

COMPUTING MODULAR POLYNOMIALS

DENIS CHARLES AND KRISTIN LAUTER

1. INTRODUCTION

The ℓ^{th} modular polynomial, $\phi_\ell(x, y)$, parameterizes pairs of elliptic curves with a cyclic isogeny of degree ℓ between them. Modular polynomials provide the defining equations for modular curves, and are useful in many different aspects of computational number theory and cryptography. For example, computations with modular polynomials have been used to speed elliptic curve point-counting algorithms ([BSS99] Chapter VII).

The standard method for computing modular polynomials consists of computing the Fourier expansion of the modular \mathbf{j} -function and solving a linear system of equations to obtain the integral coefficients of $\phi_\ell(x, y)$. According to Elkies (see [Elk98] §3) this method has a running time of $O(\ell^{4+\epsilon})$ if one uses fast multiplication. However, our analysis (given in the Appendix) shows that the running time of this method is, in fact, $\Theta(\ell^{9/2+\epsilon})$ using fast multiplication.

The object of the current paper is to compute the modular polynomial (for prime ℓ) directly modulo a prime p , without first computing the coefficients as integers. Once the modular polynomial has been computed for enough small primes, our approach can also be combined with the Chinese Remainder Theorem (CRT) approach as in [CNST98] or [ALV03] to obtain the modular polynomial with integral coefficients or with coefficients modulo a much larger prime using Explicit CRT. Our algorithm does not involve computing Fourier coefficients of modular functions. The running time of our algorithm turns out to be $O(\ell^{4+\epsilon})$ using fast multiplication. We believe our method is interesting as it is asymptotically faster; and is an essentially different approach to computing modular polynomials. Furthermore, our algorithm also yields as a corollary a fast way to compute a random ℓ -isogeny of an elliptic curve over a finite field.

The idea of our algorithm is as follows. Mestre's algorithm, Méthode des graphes [Mes86], uses the ℓ^{th} modular polynomial modulo p to navigate around the connected graph of supersingular elliptic curves over \mathbb{F}_{p^2} in order to compute the number of edges (isogenies of degree ℓ) between each node. From the graph, Mestre then obtains the ℓ^{th} Brandt matrix giving the action of the ℓ^{th} Hecke operator on modular forms of weight 2. In our algorithm we do the opposite: we compute the ℓ^{th} modular polynomial modulo p by computing all the isogenies of degree ℓ between supersingular curves modulo p via Vélu's formulae. Specifically, for a given \mathbf{j} -invariant, j (say), of a supersingular elliptic curve over \mathbb{F}_{p^2} , Algorithm 1 computes $\phi_\ell(x, j)$ modulo p by computing the $\ell + 1$ distinct subgroups of order ℓ and computing the \mathbf{j} -invariants of the $\ell + 1$ corresponding ℓ -isogenous elliptic curves. Algorithm 2 then uses the connectedness of the graph of supersingular elliptic curves over \mathbb{F}_{p^2} to move around the graph, calling Algorithm 1 for different values of j until enough information is obtained to compute $\phi_\ell(x, y)$ modulo p via interpolation.

There are several interesting aspects to Algorithms 1 and 2. Algorithm 1 does not use the factorization of the ℓ -division polynomials to produce the subgroups of order ℓ . Instead we generate independent ℓ -torsion points by picking random points with coordinates in a suitable extension of \mathbb{F}_p and taking a scalar multiple which is the group order divided by ℓ . This turns out to be more efficient than factoring the ℓ^{th} division polynomial for large ℓ .

Algorithm 2 computes $\phi_\ell(x, y)$ modulo p by doing only computations with supersingular elliptic curves in characteristic p even though $\phi_\ell(x, y)$ is a general object giving information about isogenies between elliptic curves in characteristic 0 and ordinary elliptic curves in characteristic p . The advantage that we gain by using supersingular elliptic curves is that we can show that the full ℓ -torsion is defined over an extension of degree $O(\ell)$ of the base field \mathbb{F}_{p^2} , whereas in general the field of definition can be of degree as high as $\ell^2 - 1$.

In this article we provide a running time analysis assuming fast multiplication implementation of field operations. But for small values of ℓ fast multiplication is usually not used in practice, thus we also give the running time (without the analysis) assuming a naïve implementation of field operations.

2. LOCAL COMPUTATION OF $\phi_\ell(x, j)$

The key ingredient of the algorithm is the computation of the univariate polynomial $\phi_\ell(x, j)$ modulo a prime p given a \mathbf{j} -invariant j . We describe the method to do this here.

Algorithm 1

Input: Two distinct primes p and ℓ , and j the \mathbf{j} -invariant of a supersingular elliptic curve E over a finite field \mathbb{F}_q of degree at most 2 over a prime field of characteristic p .

Output: The polynomial $\phi_\ell(x, j) = \prod_{E' \ell\text{-isogenous to } E} (x - j(E')) \in \mathbb{F}_{p^2}[x]$.

Step 1 Find the generators P and Q of $E[\ell]$:

- (a) Let n be such that $\mathbb{F}_q(E[\ell]) \subseteq \mathbb{F}_{q^n}$.
- (b) Let $S = \#E(\mathbb{F}_{q^n})$, the number of \mathbb{F}_{q^n} rational points on E .
- (c) Set $s = S/\ell^k$, where ℓ^k is the largest power of ℓ that divides S (note $k \geq 2$).
- (d) Pick two points P and Q at random from $E[\ell]$:
 - (i) Pick two points U, V at random from $E(\mathbb{F}_{q^n})$.
 - (ii) Set $P' = sU$ and $Q' = sV$, if either P' or Q' equals \mathcal{O} then repeat step (i).
 - (iii) Find the smallest i_1, i_2 such that $\ell^{i_1}P' \neq \mathcal{O}$ and $\ell^{i_2}Q' \neq \mathcal{O}$ but $\ell^{i_1+1}P' = \mathcal{O}$ and $\ell^{i_2+1}Q' = \mathcal{O}$.
 - (iv) Set $P = \ell^{i_1}P'$ and $Q = \ell^{i_2}Q'$.
- (e) Using Shanks's Baby-steps-Giant-steps algorithm check if Q belongs to the group generated by P . If so repeat step (d).

Step 2 Find the \mathbf{j} -invariants $j_1, \dots, j_{\ell+1}$ in \mathbb{F}_{p^2} of the $\ell + 1$ elliptic curves that are ℓ -isogenous to E .

- (a) Let $G_1 = \langle Q \rangle$ and $G_{1+i} = \langle P + (i-1)Q \rangle$ for $1 \leq i \leq \ell$.
- (b) For each i , $1 \leq i \leq \ell + 1$ compute the \mathbf{j} -invariant of the elliptic curve E/G_i using Vélú's formulas.

Step 3 Output $\phi_\ell(x, j) = \prod_{1 \leq i \leq \ell+1} (x - j_i)$.

Remark 2.1. In step (1e), one could alternately use the Weil pairing to check whether P and Q generate the ℓ -torsion. Doing so, however, does not lead to an asymptotic improvement in the running time of the algorithm.

The following lemma gives the possibilities for the value of n in Step (1a). We prove the following result for all elliptic curves not just supersingular ones.

Lemma 2.2. *Let E/\mathbb{F}_q be an elliptic curve, and let ℓ be a prime not equal to the characteristic of \mathbb{F}_q . Then $E[\ell] \subseteq E(\mathbb{F}_{q^n})$ where n is a divisor of either $\ell(\ell - 1)$ or $\ell^2 - 1$.*

Proof : The Weil-pairing tells us that if $E[\ell] \subseteq \mathbb{F}_{q^n}$ then $\mu_\ell \subseteq \mathbb{F}_{q^n}$ ([Sil86] Corollary 8.1.1). We seek, however, an upper bound on n , to do this we use the Galois representation coming from the ℓ -division points of E . Indeed, we have an *injective* group homomorphism ([Sil86] Chapter III, §7)

$$\rho_\ell : \text{Gal}(\mathbb{F}_q(E[\ell])/\mathbb{F}_q) \rightarrow \text{Aut}(E[\ell]) \cong \text{GL}_2(\mathbb{F}_\ell).$$

The Galois group $\text{Gal}(\mathbb{F}_q(E[\ell])/\mathbb{F}_q)$ is cyclic, thus by ρ_ℓ the possibilities for $\text{Gal}(\mathbb{F}_q(E[\ell])/\mathbb{F}_q)$ are limited to cyclic subgroups of $\text{GL}_2(\mathbb{F}_\ell)$. In other words, we are interested in the orders of the elements in $\text{GL}_2(\mathbb{F}_\ell)$. The elements of $\text{GL}_2(\mathbb{F}_\ell)$ are conjugate to one of the following types of matrices:

$$\begin{pmatrix} \alpha & 0 \\ 0 & \beta \end{pmatrix}, \begin{pmatrix} \alpha & 1 \\ 0 & \alpha \end{pmatrix}, \text{ for } \alpha, \beta \in \mathbb{F}_\ell^*, \alpha \neq \beta,$$

or those corresponding to multiplication by an element of $\mathbb{F}_{\ell^2}^*$ on the 2-dimensional \mathbb{F}_ℓ vector space \mathbb{F}_{ℓ^2} . It is easy to see that the orders of these elements all divide $\ell(\ell - 1)$ or $\ell^2 - 1$. Thus the degree of the field extension containing the ℓ -torsion points on E must divide either $\ell(\ell - 1)$ or $\ell^2 - 1$. \square

We will try step (1) with $n = \ell^2 - 1$, if steps (1d - 1e) do not succeed for some K (a constant) many trials, we repeat with $n = \ell(\ell - 1)$. The analysis that follows shows that a sufficiently large constant K will work.

For step (1b) we do not need a point counting algorithm to determine S . Since E is a supersingular elliptic curve, we have the following choices for the trace of Frobenius a_q :

$$a_q = \begin{cases} 0 & \text{if } E \text{ is over } \mathbb{F}_p \\ 0, \pm p, \pm 2p & \text{if } E \text{ is over } \mathbb{F}_{p^2}. \end{cases}$$

Not all the possibilities can occur for certain primes, but we will not use this fact here (see [Sch87]). If the curve is over \mathbb{F}_{p^2} we can determine probabilistically the value of a_q as follows. Pick a point P at random from $E(\mathbb{F}_q)$ and check if $(q + 1 + a_q)P = \mathcal{O}$. Since the pairwise gcd's of the possible group orders divide 4, with high probability only the correct value of a_q will annihilate the point. Thus in $O(\log^{2+o(1)} q)$ time we can determine with high probability the correct value of a_q . Once we know the correct trace a_q , we can find the roots (in $\overline{\mathbb{Q}}$), π and $\bar{\pi}$, of the characteristic polynomial of the Frobenius $\phi^2 - a_q\phi + q$. Then the number of points lying on E over the field \mathbb{F}_{q^n} is given by $q^n + 1 - \pi^n - \bar{\pi}^n$, this gives us S .

Note: We could have used a deterministic point counting algorithm to find $\sharp E(\mathbb{F}_q)$ but this would have cost $O(\log^6 q)$ field operations.

A lower bound on the probability that step (1d) succeeds is given by the following lemma whose proof is straightforward.

Lemma 2.3. *For a random choice of the points U and V in step (1d i) the probability that step (1d ii) succeeds is at least*

$$\left(1 - \frac{1}{\ell^2}\right)^2.$$

At the end of step (1d) we have two random ℓ -torsion points of E namely, P and Q . The probability that Q belongs to the cyclic group generated by P is $\frac{\ell}{\ell^2} = \frac{1}{\ell}$. Thus with high probability we will find in step (1e) two generators for $E[\ell]$.

Lemma 2.4. *The expected running time of Step 1 is $O(\ell^{4+o(1)} \log^{2+o(1)} q)$ and $O(\ell^6 \log^3 q)$ if fast multiplication is not used.*

Proof : The finite field \mathbb{F}_{q^n} can be constructed by picking an irreducible polynomial of degree n . A randomized method that requires on average $O((n^2 \log n + n \log q) \log n \log \log n)$ operations over \mathbb{F}_q is given in [Sho94]. Thus the field can be constructed in $O(\ell^{4+o(1)} \log^{2+o(1)} q)$ time since $n \leq \ell^2$. Step (1d) requires picking a random point on E . We can do this by picking a random element in \mathbb{F}_{q^n} treating it as the x -coordinate of a point on E and solving the resulting quadratic equation for the y -coordinate. Choosing a random element in \mathbb{F}_{q^n} can be done in $O(\ell^2 \log q)$ time. Solving the quadratic equation can be done probabilistically in $O(\ell^2 \log q)$ field operations. Thus to pick a point on E can be done in $O(\ell^{4+o(1)} \log^{2+o(1)} q)$ time. The computation in steps (1d i – iv) computes integer multiples of a point on the elliptic curve, where the integer is at most q^n , and this can be done in $O(\ell^{4+o(1)} \log^{2+o(1)} q)$ time using the repeated squaring method and fast multiplication. Shanks's Baby-steps-giant-steps algorithm for a cyclic group G requires $O(\sqrt{|G|})$ group operations. Thus step (1e) runs in time $O(\ell^{\frac{5}{2}+o(1)} \log^{1+o(1)} q)$, since the group is cyclic of order ℓ . \square

Let C be a subgroup of E , Vélu in [Vel71] gives explicit formulas for determining the equation of the isogeny $E \rightarrow E/C$ and the Weierstrass equation of the curve E/C . We give the formulas when ℓ is an odd prime. Let E is given by the equation

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6.$$

We define the following two functions in $\mathbb{F}_q(E)$ for $Q = (x, y)$ a point on $E - \{\mathcal{O}\}$ define

$$\begin{aligned} g^x(Q) &= 3x^2 + 2a_2x + a_4 - a_1y \\ g^y(Q) &= -2y - a_1x - a_3, \end{aligned}$$

and set

$$\begin{aligned} t(Q) &= 2g^x(Q) - a_1g^y(Q) \\ u(Q) &= (g^y(Q))^2 \\ t &= \sum_{Q \in (C - \{\mathcal{O}\})} t(Q) \\ w &= \sum_{Q \in (C - \{\mathcal{O}\})} (u(Q) + x(Q)t(Q)). \end{aligned}$$

Then the curve E/C is given by the equation

$$Y^2 + A_1XY + A_3Y = X^3 + A_2X^2 + A_4X + A_6$$

where

$$\begin{aligned} A_1 &= a_1, A_2 = a_2, A_3 = a_3, \\ A_4 &= a_4 - 5t, A_6 = a_6 - (a_1^2 + 4a_2)t - 7w. \end{aligned}$$

From the Weierstrass equation of E/C we can easily determine the \mathbf{j} -invariant of E/C . It is clear that this procedure can be done using $O(\ell)$ elliptic curve operations for each of the groups G_i , $1 \leq i \leq \ell + 1$. Thus step 2 can be done in $O(\ell^{4+o(1)} \log^{1+o(1)} q)$ time steps. Step 3 requires only $O(\ell)$ field operations and so the running time of the algorithm is dominated by the running time of steps 1 and 2. Note that the polynomial obtained at the end of Step 3 $\phi_\ell(x, j)$ has coefficients in $\mathbb{F}_{p^2}[x]$ since all the curves ℓ -isogenous to E are supersingular and hence their \mathbf{j} -invariants belong to \mathbb{F}_{p^2} . In summary, we have the following:

Theorem 2.5. *Algorithm 1 computes $\phi_\ell(x, j) \in \mathbb{F}_{p^2}[x]$, in fact, the list of roots of $\phi_\ell(x, j)$, and has an expected running time of $O(\ell^{4+o(1)} \log^{2+o(1)} q)$ and $O(\ell^6 \log^3 q)$ without fast multiplication.*

For our application of Algorithm 1 we will need the dependence of the running time in terms of the quantity n . We make the dependence explicit in the next theorem.

Theorem 2.6. *With notation as above, Algorithm 1 computes $\phi_\ell(x, j) \in \mathbb{F}_{p^2}[x]$ together with the list of its roots and has an expected running time of $O(n^{2+o(1)} \log^{2+o(1)} q + \sqrt{\ell} n^{1+o(1)} \log^{1+o(1)} q + \ell^2 n^{1+o(1)} \log q)$. If fast multiplication is not used then Algorithm 1 has an expected running time of $O(n^3 \log^3 q + \sqrt{\ell} n^2 \log^2 q + \ell^2 n^2 \log^2 q)$.*

In the case of ordinary elliptic curve, step (1) of Algorithm 1 can still be used, once the number of points on E/\mathbb{F}_q has been determined, by Lemma 2.2 the degree of the extension, n , is still $O(\ell^2)$. This leads to the following two results:

Corollary 2.7. *If E/\mathbb{F}_q is an elliptic curve, we can pick a random ℓ -torsion point on $E(\overline{\mathbb{F}}_q)$ in time $O(\ell^{4+o(1)} \log^{2+o(1)} q + \log^{6+o(1)} q)$ and $O(\ell^6 \log^3 q + \log^8 q)$ without fast multiplication.*

Corollary 2.8. *If E/\mathbb{F}_q is an elliptic curve, we can construct a random ℓ -isogenous curve to E in time $O(\ell^{4+o(1)} \log^{2+o(1)} q + \log^{6+o(1)} q)$ and $O(\ell^6 \log^3 q + \log^8 q)$ without fast multiplication.*

Factoring a degree d polynomial over a finite field \mathbb{F}_{q^n} requires $O(d^2(n \log q))$ operations over \mathbb{F}_{q^n} . To factor the ℓ -division polynomial we need an extension of degree roughly ℓ^2 . Thus, if we were to factor the ℓ -division polynomial to generate the isogeny, we would need $O(\ell^6 \log q)$ operations over a field of degree ℓ^2 over \mathbb{F}_q which translates to $O(\ell^{8+o(1)} \log^{2+o(1)} q)$ bit operations *even* with fast multiplication.

3. COMPUTING $\phi_\ell(x, y) \pmod p$

In characteristic $p > 2$ there are exactly

$$S(p) = \left\lfloor \frac{p}{12} \right\rfloor + \epsilon_p$$

supersingular \mathbf{j} -invariants where

$$\epsilon_p = 0, 1, 1, 2 \text{ if } p \equiv 1, 5, 7, 11 \pmod{12}.$$

In this section we provide an algorithm for computing $\phi_\ell(x, y) \pmod p$ provided $S(p) \geq \ell + 1$. The description of the algorithm follows:

Algorithm 2

Input: Two distinct primes ℓ and p with $S(p) \geq \ell + 1$.

Output: The polynomial $\phi_\ell(x, y) \in \mathbb{F}_p[x, y]$.

- (1) Find the smallest (in absolute value) discriminant $D < 0$ such that $\left(\frac{D}{p}\right) = -1$.
- (2) Compute the Hilbert Class polynomial $H_D(x) \pmod{p}$.
- (3) Let j_0 be a root of $H_D(x)$ in \mathbb{F}_{p^2} .
- (4) Set $i = 0$.
- (5) Compute $\phi_i = \phi_\ell(x, j_i) \in \mathbb{F}_{p^2}$ using Algorithm 1.
- (6) Let j_{i+1} be a root of ϕ_k for $k \leq i$ which is not one of j_0, \dots, j_i .
- (7) If $i < \ell$ then set $i = i + 1$ and repeat Step 5.
- (8) Writing $\phi_\ell(x, y) = x^{\ell+1} + \sum_{0 \leq k \leq \ell} p_k(y)x^k$, we have $\ell + 1$ systems of equations of the form $p_k(j_i) = v_{ki}$ for $0 \leq k, i \leq \ell$. Solve these equations for each $p_k(y)$, $0 \leq k \leq \ell$.
- (9) Output $\phi_\ell(x, y) \in \mathbb{F}_p[x, y]$.

We argue that the above algorithm is correct and analyze the running time. For step 1, we note that if $p \equiv 3 \pmod{4}$, then $D = -4$ works. Otherwise, -1 is a quadratic residue and writing (without loss of generality) D as $-4d$, we are looking for the smallest d such that $\left(\frac{d}{p}\right) = -1$. A theorem of Burgess ([Bur62]) tells us that $d \ll p^{\frac{1}{4\sqrt{e}}}$, and under the assumption of GRH the estimate of Ankeny ([Ank52]) gives $d \ll \log^2 p$. Computing $H_D(x) \pmod{p}$ can be done in $O(d^2(\log d)^2)$ time [LL90] §5.10. Thus step 2 requires $O(\sqrt{p} \log^2 p)$ time, and under the assumption of GRH requires $O(\log^4 p (\log \log p)^2)$ time. Since $\left(\frac{D}{p}\right) = -1$ all the roots of $H_D(x)$ are supersingular \mathbf{j} -invariants in characteristic p . $H_D(x)$ is a polynomial of degree $h(\sqrt{-D})$, the class number of the order of discriminant D , and this is $\ll |D|^{\frac{1}{2}+\epsilon}$. Finding a root of $H_D(x) \in \mathbb{F}_{p^2}$ can be done in $O(d^{1+\epsilon} \log^{2+o(1)} p)$ time using probabilistic factoring algorithms, where $d = |D|$. The graph with supersingular \mathbf{j} -invariants over characteristic p as vertices and ℓ -isogenies as edges is connected (see [Mes86]), consequently, we will always find a \mathbf{j} -invariant in step 6 that is not one of j_0, \dots, j_i . Thus the loop in steps (5) \dots (7) is executed exactly $\ell + 1$ times under the assumption that $S(p) \geq \ell + 1$. Even though Algorithm 1 requires $\tilde{O}(\ell^4 \log^2 q)$ time¹ in the worst case, we will argue that, in fact, for *all* of the iterations of the loop it actually runs in $\tilde{O}(\ell^3 \log^2 q)$ time.

Lemma 3.1. *Let E be a supersingular elliptic curve defined over \mathbb{F}_{p^2} . Then the extension degree*

$$[\mathbb{F}_{p^2}(E[\ell]) : \mathbb{F}_{p^2}]$$

divides $6(\ell - 1)$.

Proof: Let E/\mathbb{F}_{p^2} be a supersingular curve and let t be the trace of Frobenius. Then the Frobenius map ϕ satisfies

$$\phi^2 - t\phi + p^2 = 0$$

with $t = 0$ or $\pm p$ or $\pm 2p$. Suppose $t = \pm 2p$, then the Frobenius acts as multiplication by $\pm p$ on the curve E . Thus $\phi^{\ell-1}$ acts trivially on $E[\ell]$, and the ℓ -torsion points are defined over an extension of degree dividing $\ell - 1$. If $t = 0$, then $\phi^2 = -p^2$ and so $\phi^{2(\ell-1)}$ acts trivially on the ℓ -torsion. Thus $E[\ell]$ is defined over an extension of degree dividing $2(\ell - 1)$. If $t = \pm p$, then $\phi^3 = \pm p^3$ and consequently $\phi^{3(\ell-1)}$ acts trivially on the ℓ -torsion of the curve. Thus the ℓ -torsion is defined over an extension of degree dividing $3(\ell - 1)$. Thus in all cases the ℓ -torsion of the curve is defined over an extension of degree dividing $6(\ell - 1)$. \square

¹We use the soft-Oh \tilde{O} notation when we ignore factors of the form $\log \ell$ or $\log \log p$.

Thus, the loop in steps (5) \cdots (7), Algorithm 1 can be run with the quantity $n = 6(\ell - 1)$. For this value of n Algorithm 1 runs in expected time $O(\ell^{3+o(1)} \log^{2+o(1)} p)$ and so the loop runs in expected time $O(\ell^{4+o(1)} \log^{2+o(1)} p)$.

Writing the modular polynomial $\phi_\ell(x, y)$ as $x^{\ell+1} + \sum_{0 \leq k \leq \ell} p_k(y)x^k$, we know that $p_0(y)$ is monic of degree $\ell + 1$ and $\deg(p_k(y)) \leq \ell$ for $1 \leq k \leq \ell$. Thus at the end of the loop in steps (5) \cdots (7) we have enough information to solve for the $p_k(y)$ in step (8). We are solving $\ell + 1$ systems of equations, each requiring an inversion of a matrix of size $(\ell + 1) \times (\ell + 1)$. This can be done in $O(\ell^4 \log^{1+o(1)} p)$ time. Since the polynomial $\phi_\ell(x, y) \pmod p$ is the reduction of the classical modular polynomial, a polynomial with integer coefficients, the polynomial that we compute has coefficients in \mathbb{F}_p . Thus we have proved the following theorem:

Theorem 3.2. *Given ℓ and p distinct primes such that $S(p) \geq \ell + 1$, Algorithm 2 computes $\phi_\ell(x, y) \in \mathbb{F}_p[x, y]$ in expected time $O(\ell^{4+o(1)} \log^{2+o(1)} p + \log^4 p \log \log p)$ under the assumption of GRH and in $O(\ell^5 \log^3 p)$ time without fast multiplication.*

Hence, we can compute $\phi_\ell(x, y)$ modulo a prime p in $\tilde{O}(\ell^4 \log^2 p + \log^4 p)$ time if $p \geq 12\ell + 13$. If $p < 12\ell + 13$, we could still use the algorithm with ordinary elliptic curves and this would lead to a running time with the dependence on ℓ being ℓ^5 . Furthermore, we would not need the GRH since it was needed only to determine a supersingular curve in characteristic p . However, this is not very efficient.

Remark 3.3. If one is allowed to pick the prime p , such as would be the case if we are computing ϕ_ℓ over the integers using the Chinese Remainder Theorem combined with this method, then one can eliminate the assumption of GRH in the above theorem. For example, for primes $p \equiv 3 \pmod 4$ the j -invariant 1728 is supersingular. Thus in step (3) of Algorithm 2, we can start with $j_0 = 1728$ for any such prime. Hence we do not need the GRH to bound D in the analysis of the running time of the algorithm.

Acknowledgements: We would like to thank Igor Shparlinski and Steven Galbraith for useful suggestions and comments on an earlier draft of the paper. We would like to thank the anonymous referee for suggestions that improved the presentation of the paper. The first author would like to thank Microsoft Research for providing a stimulating research environment for this work.

REFERENCES

- [ALV03] Agashe, A.; Lauter, K.; Venkatesan, R.; *Constructing Elliptic Curves with a known number of points over a prime field*, in Lectures in honour of the 60th birthday of Hugh Cowie Williams, Fields Institute Communications Series, **42**, 1-17, 2003.
- [Ank52] Ankeny, N., C.; *The least quadratic non-residue*, Annals of Math., (2), **55**, 65-72, 1952.
- [BSS99] Blake, I.; Seroussi, G.; Smart, N.; *Elliptic Curves in Cryptography*, Lond. Math. Soc., Lecture Note Series, **265**, Cambridge University Press, 1999.
- [Bur62] Burgess, D., A.; *On character sums and primitive roots*, Proc. London Math. Soc., **III**, **12**, 179-192, 1962.
- [CNST98] Chao, J.; Nakamura, O.; Sobataka, K.; Tsujii, S.; *Construction of secure elliptic cryptosystems using CM tests and liftings*, Advances in Cryptology, ASIACRYPT'98 (Beijing), Lecture Notes in Computer Science, 1514, Springer-Verlag, Berlin, 1998.
- [Coh84] Cohen, P.; *On the coefficients of the transformation polynomials for the elliptic modular function*, Math. Proc. Cambridge Philos. Soc., **95**, 389-402, 1984.
- [Elk98] Elkies, Noam; *Elliptic and modular curves over finite fields and related computational issues*, in Computational Perspectives on Number Theory: Proceedings of a Conference in Honor of A.O.L. Atkin (D.A. Buell and J.T. Teitelbaum, eds.), AMS/International Press, 21-76, 1998.
- [LL90] Lenstra, A., K.; Lenstra, H., W., Jr.; *Algorithms in Number Theory*, Handbook of Theoretical Computer Science, Vol. **A**. Elsevier, 673-715, 1990.
- [Mah74] Mahler, K.; *On the coefficients of transformation polynomials for the modular function*. Bull. Austral. Math. Soc., **10**, 197-218, 1974.

- [Mes86] Mestre, J.-F.; *La méthode des graphes. Exemples et applications*, Proceedings of the international conference on class numbers and fundamental units of algebraic number fields, Nagoya Univ., Nagoya, 217-242, 1986.
- [Pet32] Petersson, H.; *Über die Entwicklungskoeffizienten der automorphen formen*, Acta Math., **58**, 169-215, 1932.
- [Sch87] Schoof, R.; *Nonsingular Plane Cubic Curves over Finite Fields*, J. Combinatorial Theory, **46**, 2, 183-208, 1987.
- [Sho94] Shoup, V.; *Fast construction of irreducible polynomials over finite fields*, J. Symbolic Computation, **17**, 371-391, 1994.
- [Sil86] Silverman, J. H.; *The Arithmetic of Elliptic Curves*, Graduate Texts in Mathematics, **106**, Springer-Verlag, 1986.
- [Vel71] Vélu, J.; *Isogénies entre courbes elliptiques*, C. R. Acad. Sc. Paris, **273**, 238-241, 1971.

APPENDIX

Elkies in [Elk98] (see §3 of that paper) claims that the usual method of computing the ℓ -th modular polynomial runs in time $O(\ell^{4+\epsilon})$. However, there is a subtle error in the analysis. We argue that in fact, the running time of this algorithm is $\Theta(\ell^{\frac{9}{2}+\epsilon})$. The first stage of the algorithm involves computing the first $\ell^2 + O(\ell)$ Fourier coefficients of the powers of the \mathbf{j} -function, namely $\mathbf{j}, \dots, \mathbf{j}^\ell$. This (as Elkies points out) can be done in $O(\ell^{3+\epsilon})$ arithmetic operations. The problem comes when we study the running time in terms of *bit-operations*. To analyze this we need to study the bit-sizes of the numbers that are handled by the algorithm. While it is true that the n -th Fourier coefficient of \mathbf{j} grows as $e^{O(\sqrt{n})}$, we need to also compute the Fourier coefficients of powers of \mathbf{j} and they grow faster (essentially because they have a higher order pole at ∞). In [Mah74] an upper bound of the form $\exp(4\pi(\sqrt{(n+k)k}))$ is proven for the n -th Fourier coefficient of \mathbf{j}^k . We show that the upper bound is quite close to the true magnitude of the Fourier coefficients below. Let $c(n)$ denote the n -th Fourier coefficient of the \mathbf{j} -function. It is well known that $c(n)$ are all positive integers and that (see [Pet32])

$$c(n) \sim \frac{e^{4\pi\sqrt{n}}}{\sqrt{2}n^{3/4}}.$$

The n -th Fourier coefficient of \mathbf{j}^k is given by

$$\sum_{a_1+a_2+\dots+a_k=n} c(a_1)c(a_2)\dots c(a_k).$$

Clearly, there is at least one partition of n (into k parts) where each of the parts, a_i , are $\geq \frac{n}{ck}$, where $c > 1$ is a constant. The asymptotic formula for $c(n)$ gives us that

$$\sum_{a_1+a_2+\dots+a_k=n} c(a_1)c(a_2)\dots c(a_k) \gg \left(e^{\sqrt{n/ck}-O(\log n)}\right)^k = e^{\Omega(\sqrt{nk})},$$

as long as k and n vary such that the ratio n/k goes to infinity. Thus a lower bound for the rate of growth of the n -th Fourier coefficient of $\mathbf{j}^k(z)$ is $e^{\Omega(\sqrt{kn})}$. Hence to compute the first $\ell^2 + O(\ell)$ Fourier coefficients of the powers \mathbf{j}^k , for $1 \leq k \leq \ell$, the bit length of the numbers involved is $O(\ell\sqrt{\ell})$ rather than the estimate $O(\ell \log \ell)$ used in [Elk98]. Thus this stage of the algorithm already requires $\Theta(\ell^{\frac{9}{2}+\epsilon})$ time using fast multiplication. The dependence of the running time on ℓ for our method, on the other hand, is $O(\ell^{4+\epsilon})$ if fast multiplication is used. The error in [Elk98] stems from the fact that a bound on the size of the coefficients of $\phi_\ell(x, y)$ was used to bound the bit sizes. These coefficients are somewhat smaller, their absolute value being bounded by $e^{O(\ell \log \ell)}$ (see [Coh84]).

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF WISCONSIN-MADISON, MADISON, WI - 53706.
E-mail address: cdx@cs.wisc.edu

MICROSOFT RESEARCH, ONE MICROSOFT WAY, REDMOND, WA - 98052.
E-mail address: klauter@microsoft.com