22.1 The Hilbert class polynomial

We now turn our attention back to the Hilbert class polynomial introduced in Lecture 20. Recall that for each imaginary quadratic order $\mathcal{O}$, we define the set

$$\mathrm{Ell}_{\mathcal{O}}(\mathbb{C}) = \{j(E) \in \mathbb{C} : \mathrm{End}(E) \simeq \mathcal{O}\}$$

of equivalence classes of elliptic curves with endomorphism ring $\mathcal{O}$ (we say such elliptic curves *have CM by* $\mathcal{O}$). By Theorem 19.2, we can uniquely identify $\mathcal{O}$ by its discriminant D.

Definition 22.1. The polynomial

$$H_D(X) = \prod_{j(E)\in \mathrm{Ell}_{\mathcal{O}}(\mathbb{C})} \big(X - j(E)\big)$$

is the *Hilbert class polynomial* (of discriminant D).

The appellation "Hilbert" is sometimes reserved for cases where D is a fundamental discriminant (in which case $H_D(X)$ is more generally called a *ring class polynomial*), but we shall use the term Hilbert class polynomial to refer to $H_D(X)$ in general. Our first objective is to use the fact that $\Phi_N \in \mathbb{Z}[X,Y]$ to prove that $H_D \in \mathbb{Z}[X]$. We require the following lemma.

Lemma 22.2. *If N is prime then the leading coefficient of $\Phi_N(X,X)$ is -1.*

Proof. We have

$$\Phi_N\big(j(\tau), j(\tau)\big) = \Big(j(\tau) - j(N\tau)\Big) \prod_{k=0}^{N-1} \Big(j(\tau) - j\Big(\frac{\tau+k}{N}\Big)\Big).$$

Recall from the proof of Theorem 21.13 that

$$j(N\tau) = \frac{1}{q^N} + \cdots,$$
$$j\Big(\frac{\tau+k}{N}\Big) = \frac{\zeta_N^{-k}}{q^{1/N}} + \cdots,$$

where $q = e^{2\pi i\tau}$, $\zeta_N = e^{2\pi i/N}$, and each ellipsis denotes terms with positive powers of q. Thus

$$j(\tau) - j(N\tau) = -\frac{1}{q^N} + \frac{1}{q} + \cdots,$$
$$j(\tau) - j\Big(\frac{\tau+k}{N}\Big) = \frac{1}{q} - \frac{\zeta_N^{-k}}{q^{1/N}} + \cdots,$$

which implies that the q-expansion of $f(\tau) = \Phi_N\big(j(\tau), j(\tau)\big)$ is $-\frac{1}{q^{2N}} + \cdots$. Since $f(\tau)$ is a polynomial in $j(\tau) = \frac{1}{q} + \cdots$, the leading term of $\Phi_N(X,X)$ must be $-X^{2N}$. □

Andrew V. Sutherland

Remark 22.3. Lemma 22.2 does not hold for composite N; in particular, when N is square $\Phi_N(X,X)$ is not even primitive (its coefficients have a non-trivial common divisor).

Before proving that $H_D \in \mathbb{Z}[X]$, we note the following classical number-theoretic result, which is a consequence of the Chebotarev[1] density theorem (the result stated here actually follows from earlier work of Dirichlet and Weber, see [2, p. 190]).

Theorem 22.4. *Let $\mathcal{O}$ be an imaginary quadratic order. Every ideal class in* $\mathrm{cl}(\mathcal{O})$ *contains infinitely many ideals of prime norm.*

Proof. This follows from Theorems 7.7 and 9.12 in [2]. □

Theorem 22.5. *The coefficients of the Hilbert class polynomial $H_D(X)$ are integers.*

Proof. Let $\mathcal{O}$ be the imaginary quadratic order of discriminant D, let $E/\mathbb{C}$ be an elliptic curve wih CM by $\mathcal{O}$, and let $\mathfrak{p}$ be a principal $\mathcal{O}$-ideal of prime norm p (the existence of $\mathfrak{p}$ is guaranteed by Theorem 22.4). Then $[\mathfrak{p}]$ is the identity in $\mathrm{cl}[\mathcal{O}]$ and therefore acts the acts trivially on $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$. Thus the elliptic curve $\mathfrak{p}E = E_{\mathfrak{p}^{-1}}$ corresponding to the torus $\mathbb{C}/\mathfrak{p}^{-1}$ is isomorphic to E. It follows that there exists a p-isogeny from E to itself. Such an isogeny is necessarily cyclic, since it has prime degree, so we must have $\Phi_p\big(j(E), j(E)\big) = 0$. Thus $j(E)$ is the root of the polynomial $-\Phi_p(X,X)$, which has integer coefficients and is also monic, by Lemma 22.2. Therefore $j(E)$ is an algebraic integer, and E can be defined by a Weierstrass equation $y^2 = x^3 + Ax + B$ whose coefficients lie in the number field $\mathbb{Q}(j(E))$.

The group $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ acts on elliptic curves defined over number fields via its action on the Weierstrass coefficients A and B: for $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ the curve E^σ is defined by the equation $y^2 = x^3 + \sigma(A)x + \sigma(B)$. Similarly, σ acts on isogenies between such curves via its action on the coefficients of the rational map defining the isogeny. If $\phi\colon E \to E$ is an endomorphism, then so is $\phi^\sigma\colon E^\sigma \to E^\sigma$. Note that for any $\phi, \psi \in \mathrm{End}(E)$ we have $(\phi+\psi)^\sigma = \phi^\sigma + \psi^\sigma$ and $(\phi \circ \psi)^\sigma = \phi^\sigma \circ \psi^\sigma$, thus we have a ring homomorphism from $\mathrm{End}(E)$ to $\mathrm{End}(E^\sigma)$, and it is invertible (apply σ^{-1} to $\mathrm{End}(E^\sigma)$), so $\mathrm{End}(E) \simeq \mathrm{End}(E^\sigma)$.

It follows that for any $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ we have $j(E^\sigma) = j(E)^\sigma \in \mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$. Thus the set of roots of $H_D(X)$ is fixed by $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, therefore $H_D \in \mathbb{Q}[X]$. Every root of $H_D(X)$ is a root of $\Phi_p(X,X)$, thus $H_D(X)$ divides $\Phi_p(X,X)$ in $\mathbb{Q}[X]$. But $\Phi_p(X,X)$ has integer coefficients and it is primitive, by Lemma 22.2, so by Gauss's lemma its divisors in $\mathbb{Q}[X]$ all lie in $\mathbb{Z}[X]$. Therefore $H_D \in \mathbb{Z}[X]$. □

Corollary 22.6. *Let $E/\mathbb{C}$ be an elliptic curve with complex multiplication. Then $j(E)$ is an algebraic integer.*

We now turn to our main goal for this lecture. We wish to prove the first main theorem of complex multiplication, which states that the Galois group of the splitting field L of $H_D(X)$ over $K = \mathbb{Q}(\sqrt{D})$ is isomorphic to the class group $\mathrm{cl}(\mathcal{O})$, and moreover, that the CM action of $\mathrm{cl}(\mathcal{O})$ on $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ is precisely the Galois action of $\mathrm{Gal}(L/K)$ on the roots of $H_D(X)$. Note that $\mathrm{cl}(\mathcal{O})$ acts transitively on $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$, so this result implies that $H_D(X)$ is irreducible over K and is therefore the minimal polynomial of each $j(E) \in \mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ over K (and over $\mathbb{Q}$).

Let $\mathcal{O}$ be the imaginary quadratic order of discriminant D, and fix an elliptic curve E_1 with CM by $\mathcal{O}$. As in the proof of Theoerem 22.5, if $\sigma \in \mathrm{Gal}(\overline{K}/K)$, then E_1^σ has CM by $\mathcal{O}$,

[1] Many different transliterations of Chebotarev's name appears in the literature, including Chebotaryov Čebotarev, Chebotarëv, Čhebotarëv, Tchebotarev, and Tschebotaröw. In Russian, his name is Чеботарёв.

and therefore $E_1^\sigma \simeq \mathfrak{a}E_1$ for some proper $\mathcal{O}$-ideal $\mathfrak{a}$. If $E_2 \simeq \mathfrak{b}E_1$ is any other elliptic curve with CM by $\mathcal{O}$, we have

$$E_2^\sigma \simeq (\mathfrak{b}E_1)^\sigma = \mathfrak{b}^\sigma E_1^\sigma = \mathfrak{b}E_1^\sigma \simeq \mathfrak{b}\mathfrak{a}E_1 = \mathfrak{a}\mathfrak{b}E_1 \simeq \mathfrak{a}E_2. \tag{1}$$

Two comments are in order. First, the innocent looking identity $(\mathfrak{b}E_1)^\sigma = \mathfrak{b}^\sigma E_1^\sigma$ used in (1) is not immediate; see [6, Prop. II.2.5] for a proof. Second, the identity $\mathfrak{b}^\sigma = \mathfrak{b}$ *is* immediate, because $\mathfrak{b} \subset K$ and $\sigma \in \mathrm{Gal}(\overline{K}/K)$ fixes every element of K; but this would not be true if we had instead used $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$.

Since our choice of E_2 was arbitrary, it follows from (1) that the action of σ on $\mathrm{Ell}_\mathcal{O}(\mathbb{C})$ is the same as the action of $\mathfrak{a}$ on $\mathrm{Ell}_\mathcal{O}(\mathbb{C})$. Because $\mathrm{Ell}_\mathcal{O}(\mathbb{C})$ is a $\mathrm{cl}(\mathcal{O})$-torsor, the map that sends each $\sigma \in \mathrm{Gal}(\overline{K}/K)$ to the corresponding class $[\mathfrak{a}]$ for which $E_1^\sigma = \mathfrak{a}E_1$ defines a group homomorphism from $\mathrm{Gal}(\overline{K}/K)$ to $\mathrm{cl}(\mathcal{O})$. Restricting this homomorphism to the splitting field L of $H_D(X)$ over K yields an injective homomorphism

$$\Psi\colon \mathrm{Gal}(L/K) \to \mathrm{cl}(\mathcal{O}).$$

To show injectivity, note that if $\Psi(\sigma)$ acts trivially on $\mathrm{Ell}_\mathcal{O}(\mathbb{C})$ then $\Psi(\sigma)$ is the identity in $\mathrm{cl}(\mathcal{O})$, and σ must fix every root of $H_D(X)$ and is therefore the identity in $\mathrm{Gal}(L/K)$.

We summarize this discussion with the following theorem.

Theorem 22.7. *Let $\mathcal{O}$ be an imaginary quadratic order of discriminant D and let L be the splitting field of $H_D(X)$ over $K = \mathbb{Q}(\sqrt{D})$. The map $\Psi : \mathrm{Gal}(L/K) \to \mathrm{cl}(D)$ that sends σ to the unique $\alpha \in \mathrm{cl}(\mathcal{O})$ for which $j(E)^\sigma = \alpha j(E)$ for all $j(E) \in \mathrm{Ell}_\mathcal{O}(E)$ is well-defined and is an injective group homomorphism.*

Thus we have embedded $\mathrm{Gal}(L/K)$ in $\mathrm{cl}(\mathcal{O})$ in a way that is compatible with each group's action on $\mathrm{Ell}_\mathcal{O}(\mathbb{C})$. It remains only to prove that Ψ is surjective. To do this we need to introduce the Artin map, which will allow us to associate to each $\mathcal{O}$-ideal $\mathfrak{p}$ of prime norm (subject to certain constraints), an element of $\sigma \in \mathrm{Gal}(L/K)$ whose action on $\mathrm{Ell}_\mathcal{O}(\mathbb{C})$ corresponds to the action of $[\mathfrak{p}]$. In order to define the Artin map we need to briefly delve into some algebraic number theory, but we will restrict ourselves to the absolute minimum we need; those who want to learn more may wish to consult [2] or [4]. Those who prefer to simply treat the Artin map as a "black box" are welcome to do so.

22.2 The Artin map

Let L be a finite abelian extension of a number field K (this means L/K is Galois and $\mathrm{Gal}(L/K)$ is a finite abelian group). Let $\mathfrak{p}$ be a prime ideal of K (an $\mathcal{O}_K$-ideal). We can factor the $\mathcal{O}_L$-ideal $\mathfrak{p}\mathcal{O}_L$ as a product of prime $\mathcal{O}_L$-ideals. When these prime ideals are all distinct, we say that $\mathfrak{p}$ is unramified in L. This holds for all but a finite set of prime ideals $\mathfrak{p}$, and we now assume that this is the case. Let $\mathfrak{P}$ be a prime ideal of L in the prime factorization of $\mathfrak{p}\mathcal{O}_L$; this means $\mathfrak{P}$ contains $\mathfrak{p}\mathcal{O}_L$, and we say that $\mathfrak{P}$ *lies above* $\mathfrak{p}$.

The subgroup $D_\mathfrak{P} = \{\sigma \in \mathrm{Gal}(L/K) : \mathfrak{P}^\sigma = \mathfrak{P}\}$ is called the *decomposition group* of $\mathfrak{P}$. Each $\sigma \in D_\mathfrak{P}$ induces an automorphism $\bar\sigma$ of the finite field $\mathbb{F}_\mathfrak{P} = \mathcal{O}_L/\mathfrak{P}$ that fixes the subfield $\mathbb{F}_\mathfrak{p} = \mathcal{O}_K/\mathfrak{p}$. Thus there is a homomorphism from $D_\mathfrak{P}$ to $\mathrm{Gal}(\mathbb{F}_\mathfrak{P}/\mathbb{F}_\mathfrak{p})$. This homomorphism is surjective [4, Prop. I.9.4], and our assumption that $\mathfrak{p}$ is unramified means that it is also injective [4, Prop. I.9.5], and therefore an isomorphism.

The group $\mathrm{Gal}(\mathbb{F}_\mathfrak{P}/\mathbb{F}_\mathfrak{p})$ is cyclic, generated by the Frobenius automorphism $x \to x^q$, where $q = \#\mathbb{F}_\mathfrak{p} = N(\mathfrak{p})$. The unique $\sigma_\mathfrak{P} \in D_\mathfrak{P} \subseteq \mathrm{Gal}(L/K)$ for which $\bar\sigma_\mathfrak{P}$ is the Frobenius automorphism is called the *Frobenius element*. In general, for any given $\mathfrak{p}$ the Frobenius element $\sigma_{\mathfrak{P}}$ depends on our choice of $\mathfrak{P}$. But the $\sigma_{\mathfrak{P}}$ are all conjugate in $\mathrm{Gal}(L/K)$, and in our situation $\mathrm{Gal}(L/K)$ is abelian, so they must all be equal. Thus there is a unique Frobenius element $\sigma_{\mathfrak{p}}$ that does not depend on our choice of $\mathfrak{P}$. The map $\mathfrak{p} \mapsto \sigma_{\mathfrak{p}}$ is known as the *Artin map* (it extends multiplicatively to a map defined on all $\mathcal{O}_K$-ideals, but this is irrelevant to us). The automorphism $\sigma_{\mathfrak{p}}$ is uniquely characterized by the fact that

$$\sigma_{\mathfrak{p}}(x) \equiv x^{N(\mathfrak{p})} \bmod \mathfrak{P}, \tag{2}$$

for all $x \in \mathcal{O}_L$ and primes $\mathfrak{P}$ that lie above $\mathfrak{p}$.

22.3 The first main theorem of complex multiplication

We are now ready to prove that $\Psi\colon \mathrm{Gal}(L/K) \to \mathrm{cl}(\mathcal{O})$ is an isomorphism. Note that we have already shown that it is injective, and this implies that $\mathrm{Gal}(L/K)$ is abelian, so we have the desired setup for applying the Artin map.

Since we have proved that the roots of $H_D(X)$ are all algebraic integers that lie in its splitting field L over $K = \mathbb{Q}(\sqrt{D})$, we now write $\mathrm{Ell}_{\mathcal{O}}(L)$ in place of $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ to emphasize that we are working with j-invariants that lie in a number field. Any elliptic curve $E/\mathbb{C}$ with CM by $\mathcal{O}$ can thus be defined over L, and we can further assume that the coefficients of the equation defining E lie in the ring of integers $\mathcal{O}_L$ (by clearing denominators). If $\mathfrak{P}$ is any prime of L (a prime $\mathcal{O}_L$-ideal), then it makes sense to reduce elements of $\mathcal{O}_L$ modulo $\mathfrak{P}$ to obtain elements of the finite field $\mathbb{F}_{\mathfrak{P}} = O_L/\mathfrak{P}$. Thus for an elliptic curve E/L we may speak of the reduction $E \bmod \mathfrak{P}$, the elliptic curve $\bar{E}/\mathbb{F}_{\mathfrak{P}}$ obtained by reducing the coefficients of E modulo $\mathfrak{P}$. We say that E has good reduction at $\mathfrak{P}$ if the discriminant of $\bar{E}$ is not zero.

Theorem 22.8. *Let $\mathcal{O}$ be an imaginary quadratic order of discriminant D and let L be the splitting field of $H_D(X)$ over $K = \mathbb{Q}(\sqrt{D})$. The map $\Psi\colon \mathrm{Gal}(L/K) \to \mathrm{cl}(\mathcal{O})$ given by Theorem 22.7 is a group isomorphism that commutes with the group actions of $\mathrm{Gal}(L/K)$ and $\mathrm{cl}(\mathcal{O})$ on $\mathrm{Ell}_{\mathcal{O}}(L)$.*

Proof. In view of Theorem 22.7, we just need to show that Ψ is surjective. So let α be an arbitrary element of $\mathrm{cl}(\mathcal{O})$. We will show that α is in the image of Ψ.

Let us fix an elliptic curve E/L with CM by $\mathcal{O}$, and let $\mathfrak{p}$ be an $\mathcal{O}_K$-ideal of prime norm p such that

(i) $\mathfrak{p} \cap \mathcal{O}$ is a proper $\mathcal{O}$-ideal contained in α.

(ii) p is unramified in L;

(iii) The elliptic curves E, $\mathfrak{p}^*E$, and $\bar{\mathfrak{p}}^*E$ have good reduction modulo every prime $\mathfrak{P}$ of L lying above p.

(iv) The elements of $\mathrm{Ell}_{\mathcal{O}}(L)$ are distinct modulo every prime $\mathfrak{P}$ of L lying above p.

The existence of such a $\mathfrak{p}$ is guaranteed by Theorem 22.4; there are infinitely many $\mathfrak{p}$ for which (i) holds, and conditions (ii)-(iv) prohibit only finitely many primes. To ease the notation, we will also use $\mathfrak{p}$ to denote the $\mathcal{O}$-ideal $\mathfrak{p} \cap \mathcal{O}$; it will be clear from context whether we are viewing $\mathfrak{p}$ as a prime of K or as an $\mathcal{O}$-ideal.

Let us now fix a prime $\mathfrak{P}$ of L that lies above $\mathfrak{p}$, and let $\bar{E}/\mathbb{F}_{\mathfrak{P}}$ be the reduction of E modulo $\mathfrak{P}$. It follows from (2) that the action of $\sigma_{\mathfrak{p}}$ on E corresponds to the action of the p-power Frobenius map π on $\bar{E}$, which gives an inseparable p-isogeny from $\bar{E}$ to $\bar{E}^{\bar{\sigma}_{\mathfrak{p}}}$. The

CM action of the $\mathcal{O}$-ideal $\mathfrak{p}$ corresponds to an isogeny of degree $N(\mathfrak{p}) = p$ from E to $\mathfrak{p}E$, and induces an isogeny ϕ from $\bar{E}$ to $\overline{\mathfrak{p}E}$. Let us now consider the possibilities for ϕ.

If ϕ is inseparable, then $\phi = \phi_{\text{sep}} \circ \pi$, by Corollary 5.16, and $\deg \phi = \deg \pi$ implies $\deg \phi_{\text{sep}} = 1$, which means that ϕ and π are isomorphic; thus $\overline{\mathfrak{p}E} \simeq \bar{E}^{\bar{\sigma}_\mathfrak{p}}$. We must then have $j(\overline{\mathfrak{p}E}) = j(\bar{E}^{\bar{\sigma}_\mathfrak{p}})$ and therefore $j(\mathfrak{p}E) = j(E^\sigma)$, by (iv). It follows that $\Psi(\sigma_\mathfrak{p}) = [\mathfrak{p}] = \alpha$, since each element of $\text{cl}(\mathcal{O})$ is determined by its action on any element of the $\text{cl}(\mathcal{O})$-torsor $\text{Ell}_\mathcal{O}(L)$.

So now suppose ϕ is separable. Then the reduction of any isogeny induced by the action of $\mathfrak{p}$ on an elliptic curve with CM by $\mathcal{O}$ must also be separable, since we get an inseparable isogeny if and only if $\Psi(\sigma_\mathfrak{p}) = [\mathfrak{p}]$, and this does not depend on the choice of E. In characteristic p, the dual of a separable p-isogeny must be inseparable, since the order of $E[p]$ is at most p. Thus the isogenies induced by $\bar{\mathfrak{p}}$, which are always dual to those induced by $\mathfrak{p}$, must have inseparable reductions. Therefore $\Psi(\sigma_\mathfrak{p}^{-1}) = \alpha$.[2] □

Corollary 22.9. *The Hilbert class polynomial $H_D(x)$ is irreducible over $K = \mathbb{Q}(\sqrt{D})$ and each of its roots $j(E)$ generates an abelian extension $K(j(E))/K$ with Galois group isomorphic to* $\text{cl}(\mathcal{O})$.

Proof. The class group $\text{cl}(\mathcal{O})$ acts transitively on the roots of $H_D(X)$ (the set $\text{Ell}_\mathcal{O}(\mathbb{C})$). By Theorem 22.8, the splitting field L of $H_D(x)$ over K must also act transitively on the roots of $H_D(X)$, which implies that $H_D(X)$ is irreducible over K. Thus each root $j(E)$ of $H_D(X)$ is an algebraic integer of degree $h(D) = |\text{cl}(\mathcal{O})| = |\text{Gal}(L/K)| = [L : K]$, and therefore generates L, and we have $\text{Gal}(L/K) \simeq \text{cl}(\mathcal{O})$, which is abelian. □

Theorem 22.10. *Let $\mathcal{O}$ be an imaginary quadratic order with discriminant D and ring class field L. Let p be a prime that is unramified in L. The following are equivalent:*

(i) *p is the norm of a principal $\mathcal{O}$-ideal;*

(ii) $\left(\frac{D}{p}\right) = 1$ *and $H_D(X)$ splits completely in $\mathbb{F}_p[X]$;*

(iii) *p splits completely in L;*

(iv) *$4p = t^2 - v^2 D$ for some integers t and v.*

When we say that p splits completely in L, we mean that the the principal $\mathcal{O}_L$-ideal (p) factors into a product of prime $\mathcal{O}_L$-ideals of norm p (degree-1 primes of L).

Proof. If $\mathfrak{p}$ is a principal $\mathcal{O}$-ideal of norm p, then $[\mathfrak{p}]$, and therefore $\sigma_\mathfrak{p}$, acts trivially on the roots of $H_D(X)$, which means that $H_D(X)$ splits into linear factors over $\mathbb{F}_\mathfrak{p} = \mathbb{F}_p$. The converse also holds, thus (i) and (ii) are equivalent.

If $\left(\frac{D}{p}\right) = 1$, then $p = \mathfrak{p}\bar{\mathfrak{p}}$ splits into degree-1 primes in K, and if $H_D(X)$ splits completely over $\mathbb{F}_p$, then its roots are all fixed by $\sigma_\mathfrak{p}$. But then $[\mathbb{F}_\mathfrak{P} : \mathbb{F}_\mathfrak{p}] = 1$, and we therefore have $N(\mathfrak{P}) = [\mathcal{O}_L : \mathfrak{P}] = [\mathcal{O}_K : \mathfrak{p}] = p$ for every prime $\mathfrak{P}$ of L lying above $\mathfrak{p}$. So p splits completely in L. The converse also holds, thus (ii) and (iii) are equivalent.

Write $D = f^2 D_K$, where $f = [\mathcal{O}_K : \mathcal{O}]$ and $D_K = \text{disc}(\mathcal{O}_K)$. Then $\mathcal{O}_K = [1, \omega_K]$, where $\omega_K = (D_K + \sqrt{D_K})/2$, and $\mathcal{O} = [1, f\omega_K]$. If (α) is a principal $\mathcal{O}$-ideal of norm p, then $\alpha = a + bf\omega_K$, for some $a, b \in \mathbb{Z}$, and

$$4p = 4N(\alpha) = 4\alpha\bar{\alpha} = 4(a + bf\omega_K)(a + bf\bar{\omega}_K) = (2a + bfD_K)^2 - b^2 D.$$

[2] In fact this never happens; we defined $\mathfrak{p}E = E_{\mathfrak{p}^{-1}}$ rather than $\mathfrak{p}E = E_\mathfrak{p}$ precisely so that we would always have $\Psi(\sigma_\mathfrak{p}) = [\mathfrak{p}]$, but we haven't actually proved this and don't need to.

Thus $4p = t^2 - v^2D$ holds for the integers $t = 2a + bfD_K$ and $v = b$. Conversely, if $4p = t^2 - v^2D$, then let $a = (t - vfD_K)/2$ and $b = v$, and set $\alpha = a + bf\omega_K$. If D is odd then $t \equiv v \bmod 2$, and if D is even then $t \equiv fD_K \bmod 2$. In either case, $a \in \mathbb{Z}$, so $\alpha \in \mathcal{O}$ generates a $\mathcal{O}$-principal ideal of norm $N(\alpha) = p$. Thus (i) and (iv) are equivalent. □

22.4 Ring class fields

The theory of complex multiplication was originally motivated not by the study of elliptic curves, but as a way to construct abelian Galois extensions. A famous theorem of Kronecker and Weber states that every finite abelian extension of $\mathbb{Q}$ lies in a cyclotomic field (a field of the form $\mathbb{Q}(\zeta_n)$, for some nth root of unity ζ_n). The effort to generalize this result to fields other than $\mathbb{Q}$ led to the development of *class field theory*, a branch of algebraic number theory that represents one of the major advances of early 20th century number theory.

In 1898 Hilbert conjectured that every number field K has a unique maximal abelian extension L/K that is unramified at every prime[3] of K, and it satisfies $\mathrm{Gal}(L/K) \simeq \mathrm{cl}(\mathcal{O}_K)$. This conjecture was proved shortly thereafter by Furtwängler, and the field L is known as the *Hilbert class field* of K. While its existence was proved, the problem of explicitly constructing L, say, by specifying a generator for L in terms of its minimal polynomial over K, remained an open problem (and for general K it still is).

After $\mathbb{Q}$, the simplest fields K to consider are imaginary quadratic fields. As a generalization of the Hilbert class field, rather than requiring L/K to be unramified at every prime $\mathcal{O}_K$-ideal, we might instead only require L/K to be unramified at every prime that is a proper $\mathcal{O}$-ideal, for some order $\mathcal{O} \subseteq \mathcal{O}_K$. This leads to the definition of the *ring class field* $L_{\mathcal{O}}$ of the order $\mathcal{O}$. The ring class field of $\mathcal{O}_K$ is then the Hilbert class field.

The ring class field $L_{\mathcal{O}}$ is uniquely characterized by the infinite set $\mathcal{S}_{L_{\mathcal{O}}/\mathbb{Q}}$ of rational primes p that split completely in $L_{\mathcal{O}}$, and with finitely many exceptions, these are precisely the primes that satisfy the equation $4p = t^2 - v^2D$ for some $t, v \in \mathbb{Z}$, where $D = \mathrm{disc}(\mathcal{O})$; see [2, Thm. 9.2, Ex, 9.3]. The Chebotarev density theorem implies that any extension M/K for which the set $\mathcal{S}_{M/\mathbb{Q}}$ matches $\mathcal{S}_{L_{\mathcal{O}}/\mathbb{Q}}$ with only finitely many exceptions must in fact be equal to $L_{\mathcal{O}}$, by [2, Thm. 8.19]. Thus we have the following corollary of Theorem 22.10, which completely solves the problem of explicitly constructing the Hilbert class field, and ring class fields, in the case that K is an imaginary quadratic field.

Corollary 22.11. *Let $\mathcal{O}$ be an imaginary quadratic order with discriminant D and let $K = \sqrt{D}$. The splitting field of $H_D(X)$ over K is the ring class field of the order $\mathcal{O}$.*

22.5 The CM method

The equation

$$4p = t^2 - v^2D$$

in part (iv) of Theorem 22.10 is known as the *norm equation*, since it arises from the principal ideal of norm p given by part (i). For $D < -4$, the integers t^2 and v^2 are uniquely determined by p and D. If the norm equation is satisfied and $j(E)$ is a root of $H_D(X)$ over $\mathbb{F}_p$, then the Frobenius endomorphism π of $E/\mathbb{F}_p$ satisfies the characteristic polynomial

[3] This includes not only all prime $\mathcal{O}_K$-ideals, but also the infinite primes of K (embeddings of K into $\mathbb{C}$). Only real infinite primes (embeddings of K into $\mathbb{R}$) can ramify, so for imaginary quadratic fields K we can safely ignore the infinite primes.

$x^2 - \mathrm{tr}(\pi)x + N(\pi)$. Viewing π as an element of $\mathrm{End}(E) \simeq \mathcal{O}$, we can apply the quadratic formula to compute

$$\pi = \frac{\mathrm{tr}(\pi) \pm \sqrt{\mathrm{tr}(\pi)^2 - 4p}}{2},$$

where $\sqrt{\mathrm{tr}(\pi)^2 - 4p}$ lies in $\mathcal{O}$ and can written as $v\sqrt{D}$ for some integer v. It follows that $\mathrm{tr}(\pi) = \pm t$. The two possible signs correspond to quadratic twists of E.

Thus given the Hilbert class polynomial $H_D(X)$ and a prime p for which the norm equation holds, we can find a root j_0 of $H_D(X)$ over $\mathbb{F}_p$ and then write down the equation $y^2 = x^3 + Ax + B$ of an elliptic curve E with $j(E) = j_0$, using $A = 3j(1728 - j)$ and $B = 2j(1728 - j)^2$. The Frobenius endomorphism π_E then satisfies $\mathrm{tr}(\pi_E) = \pm t$, and by Hasse's theorem we have

$$\#E(\mathbb{F}_p) = p + 1 - \mathrm{tr}(\pi_E).$$

The sign of $\mathrm{tr}(\pi_E)$ can be uniquely determined using the formulas in [5]. A more expedient method is to simply pick a random point $P \in E(\mathbb{F}_p)$ and check whether $(p+1-t)P = 0$ or $(p+1+t)P = 0$ both hold (at least one must). If only one of these equations is satisfied, then $\mathrm{tr}(\pi)$ is determined. By Mestre's theorem (see Lecture 8), for $p > 229$ we are guaranteed that this will work either for E or its quadratic twist, for most of the random points P we pick (when p is large the first random point P that we try is almost certain to work).

This method of constructing an elliptic curve $E/\mathbb{F}_p$ using a root of the Hilbert class polynomial is known as the *CM method*. Its key virtue is that $\#E(\mathbb{F}_p) = p + 1 - t$ is known in advance. This has many applications, one of which is an improved version of elliptic curve primality proving developed by Atkin and Morain [1], which is explored in Problem Set 11.

The main limitation of the CM method is that it requires computing (or having precomputed) the Hilbert class polynomial $H_D(X)$, which becomes very difficult when $|D|$ is large. The degree of $H_D(X)$ is the class number $h(D)$, which is asymptotically on the order of $\sqrt{|D|}$, and the size of its largest coefficient is on the order of $\sqrt{|D|}\log|D|$ bits.[4] Thus the total size of $H_D(X)$ is on the order of $|D|\log|D|$ bits, which makes it impractical to even write down if $|D|$ is large (in general, $|D|$ may be as large as the prime p we are working with). An efficient algorithm for computing $H_D(X)$ is outlined in Problem Set 11, and with a suitable implementation, it can practically handle $|D| > 10^{13}$, where the size of $H_D(X)$ is several terabytes [7]. Using class polynomials associated to alternative modular functions (which may be smaller by a large constant factor), discriminants as large as $|D| \approx 10^{15}$ can be addressed [3]; with more advanced techniques even $|D| \approx 10^{16}$ is possible [8].

22.6 Summing up the theory of complex multiplication

E	L	$\mathfrak{a}$	$ax^2 + bxy + cy^2$
↓ isomorphism	↓ homethety	↓ mod principal ideals	↓ $\mathrm{SL}_2(\mathbb{Z})$-equivalence
$j(E)$	$j(L)$	$[\mathfrak{a}]$	reduced form
↓	↓	↓	↓
$\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$	$\{j(L) : \mathcal{O}(L) = \mathcal{O}\}$	$\mathrm{cl}(\mathcal{O})$	$\mathrm{cl}(D)$

[4]Under the Generalized Riemann Hypothesis, these bounds are accurate to within an $O(\log\log|D|)$ factor.

The figure above illustrates four different objects that have been our focus of study for the last several weeks:

1. Elliptic curves $E/\mathbb{C}$ with CM by $\mathcal{O}$.
2. Lattices L (which define tori $\mathbb{C}/L$ that correspond to elliptic curves).
3. Proper $\mathcal{O}$-ideals $\mathfrak{a}$ (which may be viewed as lattices).
4. Primitive positive definite binary quadratic forms $ax^2 + bxy + cy^2$ of discriminant D (which correspond to proper $\mathcal{O}$-ideals of norm a).

Here $\mathcal{O}$ is an imaginary quadratic order of discriminant D.

In each case we have defined a notion of equivalence: isomorphism, homethety, equivalence modulo prinicipal ideals, and equivalence modulo an $\mathrm{SL}_2(\mathbb{Z})$-action, respectively, and modulo this equivalence we obtain a finite set of objects with the same cardinality $h(\mathcal{O}) = h(D)$ in each case. The two sets on the right, $\mathrm{cl}(\mathcal{O})$ and $\mathrm{cl}(D)$, are finite abelian groups that on the two sets on the left, both of which are equal to $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$. This action is free and transitive, so that $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ is a $\mathrm{cl}(\mathcal{O})$-torsor.

The integer polynomials $H_D(X)$ and $\Phi_N(X, Y)$ allow us to realize the CM torsor over any field k containing $\sqrt{D}$ where $H_D(X)$ splits completely: the roots of $H_D(X)$ form the set $\mathrm{Ell}_{\mathcal{O}}(k)$, and the action of $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O})$ sends $j(E) \in \mathrm{Ell}_{\mathcal{O}}(k)$ to a root of $\Phi_{N(\mathfrak{a})}(j(E), Y)$ that also lies in $\mathrm{Ell}_{\mathcal{O}}(k)$, via a cyclic isogeny of degree $N(\mathfrak{a})$.

References

[1] A. O. L. Atkin and Francois Morain, *Elliptic curves and primality proving*, Mathematics of Computation **61** (1993), 29–68.

[2] David A. Cox, *Primes of the form $x^2 + ny^2$: Fermat, class field theory, and complex multiplication*, Wiley, 1989.

[3] Andreas Enge and Andrew V. Sutherland, *Class invariants by the CRT method*, ANTS IX, LNCS 6197, Springer, 2010, pp. 142-156.

[4] Jürgen Neukirch, *Algebraic Number Theory*, Springer, 1999.

[5] Karl Rubin and Alice Silverberg, *Choosing the correct elliptic curve in the CM method*, Mathematics of Computation **79** (2010), 545–561.

[6] Joseph H. Silverman, *Advanced topics in the arithmetic of elliptic curves*, Springer, 1994.

[7] Andrew V. Sutherland, *Computing Hilbert class polynomials with the Chinese Remainder Theorem*, Mathematics of Computation **80** (2011), 501–538.

[8] Andrew V. Sutherland, *Accelerating the CM method*, LMS Journal of Computation and Mathematics **15** (2012), 172–204.

18.783 Elliptic Curves
Spring 2013