

p -adic class invariants

Reinier Bröker

ABSTRACT

We develop a new p -adic algorithm to compute the minimal polynomial of a class invariant. Our approach works for virtually any modular function yielding class invariants. The main algorithmic tool is modular polynomials, a concept which we generalize to functions of higher level.

1. Introduction

Let K be an imaginary quadratic number field and let \mathcal{O} be an order in K . The ring class field $H_{\mathcal{O}}$ of the order \mathcal{O} is an abelian extension of K , and the Artin map gives an isomorphism

$$\mathrm{Pic}(\mathcal{O}) \xrightarrow{\sim} \mathrm{Gal}(H_{\mathcal{O}}/K)$$

of the Picard group of \mathcal{O} with the Galois group of $H_{\mathcal{O}}/K$. In this paper we are interested in *explicitly computing* ring class fields. Complex multiplication theory provides us with a means of doing so. Letting Δ denote the discriminant of \mathcal{O} , it states that we have

$$H_{\mathcal{O}} = K[X]/(P_{\Delta}^j),$$

where P_{Δ}^j is the minimal polynomial over \mathbf{Q} of the j -invariant of the complex elliptic curve \mathbf{C}/\mathcal{O} . This polynomial is called the *Hilbert class polynomial*. It is a non-trivial fact that P_{Δ}^j has *integer* coefficients.

The fact that ring class fields are closely linked to j -invariants of elliptic curves has its ramifications outside the context of explicit class field theory. Indeed, if we let p denote a prime that is not inert in \mathcal{O} , then the observation that the roots in $\overline{\mathbf{F}}_p$ of $P_{\Delta}^j \in \mathbf{F}_p[X]$ are j -invariants of elliptic curves over $\overline{\mathbf{F}}_p$ with endomorphism ring \mathcal{O} made computing P_{Δ}^j a key ingredient in the elliptic curve primality proving algorithm [12]. Fast algorithms to compute P_{Δ}^j are also desirable from a cryptographic point of view. For example, computing P_{Δ}^j allows us to efficiently construct elliptic curves for which the discrete logarithm problem is presumed to be hard; cf. [6, Chapter 23].

There are currently three known algorithms to compute $P_{\Delta}^j \in \mathbf{Z}[X]$: a complex analytic [8], a p -adic [4, 7] and a ‘multi prime’ approach [1, 2, 22]. If the Generalized Riemann Hypothesis (GRH) holds true, the run time of all three algorithms is $\tilde{O}(|\Delta|)$; see [2]. Here the \tilde{O} -notation indicates that factors that are of logarithmic order in the main term have been disregarded. These three algorithms are efficient in the sense that the bottleneck for each algorithm is the size of the output. We are therefore inherently limited to ‘small’ discriminants. However, also for small discriminants, the coefficients of P_{Δ}^j are *huge*. For $\Delta = -23$, we obtain

$$P_{-23}^j = X^3 + 3491750X^2 - 5151296875X + 12771880859375$$

for instance. History tells us we should be able to do better. In his *Lehrbuch der Algebra* [24] from 1908, Weber introduces a modular function \mathfrak{f} from the upper half-plane \mathbf{H} to \mathbf{C} with the property that, for a suitable choice of ω , the value $\mathfrak{f}(\omega)$ generates the ring class field of $\mathbf{Z}[\omega]$ for all imaginary quadratic orders $\mathbf{Z}[\omega]$ in which 3 is unramified and 2 splits completely.

Received 11 June 2009; revised 16 April 2010.

2000 *Mathematics Subject Classification* 11G15 (primary), 11Y40 (secondary).

For $\Delta = -23$, a root of the polynomial

$$P_{-23}^f = X^3 - X^2 + 1$$

generates the Hilbert class field of $\mathbf{Q}(\sqrt{-23})$. Weber’s function f is related to the j -function via $(f^{24} - 16)^3 - jf^{24} = 0$, so computing P_{Δ}^f has the same cryptographic applications as computing P_{Δ}^j .

Following Weber, we call a function value $f(\omega)$ of a modular function $f : \mathbf{H} \rightarrow \mathbf{C}$ a *class invariant* if we have

$$K(f(\omega)) = H_{\mathbf{Z}[\omega]},$$

that is, if it generates the ring class field over $K = \mathbf{Q}(\omega)$. The logarithmic height of the coefficients of the minimal polynomial P_{Δ}^f of a class invariant is a *constant* factor smaller than the coefficients of P_{Δ}^j . This means that from a purely asymptotic point of view, there is no advantage in computing P_{Δ}^f instead of P_{Δ}^j : the difference in run time is absorbed in the O -constant. The example above shows, however, that from a more practical point of view, class invariants give a big improvement.

The theory of class invariants is well understood *in the complex analytic setting*. Using complex analytic techniques, it is now a rather mechanical process [21] to decide if $f(\omega)$ is a class invariant, and if so compute its minimal polynomial P_{Δ}^f .

In this paper we explain how to work with class invariants over *non-archimedean* fields \mathbf{Q}_p . The functions we will use are integral over $\mathbf{Z}[j]$, and most of our computations will take place in the ring \mathbf{Z}_p . Computing over \mathbf{Z}_p instead of over \mathbf{C} has the advantage that rounding errors cannot occur when computing the minimal polynomial P_{Δ}^f of a class invariant. This gives our approach an edge over the complex analytic approach. Our computer experiments indicate that our algorithm is also reasonably fast in practice. Without trying to write highly optimized code, we computed the polynomial

$$P_{-92806391}^f \in \mathbf{Z}[X]$$

for the Weber- f function and an order of discriminant roughly -10^8 in roughly 15 minutes on our standard 32-bit, 2.8 GHz PC. In comparison, it took the computer algebra package Magma many hours to compute this polynomial using a complex analytic algorithm on the same PC. We expect that with optimized code, our algorithm has the potential to obtain similar timings as in [8] where an *optimized* version of the complex analytic approach is used.

As a by-product of our algorithm, we obtain an algorithm to compute P_{Δ}^f modulo a prime p and we can in turn use this information (for certain functions f) to compute P_{Δ}^f using a ‘Chinese remainder theorem’ (CRT)-approach as in [2]. The CRT-approach is currently the fastest method [22] for computing the Hilbert class polynomial P_{Δ}^j , and we expect that the ‘one p -adic digit version’ of our algorithm can be used to give a fast CRT-approach to compute P_{Δ}^f . This paper solely focuses on computing P_{Δ}^f over the p -adics, however.

Our treatment of class invariants is of a more geometric nature than the complex analytic treatment. To keep the geometry manageable, we will mostly restrict ourselves to modular functions f of level N with the property that the natural map

$$f : \Gamma(f) \backslash \mathbf{H} \rightarrow \mathbf{A}^1$$

induced by the inclusion $\Gamma(f) \subseteq \mathrm{SL}_2(\mathbf{Z})$ has *degree one*. Here, $\Gamma(f) \subseteq \mathrm{SL}_2(\mathbf{Z})$ denotes the stabilizer of f inside the special linear group $\mathrm{SL}_2(\mathbf{Z})$. Examples of such functions include the aforementioned Weber function f and a cube root γ_2 of the j -function. If f has degree one, then we can rigorously prove that our approach works. If f has larger degree, then we need to rely on *heuristics* to prove the correctness of our algorithm. We make the heuristics precise in Section 5.

The main algorithmic ingredient of our algorithm is the *modular polynomial* Φ_l^f relating the complex analytic functions $f(z)$ and $f(lz)$ for a prime l not dividing the level N . These

polynomials Φ_l^f are a generalization of the classical modular polynomial for the j -function. In Section 5 we give the geometric interpretation of Φ_l^f and prove reduction properties of the curve $\Phi_l^f = 0$.

Our algorithm is an extension of the p -adic algorithm for the j -function [4, 7], which we briefly recall in Section 2. In Section 3 we recall properties of the modular function field and give a ‘weak version’ of Shimura reciprocity linking modular functions and ring class fields. The geometric approach to class invariants is developed in Sections 4 and 5, and the resulting algorithm is stated in Section 6. We illustrate the algorithm with a detailed example in Section 7.

2. p -adic algorithm for the j -function

In this section we explain the p -adic algorithm to compute, on input of a discriminant $\Delta < -4$, the Hilbert class polynomial $P_\Delta^j \in \mathbf{Z}[X]$. For more details, proofs, and examples, see [4, 7].

As before, let $\mathcal{O} \subset K = \mathbf{Q}(\sqrt{\Delta})$ be the imaginary quadratic order of discriminant Δ . Let p be a ‘small’ prime that splits completely in the ring class field $H_{\mathcal{O}}$. Since a prime splits completely in the ring class field if and only if it splits into principal primes in \mathcal{O} , we can find such p by looking for an integer solution to the equation

$$4p = x^2 - \Delta y^2 \tag{2.1}$$

with p prime. Under GRH, we may take p of size $\tilde{O}(|\Delta|)$ by [4, Lemma 3.1].

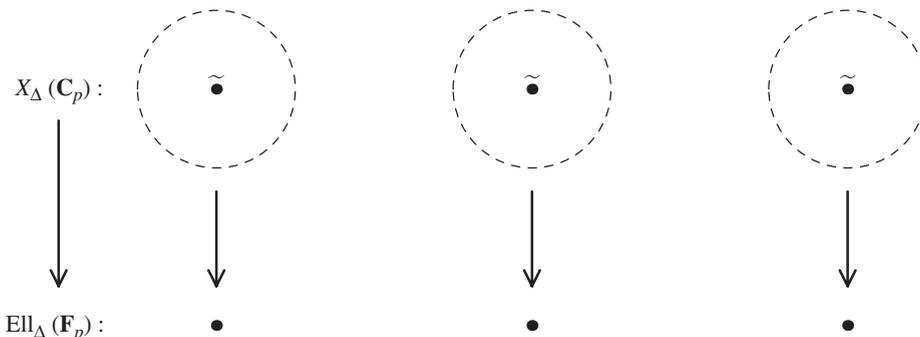
The set

$$\text{Ell}_\Delta(\mathbf{Q}_p) = \{j \in \mathbf{Q}_p \mid \exists E/\mathbf{Q}_p \text{ with } j(E) = j \text{ and } \text{End}(E) = \mathcal{O}\}$$

is a finite set of cardinality $\#\text{Pic}(\mathcal{O})$, and consists of the roots of $P_\Delta^j \in \mathbf{Q}_p[X]$. The set $\text{Ell}_\Delta(\mathbf{F}_p)$ is defined similarly. Its elements are the $\#\text{Pic}(\mathcal{O})$ j -invariants of elliptic curves over \mathbf{F}_p with endomorphism ring \mathcal{O} , or equivalently, the roots of $P_\Delta^j \in \mathbf{F}_p[X]$.

We first construct an (ordinary) elliptic curve \bar{E}/\mathbf{F}_p with $j(\bar{E}) \in \text{Ell}_\Delta(\mathbf{F}_p)$. Write $N = p + 1 - x$ with x as in equation (2.1). We try *random* curves over \mathbf{F}_p until we find a curve \bar{E} with N points. The subring $\mathbf{Z}[\text{Frob}] \subseteq \mathcal{O}$ generated by the Frobenius morphism then has index $y \geq 1$. For cryptographic applications, we are mostly interested in the case that \mathcal{O} is the maximal order. If we have $y = 1$ in this case, we know that the curve \bar{E} has endomorphism ring \mathcal{O} . For the general case, the equality $\#\bar{E}(\mathbf{F}_p) = N$ does not imply that \bar{E} has endomorphism ring \mathcal{O} . We then compute the endomorphism ring $\text{End}(\bar{E})$ using [16] and apply an isogeny of degree dividing $[\mathcal{O}_K : \mathbf{Z}[\text{Frob}]$ to find an elliptic curve with endomorphism ring \mathcal{O} . We refer to [4, Section 3] for details, and fix an elliptic curve \bar{E}/\mathbf{F}_p with $\text{End}(\bar{E}) = \mathcal{O}$ for the remainder of this section.

We want to lift $j(\bar{E}) \in \mathbf{F}_p$ to $\tilde{j} \in \text{Ell}_\Delta(\mathbf{Q}_p)$. Let \mathbf{C}_p be the completion of an algebraic closure of \mathbf{Q}_p , and put $X_\Delta(\mathbf{C}_p) = \{j \in \mathbf{C}_p \mid j \bmod p \in \text{Ell}_\Delta(\mathbf{F}_p)\}$. The set $X_\Delta(\mathbf{C}_p)$ consists of $\#\text{Pic}(\mathcal{O})$ discs of p -adic radius 1. Each disc contains exactly one element of the set $\text{Ell}_\Delta(\mathbf{Q}_p)$ that we want to compute.



Let E/\mathbf{Q}_p be an elliptic curve with endomorphism ring \mathcal{O} . For an invertible ideal $I \subset \mathcal{O}$, there exists an elliptic curve E^I/\mathbf{Q}_p and a separable isogeny $E \rightarrow E^I$ which has the subgroup $E[I]$ of I -torsion points as kernel. We obtain a bijection $\rho_I : \text{Ell}_\Delta(\mathbf{Q}_p) \rightarrow \text{Ell}_\Delta(\mathbf{Q}_p)$ that sends $j(E)$ to $j(E^I)$. The fundamental idea in [7] is that the map ρ_I has a natural extension to a map

$$\rho_I : X_\Delta(\mathbf{C}_p) \rightarrow X_\Delta(\mathbf{C}_p)$$

for invertible \mathcal{O} -ideals I that are coprime to p . For principal ideals I , the map ρ_I has the set $\text{Ell}_\Delta(\mathbf{Q}_p)$ as unique fixed points.

To define $\rho_I(j)$ for $j \in X_\Delta(\mathbf{C}_p)$, choose an elliptic curve E/\mathbf{C}_p with $j(E) = j$ that has good reduction modulo p . Assume that I is coprime to p , and let $l \in \mathbf{Z}_{>0}$ be its norm. The reduced curve E'/\mathbf{F}_p has endomorphism ring \mathcal{O} , and the subgroup $E'[I] \subset E'[l]$ lifts canonically to a subgroup $S \subset E[l]$. We put $\rho_I(j) = j(E/S)$. Note that for $j \in \text{Ell}_\Delta(\mathbf{Q}_p)$ we have $S = E[I]$, so ρ_I is indeed an extension of the map defined on $\text{Ell}_\Delta(\mathbf{Q}_p)$.

One proves [4, Theorem 4.2] that (for $\Delta < -4$) for principal ideals $I = (\alpha) \not\subset \mathbf{Z}$, the map $\rho_\alpha : X_\Delta(\mathbf{C}_p) \rightarrow X_\Delta(\mathbf{C}_p)$ is analytic, that is, it can locally be given by a power series. In this case, the derivative at $\tilde{j} \in \text{Ell}_\Delta(\mathbf{Q}_p)$ is given by $\alpha/\bar{\alpha}$; cf. [4, Lemma 4.3]. Here, $\bar{\alpha}$ denotes the complex conjugate of α . If $j_1 \in \mathbf{C}_p$ denotes any integral lift of $j(\bar{E}) \in \text{Ell}_\Delta(\mathbf{F}_p)$, the ‘Newton process’

$$j_{k+1} = j_k - \frac{\rho_\alpha(j_k) - j_k}{(\alpha/\bar{\alpha}) - 1} \quad \text{for } k \in \mathbf{Z}_{\geq 1} \tag{2.2}$$

converges to the canonical lift $\tilde{j} \in \text{Ell}_\Delta(\mathbf{Q}_p)$ if $\alpha/\bar{\alpha} - 1$ is a p -adic unit. For $k = 1$ the computation is performed with two p -adic digits of precision, and the precision is doubled at each step. The accuracy required for the computation of P_Δ^j can be explicitly bounded [4, Section 7].

The run time of the lifting phase depends heavily on the choice of α . The equality $\rho_{IJ} = \rho_I \rho_J$ shows that we want α to be *smooth*, that is, only divisible by primes of ‘small’ norm. Then ρ_α factors as a product of maps, corresponding to the prime divisors of (α) , that are quicker to compute. The smoothness properties are ‘in practice’ a lot better than what can be rigorously proved [7, Lemma 2]. At the end of this section we give more details on the explicit computation of ρ_α .

Once we have computed the canonical lift with a high enough accuracy of n p -adic digits, we need to compute its conjugates under the action of the Picard group $\text{Pic}(\mathcal{O})$. This can be done using the same techniques as before, since the action of an ideal class $[I] \in \text{Pic}(\mathcal{O})$ is given by

$$j(E) \mapsto j(E^I) = \rho_I(j(E)).$$

We compute small generators of the Picard group, and compute the Galois conjugates of \tilde{j} . In the end we expand the Hilbert class polynomial

$$P_\Delta^j = \prod_{[I] \in \text{Pic}(\mathcal{O})} (X - \rho_I(\tilde{j})) \in (\mathbf{Z}_p/p^n)[X]$$

and lift the coefficients to integers between $-p^n/2$ and $p^n/2$.

We explain how to explicitly compute the map ρ_I . It suffices to show how to treat the case that I is a prime ideal, and we let $l \neq p$ be its norm. For $j(\bar{E}) \in \text{Ell}_\Delta(\mathbf{F}_p)$, the isogeny $\bar{E} \rightarrow \bar{E}^I$ has degree l . Let $\Phi_l(X, Y) \in \mathbf{Z}[X, Y]$ be the classical *modular polynomial* of level l . It is a singular model for the modular curve $X_0(l)$ parametrizing (cyclic) l -isogenies. This means that $j(\bar{E}^I) \in \mathbf{F}_p$ is a root of $\Phi_l(X, j(\bar{E})) \in \mathbf{F}_p[X]$. Under the mild condition that the l -torsion $\bar{E}[l]$ is not \mathbf{F}_p -rational, this polynomial has only two roots in \mathbf{F}_p by [4, Theorem 5.1], namely $j(\bar{E}^I)$ and $j(\bar{E}^{\bar{I}})$. Choose a root h .

We need to decide if we have $\rho_I(j(\bar{E})) = h$ or not. Let \bar{E}/S have j -invariant h , corresponding to a cyclic subgroup $S \subset \bar{E}[l]$ of order l , that is, S is the kernel of the isogeny $\bar{E} \rightarrow \bar{E}/S$.

Choosing a Weierstraß equation

$$Y^2 = X^3 + aX + b$$

for \bar{E} , the techniques that Elkies used to improve Schoof’s original point counting algorithm [19, Sections 7 and 8] allow us to compute a polynomial $f_S \in \mathbf{F}_p[X]$ that vanishes exactly on the x -coordinates of the points in S .

Write $I = (l, c + d\pi_p)$ with $\pi_p \in \mathcal{O}$ an element of norm p . Then the group $\bar{E}[I]$ is an eigenspace for the action of Frobenius with eigenvalue $-c/d \in \mathbf{F}_l$. We now test if $(X^p, Y^p) = -c/d \cdot (X, Y)$ holds for the points in S , that is, we compute both (X^p, Y^p) and $(-c/d) \cdot (X, Y)$ in the ring

$$\mathbf{F}_p[X, Y]/(f_S(X), Y^2 - X^3 - aX - b).$$

If they are equal, we have $h = \bar{\rho}_I(j(\bar{E}))$. Otherwise, we need to take the other root of $\Phi_l(j(\bar{E}), X)$.

We have ‘decomposed’ the map $\bar{\rho}_\alpha : \text{Ell}_\Delta(\mathbf{F}_p) \rightarrow \text{Ell}_\Delta(\mathbf{F}_p)$ as a cycle of isogenies. Using modular polynomials, it is a simple matter to lift this cycle of maps

$$j(\bar{E}) \xrightarrow{I_1} j(\bar{E}^{I_1}) \longrightarrow \dots \xrightarrow{I_n} j(\bar{E}^{(\alpha)}) = j(\bar{E})$$

over \mathbf{F}_p to a ‘cycle’ $j_k \longrightarrow \dots \longrightarrow \rho_\alpha(j_k)$ over \mathbf{Q}_p . Indeed, we know that $\Phi_l(j_k, X) \in \mathbf{Z}_p[X]$ only has two roots in \mathbf{Z}_p , and since we know $j(\bar{E}^{I_1}) \in \mathbf{F}_p$, we know which root is $\rho_I(j_k)$. This enables us to compute the map ρ_α on $X_\Delta(\mathbf{C}_p)$.

3. Shimura reciprocity over the ring class field

Let $f : \mathbf{H} \rightarrow \mathbf{C}$ be a modular function. If we evaluate f at a generator ω of the \mathbf{Z} -algebra $\mathcal{O} = \mathbf{Z}[\omega]$, then the result $f(\omega)$ will typically lie in an extension field of the ring class field $H_\mathcal{O}$; see Theorem 3.2 below. However, in special cases it turns out that $f(\omega)$ does lie in $H_\mathcal{O}$ and generates the extension $H_\mathcal{O}/K$. Following Weber, we call $f(\omega)$ a class invariant in this case.

The example in the introduction shows that the minimal polynomial of a class invariant $f(\omega)$ can be a lot smaller than the Hilbert class polynomial. To quantify the improvement we obtain by using the modular function f instead of j , we define the reduction factor

$$r(f) = \frac{\deg_f(\Psi_f)}{\deg_j(\Psi_f)},$$

where Ψ_f is an irreducible polynomial with $\Psi_f(j, f) = 0$. By [14, Proposition B.3.5], the value $r(f)$ is, asymptotically, the inverse of the quotient

$$\lim_{h(j(\tau)) \rightarrow \infty} \frac{h(f(\tau))}{h(j(\tau))}.$$

Here, h is the absolute logarithmic height, and we take the limit over all complex multiplication points $\text{SL}_2(\mathbf{Z}) \cdot \tau \in \mathbf{H}$. We see that $r(f)$ is a good measure for the improvement we obtain by computing the minimal polynomial of a class invariant $f(\omega)$. We have $r(\gamma_2) = 3$, and the value $r(f) = 72$ is close to optimal in view of the upper bound $r(f) \leq 101$ proved in [5]. We refer to [9] for an overview of the ‘available functions’ and their reduction factors. In this section we explain a method, due to Shimura, that enables us to decide if a modular function f yields class invariants for a given imaginary quadratic order \mathcal{O} .

For an integer $N > 0$, let

$$\Gamma(N) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{SL}_2(\mathbf{Z}) \mid a \equiv d \equiv 1 \pmod{N}, b \equiv c \equiv 0 \pmod{N} \right\}$$

be the full congruence subgroup of level N . The modular group $\text{SL}_2(\mathbf{Z})$ acts on the complex upper half-plane \mathbf{H} and its completion $\bar{\mathbf{H}} = \mathbf{H} \cup \mathbf{P}^1(\mathbf{Q})$ by fractional linear transformations.

The quotient $\Gamma(N)\backslash\overline{\mathbf{H}}$ has the structure of a compact Riemann surface, and as such, it is isomorphic to the modular curve $X(N)$ over \mathbf{C} .

It is well known that the modular curve $X(N)$ can be defined over the cyclotomic field $\mathbf{Q}(\zeta_N)$, where ζ_N is a primitive N th root of unity. Let F_N be the function field of $X(N)$ over $\mathbf{Q}(\zeta_N)$. We have $F_1 = \mathbf{Q}(j)$. Elements of F_N are called *modular functions of level N* . Explicitly, a function $f : \overline{\mathbf{H}} \rightarrow \mathbf{C}$ is called modular if it is invariant under $\Gamma(N)$ and if the coefficients of its Fourier expansion (which it has because $\begin{pmatrix} 1 & \\ 0 & 1 \end{pmatrix}$ is an element of $\Gamma(N)$) lie in $\mathbf{Q}(\zeta_N)$.

Define the function m by

$$m(w, \tau) = -2^7 3^5 \cdot \frac{g_2(\tau)g_3(\tau)}{\Delta(\tau)} \wp(w; \langle 1, \tau \rangle)$$

for $w \in \mathbf{C}$ and $\tau \in \mathbf{H}$. Here, $\wp(\cdot; \langle 1, \tau \rangle)$ is the Weierstraß \wp -function associated to the lattice $\mathbf{Z} + \mathbf{Z} \cdot \tau$. For $r, s \in (1/N)\mathbf{Z}/\mathbf{Z}$, not both 0, define the *Fricke function* $f_{r,s}$ of level N by

$$f_{r,s}(\tau) = m(rN + s, \tau).$$

The Fourier coefficients of $f_{r,s}$ are contained in $\mathbf{Q}(\zeta_N)$. If we fix τ and let r, s vary over $(1/N)\mathbf{Z}/\mathbf{Z}$, not both equal to 0, we obtain the normalized x -coordinates of the $N^2 - 1$ non-trivial points of order N of the complex elliptic curve $\mathbf{C}/(\mathbf{Z} + \mathbf{Z} \cdot \tau)$.

THEOREM 3.1. *We have*

$$F_N = \mathbf{Q}\left(j, f_{r,s} \mid r, s \in \frac{1}{N}\mathbf{Z}/\mathbf{Z}, \text{ not both } 0\right).$$

Proof. See [17, Theorem 6.2 and the beginning of § 6.3]. □

The extension F_N/F_1 is Galois with group $\text{GL}_2(\mathbf{Z}/N\mathbf{Z})/\{\pm 1\}$. This combines the geometric $\text{SL}_2(\mathbf{Z}/N\mathbf{Z})/\{\pm 1\}$ -action coming from the Galois cover $X(N)_{\mathbf{C}}/X(1)_{\mathbf{C}}$ with the arithmetic $(\mathbf{Z}/N\mathbf{Z})^*$ -action coming from $\mathbf{Q}(\zeta_N)/\mathbf{Q}$. Here, $\sigma_d : \zeta_N \mapsto \zeta_N^d$ acts on a modular function $f = \sum_k c_k \cdot q^{k/N}$ via

$$f^{\sigma_d} = \sum_k \sigma_d(c_k) \cdot q^{k/N}. \tag{3.1}$$

Let $\widehat{\mathcal{O}} = \mathcal{O} \otimes_{\mathbf{Z}} \widehat{\mathbf{Z}}$ be the profinite completion of the imaginary quadratic order $\mathcal{O} \subset K$. Class field theory tells us that the Galois group $\text{Gal}(K^{\text{ab}}/H_{\mathcal{O}})$ of the maximal abelian extension of $H_{\mathcal{O}}$ is given by the exact sequence

$$1 \longrightarrow \mathcal{O}^* \longrightarrow \widehat{\mathcal{O}}^* \xrightarrow{\text{Artin}} \text{Gal}(K^{\text{ab}}/H_{\mathcal{O}}) \longrightarrow 1.$$

We obtain K^{ab} as the union of finite extensions $H_{N,\mathcal{O}}$ corresponding to the finite quotients

$$\widehat{\mathcal{O}}^* \twoheadrightarrow (\widehat{\mathcal{O}}/N\widehat{\mathcal{O}})^* = (\mathcal{O}/N\mathcal{O})^*.$$

The field $H_{N,\mathcal{O}}$ is called the *ray class field of conductor N for the order \mathcal{O}* , and the Artin map gives an isomorphism

$$(\mathcal{O}/N\mathcal{O})^*/\text{Im}[\mathcal{O}^*] \xrightarrow{\sim} \text{Gal}(H_{N,\mathcal{O}}/H_{\mathcal{O}}).$$

If \mathcal{O} is the maximal order of K , the field $H_{N,\mathcal{O}}$ is the ray class field of conductor N of K .

THEOREM 3.2. *Let $f \in F_N$ be modular of level N , and write $\mathcal{O} = \mathbf{Z}[\omega]$. If $f(\omega)$ is finite, we have $f(\omega) \in H_{N,\mathcal{O}}$.*

Proof. See [17, Chapter 10]. □

Let the notation be as in Theorem 3.2 above. Then we have $f(\omega) \in H_{\mathcal{O}} \subseteq H_{N,\mathcal{O}}$ if $f(\omega)$ is a class invariant. To decide if this is the case for a given value $f(\omega)$, we need to know the Galois action of $\text{Gal}(H_{N,\mathcal{O}}/H_{\mathcal{O}})$ on values of modular functions. A variant of Shimura reciprocity, described below, enables us to compute this action.

Let $g_{\omega} : (\mathcal{O}/N\mathcal{O})^* \rightarrow \text{GL}_2(\mathbf{Z}/N\mathbf{Z})$ be the map that sends x to the transpose of the matrix representing multiplication by x on the $(\mathbf{Z}/N\mathbf{Z})$ -module $\mathbf{Z}/N\mathbf{Z}\omega + \mathbf{Z}/N\mathbf{Z}$ with respect to the basis $[\omega, 1]$. The map g_{ω} connects the rows in the following diagram.

$$\begin{array}{ccccccc} \mathcal{O}^* & \longrightarrow & (\mathcal{O}/N\mathcal{O})^* & \xrightarrow{\text{Artin}} & \text{Gal}(H_{N,\mathcal{O}}/H_{\mathcal{O}}) & \longrightarrow & 1 \\ & & \downarrow g_{\omega} & & & & \\ \{\pm 1\} & \longrightarrow & \text{GL}_2(\mathbf{Z}/N\mathbf{Z}) & \longrightarrow & \text{Gal}(F_N/\mathbf{Q}(j)) & \longrightarrow & 1 \end{array}$$

Explicitly, if ω has minimal polynomial $X^2 + bX + c \in \mathbf{Z}[X]$, then we have

$$g_{\omega} : x = s\omega + t \mapsto \begin{pmatrix} t - bs & -cs \\ s & t \end{pmatrix}.$$

The content of Shimura’s reciprocity law is that the Galois conjugate $f(\omega)^x$ of $f(\omega)$ under the Artin symbol $\text{Artin}(x) \in \text{Gal}(H_{N,\mathcal{O}}/H_{\mathcal{O}})$ may be computed via the reciprocity relation

$$(f(\omega))^x = (f^{g_{\omega}(x^{-1})})(\omega);$$

cf. [20, Theorem 6.31]. If the extension $F_N/\mathbf{Q}(f)$ is Galois, we have the fundamental equivalence

$$(f(\omega))^x = f(\omega) \iff f^{g_{\omega}(x)} = f.$$

The implication \Leftarrow is immediate from the reciprocity relation. The other implication requires the hypothesis and an additional argument [20, Proposition 6.33].

We compute generators x_1, \dots, x_k for $(\mathcal{O}/N\mathcal{O})^*$ and map them to $\text{GL}_2(\mathbf{Z}/N\mathbf{Z})$ using the map g_{ω} . The value $f(\omega)$ is contained in the ring class field $H_{\mathcal{O}}$ if and only if $g_{\omega}(x_1), \dots, g_{\omega}(x_k)$ act trivially on f . If we, for instance, also know that there is an inclusion $\mathbf{Q}(j) \subseteq \mathbf{Q}(f)$, then $f(\omega)$ is also a class invariant if $g_{\omega}(x_1), \dots, g_{\omega}(x_k)$ act trivially on f . We refer to [13] for examples.

4. Class invariants over \mathbf{Q}_p

In this section we extend the p -adic algorithm from Section 2 to work with modular functions other than the j -function. The description we present in this section is not ideally suited for explicit computations yet, and serves as a stepping stone for the more practical version in Section 5. Throughout this section, we fix a modular function f of level $N \geq 1$ that is integral over the ring $\mathbf{Z}[j]$, that is, f is a root of some monic irreducible polynomial $\Psi_f(X) \in (\mathbf{Z}[j])[X]$. All known modular functions yielding class invariants are integral. We do not assume that f has degree one in this section.

As before, let p be a prime that splits completely in the ring class field $H_{\mathcal{O}}$ for the order \mathcal{O} of discriminant $\Delta < -4$. For a j -value $j(\tilde{E}) \in \text{Ell}_{\Delta}(\mathbf{Q}_p)$, the roots of the polynomial $\Psi_f(X, j(\tilde{E})) \in H_{\mathcal{O}}[X]$ lie in the ray class field of conductor $H_{N,\mathcal{O}}$ of conductor N for the order \mathcal{O} ; cf. Theorem 3.2. If we know that f yields class invariants, for instance by using Shimura reciprocity, we know that some of these roots actually lie in the ring class field $H_{\mathcal{O}}$. We need to decide which ones, and compute the action of the Galois group $\text{Gal}(H_{\mathcal{O}}/K) \cong \text{Pic}(\mathcal{O})$ on such roots.

The key observation is that f is an element of the function field

$$F_N = \mathbf{Q}\left(j, f_{r,s} \mid r, s \in \frac{1}{N}\mathbf{Z}/\mathbf{Z}, \text{ not both } 0\right)$$

of the modular curve $X(N)$ over $\mathbf{Q}(\zeta_N)$; cf. Theorem 3.1. The Fricke functions $f_{r,s}$ are normalized x -coordinates of N -torsion of points on the elliptic curve $\mathbf{C}/(\mathbf{Z} + \mathbf{Z} \cdot \tau)$, and we can write f as a \mathbf{Q} -rational function in j and the functions $f_{r,s}$.

Fix a primitive N th root of unity $\zeta_N \in \overline{\mathbf{Q}_p}$. For $a \in (\mathbf{Z}/N\mathbf{Z})^*$, let $Y(N)_a$ be the modular curve parametrizing isomorphism classes of triples

$$(E, P, Q)$$

where $P, Q \in E[N]$ form a basis that maps to ζ_N^a under the Weil pairing e_N . Then f is an element of the function field of $Y(N)_a(\overline{\mathbf{Q}_p})$ for every $a \in (\mathbf{Z}/N\mathbf{Z})^*$. Let $x \in \mathbf{Q}_p$ be a root of $\Psi_f(X, j(\tilde{E})) \in \mathbf{Q}_p[X]$. There exist $a \in (\mathbf{Z}/N\mathbf{Z})^*$ and $(\tilde{E}, P, Q) \in Y(N)_a(\overline{\mathbf{Q}_p})$ with $f(\tilde{E}, P, Q) = x \in \mathbf{Q}_p$.

The group of invertible \mathcal{O} -ideals acts on $\text{Ell}_\Delta(\mathbf{Q}_p)$ via $j(E) \mapsto j(E^I)$, and the action of the Artin symbol $[I, H_{\mathcal{O}}/K]$ of I for the extension $H_{\mathcal{O}}/K$ satisfies $j(\tilde{E})^{[I, H_{\mathcal{O}}/K]} = j(\tilde{E}^I)$. If N is coprime to the norm l of I , the isogeny

$$\varphi_I : \tilde{E} \longrightarrow \tilde{E}^I$$

extends to a natural isomorphism

$$\varphi_I : \tilde{E}[N] \xrightarrow{\sim} \tilde{E}^I[N].$$

A basis $\langle P, Q \rangle$ for $\tilde{E}[N]$ gets mapped to a basis $\langle P^I, Q^I \rangle$ for $\tilde{E}^I[N]$. We compute

$$e_N(P^I, Q^I) = e_N(P, \hat{\varphi}_I(Q^I)) = e_N(P, lQ) = \zeta_N^l$$

and conclude that we have $(\tilde{E}^I, P^I, Q^I) \in Y(N)_{la}(\overline{\mathbf{Q}_p})$. We have the fundamental equality

$$f(\tilde{E}, P, Q)^{[I, H_{N, \mathcal{O}}/K]} = f(\tilde{E}^I, P^I, Q^I).$$

We can explicitly compute the isogeny φ_I : first we compute the kernel polynomial $g_I \in \mathbf{Q}_p[X]$ corresponding to I using the ‘Atkin–Elkies techniques’ alluded to in Section 2 and then we compute the isogeny using Vélú’s formulas [23]. Hence, we have a way of computing $f(\tilde{E}, P, Q)^{[I, H_{N, \mathcal{O}}/H_{\mathcal{O}}]}$.

A root $x \in \mathbf{Q}_p$ of $\Psi_f(X, j(\tilde{E})) \in \mathbf{Q}_p[X]$ lies in $H_{\mathcal{O}}$ if and only if it is invariant under the action of

$$\text{Gal}(H_{N, \mathcal{O}}/H_{\mathcal{O}}) \cong (\mathcal{O}/N\mathcal{O})^*/\text{Im}(\mathcal{O}^*).$$

We write $x = f(\tilde{E}, P, Q)$ for some choice of basis $P, Q \in \tilde{E}[N]$ and test whether $x^{[y, H_{N, \mathcal{O}}/H_{\mathcal{O}}]} = x$ holds for all generators y of $(\mathcal{O}/N\mathcal{O})^*/\text{Im}(\mathcal{O}^*)$.

Once we have found that a certain root $x \in \mathbf{Q}_p$ lies in the ring class field $H_{\mathcal{O}}$, we need to compute its conjugates under $\text{Gal}(H_{\mathcal{O}}/K) \cong \text{Pic}(\mathcal{O})$. This proceeds exactly as before, since we have

$$x^{[I, H_{\mathcal{O}}/K]} = f(\tilde{E}, P, Q)^{[I, H_{\mathcal{O}}/K]} = f(\tilde{E}^I, P^I, Q^I) \in \mathbf{Q}_p$$

for invertible \mathcal{O} -ideals I that are coprime to the level N . If the minimal polynomial of x has integer coefficients, then the coefficients of

$$P_\Delta^f = \prod_{[I] \in \text{Pic}(\mathcal{O})} (X - f(\tilde{E}, P, Q)^{[I, H_{\mathcal{O}}/K]}) \in \mathbf{Q}_p[X]$$

lie in the subring $\mathbf{Z} \subset \mathbf{Q}_p$. If we know an upper bound on the logarithmic height of the coefficients, then we can compute all conjugates of $x = f(\tilde{E}, P, Q) \in \mathbf{Z}_p/p^n$ with high enough p -adic accuracy and lift the coefficients of $P_\Delta^f \in (\mathbf{Z}_p/p^n)[X]$ to integers between $-p^n/2$ and $p^n/2$, just like we did for the j -function in Section 2.

EXAMPLE. Let $\gamma_2 : \mathbf{H} \rightarrow \mathbf{C}$ be the holomorphic cube root of j with integral Fourier expansion. It is a classical fact that γ_2 is modular of level 3. It yields class invariants for all imaginary quadratic orders \mathcal{O} in which 3 is unramified [24, Section 125].

Let $E : Y^2 = X^3 + aX + b$ be an elliptic curve over \mathbf{Q}_p , with $p > 3$. Let $c_1, \dots, c_4 \in \overline{\mathbf{Q}_p}$ be the roots of the 3-division polynomial of degree $(3^2 - 1)/2 = 4$. Then

$$\frac{-48a}{2a - 3(c_1c_2 + c_3c_4)} \tag{4.1}$$

is a cube root of $j(E)$, as may be checked by using the Fourier expansion of the Fricke functions. Expression (4.1) nicely illustrates that γ_2 is not a function of an elliptic curve alone: some ordering on the 3-torsion is also required. We indeed get three distinct cube roots of $j(E)$. From a geometric point of view, there is no way to single out a root ‘corresponding’ to γ_2 .

We illustrate how we can use this ‘geometric γ_2 ’ to compute the polynomial $P_{-31}^{\gamma_2} \in \mathbf{Z}[X]$ for the order \mathcal{O} of discriminant $\Delta = -31$ using p -adic methods. The primes $47 = 4^2 + 31$ and $67 = 6^2 + 31$ both split completely in the Hilbert class field $H = H_{\mathcal{O}}$ of $K = \mathbf{Q}(\sqrt{-31})$. The case $p = 67$ best illustrates our techniques, since $\tilde{j} \in \text{Ell}_{\Delta}(\mathbf{Q}_p)$ then has three cube roots in \mathbf{Q}_p .

First we compute a curve \tilde{E}/\mathbf{Q}_p with $\text{End}(\tilde{E}) \cong \mathcal{O}$. Since we have $r(\gamma_2) = 3$ for the reduction factor of γ_2 , the accuracy needed is only *one third* of the required nine 67-adic digits accuracy for the computation of the Hilbert class polynomial P_{-31}^j . Using the algorithm from Section 2 we find that we may take

$$j(\tilde{E}) = 3 + 33p - 16p^2 + O(p^3) \in \mathbf{Q}_p$$

as j -invariant. The three cube roots of $j(\tilde{E})$ are

$$\begin{aligned} \eta_1 &= 18 + 26p + 38p^2 + O(p^3) \\ \eta_2 &= 53 + 3p + 30p^2 + O(p^3) \\ \eta_3 &= 63 + 36p + 65p^2 + O(p^3). \end{aligned}$$

Only one of them lies in the Hilbert class field $H \subset \mathbf{Q}_p$. Indeed, if two roots lay in H , then ζ_3 would be contained in H as well and 3 would ramify in H .

We fix a Weierstraß equation

$$Y^2 = X^3 + aX + b$$

for \tilde{E}/\mathbf{Q}_p . Let $c_1, \dots, c_4 \in \overline{\mathbf{Q}_p}$ be the four roots of the 3-division polynomial for \tilde{E} . We compute 3-torsion points P_i with x -coordinates c_i . The points P_i are defined over the unramified extension of degree four of \mathbf{Q}_p .

Let I be an \mathcal{O} -ideal that is coprime to 3. The isogeny $\varphi_I : \tilde{E} \rightarrow \tilde{E}^I$ extends to a natural isomorphism

$$\varphi_I : \tilde{E}[3] \xrightarrow{\sim} \tilde{E}^I[3].$$

Hence, we get a natural bijection

$$\varphi_I : \{\eta_1, \eta_2, \eta_3\} \xrightarrow{\sim} \{\text{cube roots of } j(\tilde{E}^I)\}.$$

For a cube root

$$\eta = \frac{-48a}{2a - 3(c_1c_2 + c_3c_4)}$$

we have

$$\eta^{[I, H_3/H]} = \frac{-48a'}{2a' - 3(c'_1c'_2 + c'_3c'_4)}.$$

Here, c'_i is the x -coordinate of $\varphi_I(P_i) \in \tilde{E}^I[3]$ and \tilde{E}^I has Weierstraß equation $Y^2 = X^3 + a'X + b'$. The group $(\mathcal{O}/3\mathcal{O})^*/\mathcal{O}^* \cong \mathbf{Z}/4\mathbf{Z}$ is generated by $\alpha = (-1 + \sqrt{-31})/2$ of norm 8. We

compute $\eta_i^{[I, H_3/H]}$ for $I = (\alpha) = \mathfrak{p}_2^3$ and obtain

$$\begin{aligned} \eta_1 &\xrightarrow{\varphi_I} \eta_1 \\ \eta_2 &\xrightarrow{\varphi_I} \eta_3 \\ \eta_3 &\xrightarrow{\varphi_I} \eta_2. \end{aligned}$$

Hence, $\eta_1 = 18 + O(p)$ is a class invariant. Note that $\varphi_{\mathfrak{p}_2}$ is just a 2-isogeny, so we do not actually need the ‘Atkin–Elkies’ techniques from [19, Sections 7 and 8].

Computing the conjugates of $\eta_1 \in H_{\mathcal{O}}$ under $\text{Gal}(H_{\mathcal{O}}/K) \cong \text{Pic}(\mathcal{O})$ proceeds similarly. We have $\text{Pic}(\mathcal{O}) \cong \mathbf{Z}/3\mathbf{Z} \cong \langle [\mathfrak{p}_2] \rangle$ and

$$\eta_1^{[\mathfrak{p}_2, H_{\mathcal{O}}/K]} = \varphi_{\mathfrak{p}_2}(\eta_1).$$

We compute the Galois conjugates of η_1 up to three 67-adic digits accuracy and expand

$$P_{-31}^{\gamma_2} = \prod_{i=1}^3 (X - \varphi_{\mathfrak{p}_2}^i(\eta_1)) = X^3 + 342X^2 + 837X + 116127 \in \mathbf{Z}[X].$$

5. Computing the action of invertible ideals

The theory developed in Section 4 is not directly suited for explicit computations. If we are given a modular function f of level N that is integral over $\mathbf{Z}[j]$ as a Fourier expansion, it is not clear how to write this as a rational function in j and the Fricke functions. Secondly, we have to partially factor the N -division polynomial to use the approach from the previous section. The degree of this polynomial is roughly N^2 , and factoring it annihilates the improvement gained by working with a ‘smaller’ function f . In this section we explain how to circumvent these problems if we *restrict* ourselves to functions f for which the natural map

$$f : \Gamma(f) \backslash \mathbf{H} \rightarrow \mathbf{A}^1$$

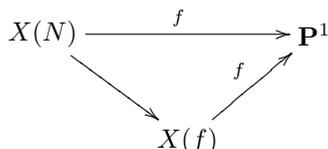
induced by the inclusion $\Gamma(f) \subseteq \text{SL}_2(\mathbf{Z})$ has degree one. Here, $\Gamma(f) \subseteq \text{SL}_2(\mathbf{Z})$ denotes the stabilizer of f inside $\text{SL}_2(\mathbf{Z})$. If f has larger degree, then we have to rely on the heuristics given at the end of the section to prove that our algorithm works.

The crucial observation is that it suffices to compute x^I , where $x \in \mathbf{Q}_p$ is a root of $\Psi_f(X, j(E)) \in \mathbf{Q}_p[X]$ and I is an invertible \mathcal{O} -ideal of norm coprime to N . Indeed, if we want to know which root x of $\Psi_f(X, j(E)) \in \mathbf{Q}_p[X]$ is a class invariant, we need to check which root is invariant under the action of $(\mathcal{O}/N\mathcal{O})^*/\text{Im}(\mathcal{O}^*)$. This amounts to computing x^I for the principal ideals I generated by generators of $(\mathcal{O}/N\mathcal{O})^*$. Once we know that $x \in \mathbf{Q}_p$ is a class invariant, we need to compute $x^I \in \mathbf{Q}_p$ for some choice of generators I of $\text{Pic}(\mathcal{O})$ that are coprime to N .

Before showing how to compute x^I we give some theory regarding modular curves and modular polynomials.

5.1. Modular curves

Let f and $\Gamma(f)$ be as above. We have $\Gamma(N) \subseteq \Gamma(f) \subseteq \text{SL}_2(\mathbf{Z})$, by the assumption that f is modular of level N . Write $X(f)$ for the modular curve corresponding to the congruence subgroup $\Gamma(f)$. The complex points of this curve are $\Gamma(f) \backslash \overline{\mathbf{H}}$. The curve $X(f)$ is a quotient of the modular curve $X(N)$ by a subgroup of $\text{SL}_2(\mathbf{Z}/N\mathbf{Z})$, and can be defined over $\mathbf{Q}(\zeta_N)$. We have a commutative diagram



and $f : X(N) \rightarrow \mathbf{P}^1$ factors through the quotient $X(f)$. Likewise, there exists a curve $X(f)_a$ for every $a \in (\mathbf{Z}/N\mathbf{Z})^*$ such that f factors through $X(N)_a \rightarrow \mathbf{P}^1$. As a complex curve, we have $X(f)_a = X(f)$.

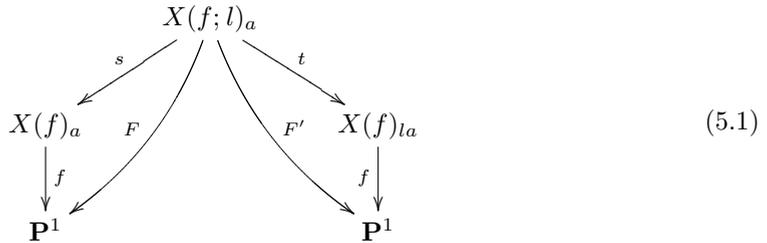
For ease of notation, we simply denote an affine point on $X(f)_a$ by a triple (E, P, Q) instead of $(\overline{E}, \overline{P}, \overline{Q})$. Here, P, Q form a basis for the N -torsion $E[N]$ of E with $e_N(P, Q) = \zeta_N^a$ for a fixed choice of ζ_N .

Let l be a prime not dividing the level N . Writing $\Gamma(f; l) = \Gamma(f) \cap \Gamma_0(l)$, we have inclusions

$$\Gamma(lN) \subseteq \Gamma(f; l) \subseteq \Gamma(f).$$

Let $Y(f; l)$ and $X(f; l)$ be the affine and projective curves corresponding to $\Gamma(f; l)$, respectively. They can be defined over $\mathbf{Q}(\zeta_{lN})$. Just as we have curves $X(f)_a$, we also have curves $X(f; l)_a$. Affine points on $X(f; l)_a$ (or points on $Y(f; l)_a$) are quadruples (E, P, Q, G) , with $(E, P, Q) \in X(f)_a$ and $G \subset E[l]$ a (cyclic) subgroup of order l .

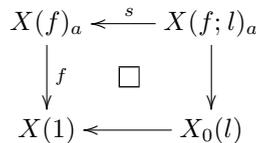
There is a natural map $s : X(f; l)_a \rightarrow X(f)_a$ and a natural map $t : X(f; l)_a \rightarrow X(f)_{la}$. The map s sends $(E, P, Q, G) \in X(f; l)_a$ to $(E, P, Q) \in X(f)_a$. The map t sends $(E, P, Q, G) \in X(f; l)_a$ to $(E/G, \varphi(P), \varphi(Q))$, where $\varphi : E \rightarrow E/G$ has kernel G . The situation is as follows.



Here, F and F' are the composed maps.

LEMMA 5.1. *The maps s, t in diagram (5.1) both have degree $l + 1$.*

Proof. We will show that the diagram

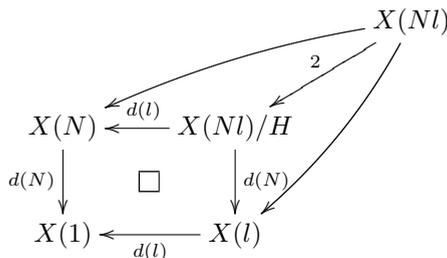


is Cartesian in the category of smooth projective curves with surjective maps. As the cover $X_0(l)/X(1)$ has degree $l + 1$, this implies that s and t have degree $l + 1$. Here, the maps on the ‘lower right part’ of the square are the forgetful maps. Instead of working over $\mathbf{Q}(\zeta_N)$, we will work over \mathbf{C} ; the same result then holds over $\mathbf{Q}(\zeta_N)$. We may then omit the subscript a in the diagram. Moreover, it is easier to work with $X(N)/X(1)$ and $X(l)/X(1)$ instead of $X(f)/X(1)$ and $X_0(l)/X(1)$, since in this case we explicitly know the Galois groups.

The fibre product of $X(l)$ and $X(N)$ is almost equal to $X(Nl)$. Indeed, writing $d(k) = \#(\mathrm{SL}_2(\mathbf{Z}/k\mathbf{Z})/\{\pm 1\})$, the degree of $X(Nl)/X(1)$ is

$$\# \mathrm{SL}_2(\mathbf{Z}/Nl\mathbf{Z})/\{\pm 1\} = 2 \cdot d(N)d(l),$$

and we obtain the diagram



where H is the subgroup $\{1\} \times \{\pm 1\} \subseteq \mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z}) \times \mathrm{SL}_2(\mathbf{Z}/l\mathbf{Z})$. Since we are working over \mathbf{C} , we know that the degrees on parallel sides of the square are equal. Hence, the square is Cartesian.

Since we have $\begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} \in \Gamma_0(l)$, the curve $X(Nl)/H$ is a cover of $X(f; l)$. Hence, the diagram

$$\begin{array}{ccc}
 X(f) & \xleftarrow{s} & X(f; l) \\
 \downarrow f & \square & \downarrow \\
 X(1) & \xleftarrow{\quad} & X_0(l)
 \end{array} \tag{5.2}$$

‘fits inside’ the bigger Cartesian diagram for $X(Nl)/H$. In particular, it is Cartesian. Since the degree of $X_0(l)/X(1)$ is $l + 1$, the same must hold for $X(f; l)/X(f)$. \square

REMARK 5.2. The curve $X_0(l)$ can be defined over \mathbf{Q} . The Cartesian diagram (5.2) shows that $X(f; l)_a$ can be defined over $\mathbf{Q}(\zeta_N)$.

5.2. Modular polynomials

Let $Y(f; l)_a$ be the affine curve defined in Subsection 5.1. We map $Y(f; l)_a$ to a curve C inside $\mathbf{A}^1 \times \mathbf{A}^1$ as in diagram (5.3) below. The map b is defined by $b(x) = (s(x), t(x))$, and the maps p_1, p_2 are the two projection maps.

$$\begin{array}{ccccc}
 & & & & F \\
 & & & & \curvearrowright \\
 Y(f; l)_a & \xrightarrow{b} & C & \xrightarrow{\quad} & \mathbf{A}^1 \times \mathbf{A}^1 & \xrightarrow{p_1} & \mathbf{A}^1 \\
 & \searrow & & & \downarrow p_2 & & \\
 & & & & \mathbf{A}^1 & & \\
 & & & & \uparrow F' & &
 \end{array} \tag{5.3}$$

The function field of C is generated by f and f_l . Here, f_l is as in (3.1) defined by $f_l(\omega) = f^{\sigma_l}(\omega)$. If f has rational Fourier coefficients, we have $f_l(\omega) = f(\omega)$. The minimal polynomial Φ_l^f of f_l over $\mathbf{Q}(\zeta_N)(f)$ is called the *modular polynomial* relating f and f_l . The coefficients of Φ_l^f need not be polynomials in f yet, but after multiplying the coefficients by the common denominator, we obtain a polynomial in $\mathbf{Q}(\zeta_N)[X, Y]$. This polynomial is a model for the curve C . As we have $\deg(F) = \deg(F')$, diagram (5.3) tells us that we have

$$\deg_X(\Phi_l^f) = \deg_Y(\Phi_l^f) = \frac{(l + 1) \deg(f)}{\deg(b)}.$$

REMARK 5.3. For $f = j$, the modular polynomial Φ_l^j is the ‘classical’ modular polynomial Φ_l that we used in Section 2.

LEMMA 5.4. *If f has rational Fourier expansion, then Φ_l^f has rational coefficients.*

Proof. It suffices to show that $X(f)$ can be defined over \mathbf{Q} . Since the algebraic closure of \mathbf{Q} inside $\mathbf{Q}(f, j)$ is \mathbf{Q} itself, the minimal polynomial Ψ_f of f over $\mathbf{Q}(j)$ is absolutely irreducible. The curve defined by $\Psi_f = 0$ is absolutely irreducible and has $\mathbf{Q}(f, j)$ as function field showing that $X(f)$ is defined over \mathbf{Q} . \square

Computing Φ_l^f is relatively easy if we know the Fourier expansion of f . We have an upper bound

$$\deg(f)(l + 1)$$

for the degrees $\deg_X(\Phi_l^f)$ and $\deg_Y(\Phi_l^f)$. By comparing the Fourier coefficients of f and f_l , we can recursively find the coefficients of Φ_l^f . The following lemma often simplifies the computations.

LEMMA 5.5. *Let f be a modular function, and let l be a prime not dividing the level of f . Suppose that the modular polynomial Φ_l^f has integer coefficients. If f is invariant under the action of either $S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \in \mathrm{SL}_2(\mathbf{Z})$ or $M = \begin{pmatrix} 0 & -l \\ 1 & 0 \end{pmatrix} \in \mathrm{GL}_2(\mathbf{Q})$, then we have*

$$\Phi_l^f(X, Y) = \Phi_l^f(Y, X),$$

that is, Φ_l^f is symmetric.

Proof. The proof is similar to the symmetry proof [17, Theorem 5.3] of the classical modular polynomial for the j -function. Assume first that f is invariant under S . If we replace z by $-1/(lz)$ in the equation $\Phi_l^f(f(z), f(lz)) = 0$, we obtain

$$\Phi_l^f(f(-1/(lz)), f(-1/z)) = 0.$$

Using the invariance of f under S , we derive

$$\Phi_l^f(f(lz), f(z)) = 0.$$

Since $\Phi_l^f(X, f)$ is irreducible in $\mathbf{C}[X, Y]$, we see that $\Phi_l^f(f, X)$ is a multiple of $\Phi_l^f(X, f)$. There exists a polynomial $g(X, Y)$ with

$$\Phi_l^f(f, X) = g(X, f)g(f, X)\Phi_l^f(f, X).$$

The Gauß lemma tells us that we have $g(X, Y) \in \mathbf{Z}[X, Y]$ and hence $g(X, Y) = \pm 1$. For $g(X, Y) = -1$, we obtain $\Phi_l^f(X, Y) = -\Phi_l^f(Y, X)$ and $\Phi_l^f(X, X) = 0$. Then $X - Y$ would be a factor of $\Phi_l^f(X, Y)$. This contradicts the irreducibility. Hence, we have $g(X, Y) = 1$ and Φ_l^f is symmetric.

The other case proceeds similarly, one replaces z by $-1/z$ in the beginning of the proof. \square

It is of great help that for many class invariants the coefficients of the modular polynomial Φ_l^f are a lot smaller than those of the classical modular polynomial for j . As an example, we consider the Weber function \mathfrak{f} from the introduction. For small primes l the coefficients of the polynomial are really small, like

$$\Phi_5^{\mathfrak{f}}(X, Y) = (X^5 - Y)(X - Y^5) + 5XY.$$

For $l = 13$ it takes at least two of these journal pages to write down the polynomial $\Phi_{13}^{\mathfrak{f}}$, but we have

$$\begin{aligned} \Phi_{13}^{\mathfrak{f}}(X, Y) &= (X^{13} - Y)(X - Y^{13}) + 5 \cdot 13XY \\ &\quad + 13(X^2Y^{12} + X^{12}Y^2 + 4X^{10}Y^4 + 4X^{10}Y^4 + 6X^6Y^8 + 6X^8Y^6). \end{aligned}$$

5.3. Computing x^I

We now show how to use the modular polynomial Φ_l^f to compute the desired value x^I . As before, we let $x \in \mathbf{Q}_p$ be a root of $\Psi_f(X, j(E)) \in \mathbf{Q}_p[X]$ and let I be an invertible \mathcal{O} -ideal of prime norm $l \nmid Np$. Let Φ_l^f be the modular polynomial defined above. From the moduli interpretation of $X(f; l)_a$, it is clear that one of the roots of $\Phi_l^f(x, X) \in \overline{\mathbf{Q}_p}[X]$ equals x^I . To see what the other roots are, we look at diagram (5.1). Above $x \in \mathbf{A}^1(\overline{\mathbf{Q}_p})$ there are $\deg(f)$ distinct points $(E_i, P_i, Q_i) \in Y(f)_a(\overline{\mathbf{Q}_p})$. Above $(E_i, P_i, Q_i) \in Y(f)_a(\overline{\mathbf{Q}_p})$, there are $l + 1$ points $(E_i, P_i, Q_i, G_j) \in Y(f; l)_a(\overline{\mathbf{Q}_p})$. Here, G_j ranges over the $l + 1$ subgroups of order l of $E_i[l]$. The points (E_i, P_i, Q_i, G_j) all map to $x \in \mathbf{A}^1(\overline{\mathbf{Q}_p})$ under F . The images under $F' : X(f; l)_a \rightarrow \mathbf{A}^1$ are exactly the roots of $\Phi_l^f(x, X)$.

REMARK 5.6. The curve $X(f; l)$ is a quotient of $X(lN)$. Since $X(N)$ has good reduction outside N , the curve $X(f; l)$ has good reduction outside lN by [15, Proposition 4.2]. Hence, the description of the roots of $\Phi_l^f(x, X)$ remains valid over $\overline{\mathbf{F}_p}$.

We need to decide which root of $\Phi_l^f(x, X)$ equals x^I . The first observation is that it suffices to look at the roots in \mathbf{Q}_p . Indeed, if x is a class invariant then we automatically have $x \in \mathbf{Q}_p$. If x is not a class invariant, then x^I need not lie in \mathbf{Q}_p . But if it does not, we have automatically proven that x is not a class invariant.

Usually, x^I is the only root of $\Phi_l^f(x, X)$ that is also a root of $\Psi_f(X, j(E^I))$. Hence, we test for all roots $\alpha \in \mathbf{Q}_p$ of $\Phi_l^f(x, X)$ whether $\Psi_f(\alpha, j(E^I)) = 0$ holds. If x is a class invariant, we find at least one such α . If we find exactly one root with this property, we have computed x^I .

LEMMA 5.7. *Let the notation be as above. Suppose that f has degree one, and let $x \in \mathbf{Q}_p$ be a class invariant. If the l -torsion of E is not \mathbf{Q}_p -rational, then there is exactly one root $\alpha \in \mathbf{Q}_p$ of both $\Phi_l^f(x, X) \in \mathbf{Q}_p[X]$ and $\Psi_f(X, j(E^I)) \in \mathbf{Q}_p[X]$.*

Proof. Let $(E, P, Q) \in Y(f)_a(\overline{\mathbf{Q}_p})$ be the unique point of $Y(f)_a$ with $f(E, P, Q) = x \in \mathbf{Q}_p$. Of the $l + 1$ points $(E, P, Q, G_i) \in Y(f; l)_a(\overline{\mathbf{Q}_p})$ lying over $(E, P, Q) \in Y(f)_a(\overline{\mathbf{Q}_p})$, only for the two points having $G_i = E[I]$ or $G_i = E[\overline{I}]$ is the value $j(t(E, P, Q, G_i))$ contained in \mathbf{Q}_p ; cf. [4, Theorem 5.1]. If both $F'(E, P, Q, E[I])$ and $F'(E, P, Q, E[\overline{I}])$ are roots of $\Phi_l^f(x, X) \in \mathbf{Q}_p[X]$, then we must have $[I] = [\overline{I}] \in \text{Pic}(\mathcal{O})$. Because x is a class invariant, we then have $F'(E, P, Q, E[I]) = F'(E, P, Q, E[\overline{I}])$. □

Assume that f has degree one, which is the case for example for the Weber- f function and the function γ_2 from the introduction. By Lemma 5.7, the polynomials $\Phi_l^f(x, X)$ and $\Psi_f(X, j(E^I))$ then have exactly one root in common: the value x^I we were after. The proof of Lemma 5.7 breaks down if f has larger degree: there are more points $(E_i, P_i, Q_i) \in Y(f)_a(\overline{\mathbf{Q}_p})$ that satisfy $f(E_i, P_i, Q_i) = x$. However, we think that it is unlikely that the conclusion of the lemma is false, leading to the following conjecture.

CONJECTURE 5.8. The conclusion of Lemma 5.7 also holds for modular functions f that yield class invariants but have degree larger than one.

Heuristics. Let f be a modular function that yields class invariants. Since we normally only use f if its minimal polynomial P_Δ^f has ‘small’ coefficients, the degree of f will typically be ‘small’. Hence, there are not many points $(E_i, P_i, Q_i) \in Y(f)_a(\overline{\mathbf{Q}_p})$ satisfying $f(E_i, P_i, Q_i) = x$.

The main reason for our heuristic is of a practical nature: we did several experiments with various modular functions (all of which had moderately small degree) and have never found a counterexample.

6. The algorithm

In this section we give the algorithm for computing the minimal polynomial of a class invariant using p -adic arithmetic. The *input* of the algorithm is a discriminant $\Delta < -4$ and a modular function f of degree one that is known to yield class invariants for the order \mathcal{O} of discriminant Δ . We assume that f is *integral* over $\mathbf{Z}[j]$, and for simplicity we assume that the polynomial P_Δ^f that we want to compute is integral. The function f needs to be specified both by its Fourier expansion and by its minimal polynomial $\Psi_f(f, X) \in \mathbf{Z}[j, X]$ over $\mathbf{C}(j)$.

The *output* of the algorithm is the polynomial $P_\Delta^f \in \mathbf{Z}[X]$.

Initialization. List the elements of the class group $\text{Pic}(\mathcal{O})$ as reduced binary quadratic forms $[a, b, c]$ and compute the precision of

$$k = \frac{\pi\sqrt{|\Delta|}}{\log 2} \sum_{[a,b,c]} \frac{1}{a} + 10$$

bits required to compute the Hilbert class polynomial for \mathcal{O} . We will use $n = \lceil k/r(f) \rceil$ bits digits precision in our computations. We note that we have no rigorous proof that this precision suffices, because the precision k for the Hilbert class polynomial is heuristic. Furthermore, the reduction factor $r(f)$ is only an *asymptotic* statement and is not proven to be correct for our particular order \mathcal{O} . Our heuristic bound of n bits has sufficed in all our experiments. If we insist on a *proven* output we should replace k by the proven bound for the Hilbert class polynomial from [2] and not divide by $r(f)$.

Step 1. Find a prime $p \nmid N$ and an elliptic curve \bar{E}/\mathbf{F}_p with $\text{End}(\bar{E}) \cong \mathcal{O}$ using the approach outlined in Section 2. Compute the zeros $x_1, \dots, x_k \in \mathbf{F}_p$ of $\Psi_f(X, j(\bar{E})) \in \mathbf{F}_p[X]$.

Step 2. We have to decide which of these zeros is the reduction of a class invariant. First we compute the structure of the group $(\mathcal{O}/N\mathcal{O})^*$. This is a well-known computation in algebraic number theory, and the standard way of doing this is by localizing the ring \mathcal{O} at the primes l dividing N , and then applying the l -adic logarithm to reduce the multiplicative problem to a computation with *additive* groups. We refer to [18, § II.5] for details.

Sieve in the set

$$S = \{a + b\pi_p \mid a, b \in \mathbf{Z}, b \neq 0, (a, b) = 1, a + b\pi_p \text{ and } pN\Delta \text{ are coprime}\}$$

for smooth elements y_1, \dots, y_t generating $(\mathcal{O}/N\mathcal{O})^*$. Here, π_p is an element of norm p .

Step 3. Write $(y_1) = \alpha_1 \cdots \alpha_s$, with $N(\alpha_i) = l_i \in \mathbf{Z}$ prime. Compute the cycle

$$j(E) \xrightarrow{\bar{\rho}_{\alpha_1}} j(\bar{E}^{\alpha_1}) \xrightarrow{\bar{\rho}_{\alpha_2}} \dots \xrightarrow{\bar{\rho}_{\alpha_t}} j(\bar{E}^{(y_1)}) = j(\bar{E})$$

of j -invariants over \mathbf{F}_p as in Section 2 using the modular polynomials for j . Using the linear algebra technique explained in Section 5, compute the modular polynomial $\Phi_{l_1}^f$ of degree l_1 for f . Next, compute all roots $\eta_i \in \mathbf{F}_p$ of $\Phi_{l_1}^f(x_1, X) \in \mathbf{F}_p[X]$ that also satisfy $\Psi_f(\eta_i, j(\bar{E}^{\alpha_1})) = 0$. By Lemma 5.7, we find either zero or one such root η_i . If we find zero roots, then x_1 is not the reduction of a class invariant. If we find one root, we have computed $x_1^{\alpha_1}$.

Continuing like this, compute a series

$$x_1 \xrightarrow{\bar{\rho}_{\alpha_1}} x_1^{\alpha_1} \xrightarrow{\bar{\rho}_{\alpha_2}} \dots \xrightarrow{\bar{\rho}_{\alpha_t}} x_1^{(y_1)}.$$

If we have $x_1^{(y_1)} = x_1$, compute $x_1^{(y_2)}$, etc. If x_1 is invariant under all generators y_1, \dots, y_t of $(\mathcal{O}/N\mathcal{O})^*$, it is the reduction of a class invariant. Otherwise, repeat this computation with x_2 , etc. As there are only finitely many x_k , this computation terminates after a finite number of steps.

Step 4. Say that $x \in \mathbf{F}_p$ is the reduction of a class invariant. Choose a smooth \mathcal{O} -ideal $(\alpha) = \alpha_1 \dots \alpha_u$ for the map ρ_α from Section 2 by sieving in the set S from Step 2.

Compute a cycle

$$j(\overline{E}) \xrightarrow{\bar{\rho}_{\alpha_1}} j(\overline{E}^{\alpha_1}) \xrightarrow{\bar{\rho}_{\alpha_2}} \dots \xrightarrow{\bar{\rho}_{\alpha_u}} j(\overline{E}^{(\alpha)}) = j(\overline{E}),$$

and use this cycle to compute the corresponding cycle

$$x \xrightarrow{\bar{\rho}_{\alpha_1}} x^{\alpha_1} \xrightarrow{\bar{\rho}_{\alpha_2}} \dots \xrightarrow{\bar{\rho}_{\alpha_u}} x^\alpha = x$$

for x , just as in Step 3.

Step 5. Lift $\overline{E}/\mathbf{F}_p$ to E_1/\mathbf{Q}_p by lifting the coefficients of the Weierstraß equation for \overline{E} arbitrarily. We use two p -adic digits accuracy in this step.

Step 6. Lift $x \in \mathbf{F}_p$ to $x_1 \in \mathbf{Z}_p/(p^2)$ as a root of $\Psi_f(X, j(E_1)) \in \mathbf{Z}_p[X]$. As in Step 5, write $(\alpha) = \alpha_1 \dots \alpha_t$. Compute x^{α_1} as the unique root of $\Phi_l^f(x_1, X) \in (\mathbf{Z}_p/(p^2))[X]$ that reduces to $x_1^{\alpha_1} = x^{\alpha_1}$ modulo p , where l is the norm of α_1 .

Lift the cycle from Step 4 to a cycle

$$x \xrightarrow{\rho_{\alpha_1}} x^{\alpha_1} \longrightarrow \dots \longrightarrow x^\alpha$$

over \mathbf{Q}_p with two p -adic digits accuracy. We will typically *not* have $x = x^\alpha$.

Compute $\rho_\alpha(j(E_1))$ as the unique root of $\Psi_f(x_1^{(\alpha)}, X) \in \mathbf{Z}_p[X]$ that reduces to $j(\overline{E})$ modulo p .

Step 7. Update $\rho_\alpha(j(E_1))$ to $j(E_2)$ according to the ‘Newton formula’ (2.2).

Step 8. Repeat Step 6 with $j(E_1)$ replaced by $j(E_2)$. We now work with four p -adic digits precision. We obtain $j(E_3)$. Continue this iteration process until we have computed the canonical lift $j(\tilde{E})$ with $n = \lceil k/r(f) \rceil$ bits or $m = n(\log 2)/(\log p)$ p -adic digits accuracy. Compute the ‘canonical lift’ $\tilde{x} \in \mathbf{Z}_p$ of x as the root of $\Psi_f(X, j(\tilde{E}))$ reducing to $x \in \mathbf{F}_p$.

Step 9. Compute the conjugates of \tilde{x} under $\text{Pic}(\mathcal{O})$ in the same fashion as before: for an invertible \mathcal{O} -ideal I of norm l coprime to N , compute $j(\overline{E}^I) \in \mathbf{F}_p$ as in Section 2. Knowing $j(\overline{E}^I)$, compute the unique root $\underline{\beta} \in \mathbf{F}_p$ of $\Phi_l^f(x, X) \in \mathbf{F}_p[X]$ that also satisfies $\Psi_f(\beta, j(\overline{E}^I)) = 0$.

Since we know the reduction $\tilde{x}^I = \beta$ of \tilde{x}^I , we know which root of $\Phi_l^f(\tilde{x}, X) \in \mathbf{Z}_p[X]$ is \tilde{x}^I .

Step 10. Expand the polynomial

$$P_\Delta^f = \prod_{[I] \in \text{Pic}(\mathcal{O})} (X - \tilde{x}^I) \in (\mathbf{Z}_p/(p^m))[X],$$

and lift the coefficients of P_Δ^f from $\mathbf{Z}_p/(p^m) = \mathbf{Z}/(p^m)$ to \mathbf{Z} , where we take the representative between $-p^m/2$ and $p^m/2$. Return $P_\Delta^f \in \mathbf{Z}[X]$.

REMARK 6.1. The algorithm presented in this section also works for modular functions of degree >1 for which Conjecture 5.8 holds.

We expect the run time of this algorithm to be $\tilde{O}(|D|)$ for fixed f , just like the run time for the p -adic algorithm for the j -function from [4]. There is a serious obstacle that prevented us from *proving* this run time. The problem is that we cannot prove a reasonable smoothness bound on the generators y_1, \dots, y_t in Step 2. If GRH holds true, we have a bound for the α we find in Step 5; see [7, Lemma 2]. Although being B -smooth and lying in a prescribed

residue class in $(\mathcal{O}/N\mathcal{O})^*$ are quite unrelated, there appears to be no hope in proving similar smoothness bounds for the generators in Step 2.

Furthermore, since the reduction factor of f only holds *asymptotically* we have no rigorous proof that the precision

$$n = \lceil k/r(f) \rceil$$

we use in the computation will suffice. For a *proven* run time we would have to use k digits, annihilating the improvement we obtain by using with a smaller function. We note that this obstacle also prevented a run time analysis of the complex analytic method to compute minimal polynomials of class invariants; see [8].

7. Example

We illustrate the p -adic algorithm by working with the function

$$f(z) = \frac{\eta(z/5)\eta(z/7)}{\eta(z)\eta(z/35)} \in \mathbf{Z}[[q^{1/35}]],$$

where $\eta(z)$ denotes the Dedekind eta function. More examples, including a large example and an example where the polynomial P_Δ^f does not have integer coefficients, can be found in [3, Chapter 7]. As displaying large numbers is not particularly pleasing to the human eye, we work with a relatively small discriminant. Let \mathcal{O} be the order of discriminant $\Delta = -1571$. By [10, Theorem 3], the value $f(\omega)$ is a class invariant for a suitable choice of generator of the \mathbf{Z} -algebra $\mathcal{O} = \mathbf{Z}[\omega]$. Furthermore, the polynomial P_Δ^f has integer coefficients, and the size of these coefficients is a factor $r(f) = 24$ smaller than for the j -function.

By explicitly computing the conjugates of f over $\mathbf{Q}(j)$, we compute the minimal polynomial Ψ_f of f . In accordance with the example in [11], we find

$$\Psi_f(j, X) = X^{48} + (-j + 708)X^{47} + \dots + 12X + 1 \in \mathbf{Z}[j, X].$$

The function f generates the function field of $X_0(35)$ over $\mathbf{C}(j)$ and therefore has degree two. Using linear algebra, we compute some modular polynomials:

$$\begin{aligned} \Phi_2^f &= X^3 + Y^3 - X^2Y^2 + 2(XY^2 + X^2Y) + XY \\ \Phi_3^f &= X^4 + Y^4 - X^3Y^3 + 3(X^2Y^3 + X^3Y^2) + 3(Y^3X + X^3Y) \\ &\quad + 6(X^2Y^2) - 3(Y^2X + X^2Y) - XY. \end{aligned}$$

The fact that these polynomials have degree $l + 1$ and not $2(l + 1)$ is due to the fact that f is invariant under the Atkin–Lehner involution (the map $b : X(f) \rightarrow C$ in diagram (5.3) has degree two). Since 5 and 7 divide the level 35 of f , we cannot use Φ_5^f and Φ_7^f . We computed all modular polynomials for primes up to 23. This should be seen as a *precomputation*.

The (heuristic) precision required to compute the Hilbert class polynomial for this order is $k = 550$ bits. As we have $r(f) = 48/2 = 24$, we will use $\lceil 550/24 \rceil = 23$ bits accuracy in our computations.

The prime $p = 449$ splits completely in the ring class field $H_{\mathcal{O}}$, and the elliptic curve

$$\overline{E} : Y^2 = X^3 + X + 16$$

of j -invariant 383 has endomorphism ring \mathcal{O} . The polynomial $\Psi_f(j(\overline{E}), X) \in \mathbf{F}_p[X]$ has the four roots $b_1 = 62, b_2 = 130, b_3 = 239$ and $b_4 = 358$ in \mathbf{F}_p . We know [10, Theorem 3] that each one of them is a reduction of a class invariant. To illustrate our p -adic techniques, we reprove this.

A root $b_i \in \mathbf{F}_p$ is the reduction of a class invariant if it is invariant under the action of the group $(\mathcal{O}/35\mathcal{O})^*/\{\pm 1\} \cong \mathbf{Z}/2\mathbf{Z} \times \mathbf{Z}/2\mathbf{Z} \times \mathbf{Z}/12\mathbf{Z} \times \mathbf{Z}/12\mathbf{Z}$. We take $\{\pi_p, 2\pi_p - 11, 2\pi_p - 19, -\pi_p - 28\}$ as a generating set for $(\mathcal{O}/35\mathcal{O})^*$, where π_p is an element of norm p . We choose this particular set of generators, because the elements have smooth norm (except π_p).

Since b_i is an element of \mathbf{F}_p , it is invariant under the action of π_p . The element $\alpha = 2\pi_p - 11$ has order 12 in $(\mathcal{O}/35\mathcal{O})^*$, and the ideal (α) of norm $1587 = 3 \cdot 23^2$ factors as

$$(\alpha) = \mathfrak{p}_3 \cdot \mathfrak{p}_{23}^2 = (3, \pi_p - 1) \cdot (23, \pi_p - 17).$$

We compute the cycle of j -invariants over \mathbf{F}_p for the map $\bar{\rho}_\alpha : \text{Ell}_\Delta(\mathbf{F}_p) \rightarrow \text{Ell}_\Delta(\mathbf{F}_p)$:

$$j(\bar{E}) = 383 \xrightarrow{\mathfrak{p}_3} 13 \xrightarrow{\mathfrak{p}_{23}} 24 \xrightarrow{\mathfrak{p}_{23}} 383.$$

The modular polynomial $\Phi_3^f(b_1, X) \in \mathbf{F}_p[X]$ has two roots, namely 64, 95 $\in \mathbf{F}_p$. We check that 64 satisfies $\Psi_f(13, 64) = 0$, where 13 is the j -invariant of $\bar{E}^{\mathfrak{p}_3}$. The other root 95 does not satisfy $\Psi_f(13, 95) = 0$, and we conclude that we have $b_1^{\mathfrak{p}_3} = 64 \in \mathbf{F}_p$. Continuing like this, we compute

$$b_1 = 62 \xrightarrow{\mathfrak{p}_3} 64 \xrightarrow{\mathfrak{p}_{23}} 34 \xrightarrow{\mathfrak{p}_{23}} 62.$$

The computation for b_2, b_3, b_4 proceeds similarly and they are also invariant under the action of (α) . The computation for the other two elements of our generating set is similar. All four elements b_i are invariant under the action of $(\mathcal{O}/35\mathcal{O})^*$, proving that they are reductions of class invariant.

We will work with $b = b_1 = 62 \in \mathbf{F}_p$. For the polynomial P_Δ^j we would have needed 62 p -adic digits accuracy. For P_Δ^f we only need three p -adic digits. As element α for the map $\rho_\alpha : X_D(\mathbf{C}_p) \rightarrow X_D(\mathbf{C}_p)$ we again take $\alpha = 2\pi_p - 11$ of norm $3 \cdot 23^2$. We lift \bar{E}/\mathbf{F}_p to the curve E_1/\mathbf{Q}_p defined by $Y^2 = X^3 + X + 16$ of j -invariant $j(E_1) = 383 + 224p \in \mathbf{Q}_p$. This leads to the lift $b_1 = 62 + 45p \in \mathbf{Q}_p$.

We compute the ‘cycle’ for $b_1 \in \mathbf{Q}_p$ corresponding to the map ρ_α :

$$b_1 = 62 + 45p \xrightarrow{\mathfrak{p}_3} 64 + 175p \xrightarrow{\mathfrak{p}_{23}} 34 + 6p \xrightarrow{\mathfrak{p}_{23}} 62 - 198p = b_1^{(\alpha)}.$$

The degree two polynomial $\Psi_f(X, b_1^{(\alpha)}) \in \mathbf{Z}_p[X]$ has roots $131 - 94p + O(p^2)$ and $383 - 119p + O(p^2)$. We conclude that we have $\rho_\alpha(j(E_1)) = 383 - 119p \in \mathbf{Q}_p$. We update this j -value according to the ‘Newton formula’ (2.2) and obtain $j(E_2) = 383 - 98p \in \mathbf{Q}_p$. This is the j -invariant of the canonical lift in two p -adic digits accuracy. We compute $\tilde{b} = 62 - 64p + O(p^2) \in \mathbf{Q}_p$. Similarly, we compute $j(E_3) = 383 - 98p + 127p^2$ and $\tilde{b} = 62 - 64p + 66p^2$. To compute the conjugates of \tilde{b} under $\text{Pic}(\mathcal{O}) \cong \mathbf{Z}/17\mathbf{Z} \cong \langle \mathfrak{p}_3 \rangle$ we use the modular polynomial Φ_3^f once more. In the end we expand the polynomial

$$P_\Delta^f = \prod_{[I] \in \text{Pic}(\mathcal{O})} (X - \tilde{b}^I) \in (\mathbf{Z}_p/p^3)[X]$$

and we lift the coefficients to integers between $-p^3/2$ and $p^3/2$ to find

$$\begin{aligned} P_{-1571}^f &= X^{17} + 21X^{16} + 918X^{15} - 11046X^{14} + 49849X^{13} - 115187X^{12} \\ &\quad + 112918X^{11} + 168294X^{10} - 275500X^9 + 361744X^8 - 403346X^7 \\ &\quad + 181066X^6 - 10143X^5 - 3403X^4 - 4290X^3 + 1422X^2 \\ &\quad - 71X + 1 \in \mathbf{Z}[X]. \end{aligned}$$

Acknowledgements. I thank Bas Edixhoven and Peter Stevenhagen for helpful discussions.

References

1. A. AGASHE, K. LAUTER and R. VENKATESAN, ‘Constructing elliptic curves with a known number of points over a prime field’, *High primes and misdemeanours: lectures in honour of the 60th birthday of H. C. Williams*, Fields Institute Communications Series 41 (American Mathematical Society, Providence, RI, 2004) 1–17.
2. J. BELDING, R. BRÖKER, A. ENGE and K. LAUTER, ‘Computing Hilbert class polynomials’, *Algorithmic number theory symposium VIII*, Lecture Notes in Computer Science 5011 (Springer, Berlin, 2008).

3. R. BRÖKER, 'Constructing elliptic curves of prescribed order', PhD Thesis, Universiteit Leiden, 2006.
4. R. BRÖKER, 'A p -adic algorithm to compute the Hilbert class polynomial', *Math. Comp.* 77 (2008) 2417–2435.
5. R. BRÖKER and P. STEVENHAGEN, 'Constructing elliptic curves of prime order', *Contemp. Math.* 463 (2008) 17–28.
6. H. COHEN, G. FREY *et al.*, *Handbook of elliptic and hyperelliptic curve cryptography* (Chapman & Hall, 2006).
7. J.-M. COUVEIGNES and T. HENOCQ, 'Action of modular correspondences around CM-points', *Algorithmic number theory symposium V*, Lecture Notes in Computer Science 2369 (Springer, Berlin, 2002) 234–243.
8. A. ENGE, 'The complexity of class polynomial computation via floating point approximations', *Math. Comp.* 78 (2009) 1089–1107.
9. A. ENGE and F. MORAIN, 'Comparing invariants for class fields of imaginary quadratic fields', *Algorithmic number theory symposium V*, Lecture Notes in Computer Science 2369 (Springer, Berlin, 2002) 252–266.
10. A. ENGE and R. SCHERTZ, 'Constructing elliptic curves over finite fields using double eta-quotients', *J. Théor. Nombres Bordeaux* 16 (2004) 555–568.
11. A. ENGE and R. SCHERTZ, 'Modular curves of composite level', *Acta Arith.* 118 (2005) 129–141.
12. J. FRANKE, T. KLEINJUNG, F. MORAIN and T. WIRTH, 'Proving the primality of very large numbers with fast ECPP', *Algorithmic number theory symposium VI*, Lecture Notes in Computer Science 3076 (Springer, Berlin, 2004) 194–207.
13. A. GEE and P. STEVENHAGEN, 'Generating class fields using Shimura reciprocity', *Algorithmic number theory*, Lecture Notes in Computer Science 1423 (Springer, Berlin, 1998) 441–453.
14. M. HINDRY and J. H. SILVERMAN, *Diophantine geometry, an introduction*, Graduate Texts in Mathematics 201 (Springer, Berlin, 2000).
15. A. J. DE JONG, 'Families of curves and alterations', *Ann. Inst. Fourier (Grenoble)* 47 (1997) 599–621.
16. D. KOHEL, 'Endomorphism rings of elliptic curves over finite fields', PhD Thesis, University of California, Berkeley, 1996.
17. S. LANG, *Elliptic functions*, Graduate Texts in Mathematics 112 (Springer, Berlin, 1987).
18. J. NEUKIRCH, *Algebraic number theory*, Grundlehren der Mathematischen Wissenschaften 322 (Springer, Berlin, 1999).
19. R. SCHOOF, 'Counting points on elliptic curves over finite fields', *J. Théor. Nombres Bordeaux* 7 (1993) 219–254.
20. G. SHIMURA, *Introduction to the arithmetic theory of automorphic forms* (Princeton University Press, Princeton, NJ, 1971).
21. P. STEVENHAGEN, 'Hilbert's 12th problem, complex multiplication and Shimura reciprocity', *Class field theory—its centenary and prospect*, Advanced Studies in Pure Mathematics 30 (ed. K. Miyake; Mathematical Society of Japan, Tokyo, 2001) 161–176.
22. A. V. SUTHERLAND, 'Computing Hilbert class polynomials with the Chinese remainder theorem' *Math. Comp.*, to appear, available at <http://arxiv.org/abs/0903.2785>.
23. J. VÉLU, 'Isogénies entre courbes elliptiques', *C. R. Acad. Sci. Paris Sér. A–B* 273 (1971) A238–A241.
24. H. WEBER, *Lehrbuch der Algebra* (Friedrich Vieweg und Sohn, 1908).

Reinier Bröker
 Brown University
 Box 1917, 151 Thayer Street
 Providence, RI
 USA

reinier@math.brown.edu