18.783 Elliptic Curves
Lecture 20

Andrew Sutherland

April 20, 2022

The modular polynomial $\Phi_N \in \mathbb{Z}[X, Y]$

In the last lecture we proved that $\mathbb{C}(\Gamma_0(N)) = \mathbb{C}(j, j_N)$, where $j_N(\tau) := j(N\tau)$.

Definition

The modular polynomial Φ_N is the minimal polynomial of j_N over $\mathbb{C}(j)$.

We may write $\Phi_N \in \mathbb{C}(j)[Y]$ as

$$\Phi_N(Y) = \prod_{i=1}^{n} (Y - j_N(\gamma_i \tau)),$$

where $\{\gamma_1, \ldots \gamma_n\}$ is a set of right coset representatives for $\Gamma_0(N)$.

The coefficients of $\Phi_N(Y)$ are symmetric polynomials in $j_N(\gamma_i \tau)$ and lie in $\mathbb{C}[j]$.

If we replace j by X we obtain a polynomial $\Phi_N(X, Y)$ whose coefficients lie in $\mathbb{Z}$.
It is a canonical plane (singular) model for the modular curve $X_0(N)$.

Isogenies

If $L_1 \subseteq L_2$ are lattices in $\mathbb{C}$, and $E_1 := E_{L_1}$ and $E_2 := E_{L_2}$ are the corresponding elliptic curves over $\mathbb{C}$, the inclusion $L_1 \subseteq L_2$ induces an isogeny $\phi\colon E_1 \to E_2$ whose kernel is isomorphic to the finite abelian group L_2/L_1.

$$\begin{array}{ccc} \mathbb{C}/L_1 & \xrightarrow{\iota} & \mathbb{C}/L_2 \\ \downarrow \simeq & & \downarrow \simeq \\ E_1(\mathbb{C}) & \xrightarrow{\phi} & E_2(\mathbb{C}) \end{array}$$

If we replace L_2 by the homothetic lattice NL_2, where $N = [L_2 : L_1] = \deg \phi$, the inclusion $NL_2 \subseteq L_1$ induces the dual isogeny $\hat{\phi}\colon E_2 \to E_1$ (up to isomorphism).

The composition $\phi \circ \hat{\phi}$ is the multiplication-by-N map on E_2, corresponding to the lattice inclusion $NL_2 \subseteq L_2$, with kernel $L_2/NL_2 \simeq \mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z} \simeq E_2[N]$.

Cyclic lattices and isogenies

Definition

If $L_1 \subseteq L_2$ with L_2/L_1 cyclic, then L_1 is a cyclic sublattice of L_2.
An isogeny $\phi\colon E_1 \to E_2$ is cyclic if its kernel is cyclic.

If ϕ is induced by $L_1 \subseteq L_2$ then ϕ is cyclic if and only if L_1 is a cyclic sublattice of L_2.

Every isogeny is a composition of cyclic isogenies (since prime degree implies cyclic).
We thus restrict our attention to cyclic sublattices of prime index.

Lemma

Let $L = [1, \tau]$ be a lattice with $\tau \in \mathcal{H}$ and let N be prime. The cyclic sublattices of L of index N are the lattice $[1, N\tau]$ and the lattices $[N, \tau + k]$, for $0 \le k < N$.

Proof: To the board!

Roots of the modular polynomial represent isogenies

Theorem

For all $j_1, j_2 \in \mathbb{C}$, we have $\Phi_N(j_1, j_2) = 0$ if and only if j_1 and j_2 are the j-invariants of elliptic curves over $\mathbb{C}$ over that are related by a cyclic isogeny of degree N.

Proof: To the board!

This theorem also applies to any field that can be embedded in $\mathbb{C}$, including all number fields. It can be extended via the Lefschetz principle to any field of characteristic zero, and as shown by Igusa, to fields of positive characteristic $p \nmid N$.

Theorem

Let $N > 1$ be an integer and let k be a field of characteristic not dividing N. For all $j_1, j_2 \in k$ we have $\Phi_N(j_1, j_2) = 0$ if and only if j_1 and j_2 are the j-invariants of elliptic curves over k that are related by a cyclic isogeny of degree N defined over k.

A few words of warning...

Remark

Over $\mathbb{C}$ we have $\Phi_N(j(E_1), j(E_2)) = 0$ if and only if E_1 and E_2 are related by a cyclic isogeny of degree N, but this is not true in general because $j(E_1) = j(E_2)$ only implies $E_1 \simeq E_2$ over algebraically closed fields. In general we may need to consider twists.

Remark

We should note that if $\phi\colon E_1 \to E_2$ is a cyclic N-isogeny, the pair of j-invariants $(j(E_1), j(E_2))$ does **not** uniquely determine ϕ, not even up to isomorphism.

Suppose $\mathrm{End}(E_1) \simeq \mathcal{O}$ and $\mathfrak{p} \neq \bar{\mathfrak{p}}$ is a proper $\mathcal{O}$-ideal of prime norm such that $[\mathfrak{p}]$ has order 2 in $\mathrm{cl}(\mathcal{O})$. Then $\mathfrak{p}E_1 \simeq \bar{\mathfrak{p}}E_1$ but $\phi_{\mathfrak{p}} : E_1 \to \mathfrak{p}E_1$ and $\phi_{\bar{\mathfrak{p}}} : E_1 \to \bar{\mathfrak{p}}E_1$ have distinct kernels and cannot be related by an isomorphism.

In this situation $\Phi_p(j(E_1), Y)$ will have $j(E_2)$ as a double root.

The polynomial $\Phi_N \in \mathbb{Z}[X, Y]$

The dual isogeny implies that $\Phi_N(j_1, j_2) = 0$ if and only if $\Phi_N(j_2, j_1) = 0$.
In fact $\Phi_N(X, Y) = \Phi_N(Y, X)$ is symmetric in the variables X and Y.

Theorem

$\Phi_N(X, Y) = \Phi_N(Y, X)$ *for all* $N > 1$.

Proof: To the board!

It follows that for prime N the polynomial $\Phi_N(X, Y)$ has degree $N + 1$ in X and Y.

Example

For $N = 2$ we have

$$\Phi_2(X, Y) = X^3 + Y^3 - X^2Y^2 + 1488(X^2Y + XY^2) - 162000(X^2 + Y^2) + 40773375XY + 8748000000(X + Y) - 157464000000000.$$

The bitsize of Φ_N is $O(N^3 \log N)$; Φ_{1009} is about 4 GB, and Φ_{10007} is about 5 TB.

Moduli spaces

In the same way that the j-function defines a bijection from $Y(1) = \mathcal{H}/\Gamma(1)$ to $\mathbb{C}$ (which we may regard as an affine curve in $\mathbb{C}^2$), the functions $j(\tau)$ and $j_N(\tau)$ define a bijection from $Y_0(N) = \mathcal{H}/\Gamma_0(N)$ to the affine curve $\Phi_N(X, Y) = 0$ via the map

$$\tau \mapsto \big(j(\tau), j_N(\tau)\big).$$

If $\{\gamma_k\}$ is a set of right coset representatives for $\Gamma_0(N)$ then for each γ_k we have

$$\gamma_k\tau \mapsto \big(j(\gamma_k\tau), j_N(\gamma_k\tau)\big) = \big(j(\tau), j_N(\gamma_k\tau)\big),$$

These points correspond to cyclic N-isogenies $E \to E'$ with $j(E) = j(\tau)$ and $j(E') = j_N(\gamma_k\tau))$. We can thus view the modular curve $Y_0(N)$, equivalently, the non-cuspidal points on $X_0(N)$, as parameterizing cyclic N-isogenies.

But recall our warning that the pair $(j(E), j(E'))$ do not uniquely determine $E \to E'$.

Moduli spaces

A cyclic N-isogeny $\phi\colon E \to E'$ is uniquely determined by a pair $(E, \langle P \rangle)$, where P is any generator for $\ker \phi$ (so P is a point of order N).

Every such pair $(E, \langle P \rangle)$ thus corresponds to a non-cuspidal point of $X_0(N)$.
Two pairs $(E, \langle P \rangle)$ and $(E', \langle P' \rangle)$ correspond to the same point if and only if there exists an isomorphism $\varphi\colon E \xrightarrow{\sim} E'$ such that $\varphi(\langle P \rangle) = \langle P' \rangle$.

The modular curve $X_0(N)$ is the moduli space of cyclic N-isogenies of elliptic curves, in which each non-cuspidal point represents an isomorphism class of pairs $(E, \langle P \rangle)$.

For $X(N)$ take isomorphism classes of triples (E, P_1, P_2), where $E[N] = \langle P_1, P_2 \rangle$.
For $X_1(N)$ take isomorphism classes of pairs (E, P), where $P \in E[N]$ has order N.
As above, these describe the non-cuspidal points, there are also cusps.

Elliptic curves with complex multiplication

Recall that for each imaginary quadratic order $\mathcal{O}$, we have the set

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

of isomorphism classes of elliptic curves with complex multiplication (CM) by $\mathcal{O}$.
Every elliptic curve $E/\mathbb{C}$ with CM by $\mathcal{O}$ is of the form $E_{\mathfrak{b}}$, where $\mathfrak{b}$ is a proper $\mathcal{O}$-ideal for which $j(\mathfrak{b}) = j(E)$ (note that $j(\mathfrak{b}) = j(E)$ depends only on the class $[\mathfrak{b}]$ in $\mathrm{cl}(\mathcal{O})$).
If $[\mathfrak{a}]$ is an element of $\mathrm{cl}(\mathcal{O})$, then $\mathfrak{a}$ acts on $E_{\mathfrak{b}}$ by the isogeny

$$\phi_{\mathfrak{a}} \colon E_{\mathfrak{b}} \to E_{\mathfrak{a}^{-1}\mathfrak{b}}$$

of degree $\mathrm{N}\mathfrak{a}$ induced by the lattice inclusion $\mathfrak{b} \subseteq \mathfrak{a}^{-1}\mathfrak{b}$. As with $E_{\mathfrak{b}}$, the isomorphism class of $E_{\mathfrak{a}^{-1}\mathfrak{b}}$ depends only on the class $[\mathfrak{a}^{-1}\mathfrak{b}]$ in $\mathrm{cl}(\mathcal{O})$, and we proved that this action is free and transitive, meaning that $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ is a $\mathrm{cl}(\mathcal{O})$-torsor.

The set $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ is finite, with cardinality equal to the class number $h(\mathcal{O}) := \#\,\mathrm{cl}(\mathcal{O})$.

The Hilbert class polynomial

Definition

Let $\mathcal{O}$ be an imaginary quadratic order of discriminant D. The polynomial

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

is the Hilbert class polynomial for $\mathcal{O}$ (and for D), a monic polynomial of degree $h(\mathcal{O})$. Its roots are the j-invariants of all elliptic curves with CM by $\mathcal{O}$.

Lemma

If N is prime then the leading term of $\Phi_N(X, X) \in \mathbb{Z}[X]$ is $-X^{2N}$.

Proof: To the board!

Remark

This lemma does not hold for general N.

The Hilbert class polynomial

Theorem

Let $\mathcal{O}$ be an imaginary quadratic order. Every ideal class in $\mathrm{cl}(\mathcal{O})$ contains infinitely many ideals of prime norm. **Proof**: See Theorems 7.7 and 9.12 in Cox.

Theorem

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

Proof: To the board!

Corollary

Let $E/\mathbb{C}$ be an elliptic curve with complex multiplication. Then $j(E) \in \overline{\mathbb{Z}}$.

The action of Galois

The groups $\mathrm{cl}(\mathcal{O})$ and $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ both act on the roots of $H_D(X)$.
How are these group actions related?

We consider $\mathrm{Gal}(L/K)$, where L is the splitting field of $H_D(X)$ over $K = \mathbb{Q}(\sqrt{D})$.
(we use K rather than $\mathbb{Q}$ because $\mathrm{Gal}(L/K)$ acts trivially on $\mathcal{O}$).

The first main theorem of complex multiplication states that $\mathrm{Gal}(L/K) \simeq \mathrm{cl}(\mathcal{O})$.

Let $\mathcal{O}$ be the imaginary quadratic order of discriminant D, and fix E_1 with CM by $\mathcal{O}$.

Each $\sigma \in \mathrm{Gal}(L/K)$ can be viewed as the restriction to L of an element of $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ that fixes K, and the elliptic curve E_1^σ also has CM by $\mathcal{O}$.

Thus $E_1^\sigma \simeq \mathfrak{a}E_1$ for some proper $\mathcal{O}$-ideal $\mathfrak{a}$, since $\mathrm{cl}(\mathcal{O})$ acts transitively on $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$.

The first main theorem of complex multiplication

If $E_2 \simeq \mathfrak{b}E_1$ is any other elliptic curve with CM by $\mathcal{O}$, then

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

(the innocent looking identity $(\mathfrak{b}E_1)^\sigma = \mathfrak{b}^\sigma E_1^\sigma$ is not immediate; see Silverman).
Thus the action of σ is the same as the action of $\mathfrak{a}$.

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 unique class $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O})$ for which $E_1^\sigma = \mathfrak{a}E_1$ defines a group homomorphism

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

This homomorphism is injective because, the only the identity in $\mathrm{Gal}(L/K)$ acts trivially on the roots of $H_D(X)$, and the same is true of $\mathrm{cl}(\mathcal{O})$. We have an embedding of $\mathrm{Gal}(L/K)$ in $\mathrm{cl}(\mathcal{O})$ that is compatible with the actions of both groups on $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$.

It remains only to prove that Ψ is surjective, equivalently, $H_D(X)$ is irreducible over K.