)P_C(t)$ when d is even.*

Proof. The following argument is from [Ked01] when d is odd. From the explicit description of A^\dagger , it follows immediately that, if $B_k = k[x]_Q$ and $B_R = R[x]_Q$, then F_p lifts to B^\dagger as a σ -linear map with $x \mapsto x^p$ and

$$A_+^\dagger \simeq B^\dagger \quad \text{and} \quad (\tilde{\Omega}_{A^\dagger/R}^1)^+ \simeq \tilde{\Omega}_{B^\dagger/R}^1$$

as F_p -modules. Thus (abbreviating $H^i(A_k; K)$ to H^i and using subscripts for the \pm eigenspaces), $H^0 = H_+^0$ and H_+^1 are F_* -isomorphic to the cohomology groups for $\text{Spec}(B_k)$. This is isomorphic to $\mathbb{P}^a := \mathbb{P}_k^1 \setminus S$, where S is the set of finite places corresponding to the irreducible factors of $Q \in k[x]$ and the point at infinity.

Essentially, the contribution to cohomology resulting from the removal of closed points from C to get to C^a is precisely matched by the removal of the corresponding points from \mathbb{P}^1 in the odd case and gives the H_+^0 component. When d is even, as well as the Weierstrass points, we are removing 2 k -rational points from C at infinity which are swapped by the hyperelliptic involution and lie over a single k -rational point of \mathbb{P}^1 . This leads to an extra eigenvalue q in each of the $+$ and $-$ components of H^1 . Formally, this follows easily from the trace formula as we now show.

The fixed-point theorem for Monsky-Washnitzer cohomology gives the following trace formula for a general affine X/k of dimension n with (finite-dimensional) cohomology groups H^i :

$$\#X(\mathbf{F}_{q^r}) = \sum_{i=0}^n (-1)^i \text{Trace}_K((q^n F_*^{-1})^r | H^i) \quad \forall r \geq 1$$

Let N_r = the number of roots of $Q(x)$ over \mathbf{F}_{q^r} and $\delta = 0$ if d is odd and 1 if d is even. The MW trace formula for C^a and \mathbb{P}^a and Weil formula for $\#C(\mathbf{F}_{q^r})$ give

$$\begin{aligned} (C^a) \quad q^r - S_r(\alpha) - N_r - \delta &= \text{Tr}((qF_*^{-1})^r | H^0) - \text{Tr}((qF_*^{-1})^r | H_+^1) - \text{Tr}((qF_*^{-1})^r | H_-^1) \\ (\mathbb{P}^a) \quad q^r - N_r &= \text{Tr}((qF_*^{-1})^r | H^0) - \text{Tr}((qF_*^{-1})^r | H_+^1) \end{aligned}$$

Subtracting gives

$$\text{Tr}((qF_*^{-1})^r | H_-^1) = S_r(\alpha) + \delta \quad \forall r \geq 1$$

which implies that the eigenvalues of qF_*^{-1} on H_-^1 are $\{\alpha_i\} \cup \{1\}_{d \text{ even}}$. Hence, the eigenvalues of F_* are $\{\alpha_i\} \cup \{q\}_{d \text{ even}}$.

Therefore the characteristic polynomial of F_* on H_-^1 is $P_C(t)$, if d is odd, or $(t - q)P_C(t)$, if d is even. \square

3.2. Differential basis choices.

We first note that Kedlaya's assertion that $\{x^i dx/y : 0 \leq i \leq d-2\} \cup \{x^i dx/y^2 : 0 \leq i \leq d-1\}$ form a basis for H^1 remains true for d even.

By Thm. 5.6 of [MW68], the natural map $H_{dR}^1(C_K^a/K) \rightarrow H^1(A_k; K)$ is an isomorphism, where C_K, C_K^a are the hyperelliptic lifts of C, C^a to K corresponding to the lift of $Q(x)$. The reductions **RednA** and **RednB** on algebraic differentials show that the above set of differentials generate $H_{dR}^1(C_K^a/K)$ and a similar argument shows that no nontrivial K -linear sum of them is of the form df for $f \in K[x, y, y^{-1}]/(y^2 - Q(x))$ [Note: any element of this algebra is a finite sum of the form $\sum_{n=0}^N a_n(x)y^{-n}$]

Remark. That the given differentials form a basis also follows easily from general de Rham theory for complete curves and their open affine subsets applied to $H_{dR}^1(C_K/K)$ and $H_{dR}^1(C_K^a/K)$.

Definition 3.2. We define two sets of differentials, B_1 and B_2 .

$$B_1 = \{dx/y, x(dx/y), \dots, x^{d-2}(dx/y)\}$$

$$B_2 = \{dx/y^3, x(dx/y^3), \dots, x^{d-2}(dx/y^3)\}$$

The classes of the differentials in B_1 give a basis for H_-^1 . B_2 is our pseudo-basis whose classes only give a basis for H_-^1 when d is odd, as we shall see.

For convenience, we also define V_2 as the $(d-1)$ -dimensional K -vector subspace of $\tilde{\Omega}_{A_K^\dagger/K}^1$ with basis B_2 and η as the class map into H_-^1

$$\eta : V_2 \longrightarrow H_-^1 \quad x^{i-1}(dx/y^3) \mapsto [x^{i-1}(dx/y^3)]$$

Lemma 3.3.

(i) (*Kedlaya*) Let $m > 2$, $S \in R[x]$ with $\deg(S) \leq d-1$. Under **RednA**, let

$$S(dx/y^m) \equiv T(x)\{(dx/y) \text{ } m \text{ odd}, (dx/y^2) \text{ } m \text{ even}\} \quad T(x) \in K[x], \deg(T) < d$$

then $p^{\lfloor \log_p(m-2) \rfloor} T \in R[x]$.

(ii) Let $S \in R[x]$ with $\deg(S) = m \geq d-1$. Under **RednB** let

$$S(dx/y) \equiv T(x)(dx/y) \quad T(x) \in K[x], \deg(T) < d-1$$

then $p^{\lfloor \log_p(2m-d+2) \rfloor} T \in R[x]$. If d is even, $p^{\lfloor \log_p(m-(d/2)+1) \rfloor} T \in R[x]$.

In either case, the $d(\sum_a^b S_r(x)y^r)$ differential giving the reduction can be chosen with $p^u S_r(x) \in R[x] \forall r$ for the same p^u .

Proof. i) is just Lemma 2 of [Ked01]. Note that in the statement of that Lemma, $\log_p(2m+1)$ should be replaced by $\log_p(2m-1)$ (with $m \geq 1$) and in the proof, every $\pm m$ as the upper or lower limit of a sum should be replaced by $\pm(m-1)$. The proof of the lemma works just as well for d even or odd and the final statement about $d(\sum_a^b S_r(x)y^r)$ above is what is actually proven in Lemma 2.

ii) This is essentially Lemma 3 of [Ked01] (or rather the corrected statement in the errata, [Ked03]). As Kedlaya notes, ii) and the statement about $d(\sum_a^b S_r(x)y^r)$ follow in the same way as part i) (but more easily). We have that $S(dx/y) - d(\sum_{r=0}^{m+d-1} 2a_r x^r y) = T(dx/y)$, $d(2x^r y) = ((d+2r)x^{d+r-1} + \dots)(dx/y)$ and the coefficient of x^s in T is zero for $s \geq d-1$. Kedlaya's argument - considering formal expansions of the differentials with respect to a local parameter at one of the points at infinity - effectively shows that the largest power of p that may occur in denominators is the largest power of p that can divide *one* of the $d+2r$ (rather than their product). When d is even, it is only necessary to consider divisibility of $(d/2) + r$ since p is odd. \square

Any element of $\tilde{\Omega}_{A_K^\dagger/K}^1$ can be written uniquely in the form $\sum_{-\infty}^{+\infty} S_n(x)y^n dx$ with $\deg(S_n) < d$, which we refer to as its *standard expansion*.

Lemma 3.4.

- (i) For all $\omega \in B_2$, the standard expansion of $F_{p*}\omega$ is of the form $\sum_{n \geq 3} B_n(x)(dx/y^n)$.
- (ii) **RednA** on the $\sum_{n \geq 1} S_n(x)(dx/y^n)$ part of the standard expansion of $F_{p*}(x^i(dx/y))$ (resp. $F_{p*}(x^i(dx/y^3))$) gives a form which is a linear combination of elements of B_1 (resp. B_2) with p -integral coefficients.
- (iii) Consider the coefficients of the B_1 expansion resulting from **RednB** on the $\sum_{n \geq 1} S_n(x)y^n dx$ part of the standard expansion of $F_{p*}(x^{i-1}(dx/y))$.
 - (a) If $d = 2g + 1$, then these coefficients are p -integral for $i \leq g$ and for $i = g + r$ have denominator bounded by $p^{-\lfloor \log_p(2r-1) \rfloor}$.
 - (b) If $d = 2g + 2$, then these coefficients are p -integral for $i \leq g + 1$ and for $i = g + r + 1$ have denominator bounded by $p^{-\lfloor \log_p(r) \rfloor}$.

By part (i), we can use **RednA** to reduce $F_{p*}\omega$ back to linear combinations of elements in B_2 rather than descending to B_1 . This is what is meant in part (ii). In this way, we get a σ_p -linear map (also denoted F_{p*}) $V_2 \rightarrow V_2$.

Proof. We have, for $1 \leq i \leq d - 1$, $k = 0$ or 1 ,

$$\begin{aligned}
 F_{p*}(x^{i-1}(dx/y^{2k+1})) &= x^{p(i-1)}y^{-(2k+1)p} \left(1 + p \left(\frac{Q^{\sigma_p}(x^p) - (Q(x)^p)}{p}\right) y^{-2p}\right)^{-(2k+1)/2} d(x^p) \\
 &= px^{pi-1}y^{-(2k+1)p} (1 + p\{a_1(x)y^{-2} + \dots + a_p(x)y^{-2p}\})^{-(2k+1)/2} dx \\
 &= px^{pi-1}y^{-(2k+1)p} \left(1 + \sum_{n=1}^{\infty} \binom{-(2k+1)/2}{n} p^n \{\dots\}^n\right) dx \\
 &= px^{pi-1} \left(\sum_{m \text{ odd} \geq (2k+1)p} p^{\lfloor \frac{m-p}{2p} \rfloor - k} b_m(x)y^{-m}\right) dx
 \end{aligned}$$

with $a_i(x), b_i(x) \in R[x]$ of degree less than d . Note that $b_{(2k+1)p}(x) = 1$ and that $\{\dots\}^n$ when expanded is then reduced to the form $A_1(x)y^{-2} + \dots + A_{pn}(x)y^{-2pn}$ with $A_i(x) \in R[x]$ of degree less than d .

When we multiply each term in the final sum by x^{pi-1} and reduce using the relation $y^2 = Q(x)$, we see that the result is

$$F_{p*}(x^{i-1}(dx/y^{2k+1})) = \sum_{m \text{ odd} \geq m_0} c_m(x)y^{-m}dx$$

where

$$m_0 \geq (2k+1)p - 2\lfloor (pi-1)/d \rfloor \quad (1)$$

and each $c_m(x) \in pR[x]$. Here we have used $b_{(2k+1)p}(x) = 1$ to get $pi-1$ rather than $pi+d-2$. Furthermore,

$$c_m(x) \in p^{\lfloor \frac{m-p}{2p} \rfloor + 1 - k} R[x] \quad \forall m \geq (2k+1)p \quad (2)$$

- (i) When $k = 1$, by (1) with $i = d - 1$, $m_0 \geq p + 2 > 3$.

(ii) First note that for $m < (2k+1)p \leq p^2$, $\log_p(m-2) < 2$. From Lemma 3.3 and (2), we see that it suffices to prove that

$$\left\lceil \frac{m-p}{2p} \right\rceil + 1 - k - \lfloor \log_p(m-2) \rfloor \geq 0 \quad \forall m \text{ odd} \geq (2k+1)p$$

For $k=0$, the inequality with the floor and ceiling brackets removed holds for $m > 2p+1$ by elementary calculus. For $p \leq m \leq 2p+1$, it is clear.

For $k=1$ and $p \geq 5$, the inequality again holds for $m \geq 5p$ by calculus and for $3p \leq m < 5p$ it is clear.

For $k=1$ and $p=3$, the inequality holds for $m \geq 3p^2+1$ by calculus and for $3p \leq m < 3p^2+1$ it is again easy to check directly.

(iii) Consider the $px^{pi-1}p^\alpha b_m(x)y^{-m}$ terms that give contributions to the $\sum_{n \geq 1}$ sum. Expressing $x^{pi-1}b_m(x)$ as $u_r(x)y^{2r} + \dots u_0(x)$ with $\deg(u_i) < d$, we must have $r \geq (m-1)/2$ and the contribution will be expressible in the form $S(dx/y)$ with $\deg(S) = pi-1 + \deg(b_m) - d(m-1)/2$. This last expression must be greater than or equal to $d-1$ for non-trivial reduction under **RednB**. For such m , writing d_m for $\deg(b_m)$, the above and Lemma 3.3 (ii) show that the power of p in the denominator of the **RednB** reduction of the contribution from the index m term is bounded above by

$$\begin{aligned} & \lfloor \log_p(2pi - md + 2d_m) \rfloor - 1 - \lceil (m-p)/2p \rceil & \text{if } d = 2g+1 \\ & \lfloor \log_p(pi - m(d/2) + d_m + 1) \rfloor - 1 - \lceil (m-p)/2p \rceil & \text{if } d = 2g+2 \end{aligned}$$

We have that $d_p = 0$ ($b_p(x) = 1$) and $d_m \leq d-1$ for $m \geq p+2$. Since $m \geq p$ is odd, the above expressions are maximal when $m=p$. (a) and (b) follow easily from this. \square

The bounds in Lemma 3.4 (iii) for denominators in the reduction of $F_{p*}(x^{i-1}(dx/y))$ are sharp. The proof shows that the first term in the power series expansion $px^{pi-1}(dx/y^p)$ is the only one that can contribute to the given maximal power of p and for a general Q it does indeed lead to denominators equal to the bounds.

Thus, as is readily confirmed in practice by computer computations, we reach the following

Conclusion: When $d = 2g+1$ and $p > 2g-1$ or $d = 2g+2$ and $p > g$, the transformation matrix M for F_{p*} w.r.t. basis B_1 for H_-^1 is p -integral. When these equalities for p do not hold however, for a general Q , entries in the lower rows of M have powers of p in the denominator given by the bounds in the last part of Lemma 3.4.

On the other hand, Lemma 3.4 shows that **RednA** applied to $F_{p*}(\omega)$ for $\omega \in B_2$ reduces back to an expression that is always an R -linear combination of the elements of B_2 , so formally leads to a p -integral transformation matrix M .

If B_2 gives a basis for H_-^1 , then this M genuinely represents F_{p*} on that space and B_2 can replace B_1 as the chosen basis for computations.

Even when B_2 doesn't give a basis, this M can still be used. The above shows that the kernel of η and its image in H_-^1 are F_{p*} - and hence also F_* -stable.

The following result will be demonstrated in the next section.

Proposition 3.5.

- (i) η is an isomorphism when $d = 2g + 1$ but has a 1-dimensional kernel and cokernel when $d = 2g + 2$.
- (ii) In the latter case, F_* is the identity on $\ker(\eta)$ and acts as multiplication by q on $H_-^1/\text{Im}(\eta)$.

This justifies the adaptation of Kedlaya's algorithm given in Section 2, which always works with a p -integral M . In summary:

New Algorithm

- $d = 2g + 1$. If $p \geq 2g$ then the algorithm is unchanged. If $p < 2g$ then the algorithm is as before, but use differential basis B_2 instead of B_1 .
- $d = 2g + 2$, $p > g$. Apply the algorithm as for odd d with differential basis B_1 . At the end, remove a factor $t - q$ from the characteristic polynomial of F_* .
- $d = 2g + 2$, $p \leq g$. Formally apply the algorithm as for odd d with pseudo-basis B_2 . At the end, remove a factor $t - 1$ from the characteristic polynomial of F_* .

Efficiency If $N_1 = \lceil (ng/2) + \log_p(2\binom{2g}{g}) \rceil$ ($q = p^n$) and $N = N_1 + \lfloor \log_p(2N_1) \rfloor + 1$, then estimates using Lemma 3.4 and the Weil bound show that it suffices to compute $(1 + (Q^\sigma(x^p) - Q(x)^p)y^{-2p})^{-(2k+1)/2}$ to accuracy p^N in order that M is of sufficient p -adic accuracy to determine $P_C(t)$. Here, $k = 0$ if we use B_1 and $k = 1$ for B_2 . Using $k = 1$ rather than $k = 0$ makes virtually no difference in computational efficiency here, and the reduction of $F_{p*}(\omega)$ back to a linear combination of basis elements is in fact slightly better when using B_2 .

However, $d = 2g + 2$ rather than $2g + 1$ does increase the size of the bases by 1 element, meaning that one extra reduction of a $F_{p*}(\omega)$ has to be performed. Also the $(d - 1) \times (d - 1)$ matrix M , which has to be σ -powered to the n th power, has an extra row and column. This does make a small difference (more so for smaller g), which makes it worth looking for a k -rational root of $Q(x)$ and moving that to ∞ to transform to $d = 2g + 1$. In general, though, no such transformation is possible.

4. PROOF OF PROPOSITION 3.5

Proposition 3.5 of the last section on the η map is proven in the following three lemmas.

Lemma 4.1. *If $d = 2g + 1$ then η is an isomorphism onto H_-^1 .*

If $d = 2g + 2$ then η has a one dimensional kernel generated by $V(dx/y^3) = d(-2S/y)$ where $V = SQ' - 2S'Q$ and $S = x^{g+1} + \dots \in K[x]$ is the unique monic degree $g + 1$ polynomial such that V is of degree $\leq 2g$.

Proof. Using the fact that B_1 is a basis for H_-^1 and the **RednA** formula, we see that an element of the kernel of η corresponds to a differential of the form $V(dx/y^3)$ with $\deg(V) \leq d-2$ and $V = SQ' - 2S'Q$.

If $V = a_r x^r + \dots$ with $r \geq 0$, $a_r \neq 0$, then the leading term of $SQ' - 2S'Q$ is $(d-2r)a_r x^{d+r-1}$, so d must equal $2r$. So, $d = 2g + 2$ and $r = g + 1$. Normalising S so that its leading coefficient is 1, it follows easily that its lower coefficients are completely determined by the condition on $\deg(V)$. Explicitly, if b_i is the coefficient of x^i in S , then the condition that the coefficient of x^{d+i-1} in $SQ' - 2S'Q$ is zero, $0 \leq i \leq g$, translates into

$$(2g + 2 - 2i)b_i = \text{some linear combination of } b_j, j \geq i + 1$$

This determines the b_i inductively and gives a unique S and V up to K -scaling. \square

Lemma 4.2. *When $d = 2g + 2$, F_* acts trivially on $\ker(\eta)$.*

Proof. From the last lemma, $\ker(\eta)$ is 1-dimensional and generated by $\omega = V(dx/y^3) = d(-2S/y)$ with $S = x^{g+1} + \dots$. As $\ker(\eta)$ is F_* -stable, ω is an eigenvector for F_* with eigenvalue λ , say. We must show that $\lambda = 1$.

Considering the images in H_-^1 and using Lemma 3.4 (ii), we get

$$F_*(\omega) = \lambda\omega - 2d(f) \quad f = \sum_{r=1}^{\infty} \frac{B_r(x)}{y^{2r+1}} \in (A_K^\dagger)^- \Rightarrow \quad d\left(F\left(\frac{S}{y}\right)\right) = \lambda d\left(\frac{S}{y}\right) + d\left(\frac{B_1}{y^3} + \frac{B_3}{y^5} \dots\right)$$

So

$$F\left(\frac{S}{y}\right) = \lambda \left(\frac{S}{y}\right) + \left(\frac{B_1}{y^3} + \frac{B_3}{y^5} \dots\right) \in (A_K^\dagger)^- \quad (3)$$

In fact, this equality is true up to addition of a constant in K , but as both sides are in the $-$ eigenspace, the constant must be zero. The B_i here have degree $< d$.

Now, as in the proof of Lemma 3.4, we see that if the standard expansion of $f \in A_K^\dagger$ is of the form $\sum_{n \geq 3} a_n(x)/y^n$, then $F_p(f)$ has the same property.

Also, expanding $S(x^p)$ as $u_r(x)Q(x)^r + \dots + u_0(x) = u_r(x)y^{2r} + \dots + u_0(x)$ with $\deg(u_i) < d$, we easily get that $r = (p-1)/2$ and $u_r(x) = x^{g+1} + \dots$.

Then, using $F_p(1/y) = y^{-p}(1 + a_2(x)/y^2 + a_4(x)/y^4 + \dots)$, we find that $F_p(S/y) = S_1(x)/y + b_3(x)/y^3 + \dots$ with $S_1(x) = x^{g+1} + \dots$. Iterating, we see that the same holds for $F(S/y)$. Then, (3) implies that $\lambda = 1$. \square

Lemma 4.3. *When $d = 2g + 2$, F_* acts on $H_-^1/\text{Im}(\eta)$ as multiplication by q .*

Proof. We already know that $\text{Im}(\eta)$ is an F_* -stable codimension 1 subspace of H_-^1 and that the eigenvalues of F_* on H_-^1 are q and the roots of $P_C(t)$, the numerator of the zeta-function of C . We need to show that the eigenvalues of F_* on $\text{Im}(\eta)$ are precisely these latter roots.

We will prove the lemma by using an isomorphism to an odd degree model over an extension \mathbf{F}_{q^r} of k where $Q \in k[x]$ has a root. In fact, replacing F by F^r corresponds to replacing the basefield $k = \mathbf{F}^q$ by $k_1 = \mathbf{F}^{q^r}$ and the roots of $P_{C/k_1}(t)$ are the r th powers of the roots of $P_{C/k}(t)$.

These latter roots have absolute value $q^{r/2}$ in every complex embedding whereas q^r obviously has absolute value q^r . So we can assume that Q has a root in k .

First note that

$$\mathrm{Im}(\eta) = \{\omega \in H_-^1 \mid \mathrm{Residue}_{\infty_1}(\omega) = \mathrm{Residue}_{\infty_2}(\omega) = 0\}$$

as both sides of the equality have codimension 1 in H_-^1 and the LHS lies in the RHS (in fact, all differentials of the form $x^i(dx/y^3)$, $i \leq d-2$ are holomorphic at both points at infinity).

We can translate a root of $Q(x)$ to zero by a $x \mapsto x - \alpha$ translation (this changes the lift of F but not $\mathrm{Im}(\eta)$), so assume that $Q(x) = x^{2g+2} + a_{2g+1}x^{2g+1} + \dots + a_1x \in k[X]$, $a_1 \neq 0$.

Let $\tilde{Q}(X) = X^{2g+1} + (a_2/a_1^2)X^{2g} + \dots + (1/a_1^{2g+2})$.

The equation $Y^2 = \tilde{Q}(X)$ defines a new smooth, odd-degree affine model for C and we have

$$B_k := \frac{k[X, Y, Y^{-1}]}{(Y^2 - \tilde{Q}(X))} \hookrightarrow A_k = \frac{k[x, y, y^{-1}]}{(y^2 - Q(X))} \quad X \mapsto 1/(a_1x), \quad Y \mapsto y/(a_1x)^{g+1}$$

[note: $1/(a_1x) = (1/(a_1y^2))(a_1 + a_2x + \dots) \in A_k$]. Letting B^\dagger be the smooth lift of B_k corresponding to the lift to $R[X]$ of \tilde{Q} with the coefficient lift compatible with that of Q , we get the corresponding commutative diagram

$$\begin{array}{ccc} B^\dagger & \longrightarrow & A^\dagger \\ F^{(1)} \downarrow & & \downarrow F^{(2)} \\ B^\dagger & \longrightarrow & A^\dagger \end{array}$$

for some choice of q -Frobenius lifts $F^{(1)}$ and $F^{(2)}$. All maps commute with the automorphisms induced by the hyperelliptic involution.

One easily sees that $A_k = B_k[1/X]$. The Main Theorem of [Mon68] implies that

$$H^1(B_k; K) \hookrightarrow H^1(A_k; K) = H^1$$

with image the K -subspace of elements with residues 0 at ∞_1 and ∞_2 , the images of points with $X = 0$ under the automorphism of C induced from $B_k \hookrightarrow A_k$. [In fact, a bit of computation verifies the residue condition directly from the explicit maps].

Thus $\mathrm{Im}(H^1(B_k; K)^-) = \mathrm{Im}(\eta)$ and as we know that the eigenvalues of F_* on $H^1(B_k; K)^-$ are the roots of $P_C(t)$ (the odd degree case), the result follows. \square

REFERENCES

- [BCP97] W. Bosma, J. Cannon, and C. Playoust, *The Magma algebra system I. The user language.*, J. Symbolic Computation **24** (1997), 235–265.
- [Car07] Robert Carls, *Canonical coordinates on the canonical lift*, J. Ramanujan Math. Soc. **22** (2007), no. 1, 1–14.
- [CDV06] W. Castryck, J. Denef, and F. Vercauteren, *Computing Zeta Functions of Nondegenerate Curves*, Int. Math. Res. Papers **12** (2006), 1–57.
- [CL09] Robert Carls and David Lubicz, *A p -adic quasi-quadratic time point counting algorithm*, Int. Math. Res. Not. IMRN (2009), no. 4, 698–735.

- [DV06a] J. Denef and F. Vercauteren, *Computing zeta functions of C_{ab} curves using Monsky-Washnitzer cohomology*, Finite fields and Their Applications **12(1)** (2006), 78–102.
- [DV06b] ———, *An extension of Kedlaya's algorithm to hyperelliptic curves in characteristic 2*, J. Cryptology **19(1)** (2006), 1–25.
- [Edi06] B. Edixhoven, *Point counting after Kedlaya*, EIDMA-Stieltjes Graduate course, Leiden, September 22–26, 2003, online copy at "http://www.math.leidenuniv.nl/~edix/mathofcrypt/carls_edixhoven/kedlaya.pdf", October 2006.
- [Ger07] Ralf Gerkmann, *Relative rigid cohomology and deformation of hypersurfaces*, Int. Math. Res. Pap. IMRP (2007), no. 1, Art. ID rpm003, 67.
- [Har77] Robin Hartshorne, *Algebraic Geometry*, GTM 52, Springer, 1977.
- [Hub08] Hendrik Hubrechts, *Point counting in families of hyperelliptic curves*, Found. Comput. Math. **8** (2008), no. 1, 137–169.
- [Ked01] K. Kedlaya, *Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology*, J. Ramanujan Math. Soc. **16** (2001), 318–330.
- [Ked03] ———, *Comments/errata for "Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology"*, J. Ramanujan Math. Soc. **18** (2003), 417–418.
- [Ked04] ———, *Computing zeta functions via p -adic cohomology*, Algorithmic Number Theory, 6th International Symposium, ANTS-VI, LNCS 3076, Springer, 2004.
- [LL06] Reynald Lercier and David Lubicz, *A quasi quadratic time algorithm for hyperelliptic curve point counting*, Ramanujan J. **12** (2006), no. 3, 399–423.
- [LW08] Alan G. B. Lauder and Daqing Wan, *Counting points on varieties over finite fields of small characteristic*, Algorithmic number theory: lattices, number fields, curves and cryptography, Math. Sci. Res. Inst. Publ., vol. 44, Cambridge Univ. Press, 2008, pp. 579–612.
- [Mon68] P. Monsky, *Formal cohomology: II. the cohomology sequence of a pair*, Ann. Math. **88(2)** (1968), 218–238.
- [Mon71] ———, *Formal cohomology: III. fixed point theorems*, Ann. Math. **93(2)** (1971), 315–343.
- [MW68] P. Monsky and G. Washnitzer, *Formal cohomology: I*, Ann. Math. **88(2)** (1968), 181–217.
- [vdB08] T. van den Bogaart, *About the choice of a basis in Kedlaya's algorithm*, extract from the author's PhD thesis available online at arXiv as arXiv:0809.1243v1[math.AG], September 2008.

SCHOOL OF MATHEMATICS AND STATISTICS F07, UNIVERSITY OF SYDNEY, NSW 2006, AUSTRALIA