

EXPLICIT COLEMAN INTEGRATION IN LARGER CHARACTERISTIC

ALEX J. BEST

ABSTRACT. We describe a more efficient algorithm to compute p -adic Coleman integrals on odd degree hyperelliptic curves for large primes p . The improvements come from using fast linear recurrence techniques when reducing differentials in Monsky-Washnitzer cohomology, a technique introduced by Harvey [Har07] when computing zeta functions. The complexity of our algorithm is quasilinear in \sqrt{p} and is polynomial in the genus and precision. We provide timings comparing our implementation with existing approaches.

1. INTRODUCTION

In 2001, Kedlaya introduced an algorithm for computing the action of Frobenius on the Monsky-Washnitzer cohomology of odd degree hyperelliptic curves over \mathbf{Q}_p [Ked01]. This has been used to compute zeta functions of the reductions modulo p of such curves, and, starting with the work of Balakrishnan-Bradshaw-Kedlaya [BBK10], to evaluate Coleman integrals between points on them. Computation of Coleman integrals requires more information to be retained throughout the execution of the algorithm than is needed to compute only the way Frobenius acts on cohomology classes, which is all that is needed to compute zeta functions.

Harvey [Har07] introduced a variant of Kedlaya's algorithm, its run time in terms of p alone is $\tilde{O}(\sqrt{p}) := O(\sqrt{p} \log^k \sqrt{p})$ for some $k \in \mathbf{Z}$. In [BBK10] the authors asked if it is possible to use Harvey's techniques when computing Coleman integrals.

Here we show that one can obtain the same efficiency improvements in Kedlaya's algorithm as Harvey did, whilst retaining enough information to compute Coleman integrals. Specifically, we obtain the following result:

Theorem 1.1. *Let X/\mathbf{Z}_p be a genus g , odd degree hyperelliptic curve. Then for the basis $\{\omega_i = x^i dx/2y\}_{i=0}^{2g-1}$ of $H_{\text{dR}}^1(X)$, let M be the matrix of Frobenius acting on this basis, and $N \in \mathbf{N}$ be such that X and $P, Q \in X(\mathbf{Q}_p)$ are known to precision p^N , assume $p > (2N - 1)(2g + 1)$. Then, if multiplying two $g \times g$ matrices requires $O(g^\omega)$ ring operations, the vector of Coleman integrals $(\int_P^Q \omega_i)_{i=0}^{2g-1}$ can be computed in time $\tilde{O}(g^\omega \sqrt{p} N^{5/2} + N^4 g^4 \log p)$ to absolute p -adic precision $N - v_p(\det(M - I))$.*

2010 *Mathematics Subject Classification.* Primary 11G20; Secondary 11Y16, 14F30.

Key words and phrases. Coleman integration, hyperelliptic curves, Kedlaya's algorithm.

I would like to thank Jennifer Balakrishnan, for suggesting this as something that might be possible, and for many subsequent helpful conversations and comments. Additional thanks are due to Jan Tuitman for remarking that ramified extensions should be avoided, by using Lemma 3.2. I have had many interesting conversations with Sachi Hashimoto about Coleman integration. Finally I would like to thank the reviewers for their suggestions. I am grateful for support from the Simons Foundation as part of the Simons Collaboration on Arithmetic Geometry, Number Theory, and Computation #550023.

As surveyed in [BBK10] there are many applications of Coleman integration in arithmetic geometry, notably they are central to the method of Chabauty-Coleman-Kim. This method has been made explicit in some cases, such as in [BD18] Example 2. There, and in general, when working over number fields it is useful to work only with p that split. This is an additional condition on p , which often results in having to take larger p , which gives one motivation for the current work.

In §2 and §3 we recall the set-up for Coleman integration, and, most importantly, exactly what data is needed to compute Coleman integrals on hyperelliptic curves. In §4 we examine the reduction procedure used by Harvey in more detail. We then come to our main new ideas, creating an appropriate recurrence that computes the data necessary for Coleman integration. In §5 we introduce a modification of the linear recurrence algorithm used by Harvey, which is specialised to the type of recurrences we obtained. This is useful when computing Coleman integrals between many endpoints simultaneously. In §6 we describe the main algorithm in detail. In §7 and §8 we analyse its correctness and complexity. Finally in §9 and §10 we give some timings and examples obtained with a SageMath/C++ implementation, showing its practical use.

2. SET-UP AND NOTATION

Throughout we work with a fixed prime p and an odd degree hyperelliptic curve X/\mathbf{Z}_p , of genus $g \geq 1$, given as $y^2 = Q(x)$ with $Q(x) \in \mathbf{Z}_p[x]$. Where $Q(x) = x^{2g+1} + P(x)$ with $\deg(P) \leq 2g$. We assume that the reduction of $Q(x)$ to $\mathbf{F}_p[x]$ has no multiple roots. We fix a desired p -adic precision $N \geq 1$ such that

$$p > (2N - 1)(2g + 1). \quad (2.1)$$

Let ι denote the hyperelliptic involution, given on the finite affine chart as $(x, y) \mapsto (x, -y)$; the fixed points of this involution are called *Weierstrass points*.

We will make use of several notions from rigid geometry. Points of $X(\mathbf{Q}_p)$ which reduce to the same point in $X_{\mathbf{F}_p}(\mathbf{F}_p)$ are said to lie in the same *residue disk*. A residue disk that contains a Weierstrass point is a *Weierstrass residue disk*.

3. COLEMAN INTEGRATION

Coleman integration is a p -adic (line) integration theory developed by Robert Coleman in the 1980s [Col82, CdS88, Col85]. Here we briefly summarise the set-up for this theory (for more precise details, see, for example, [Bes12]). We also recall the key inputs, which are obtained from Kedlaya's algorithm, for performing explicit Coleman integration on hyperelliptic curves, as described in [BBK10].

The setting for Coleman integration as we will be using it is via the Monsky-Washnitzer weak completion of the coordinate ring of the curve minus its Weierstrass points. So, letting $A = \mathbf{Z}_p[x, y, y^{-1}]/(y^2 - Q(x))$, its weak completion is the space A^\dagger of series $\sum_{i=-\infty}^{\infty} R_i(x)y^{-i}$ with $R_i \in \mathbf{Z}_p[x]$, $\deg R_i \leq 2g$ subject to the condition that $\liminf_{|i| \rightarrow \infty} v_p(R_i)/|i| > 0$. The p -power Frobenius on $\bar{A} = A/p$ can be lifted to a function $\phi: A^\dagger \rightarrow A^\dagger$ by sending $x \mapsto x^p$ and $y \mapsto y^{-p} \sum_{k=0}^{\infty} \binom{-1/2}{k} (\phi(Q(x)) - Q(x)^p)^k / y^{2pk}$. We will consider differentials in $\Omega_{A^\dagger}^1 = A^\dagger dx \oplus A^\dagger dy/(2y dy - Q'(x) dx)$ with d the exterior derivative

$$d: A^\dagger \rightarrow \Omega_{A^\dagger}^1; \quad \sum_{i=-\infty}^{\infty} \frac{R_i(x)}{y^i} \mapsto \sum_{i=-\infty}^{\infty} R'_i(x)y^{-i} dx - R_i(x)iy^{-i-1} dy. \quad (3.1)$$

We will say that f is a *primitive* of the exact differential df . We then define the Monsky-Washnitzer cohomology of A to be $H_{\text{MW}}^1(\overline{A}) = \Omega_{A^\dagger}^1 \otimes \mathbf{Q}_p / d(A^\dagger \otimes \mathbf{Q}_p)$. The action of Frobenius and of the hyperelliptic involution can be extended to $\Omega_{A^\dagger}^1$ and $H_{\text{MW}}^1(\overline{A})$ and the actions of ϕ and ι commute. In particular we have an eigenspace decomposition of all of these spaces under ι into *even* and *odd* parts; the odd part will be denoted with a $-$ superscript. Let $A_{\text{loc}}(X)$ denote the \mathbf{Q}_p -valued functions on $X(\mathbf{Q}_p)$ which are given by a power series on each residue disk.

Theorem 3.1 (Coleman). *There is a unique (up to a global constant of integration) \mathbf{Q}_p -linear integration map $\int : \Omega_{A^\dagger}^1 \otimes \mathbf{Q}_p \rightarrow A_{\text{loc}}(X)$ satisfying:*

- (1) *Frobenius equivariance, $\int \phi^* \omega = \phi^* \int \omega$,*
- (2) *the fundamental theorem of calculus, $d \circ \int$ is the identity on $\Omega_{A^\dagger}^1 \otimes \mathbf{Q}_p$,*
- (3) *and $\int \circ d$ is the natural map $A^\dagger \rightarrow A_{\text{loc}}(X)$ (constant functions).*

Given points $P, Q \in X(\mathbf{Q}_p)$ the definite integral $\int_P^Q \omega$ is then defined as $(\int \omega)(Q) - (\int \omega)(P)$, which is a well-defined function of P, Q .

After fixing a basis $\{\omega_i\}_{i=0}^{2g-1}$ of $H_{\text{MW}}^1(\overline{A})^- = H_{\text{dR}}^1(X)$, any 1-form of the second kind $\omega \in \Omega_{A^\dagger}^1$ can be expressed as $\omega = df + \sum_{i=0}^{2g-1} a_i \omega_i$, $f \in A^\dagger$, so by Theorem 3.1 we see that for some $a_i \in \mathbf{Q}_p$

$$\int_P^Q \omega = f(Q) - f(P) + \sum_{i=0}^{2g-1} a_i \int_P^Q \omega_i. \quad (3.2)$$

We can therefore reduce to the case of integrating only the basis differentials ω_i and evaluating the primitive f . The complexity of reducing to this case depends on how ω is presented. For example, if ω has many terms, the total run time can be dominated by finding f and evaluating $f(Q) - f(P)$ in the above. So we will focus on computing $\left\{ \int_P^Q \omega_i \right\}_{i=0}^{2g-1}$. In many applications, all that we need to integrate are \mathbf{Q}_p -linear combinations of the basis differentials.

The work of Balakrishnan-Bradshaw-Kedlaya [BBK10] describes how to explicitly compute Coleman integrals for differentials on odd degree hyperelliptic curves. They describe how to reduce the problem of computing general Coleman integrals between two points to that of finding a matrix M and $f_i \in A^\dagger$ such that

$$\phi^* \omega_i = df_i + \sum_j M_{ij} \omega_j \in \Omega_{A^\dagger}^1. \quad (3.3)$$

Before stating a form of their algorithm, we recall a useful result which allows us to deal with the difficulties arising when the endpoints of the integral are Weierstrass. This can be problematic, as we need to evaluate primitives as in (3.2); if the endpoints are in Weierstrass residue disks, these power series may not converge.

Lemma 3.2 ([BBK10] Lemma 16). *Let $P, Q \in X(\mathbf{Q}_p)$ with Q Weierstrass and let $\omega \in \Omega_{A^\dagger}^{1,-}$ be an odd differential without poles at P, Q . Then $\int_P^Q \omega = \frac{1}{2} \int_P^{\iota(P)} \omega$.*

In particular, if P is also a Weierstrass point, then the integral is zero.

Lemma 3.2 allows us to express general integrals as linear combinations of integrals between two points in non-Weierstrass residue disks and integrals between two

points in the same residue disk (known as *tiny integrals*). Evaluating tiny integrals uses formal integration of power series, see [BBK10] Algorithm 8.

Note that ∞ is a Weierstrass point so Lemma 3.2 applies with $Q = \infty$; integrals based at ∞ can be rewritten as a linear combination of a tiny integral and an integral between two non-Weierstrass points. Specifically, for a Teichmüller point P , if we know the matrix M expressing the action of Frobenius on the basis differentials ω_i , we can use the Frobenius equivariance of the Coleman integral to deduce

$$\begin{pmatrix} \int_P^\infty \omega_i \\ \vdots \\ \int_P^\infty \omega_i \\ \vdots \end{pmatrix} = \frac{1}{2} \begin{pmatrix} \int_{P^{(P)}} \omega_i \\ \vdots \\ \int_{P^{(P)}} \omega_i \\ \vdots \end{pmatrix} = \frac{(M-I)^{-1}}{2} \begin{pmatrix} f_i(P) - f_i(\iota(P)) \\ \vdots \\ f_i(P) - f_i(\iota(P)) \\ \vdots \end{pmatrix} = (M-I)^{-1} \begin{pmatrix} \vdots \\ f_i(P) \\ \vdots \end{pmatrix} \quad (3.4)$$

The last equality holds as we are using odd differentials, so the df_i must also be odd, so from the expansion of (3.1) we see that the f_i must also be odd (up to the constant term, which cancels).

So we will fix ∞ as our basepoint and compute only integrals of the form $\int_P^\infty \omega$; general integrals can be obtained by subtracting two of the above type. We will use the following algorithm, c.f. [BBK10] Remark 15:

Algorithm 3.3. *Input:* $P \in X(\mathbf{Q}_p)$, the matrix of Frobenius M , and if P is not in a Weierstrass residue disk, $\{f_i(P')\}_{i=0}^{2g-1}$ for the unique Teichmüller point P' in the same residue disk as P , and f_i as in (3.3).

Output: $\{\int_P^\infty \omega_i\}$ for $0 \leq i \leq 2g-1$.

- (1) *If P is in a Weierstrass residue disk:* Let P' be the Weierstrass point in the same residue disk, so that $\int_{P'}^\infty \omega_i = 0$ for all i .
Else: Let P' be the (unique) Teichmüller point in the same residue disk as P . Then compute the vector of $\int_{P'}^\infty \omega_i$ using (3.4).
- (2) For each i , compute the tiny integral $\int_P^{P'} \omega_i$, as in [BBK10] Algorithm 8.
- (3) For each i , sum the result of Steps 1 and 2 to get $\int_P^\infty \omega_i = \int_P^{P'} \omega_i + \int_{P'}^\infty \omega_i$.

Variants of this algorithm are possible, c.f. [BBK10] Algorithm 11. From the version stated above, it is clear that, beyond solving a linear system and computing tiny integrals, the matrix of Frobenius and evaluations of the primitives f_i at Teichmüller points in non-Weierstrass residue disks are all the input data that is needed to compute arbitrary Coleman integrals. We shall refer to this data as the *Coleman data*. To compute Coleman integrals efficiently, we require an efficient way of computing this data, possibly for several disks of interest.

Remark 3.4. We do not need to compute the f_i themselves to compute integrals, only evaluations at Teichmüller points in prescribed non-Weierstrass residue disks. This simplification is key to our ability to write down a suitable recurrence. Moreover, once the Coleman data is computed, it can be saved and will not need to be recomputed if integrals between other points in the same residue disks are required.

4. REDUCTIONS IN COHOMOLOGY

4.1. Kedlaya's algorithm and Harvey's work. Kedlaya's algorithm computes the action of Frobenius on Monsky-Washnitzer cohomology up to a specified precision. The general strategy is to begin with a finite p -adic approximation of $\phi^*\omega$ as a (Laurent) polynomial in x and y multiplied by the differential $dx/2y$. This is

reduced step-by-step via cohomologous differentials of lower polynomial degree, by subtracting appropriate exact forms dg for polynomials g . This process is continued until one is left with a \mathbf{Q}_p -linear combination of basis elements, and we have an expression of the form (3.3). For a given basis $\{\omega_i\}$ of $H_{\text{MW}}^1(\overline{A})^-$, writing each $\phi^*\omega_i$ in terms of this basis results in a matrix of Frobenius acting on $H_{\text{MW}}^1(\overline{A})^-$.

The innovation in [Har07] is to express the reduction process as a linear recurrence, where the coefficients are linear polynomials in the index of the recurrence. A term several steps later in such recurrences can then be found more efficiently than the straightforward sequential approach, via the algorithm of Bostan-Gaudry-Schost [BGS07] Theorem 15. Here we also ultimately appeal to these methods, and so we must examine in more detail the polynomials g used in the reduction steps. We will describe the sum of the evaluations of these g at points of interest as a linear recurrence, so that they may be computed along with the reductions.

We use the basis of $H_{\text{MW}}^1(\overline{A})^-$ consisting of $\omega_i = x^i dx/2y$ for $0 \leq i \leq 2g - 1$. This differs by a factor of 2 from the basis used by Harvey and Kedlaya; this choice reduces the number of 2's appearing in our formulae and so appears more natural here. Changing the basis by a scalar multiple has no effect on the matrix of Frobenius, only the exact differentials. An approximation to $\phi^*\omega_i$ is given in [Har07] (4.1) by letting $C_{j,r}$ be the coefficient of x^r in $Q(x)^j$ and $B_{j,r} = p\phi(C_{j,r}) \sum_{k=j}^{N-1} (-1)^{k+j} \binom{-1/2}{k} \binom{k}{j} \in \mathbf{Z}_p$ so that

$$\phi^*\omega_i \equiv \sum_{j=0}^{N-1} \sum_{r=0}^{(2g+1)j} B_{j,r} x^{p(i+r+1)-1} y^{-p(2j+1)+1} \frac{dx}{2y} \pmod{p^N}. \quad (4.1)$$

In (4.1), there are only $(2g+1) \frac{N(N-1)}{2} + N$ terms in total and the exponents of x and y that appear are always congruent to -1 or $1 \pmod{p}$ respectively.

As in [Har07] Section 5, we work with finite-dimensional vector spaces over \mathbf{Q}_p

$$W_{s,t} = \left\{ f(x) x^s y^{-2t} \frac{dx}{2y} : \deg f \leq 2g \right\} = \left\langle x^i x^s y^{-2t} \frac{dx}{2y} \right\rangle_{i=0}^{2g}, \quad (4.2)$$

for $s \geq -1$, $t \geq 0$, where, in addition, we restrict $W_{-1,t}$ to be the subspace of the above for which the coefficient of x^{-1} is zero (i.e. for which $f(0) = 0$).

Notice that $W_{-1,0}$ is naturally identified with $H_{\text{MW}}^1(\overline{A})^-$ with the basis chosen above, so that ω_i is the i th basis element of $W_{-1,0}$. In order to derive an expression for $\phi^*\omega_i$ as a linear combination of the other basis elements, we begin with the approximation of $\phi^*\omega_i$ from (4.1). Then starting with the terms of highest degree in x , which are each inside of some $W_{s,t}$ we reduce ‘‘horizontally’’, finding a cohomologous element of $W_{s-1,t}$ by subtracting an appropriate exact differential. This process is repeated until $s = -1$, but whenever we reach a space $W_{s,t}$ containing a term from (4.1), we add it to the current differential under consideration. We do this for each t appearing as an exponent for a monomial in the original approximation, and for each such t we obtain an element of $W_{-1,t}$. We then reduce ‘‘vertically’’, beginning with the largest t we have, we subtract appropriate exact differentials to reduce the element of each $W_{-1,t}$ to a cohomologous one in $W_{-1,t-1}$ while $t \geq 1$. This is continued until we have reduced everything to the space $W_{-1,0}$, and we

have obtained a linear combination of the basis differentials that is cohomologous to $\phi^*\omega_i$ up to the specified precision.

Note that many horizontal *rows* will not be considered at all. When p is large enough, most steps simply involve reducing terms we already have, as there are comparatively few terms in the (4.1) compared to the total degree. Doing multiple reduction steps quickly will therefore improve the run time of this procedure, even though we have to add new terms occasionally. This is where Harvey applies linear recurrence techniques to speed up this reduction process. We now state the reductions we will use; compared to [Har07] (5.2) and (5.3) we must be more explicit about the exact form we are subtracting, as this data is important for us.

4.2. Horizontal reduction. To reduce horizontally from $W_{s,t}$ to $W_{s-1,t}$, we express the highest order basis element $x^{2g}x^s y^{-2t} dx/2y \in W_{s,t}$ as a cohomologous term in $W_{s-1,t}$. The other basis elements are naturally basis elements for $W_{s-1,t}$ just with their indices shifted by 1.

Lemma 4.1 (Horizontal reduction). *We have*

$$\begin{aligned} & x^{2g}x^s y^{-2t} \frac{dx}{2y} - \frac{-1}{(2t-1)(2g+1)-2s} d(x^s y^{-2t+1}) \\ &= \frac{2sP(x) - (2t-1)xP'(x)}{(2t-1)(2g+1)-2s} x^{s-1} y^{-2t} \frac{dx}{2y} \in W_{s-1,t}. \end{aligned} \quad (4.3)$$

Proof. We directly compute

$$\begin{aligned} d(x^s y^{-2t+1}) &= s x^{s-1} y^{-2t+1} dx + (-2t+1) x^s y^{-2t} dy \\ &= \left(s x^{s-1} y^{-2t+1} + \frac{1}{2} (-2t+1) x^s y^{-2t-1} Q'(x) \right) dx \\ &= (2sQ(x) - (2t-1)xQ'(x)) x^{s-1} y^{-2t} \frac{dx}{2y} \\ &= (2s - (2t-1)(2g+1)) x^{2g+1} x^{s-1} y^{-2t} \frac{dx}{2y} \\ &\quad + (2sP(x) - (2t-1)xP'(x)) x^{s-1} y^{-2t} \frac{dx}{2y}. \end{aligned} \quad (4.4)$$

Therefore, by subtracting $\frac{1}{2s-(2t-1)(2g+1)} d(x^s y^{-2t+1})$ from $x^{2g}x^s y^{-2t} dx/2y$, the remaining terms are all as stated, and of lower degree. \square

4.3. Vertical reduction. To reduce vertically from $W_{-1,t}$ to $W_{-1,t-1}$, we express the $2g$ basis elements $x^i y^{-2t} dx/2y \in W_{-1,t}$ as cohomologous terms in $W_{-1,t-1}$.

Lemma 4.2 (Vertical reduction). *Let $R_i(x), S_i(x) \in \mathbf{Z}_p(x)$ be such that $x^i = R_i(x)Q(x) + S_i(x)Q'(x)$ with $\deg R_i \leq 2g-1$, $\deg S_i \leq 2g$. Then*

$$x^i y^{-2t} \frac{dx}{2y} - \frac{-1}{2t-1} d(S_i(x) y^{-2t+1}) = \frac{(2t-1)R_i(x) + 2S_i'(x)}{2t-1} y^{-2(t-1)} \frac{dx}{2y} \in W_{-1,t-1}.$$

Proof. We have that

$$x^i y^{-2t} \frac{dx}{2y} = (R_i(x)Q(x) + S_i(x)Q'(x)) y^{-2t} \frac{dx}{2y} = R_i(x) y^{-2t+2} \frac{dx}{2y} + S_i(x) y^{-2t} dy,$$

and also that $d(S_i(x)y^{-2t+1}) = S'_i(x)y^{-2t+1} dx + (-2t+1)S_i(x)y^{-2t} dy$. Therefore by subtracting $\frac{1}{-2t+1} d(S_i(x)y^{-2t+1})$ from $x^i y^{-2t} dx/2y$, we see that

$$\begin{aligned} x^i y^{-2t} \frac{dx}{2y} &\sim R_i(x)y^{-2t+2} \frac{dx}{2y} + \frac{1}{2t-1} S'_i(x)y^{-2t+1} dx \\ &= \frac{(2t-1)R_i(x) + 2S'_i(x)}{2t-1} y^{-2(t-1)} \frac{dx}{2y}. \end{aligned} \quad (4.5)$$

□

4.4. Towards a faster algorithm. In order to make use of the same linear recurrence techniques as Harvey, we express the reduction process as we descend through the indices s, t as a linear recurrence with coefficients linear polynomials in s, t . We describe such a recurrence that retains enough information to compute Coleman integrals. By working with a number of evaluations of the primitives on prescribed points on the curve, rather than the primitives themselves as power series, we only have to deal with a vector of fixed size at each step. This is preferable to maintaining a power series as we reduce, adding terms at each step.

We will now give an idea of the approach, giving the details in the next section. Let us first consider the end result of one row of the horizontal reduction process. Fixing a row t , after the reduction we have an equality of the form

$$\sum_{s \geq 0} a_s x^s y^{-2t} \frac{dx}{2y} - d \left(\sum_{s \geq 0} c_s x^s y^{-2t+1} \right) = \sum_{i=0}^{2g-1} m_i x^i y^{-2t} \frac{dx}{2y} \in W_{-1,t} \quad (4.6)$$

in which the terms of the exact differential were found in decreasing order as the reductions are performed. Unfortunately, adding each new term as it is obtained is not a *linear* recurrence in the index s , as we have s appearing in the exponent of x in each term. Instead we observe that we can express the exact differential as

$$d((c_0 + x(c_1 + x(\cdots + x(c_r))))y^{-2t+1}). \quad (4.7)$$

In essence, we are applying the *Horner scheme* for polynomial evaluation.

Now we specialise to the case of computing the evaluation $f_i(P)$ of the primitive for some point $P = (x(P), y(P))$. We can, at each step, compute a further bracketed term starting from the innermost; using the given x, y values, we get a recurrence whose final term is the same as the original evaluation. So we can compute the terms of a recurrence of the form

$$f_{i,0} = 0, f_{i,n} = x(P)f_{i,n-1} - \frac{1}{(2t-1)(2g+1) - 2s} d_{i,n} \quad (4.8)$$

where $s = s_{\max} - n$ decreases from its maximum value, and $d_{i,n}$ is the coefficient of the monomial removed in the n th step of the reduction process. Multiplying the result of this recurrence by the factor y^{-2t+1} (which is constant along the row) will result in the evaluation of the primitive for the row. At each step we will no longer have an evaluation of the primitive so far, it is only after completing all the reduction steps that each term will have the correct power of x .

We may use the same technique for the vertical reductions; here we have

$$\sum_{t \geq 0} \sum_{i=0}^{2g-1} m_i x^i y^{-2t} \frac{dx}{2y} - d \left(\sum_{t \geq 1} \sum_{i=0}^{2g} d_{ti} S_i(x) y^{-2t+1} \right) = \sum_{i=0}^{2g-1} M_i x^i \frac{dx}{2y} \in W_{-1,0},$$

where now writing $d_t = \sum_{i=0}^{2g-1} d_{t,i} S_i(x)$, the exact differential can be expressed as

$$d(y^{-1}(d_1 + y^{-2}(d_2 + y^{-2}(\cdots (d_{r-1} + y^{-2}(d_r)) \cdots))).) \quad (4.9)$$

Remark 4.3. The factor y^{-2t+1} appears in every term in the primitive in row t . It is the same factor in the primitive for the vertical reduction from row t to row 0. So we can initialise the vertical recurrence from $W_{-1,t}$ with both the differential and the evaluations obtained from horizontal reduction along row t , and let the vertical reduction steps multiply the evaluation of the row primitives by this factor.

Now we write down the recurrences for both horizontal and vertical reductions precisely using matrices acting on appropriate vector spaces.

4.5. The recurrence. We will now switch to working with a \mathbf{Q}_p vector $h^t(s) \in W_{s,t} \times \mathbf{Q}_p^L$ (resp. $v(t) \in W_{-1,t} \times \mathbf{Q}_p^L$); these are of length $2g+1+L$ (resp. $2g+L$) in the horizontal case (resp. vertical case). The first entries represent the current differential we have reduced to, with respect to the basis given in (4.2). The last L entries will contain the evaluations of the terms of the primitive picked up so far, one for each of the L points $P_1, \dots, P_L \in X(\mathbf{Q}_p)$ we want evaluations at.

When we horizontally reduce, using the result of Lemma 4.1, the two terms we are interested in, the exact differential and the reduction, have a common denominator of $D_H^t(s) = (2t-1)(2g+1) - 2s$. Similarly, in the vertical case, the two terms of interest in Lemma 4.2 have a common denominator of $D_V(t) = 2t-1$.

Writing out the result of a single reduction step in terms of these vectors, we see that we need to compute the terms of the recurrence given by $h^t(s) = R_H^t(s+1)h^t(s+1)$ in the horizontal case, for $R_H^t(s)$ defined by

$$D_H^t(s)R_H^t(s) = M_H^t(s) = \left(\begin{array}{cccc|cccc} 0 & \cdots & 0 & p_0^t & & & & \\ D_H^t(s) & \cdots & 0 & p_1^t & & & & \\ \vdots & \ddots & \vdots & \vdots & & & & \\ 0 & \cdots & D_H^t(s) & p_{2g}^t & & & & \\ \hline 0 & \cdots & 0 & -1 & x(P_1)D_H^t(s) & \cdots & 0 & \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \\ 0 & \cdots & 0 & -1 & 0 & \cdots & x(P_L)D_H^t(s) & \end{array} \right), \quad (4.10)$$

where p_i^t is the linear function of s obtained as the coefficient of x^i in $2sP(x) - (2t-1)xP'(x)$. To divide through by $D_H^t(s)$ we must multiply some terms by $D_H^t(s)$.

For the vertical reductions we use $R_V(t)$ defined by

$$D_V(t)R_V(t) = M_V(t) = \left(\begin{array}{ccc|cccc} (2t-1)r_{0,0} + 2s'_{0,0} & \cdots & (2t-1)r_{2g-1,0} + 2s'_{2g-1,0} & & & \\ \vdots & \ddots & \vdots & & & \\ (2t-1)r_{0,2g-1} + 2s'_{0,2g-1} & \cdots & (2t-1)r_{2g-1,2g-1} + 2s'_{2g-1,2g-1} & & & \\ \hline -S_0(x(P_1)) & \cdots & -S_{2g-1}(x(P_1)) & y(P_1)^{-2}D_V(t) & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ -S_0(x(P_L)) & \cdots & -S_{2g-1}(x(P_L)) & 0 & \cdots & y(P_L)^{-2}D_V(t) \end{array} \right), \quad (4.11)$$

where $r_{i,j}$ is the coefficient of x^j in $R_i(x)$ and $s'_{i,j}$ is the coefficient of x^j in $S'_i(x)$. Once again we have multiplied the rightmost block by $D_V(t)$ to extract the common denominator. We do this to express the reduction steps as linear recurrences with linear polynomial coefficients, rather than rational function coefficients.

Introducing the notation $M_H^t(a,b) = M_H^t(a+1) \cdots M_H^t(b-1)M_H^t(b)$ (and the analogous $M_V(a,b)$), we can write the upshot of the above as

Theorem 4.4. *Let $h^t(s) = (\omega, 0) \in W_{s,t} \times \mathbf{Q}_p^L$, and $f: X(\mathbf{Q}_p) \rightarrow \mathbf{Q}_p$, write $c(f)$ for the correction factor, the linear endomorphism of $W_{s,t} \times \mathbf{Q}_p^L$ that is the identity on $W_{s,t}$ and scales each component of \mathbf{Q}_p^L by $f(P_\ell)$ for the corresponding P_ℓ . Then the reduced vector $c(y^{-1})R_V(0,t)c(y^2)R_H^t(-1,s)h^t(s) \in W_{-1,0} \times \mathbf{Q}_p^L$ is such that the projection onto $W_{-1,0}$ is some $\tilde{\omega}$ with $\tilde{\omega} = \omega - d(g)$ for some $g \in A^\dagger$, and the projection onto \mathbf{Q}_p^L is $(g(P_1), \dots, g(P_L))$.*

As the approximation in (4.1) has summands that occur in several different $W_{s,t}$'s, we cannot simply find the product matrix and apply it to a single vector. Instead, we must work through the various subspaces doing as many reductions as possible before we reach a new monomial from the original approximation. As D_H and D_V are scalar matrices, we can commute them past the M_V 's and M_H 's. This separates out the components so we can work just with products of matrices of linear polynomials. This reduces the problem to finding several of the products $M_V(a,b)$ and $M_H^t(a,b)$. In practice, to use as little p -adic precision as we can, we must not allow too many runs of multiplications by p and then divisions by p , so that the relative precision stays as large as possible. This will be addressed in §7.

5. LINEAR RECURRENCE ALGORITHMS

In this section, we recall and adapt some methods for finding subsequent terms of linear recurrences with linear polynomial coefficients. The set-up is that we are given an $m \times m$ matrix $M(x)$ with entries that are linear polynomials over a ring R and wish to obtain several products $M(x,y) = M(y)M(y-1) \cdots M(x+1)$ for $x < y$ integers. We let $\text{MM}(m,n)$ be the number of ring operations used when multiplying an $n \times m$ matrix by an $m \times m$ matrix, both with entries in R . Then $\text{MM}(m) = \text{MM}(m,m)$ is the cost of multiplying two $m \times m$ matrices. We will not say much about these functions here, as modern theoretical bounds for these functions do not affect the point of our main result; however see [LG12] for some recent work on the topic. Using naive matrix multiplication, we have $\text{MM}(m,n) = O(m^2n)$, which if $m^2 = o(n)$, cannot be improved upon asymptotically. Whenever $n \geq m$ we can partition an $n \times m$ matrix into roughly n/m blocks each of size $m \times m$. These blocks can then be multiplied individually for a run time of $\text{MM}(m,n) = O(\text{MM}(m) \frac{n}{m})$. We will also let $M(n)$ be the number of ring operations needed to multiply two polynomials of degree n with coefficients in R .

The method of Bostan-Gaudry-Schost requires that certain elements of R be invertible. Moreover, they assume as input a product $D(\alpha, \beta, k)$ of several of these inverses. We will apply these methods in $\mathbf{Z}/p^N\mathbf{Z}$ where the cost of computing inverses is negligible compared to the rest of the algorithm, so we will take this step for granted; see [BGS07] for more details.

With the above set-up, Harvey, [Har07] Theorem 6.2, adjusts the algorithm of Bostan-Gaudry-Schost, [BGS07] Theorem 15, to prove the following theorem:

Theorem 5.1. *Let $M(x)$ be a $m \times m$ matrix with entries that are linear polynomials in $R[x]$, let $0 \leq K_1 < L_1 \leq K_2 < L_2 \leq \cdots \leq K_r < L_r \leq K$ be integers, and let $s = \lfloor \log_4 K \rfloor$. Suppose that $2, 3, \dots, 2^s + 1$ are invertible in R . Suppose also that $r < K^{\frac{1}{2}-\epsilon}$, with $0 < \epsilon < 1/2$. Then $M(K_1, L_1), \dots, M(K_r, L_r)$ can be computed using $O(\text{MM}(m)\sqrt{K} + m^2M(\sqrt{K}))$ ring operations in R .*

In order to apply this theorem to the above recurrences for computing the Coleman data, we introduce a variant better suited to the recurrences we obtained in

§4.4. If we simply applied the same algorithm/result as Harvey naively, we would not get as good a run time in general.

Theorem 5.2. *With the same set-up as Theorem 5.1, except that now let $M(x)$ be instead an $(m+n) \times (m+n)$ block lower triangular matrix with 4 blocks, with top left block an $m \times m$ matrix and bottom right block a diagonal matrix:*

$$\left(\begin{array}{c|ccc} A & & & 0 \\ \hline B & d_1 & & \\ & & \ddots & \\ & & & d_n \end{array} \right). \quad (5.1)$$

Then the interval products $M(K_1, L_1), \dots, M(K_r, L_r)$ can be computed using only $O\left((\text{MM}(m) + \text{MM}(m, n))\sqrt{K} + (m^2 + mn)\text{M}(\sqrt{K})\right)$ ring operations in R .

Proof. The algorithm to do this is the same as the one given for Theorem 5.1 in [Har07] Theorem 6.2, only adjusted to take advantage of the fact that the matrices used are of a more restricted form as follows:

First, note that a product of matrices of the assumed form is again of the same shape, so one can work only with matrices of this form throughout. Such matrices should then be stored without keeping track of the entries that are always 0, as a pair of matrices A, B of size $m \times m$ and $n \times m$ respectively, and a list containing the n bottom right diagonal entries. Now the algorithm of Harvey and Bostan-Gaudry-Schost should be applied using this fixed representation.

The complexity of this algorithm is dominated by two main subtasks: shifting evaluations of the matrices and matrix multiplication. During the shifting step, we need only interpolate the non-zero entries; there are $(m+n)m+n$ of these. The number of ring operations required for this is then $O((m^2 + mn)\text{M}(\sqrt{K}))$.

For the matrix multiplication steps, the restricted form of the matrix once again allows us to use a specialised matrix multiplication routine. Here we can evaluate the block matrix product more efficiently, multiplying only the non-zero blocks, and using the fact that multiplying an $n \times m$ matrix on the right by a square diagonal matrix stored as a list uses only $O(nm)$ operations. Therefore the total complexity of multiplying two matrices of this form is $O(\text{MM}(m, n) + \text{MM}(m))$. As we do not modify the algorithm in any other way, the result follows. \square

The conditions on the matrix in Theorem 5.2 are precisely those satisfied by the matrices $M_H^t(s)$ and $M_V(t)$ from §4. So we may use this algorithm for computing block horizontal and vertical reductions for certain intervals.

Remark 5.3. As well as utilising the polynomial structure of our matrices, for any row with sufficiently many terms compared to the desired precision, it is also possible to interpolate p -adically. This idea is due to Kedlaya and is explained in [Har07] Section 7.2.1. Using this allows us to compute fewer interval products using Theorem 5.2 by interpolating the remaining ones.

If we could compute to infinite precision, it would be optimal to reduce as far as possible at each reduction step, i.e., until we get to index of a new term that needs adding. However, in practice, we should divide by p as soon as possible, in order to reduce the number of extra p -adic digits needed throughout. Therefore analysing when divisions by p occur informs which interval products are found.

6. THE ALGORITHM

In this section we describe the complete algorithm derived in the previous sections. The flow of the algorithm is the same as that of Harvey, only we use our larger matrices throughout and have to make some small adjustments to the evaluations. Care should be taken in all steps where division occurs, see §7.

Algorithm 6.1 (Computation of Coleman data). *Input:* A list of points $\{P_\ell\}_{1 \leq \ell \leq L}$ in non-Weierstrass residue disks, precision N .

Output: Matrix of Frobenius M , modulo p^N , such that $\omega_i = df_i + \sum_j M_{ij}\omega_j$, evaluations $f_i(P_\ell)$ modulo p^N for all i, ℓ also.

- (1) For each row index $t = (p(2j+1) - 1)/2$ for $0 \leq j \leq N-1$ do:
 - (a) Compute the horizontal reduction matrices $M_H^t((k-1)p, kp - 2g - 2)$ and $D_H^t((k-1)p, kp - 2g - 2)$ for $0 \leq k \leq (2g+1)(j+1) - 1$ using Theorem 5.2, and the p -adic interpolation outlined in [Har07] 7.2.1, for $k > N$.
 - (b) For each basis differential ω_i , $0 \leq i \leq 2g-1$ do:
 - (i) Initialise a vector $h_{ij} \in (\mathbf{Z}/p^{N+1}\mathbf{Z})^{2g+1+L}$
 - (ii) For each column index $s = p(i+r+1) - 1$ for $r = (2g+1)j$ down to 0 do:
 - (A) Add the $x^s y^{-2t}$ term of (4.1) to h_{ij} .
 - (B) Set $h_{ij} = R_H^t(kp - 2g - 2, kp)h_{ij}$ by doing $2g+2$ matrix-vector products.
 - (C) Set $h_{ij} = R_H^t((k-1)p, kp - 2g - 2)h_{ij}$.
 - (D) Set $h_{ij} = R_H^t((k-1)p)h_{ij}$.
- (2) Initialise a $2g \times L$ matrix for the evaluations E and a $2g \times 2g$ matrix for the action of Frobenius M .
- (3) Compute the vertical reduction matrices $M_V(0, (p-1)/2)$, $M_V((p-1)/2 + jp, (p-1)/2 + (j+1)p)$ for $1 \leq j < N$ and the corresponding $D_V(t)$'s to precision p^{N+1} using Theorem 5.2, and divide through to obtain the corresponding R_V 's, label them R_j .
- (4) For each basis differential ω_i , $0 \leq i \leq 2g-1$:
 - (a) Initialise a zero vector $v_i \in (\mathbf{Z}/p^N\mathbf{Z})^{2g+L}$.
 - (b) For each row index $t = (p(2j+1) - 1)/2$ for $j = N-1$ down to 0 do:
 - (i) Add the last $2g+L$ entries of h_{ij} to v_i , correcting the last L entries as in Theorem 4.4.
 - (ii) Set $v_i = R_j v_i$.
 - (c) Set the i th column of M to be the first $2g$ entries of v_i .
 - (d) Set the i th row of E to be the last L entries of v_i , correcting them to be evaluations as in Theorem 4.4.
- (5) Output the matrix of Frobenius M and the matrix of evaluations E .

Remark 6.2. We have not used the fact that in Algorithm 3.3 we only needed to evaluate at Teichmüller points. Using Teichmüller points only serves to make the description of Coleman integration a little simpler, and provides a convenient set of points corresponding to residue disks. This allows one to store the output of Algorithm 6.1 for further computations involving the same set of residue disks.

One simpler variant of this algorithm is to compute evaluations for one point at a time, re-running the whole procedure including finding the matrix of Frobenius once for each point. The advantage of this method is not needing a specialised version of

the linear recurrence algorithms as in Theorem 5.2. While this would result in the same theoretical run time if $g^2 \in o(p)$, recomputing the matrix of Frobenius would be a duplication of work and inefficient in many parameter ranges.

7. PRECISION

In this section we examine the level of p -adic precision that needs to be maintained throughout, in order to compute the matrix of Frobenius and evaluations of primitives to precision $O(p^N)$. We follow Harvey's approach in [Har07] Section 7 and prove that analogous results hold for our recurrence.

Lemma 7.1. *During horizontal reduction, the evaluations of the primitives remain integral. Moreover, if the calculations are performed with initial data known to absolute precision p^N and intermediate computations are performed with an absolute precision cap of p^{N+1} , then whenever division by p occurs, the dividend is known to absolute precision p^{N+1} , so that the quotient is known to absolute precision $O(p^N)$.*

Proof. As we begin with evaluation 0, we must show that if the evaluations are integral, they remain so after several reduction steps. Any point $P = (x, y)$ that we are evaluating at is assumed not be in a Weierstrass residue disk and in particular not in the residue disk at infinity. Hence x is integral and multiplication by it will never reduce p -adic valuation.

In the horizontal reduction matrix (4.10), the only nonzero terms in the bottom left block are the -1 s in the rightmost column which will not disturb integrality.

When $D_H^t(s) \equiv 0 \pmod{p}$, it is shown in [Har07] Claim 7.3, using the assumptions on p in (2.1), that the vector currently being reduced has its $(2g+1)$ -component divisible by p and is correct to absolute precision p^{N+1} . Thus this can be divided by $D_H^t(s)$ while keeping absolute precision p^N . Every column of $M_H^t(s)$ other than the $(2g+1)$ st has $D_H^t(s)$ as a factor, so the division can be performed.

All other steps follow directly from the work of Harvey. \square

Lemma 7.2. *During vertical reduction, the evaluations of the primitives remain integral. Moreover, if the calculations are performed with initial data known to absolute precision p^N and intermediate computations are performed with an absolute precision cap of p^{N+1} , then whenever division by p occurs, the dividend is known to absolute precision p^{N+1} , so that the quotient is known to absolute precision $O(p^N)$.*

Proof. Any point $P = (x, y)$ that we are evaluating at is assumed not to be in a Weierstrass residue disk and in particular not in the residue disk at infinity. Hence y is a unit and multiplying or dividing by it will not change p -adic valuation.

We check that the analysis in [Har07] Lemmas 7.7 and 7.9 may be adjusted to apply with our extended $M_V(t)$. Assume that $t \equiv 1/2 \pmod{p}$ so that $D_V(t) \equiv 0 \pmod{p}$, in this case $v_p(D_V(t)) = 1$ as (2.1) implies $D_V(t) < p^2$. Unlike in [Har07] Lemma 7.7, our matrix $M_V(t)$ will not have integral inverse as $D_V(t)$ appears in the bottom right block, so $M_V(t)$ is singular mod p . Instead, the inverse of the block lower triangular $M_V(t)$ has integral top left block, and the bottom two blocks have valuation at least -1 . Now letting $t_0 = (p-1)/2$ and $X = D_V(t_0, t_0 + p + 1)^{-1} M_V(t_0, t_0 + p + 1)$, the argument in [Ked01] Lemma 2 implies that pX is integral. The argument says that taking $\omega \in W_{-1, t_0 + p + 1}$ with integral coefficients, the primitive g of $X\omega - \omega$ becomes integral after multiplication by p , and hence the evaluation of pg at a point in a non-Weierstrass residue disk is

integral. The entries in the bottom left block of X are evaluations of this form up to a power of $y(P)$, which will not affect integrality. The bottom right block of X is integral already as it is simply a power of the diagonal matrix $\text{diag}((y(P_\ell)^{-2})_\ell)$. So each term of the block matrix product $(pX)M_V(t_0 + p + 1)$ is integral, and $M_V(t_0, t_0 + p) = D_V(t_0, t_0 + p + 1)X M_V(t_0 + p + 1)^{-1}$ is divisible by p . \square

Remark 7.3. Multiplying by $(M - I)^{-1}$, as in (3.4), will lose $v_p(\det(M - I))$ digits of absolute p -adic precision. As $v_p(\det(M - I)) = v_p(\text{Jac}(X)(\mathbf{F}_p)[p])$, this is at most g in the *anomalous case*, and in general we expect that it is 0, so if $g = O(N)$ the whole computation can be repeated with the extra precision required at no extra asymptotic cost.

8. RUN TIME ANALYSIS

Having described the algorithm in detail, we now analyse its run time, in order to prove Theorem 1.1. First of all we analyse each step of Algorithm 6.1.

The main step is the computation of the reduction matrices via Theorem 5.2. In this case, we have $m = 2g$ (+1 in the horizontal case) and $n = L$. When reducing horizontally, for each row the largest index is bounded by $K = O(Np)$. When reducing vertically our index is also at most $O(Np)$. As there are N rows in total, we obtain a total of

$$O\left(N((\text{MM}(g) + \text{MM}(g, L))\sqrt{Np} + (g^2 + gL)\text{M}(\sqrt{Np}))\right) \quad (8.1)$$

ring operations to compute the matrices. Using that $\text{M}(d) \in \tilde{O}(d)$, that $\text{MM}(m) = m^\omega$ for some $2 \leq \omega \leq 3$, and the above discussion of $\text{MM}(m, n)$, we simplify to $O((g^\omega + Lg^{\omega-1})\sqrt{p}N^{3/2})$ ring operations, bit complexity $\tilde{O}((g^\omega + Lg^{\omega-1})\sqrt{p}N^{5/2})$.

The remaining operations are exactly as analysed by Harvey in [Har07] Section 7.4. With our larger, but still sparse, horizontal reduction matrices, each reduction step without Theorem 5.2 uses $O(g + L)$ rather than $O(g)$ ring operations, for a total of $O(N^3g^3(g + L))$ ring operations, or $\tilde{O}(N^4g^3(g + L)\log p)$ bit operations. We then have a total time complexity of

$$\tilde{O}\left((g^\omega + Lg^{\omega-1})\sqrt{Np}N^2 + N^4g^3(g + L)\log p\right). \quad (8.2)$$

Now we turn to the algorithm for computing Coleman integrals, obtained by running Algorithm 6.1 once and then Algorithm 3.3 once for each point. The analysis here is the same as that in [BBK10] Section 4.2, where, by using Algorithm 6.1 instead of Kedlaya's algorithm, we may replace the $\tilde{O}(pN^2g^2)$ in their complexity analysis with (8.2). The remaining steps to complete the Coleman integration are logarithmic in p and are dominated by the logarithmic in p term of (8.2).

If L is fixed (for example $L = 2$ when computing integrals between two points) the complexity is as in [Har07] Theorem 1.1. This finishes the proof of Theorem 1.1.

Remark 8.1. The version of Kedlaya's algorithm used in [BBK10] Algorithm 10, seems to have an advantage in that it outputs the power series of the f_i 's. This could of course be re-used later to evaluate at further points without re-running Kedlaya's algorithm. However, for p large enough, this series has so many terms that it is faster asymptotically to recompute everything with the algorithm given here, than it is to evaluate the power series at one point.

$p \setminus N$	1	3	5	7	9
131	1.14/0.01	3.67/0.02	9.36/0.07	16.90/0.12	20.06/0.49
257	1.96/0.01	8.90/0.03	20.83/0.07	30.91/0.18	63.14/0.68
521	4.73/0.01	19.23/0.03	39.18/0.08	86.49/0.62	162.81/0.91

TABLE 1. Timings for genus 3: Sage 8.0 time/New time (sec)

9. IMPLEMENTATION

We have implemented this algorithm in C++ as an extension of David Harvey’s `hypellfrob` package. This extension has been wrapped and can be easily used from within Sage [Sag18]. The implementation is included as part of the supplementary materials to this paper. This implementation uses naive matrix multiplication (for which $\omega = 3$) and does not take into account the special form of the matrices, as in Theorem 5.2; so the run time of this implementation will not have the asymptotic behaviour stated in (8.2) for the parameter L .

In Table 1, we list some timings obtained using this implementation in genus 3, for various primes p and p -adic precision bounds N . For comparison, we also list timings for the functionality for computing Coleman integrals in Sage 8.0. The implementation in Sage is written in Python, rather than C++, so we would expect some speed-up even if a superior algorithm was not used. Specifically we have compared the time to compute the Coleman data only, and do not include any of the time spent doing the linear algebra and tiny integral steps of Coleman integration, which should be comparatively fast. As such, we only time the components that will differ between the old and new approaches. For the existing Sage code we have timed both finding the matrix of Frobenius and the primitives (by calling `monskey_washnitzer.matrix_of_frobenius_hyperelliptic`), and the time to evaluate the resulting primitive at one point. This is compared with the time taken by the new implementation, called from its Sage wrapper with one point specified, this outputs the matrix of Frobenius and the evaluations at that point. All timings and examples are on a single 16 AMD Opteron 8384 2.7GHz processor on a machine with 16 cores and 82 GB RAM. While this table is mostly intended to show practicality, in the $N = 9$ column the square root dependence on p can be seen. The large jump in the timings between $p \approx 256$ and $p \approx 512$ for $N = 7$ could be explained by the fact that this is the cut off between when an element of $\mathbf{Z}/p^N\mathbf{Z}$ is representable in one machine word.

10. EXAMPLES

In this section we give an explicit example of a computation we can perform with this technique, demonstrating how large we can feasibly take the parameters. We compare our implementation to the existing functionality for Coleman integration in Sage 8.0 for this example.

The current implementation uses the basis $x^i dx/y$, to remain consistent with Harvey’s notation. As the existing functionality for Coleman integration in Sage 8.0 uses the basis $x^i dx/2y$ for cohomology, we must divide the obtained evaluations by 2 to compare them to those returned by Sage or Algorithm 6.1.

Example 10.1. Let $C: y^2 = x^5 + \frac{33}{16}x^4 + \frac{3}{4}x^3 + \frac{3}{8}x^2 - \frac{1}{4}x + \frac{1}{16}$ be Leprévost’s curve, as in [BBK10] Example 21. Then letting $P = (-1, 1)$, $Q = (0, \frac{1}{4})$ and

$p = 2^{45} + 59 = 35184372088891$, using our implementation we can compute the matrix of Frobenius M to 1 p -adic digit of precision, and also that

$$\begin{aligned} f_0(P) - f_0(Q) &= O(p), & f_1(P) - f_1(Q) &= O(p), \\ f_2(P) - f_2(Q) &= 7147166195043 + O(p), & f_3(P) - f_3(Q) &= 9172338112529 + O(p). \end{aligned}$$

Computing this (and finding $(M - 1)^{-1}$) takes a total of 27.8 minutes (with a peak memory usage of 2.9GB). Evaluating Coleman integrals for such a large prime is far out of the range of what was possible to compute in reasonable amount of time using the previous implementation. In fact, even when $p = 2^{14} + 27$ the existing Sage functionality takes 53.2 minutes, and uses a larger volume of memory (12GB).

As we have used only 1 digit of p -adic precision, the points P and Q are congruent up to this precision to the corresponding Teichmüller point in their residue disk. So, for this example, we do not need to worry about computing tiny integrals; the vector of Coleman integrals $\int_Q^P \omega_i$ can be obtained from the above vector of evaluations by multiplying by $(M - 1)^{-1}$. Doing this gives us the vector $(O(p), O(p), 9099406574713 + O(p), 7153144612900 + O(p))$ reflecting the holomorphicity of the first two basis differentials only. We have also run the same example with precision $N = 3$; this took 22.5 hours and used a peak of 50GB of memory.

11. FUTURE DIRECTIONS

The assumptions on the size of p allow us to use at most one extra digit of p -adic precision; it should be possible to relax this assumption somewhat, using a more complicated algorithm instead. Similarly it should be possible to work over extensions of \mathbf{Q}_p , or remove the assumption that $Q(x)$ is monic.

Kedlaya’s algorithm has been generalised to other curves and varieties, e.g. [Har12, GG01, Gon15, Tui17] and Harvey’s techniques have also been generalised to some of these cases [Min10, ABC⁺18]. Moreover, explicit Coleman integration has also been carried out in some of these settings, for even degree hyperelliptic curves [Bal15], and for general curves [BT17]. It would be interesting to adapt our techniques to those contexts. Iterated Coleman integrals are also of interest and have been made computationally effective [Bal13]. Extending the algorithm presented here to compute iterated integrals is another natural next step. Harvey has also described an *average polynomial time* algorithm for dealing with for many primes at once [Har14]. The author plans to explore the feasibility of analogous techniques when computing Coleman integrals.

REFERENCES

- [ABC⁺18] Arul V., Best A.J., Costa E., Magner R., Triantafillou N. “Computing Zeta Functions of Cyclic Covers in Large Characteristic”. In “ANTS XIII—Proceedings of the Thirteenth Algorithmic Number Theory Symposium”, This volume. Math. Sci. Publ., Berkeley, CA, 2018.
- [Bal13] Balakrishnan J.S. “Iterated Coleman integration for hyperelliptic curves”. In “ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium”, volume 1 of *Open Book Ser.*, pages 41–61. Math. Sci. Publ., Berkeley, CA, 2013. doi:10.2140/obs.2013.1.41.
- [Bal15] ———. “Coleman integration for even-degree models of hyperelliptic curves”. *LMS J. Comput. Math.*, 18(1):258–265, 2015.
- [BBK10] Balakrishnan J.S., Bradshaw R.W., Kedlaya K.S. “Explicit Coleman integration for hyperelliptic curves”. In “Algorithmic number theory”, volume 6197 of *Lecture Notes in Comput. Sci.*, pages 16–31. Springer, Berlin, 2010.

- [BD18] Balakrishnan J.S., Dogra N. “Quadratic Chabauty and rational points I: p -adic heights”. *Duke Mathematical Journal*, 2018. <http://arxiv.org/abs/1601.00388v2>.
- [Bes12] Besser A. “Heidelberg lectures on Coleman integration”. In “The arithmetic of fundamental groups—PIA 2010”, volume 2 of *Contrib. Math. Comput. Sci.*, pages 3–52. Springer, Heidelberg, 2012.
- [BGS07] Bostan A., Gaudry P., Schost E. “Linear recurrences with polynomial coefficients and application to integer factorization and Cartier-Manin operator”. *SIAM J. Comput.*, 36(6):1777–1806, 2007.
- [BT17] Balakrishnan J.S., Tuitman J. “Explicit Coleman integration for curves”. 2017. [1710.01673v2](https://arxiv.org/abs/1710.01673v2).
- [CdS88] Coleman R., de Shalit E. “ p -adic regulators on curves and special values of p -adic L -functions”. *Invent. Math.*, 93(2):239–266, 1988.
- [Col82] Coleman R.F. “Dilogarithms, regulators and p -adic L -functions”. *Invent. Math.*, 69(2):171–208, 1982.
- [Col85] ———. “Torsion points on curves and p -adic abelian integrals”. *Ann. of Math. (2)*, 121(1):111–168, 1985.
- [GG01] Gaudry P., Gürel N. “An extension of Kedlaya’s point-counting algorithm to superelliptic curves”. In “Advances in cryptology—ASIACRYPT 2001 (Gold Coast)”, volume 2248 of *Lecture Notes in Comput. Sci.*, pages 480–494. Springer, Berlin, 2001.
- [Gon15] Gonçalves C. “A point counting algorithm for cyclic covers of the projective line”. In “Algorithmic arithmetic, geometry, and coding theory”, volume 637 of *Contemp. Math.*, pages 145–172. Amer. Math. Soc., Providence, RI, 2015.
- [Har07] Harvey D. “Kedlaya’s algorithm in larger characteristic”. *Int. Math. Res. Not. IMRN*, (22):Art. ID rnm095, 29, 2007.
- [Har12] Harrison M.C. “An extension of Kedlaya’s algorithm for hyperelliptic curves”. *J. Symbolic Comput.*, 47(1):89–101, 2012.
- [Har14] Harvey D. “Counting points on hyperelliptic curves in average polynomial time”. *Ann. of Math. (2)*, 179(2):783–803, 2014.
- [Ked01] Kedlaya K.S. “Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology”. *J. Ramanujan Math. Soc.*, 16(4):323–338, 2001.
- [LG12] Le Gall F. “Faster algorithms for rectangular matrix multiplication”. In “2012 IEEE 53rd Annual Symposium on Foundations of Computer Science—FOCS 2012”, pages 514–523. IEEE Computer Soc., Los Alamitos, CA, 2012.
- [Min10] Minzlaff M. “Computing zeta functions of superelliptic curves in larger characteristic”. *Math. Comput. Sci.*, 3(2):209–224, 2010.
- [Sag18] Sage Developers, The. *SageMath, the Sage Mathematics Software System (Version 8.0.0)*, 2018. <http://www.sagemath.org>.
- [Tui17] Tuitman J. “Counting points on curves using a map to \mathbf{P}^1 , II”. *Finite Fields Appl.*, 45:301–322, 2017.

111 CUMMINGTON MALL, BOSTON MA 02215
E-mail address: alex.j.best@gmail.com