

Description

These problems are related to the material covered in Lectures 22-24. Problems 4 and 5 use material on the Weil pairing that is covered in the notes for Lecture 23 that we did not cover in class.

Instructions: Solve any combination of Problems that sum to 100 points. Your solutions are to be written up in latex and submitted as a pdf-file via [Gradescope](#) by **midnight** on the date due.

Collaboration is permitted/encouraged, but you must identify your collaborators or the name of your group on [pset partners](#), as well any references you consulted that are not posted on the course website (either in the [syllabus](#) or with the [lecture notes](#)). If there are none, write “**Sources consulted: none**” at the top of your solution. The first to spot each non-trivial typo/error in the problem sets or lecture notes will receive 1-5 points of extra credit. In cases where your solution involves code, please either include your code in your write up, or (better) the name of a notebook in your 18.783 CoCalc project containing you code (use a separate notebook for each problem).

This [Sage notebook](#) contains modular polynomials and helper functions from previous problem sets that you may find useful.

Problem 1. Isogeny volcanoes (98 points)

For the purposes of this problem, an isogeny volcano is an ordinary component of an ℓ -isogeny graph $G_\ell(\mathbb{F}_q)$ that does not contain 0 or 1728, where $\ell \nmid q$. This is a bi-directed graph that we regard as an undirected graph.

(a) Use the CM method to explicitly construct isogeny volcanoes that meet each of the following sets of criteria:

- (i) $\ell = 2$, $d = 3$, V_0 is a 5-cycle;
- (ii) $\ell = 3$, $d = 2$, V_0 contains a single edge;
- (iii) $\ell = 7$, $d = 1$, V_0 contains a single vertex with two self-loops.

In your answers, specify the finite field used, the discriminant of the order \mathcal{O}_0 corresponding to V_0 , and list each bi-directed edge just once, as a pair (v_1, v_2) of j -invariants corresponding to a horizontal or descending edge.

(b) Use the CM method to construct a single ordinary elliptic curve E/\mathbb{F}_q that simultaneously satisfies all of the following criteria:

- (i) $j(E)$ is on the floor of its 2-volcano, which has depth 6.
- (ii) $j(E)$ is on the surface of its 3-volcano, which has depth 3.
- (iii) $j(E)$ is on the middle level of its 5-volcano, which has depth 2.
- (iv) $j(E)$ is on the floor of its 7-volcano, which has depth 5.
- (v) $j(E)$ is one of exactly two vertices in its 11-volcano.


(vi) $j(E)$ is the only vertex in its 13-volcano.

In your answer, specify the finite field \mathbb{F}_q , the j -invariant $j(E)$, and the discriminant D of the order $\mathcal{O} \simeq \text{End}(E)$.

- (c) Prove that the cardinality of a 2-isogeny volcano with an odd number of vertices must be a Mersenne number (an integer of the form $2^n - 1$). Give an explicit example of a 2-isogeny volcano with 15 vertices.
- (d) Prove that every ordinary elliptic curve E/\mathbb{F}_q is isogenous to an elliptic curve E'/\mathbb{F}_q for which $E'(\mathbb{F}_q)$ is a cyclic group.

Problem 2. Computing modular polynomials (98 points)

As we have seen, the modular polynomials $\Phi_\ell(X, Y)$ play a key role in many theoretical and practical applications of elliptic curves. One can compute them using the q -expansions of the modular functions $j(z)$ and $j(\ell z)$, but this approach is difficult to implement efficiently and extremely memory intensive. In this problem you will implement a more efficient algorithm that uses isogeny volcanoes. The strategy is to use a CRT approach, working modulo primes p that are carefully selected to achieve a configuration of ℓ -volcanoes similar to that depicted below:



Here we have a configuration of three ℓ -volcanoes, with $\ell = 7$, each of depth $d = 1$. There are a total of $\ell + 2$ vertices on the surface (any value greater than $\ell + 1$ suffices).

Provided we have completely “mapped” this configuration of ℓ -volcanoes, meaning that we know the j -invariants of every vertex in the figure and the edges between them, we can compute $\Phi_\ell(X, Y)$ as follows. For any particular j -invariant j_i on the surface, we know the values of all the roots of $\phi_i(Y) = \Phi_\ell(j_i, Y)$, since we know the $\ell + 1$ neighbors of j_i in $G_\ell(\mathbb{F}_p)$. We can therefore compute each ϕ_i as the product of its linear factors. If we then consider the coefficient of Y^k in ϕ_i , we know (at least) $\ell + 2$ values c_{ik} of this coefficient, corresponding to $\ell + 2$ distinct j_i . This suffices to uniquely determine the polynomial $\psi_k(X)$ of degree at most $\ell + 1$ for which $\psi_k(j_i) = c_{ik}$, via Lagrange interpolation:

$$\psi_k(X) = \sum_{i=1}^{\ell+2} c_{ik} \prod_{m \neq i} \frac{(X - j_m)}{(j_i - j_m)}$$

- (a) Prove that $\Phi_\ell(X, Y) = \sum_{k=0}^{\ell+1} \psi_k(X) Y^k$.

To make things simpler, we will use a configuration with (at least) $\ell + 2$ isomorphic ℓ -volcanoes, each with one vertex on the surface and $\ell + 1$ neighbors on the floor. This can be achieved using a fundamental discriminant D with $(\frac{D}{\ell}) = -1$ and $h(D) \geq \ell + 2$. The vertex on the surface of each ℓ -volcano will have endomorphism ring equal to the maximal order for $K = \mathbb{Q}(\sqrt{D})$, and the vertices on the floor will then have endomorphism ring equal to the order \mathcal{O}' with discriminant $\ell^2 D$ (note that \mathcal{O}' has index ℓ in \mathcal{O}). For convenience, we will choose D so that both $\text{cl}(D)$ and $\text{cl}(\ell^2 D)$ are cyclic groups generated

by prime forms of norm $\ell_0 = 3$ (so we can use ℓ_0 -volcanoes of depth 0; see part (d) of Problem 1). This idealized setup is not always achievable, but it will work for our example using $\ell = 17$ and $D = -2339$, with class number $h(D) = 19$.

The key challenge is to map our set of ℓ -volcanoes without using the polynomial Φ_ℓ . Mapping the surface is easy: the vertices on the surface of our set of ℓ -volcanoes are the roots of the Hilbert class polynomial H_D (each root constitutes the surface of its own volcano). The vertices on the floor are the roots of the Hilbert class polynomial $H_{\ell^2 D}$, but this polynomial is much larger than H_D and we don't want to compute it, since it would take time $\tilde{O}(\ell^4)$. Instead we will use Vélú's formulas from Lecture 5 to compute a descending isogeny from each vertex on the surface. The kernel of this isogeny is a cyclic subgroup of $E[\ell]$, and Vélú's formulas require us to enumerate the points in the kernel, which may lie in an extension field of degree as large as $\ell^2 - 1$ (the degree of the ℓ -division polynomial). But we will choose primes $p \equiv 1 \pmod{\ell}$ that satisfy the norm equation $4p = t^2 - \ell^2 D$. This ensures that the elliptic curves E/\mathbb{F}_p with endomorphism ring \mathcal{O}_K have rational ℓ -torsion (provided we choose the correct twist); in this situation Vélú's formulas are very efficient.

- (b) With $\ell = 17$ and $D = -2339$, find a prime $p \equiv 1 \pmod{\ell}$ that satisfies $4p = t^2 - \ell^2 D$. Note that this requires $t \equiv \pm 2 \pmod{\ell}$, and with $t \equiv 2 \pmod{\ell}$ we will have $p + 1 - t$ divisible by ℓ^2 . Use Sage to compute the Hilbert class polynomial $H_D(X)$ and find the roots of $H_D \pmod{p}$. For each of the roots j_1, \dots, j_h of H_D , construct an elliptic curve E_i with j -invariant j_i , and attempt to find a point $P_i \in E(\mathbb{F}_p)$ with order ℓ by computing random $P_i = mP$ with $m = (p + 1 - t)/\ell^2$. If you find $P_i \neq 0$ and $\ell P_i \neq 0$ then you will need to replace E_i with a quadratic twist $y^2 = x^3 + d^2 A + d^3 B$, where $d \in \mathbb{F}_p^\times$ is any non-square.

We are now ready to apply Vélú's formulas to each pair (E_i, P_i) to obtain an ℓ -isogenous curve E'_i . Since every curve E'_i that is ℓ -isogenous to E_i lies on the floor, it does not matter which P_i we choose, any point of order ℓ will work. Below is a simplified algorithm that implements Vélú's formulas for the case where we have a cyclic subgroup generated by a point P of odd order on an elliptic curve given in short Weierstrass form $y^2 = x^3 + Ax + B$ over a finite field \mathbb{F}_p with $p > 3$.

1. Set $t \leftarrow 0$, $w \leftarrow 0$, and $Q \leftarrow P$.
2. Repeat $(\ell - 1)/2$ times:
 - a. Set $s \leftarrow 6Q_x^2 + 2A$, and then set $u \leftarrow 4Q_y^2 + sQ_x$.
 - b. Set $t \leftarrow t + s$, $w \leftarrow w + u$, and $Q \leftarrow Q + P$.
3. Set $A' = A - 5t$ and $B' = B - 7w$.
4. Output the curve E'/\mathbb{F}_p defined by $y^2 = x^3 + A'x + B'$.

In the description above Q_x and Q_y are the affine coordinates (x, y) of the point Q .

- (c) Implement the above algorithm and use it to compute elliptic curves E'_i that are ℓ -isogenous to the curves E_i you computed in step 2. Let j'_1, \dots, j'_h be the corresponding j -invariants.

Now comes the interesting part. We want to enumerate the vertices on the floor of our ℓ -volcano, but there are no horizontal ℓ -isogenies between vertices on the floor! Instead, we must go up to the surface and back down, which amounts to computing an isogeny of degree ℓ^2 . If we return to the same vertex this is just the multiplication-by- ℓ map (the composition of an ℓ -isogeny with its dual), but otherwise it is a cyclic isogeny of degree ℓ^2 , corresponding to the CM action of a proper \mathcal{O}' -ideal of norm ℓ^2 .

- (d) For $(\frac{D}{\ell}) = -1$, show that there are ℓ inequivalent integral primitive positive definite binary quadratic forms (ℓ^2, b, c) of discriminant $\ell^2 D$ (in our example these will all be reduced forms). These forms generate a cyclic subgroup G of $\text{cl}(\ell^2 D)$ of order $\ell + 1$. For $\ell = 17$ and $D = -2339$, determine a generator $f = (a, b, c)$ for G .

We don't want to use Φ_{ℓ^2} to compute the action of f (we don't even know Φ_{ℓ} yet!). But as in problem 1 of Problem Set 11, we can compute the action of an \mathcal{O}' -ideal of large norm using the action \mathcal{O}' -ideals of much smaller norm. In our example, we can use an \mathcal{O}' -ideal of norm $\ell_0 = 3$ to enumerate all the vertices on the floor of our set of volcanoes, and then determine the action of f by computing a discrete logarithm in $\text{cl}(\ell^2 D)$. Recall that we chose D so that a prime form of norm 3 generates $\text{cl}(\ell^2 D)$, so this is easy.

- (e) Use $\Phi_{\ell_0} = \Phi_3$ to enumerate all the vertices on the floor as a cycle of 3-isogenies.
- (f) Compute the discrete logarithm k of the form f from part (d) with respect to a prime form of norm $\ell_0 = 3$ in $\text{cl}(\ell^2 D)$. There is no need to distinguish inverses, and you should find that $(\ell + 1)k \equiv 0 \pmod{h(\ell^2 D)}$. Feel free to use brute force (a linear search); the time will be dominated by later steps in any case. Knowing k , you can now identify the subsets in the enumeration of part (e) that correspond to cosets of G . Each of these subsets will contain exactly one the j -invariants j'_i that you computed in step 3 and corresponds to the $\ell + 1$ "children" of j_i (its neighbors on the floor).
- (g) For each of $\ell + 2$ vertices j_i on the surface, compute the polynomial $\phi_i(Y) = \Phi_{\ell}(j_i, Y) = \prod_n (Y - j_{im})$, where the j_{im} range over the $\ell + 1$ children of j_i that you identified in part (f). Then, for k ranging from 0 to $\ell + 1$, interpolate the unique polynomial $\psi_k(X)$ of degree at most $\ell + 1$ for which $\psi_k(j_i)$ is equal to the coefficient of Y^k in $\phi_i(Y)$. You can do this with Sage: first create the polynomial ring `R.<X>=PolynomialRing(GF(p))`, and then use

```
R.lagrange_polynomial([(x0, y0), (x1, y1), ..., (xn, yn)])
```

to compute the unique polynomial $f(X)$ of degree at most n for which $f(x_i) = y_i$. Note that $\psi_{\ell+1}(X)$ must be the constant polynomial 1.

Finally, compute $\Phi_{\ell}(X, Y) = \sum_{k=0}^{\ell+1} \psi_k(X) Y^k \pmod{p}$. As a sanity check, verify that the coefficients are symmetric: $\Phi_{\ell}(X, Y) = \Phi_{\ell}(Y, X)$.

If you need to debug your algorithm, you may find it helpful to compute the Hilbert class polynomial $H_{\ell^2 D}(X)$ and then verify that the j -invariants j'_i computed in step 3 are actually roots of $H_{\ell^2 D} \pmod{p}$.

Provided that $D = O(\ell^2)$ and $\ell_0 = O(\log \ell)$, one can show that the algorithm you have implemented takes time $O(\ell^2 \log^3 p \log \log p)$, which is nearly optimal, since it is

quasi-linear in the size of $\Phi_\ell \bmod p$. By applying the same algorithm to a sufficiently large set of suitable primes p_i (it suffices to have $\sum \log p_i > 6\ell \log \ell + 18\ell$), one can then use the Chinese remainder theorem (as in problem 2 of Problem Set 11) to compute the coefficients of $\Phi_\ell \in \mathbb{Z}[X, Y]$. Under the GRH, the total time to compute Φ_ℓ over \mathbb{Z} is $O(\ell^3 \log^3 \ell \log \log \ell)$; see [1]. In practical terms, this algorithm can be used to compute Φ_ℓ even when ℓ is well into the thousands and Φ_ℓ is hundreds of gigabytes.

To convince ourselves that $\Phi_{17} \bmod p$ is correct, let's use it to compute a 17-volcano.

- (h) Using the same prime p , pick a different discriminant D for which $4p = t^2 - v^2D$, with $17 \nmid v$ and $\left(\frac{D}{17}\right) = 1$, such that $h(D) \geq 10$. Use Sage to find a root $j_0 \in \mathbb{F}_p$ of the Hilbert class polynomial $H_D(X) \bmod p$. Then use the polynomial $\Phi_{17}(X, Y) \bmod p$ to enumerate the vertices in the 17-volcano containing j_0 , which has depth 0 and degree 2 (since $\ell \nmid v$ and $\left(\frac{D}{17}\right) = 1$) and therefore consists of a single cycle. List the j -invariants of this cycle in order.
- (i) Let \mathfrak{a} be a prime ideal of norm 17 in the order \mathcal{O} of discriminant D . Compute the order of $[\mathfrak{a}]$ in $\text{cl}(\mathcal{O})$ and verify that it matches the length of the cycles you computed in part (h). You can construct the the order \mathcal{O} in Sage via `O=QuadraticField(D).maximal_order()` to create the order \mathcal{O} , then use `a=O.ideal(17).factor()[0][0]` to construct \mathfrak{a} .

Problem 3. Supersingular isogeny graphs (49 points)

Let p be and ℓ be distinct primes. Recall from Theorem 14.16 that the j -invariant of every supersingular elliptic curve over $\overline{\mathbb{F}}_p$ lies in \mathbb{F}_{p^2} . In this problem you will explore some properties of the supersingular components of $G_\ell(\mathbb{F}_{p^2})$.¹

- (a) Compute the graph of the component of $G_2(\mathbb{F}_{97^2})$ containing the supersingular j -invariant 1. You may wish to draw the graph on paper, but in your write-up just give a complete list of directed edges.
- (b) Prove that every supersingular vertex in $G_\ell(\mathbb{F}_{p^2})$ has out-degree $\ell + 1$, and conclude that no supersingular component of $G_\ell(\mathbb{F}_{p^2})$ is an ℓ -volcano. Show by example that the in-degree need not be $\ell + 1$.
- (c) Design an efficient *Las Vegas* algorithm that, given an arbitrary j -invariant in \mathbb{F}_{p^2} , determines whether it lies in an ordinary or supersingular component of $G_\ell(\mathbb{F}_{p^2})$ by detecting the difference between these components as abstract graphs. Prove that if $\ell = O(1)$ then the expected running time of your algorithm is $\tilde{O}(n^3)$, where $n = \log p$.²

The fastest known algorithms for computing the trace of Frobenius all have complexity $\Omega(n^4)$, so your algorithm provides a way to determine whether a given elliptic curve over a finite field is ordinary or supersingular that is asymptotically more efficient than checking whether the trace of Frobenius is divisible by p , and in practice, it should be *much* faster.

¹There is in fact only one supersingular component of $G_\ell(\mathbb{F}_{p^2})$, see [3, Cor. 78], but won't use this.

²As usual, the soft \tilde{O} -notation ignores factors that are polylogarithmic in n .

(d) By applying your algorithm to $G_2(\mathbb{F}_{p^2})$, determine which of the following j -invariants is supersingular. List the running time of your algorithm in each case.

(i) $p = 2^{64} + 81$:

```
p=2^64+81
R.<t> = PolynomialRing(GF(p))
F.<a> = GF(p^2, modulus=t^2+5)
j1=8326557536028784306*a + 13186271742734526835
j2=17095442389470987916*a + 5391379569813173462
j3=8201451720284342414*a + 1239990603471114829
j4=3832397532494683106*a + 3456346199771023610
j5=6995663267023152807*a + 5118305496003400382
```

(ii) $p = 2^{498}(2^{17} - 1) + 5^2 \cdot 11^2$:

```
p=2^498*(2^17-1)+5^2*11^2
F.<a>=GF(p^2)
j1=F(1068730309040382537178579357918315740437237673601\
46365282990696994391226239701748935923381766723513633\
617314116677847252974815762274295992015602852450016138)
j2=F(9307837638889485802864130889597342112431240717617\
79743203146570670576874073881819468942290046762690325\
81122360838583736151525289450839654218958090187901480)
```

Be patient, it may take a while for your program to run on the last two examples (but it should not take more than an hour).

Problem 4. Pairing attack on the discrete logarithm problem (49 points)

In the early days of elliptic curve cryptography supersingular curves were initially considered ideal candidates for discrete logarithm based cryptography (using a prime order subgroup of the rational points) because for these curves it is easy to determine the group order ($p + 1$ over prime fields for $p > 3$) and there are special techniques to speed up scalar multiplication. However, supersingular curves were quickly ruled out once it was discovered by Menezes, Okamoto, and Vanstone [4] that one can use the Weil pairing to reduce the computation of a discrete logarithm in an order n subgroup of $E(\mathbb{F}_q)$ to the computation of a discrete logarithm in a finite field \mathbb{F}_{q^k} that contains the group $\mu_n \subseteq \overline{\mathbb{F}}_q$ of n th roots of unity. As we saw in Lecture 10, there are subexponential-time algorithms to solve the discrete logarithm problem in a finite field, whereas no such algorithm is known for the discrete logarithm problem on an elliptic curve. In general k will be very large (exponential in $\log q$) and this reduction does not make the problem of computing discrete logarithms in $E(\mathbb{F}_q)$ any easier. But for supersingular curves this is not the case.

Let $p > 3$ be a prime, let E/\mathbb{F}_p be a supersingular curve, let $n > 2$ be a divisor of $p + 1$, and let μ_n denote the multiplicative group of n th roots of unity in $\overline{\mathbb{F}}_p$.

(a) Prove that $\mu_n \not\subseteq \mathbb{F}_p^\times$ and $E[n] \not\subseteq E(\mathbb{F}_p)$, but $\mu_n \subseteq \mathbb{F}_{p^2}^\times$ and $E[n] \subseteq E(\mathbb{F}_{p^2})$.

Let $P \in E(\mathbb{F}_p)$ be a point of prime order n , and let $Q \in \langle P \rangle$.³ Consider the following algorithm Las Vegas algorithm to compute $\log_P Q$:

³For composite n one uses the Pohlig-Hellman algorithm (see Lecture 9) to reduce to the prime case.

1. Generate a random point $R \in E(\mathbb{F}_{p^2})$ and compute $S = mR$ with $m := \#E(\mathbb{F}_{p^2})/n$.
2. Compute $a = e_n(S, P)$, and if $a = 1$ then return to step 1.
3. Compute $b = e_n(S, Q)$.
4. Compute $\log_a b$ in $\mathbb{F}_{p^2}^\times$ and output the result.

(b) Prove that the expected number of times the algorithm executes step 1 is $1 + o(1)$.

(c) Prove that the algorithm outputs $\log_P Q$.

The expected running time the algorithm is completely dominated by the time to compute $\log_a b$ in $\mathbb{F}_{p^2}^\times$, which is heuristically $L[1/3, c]$.

(d) Let E be an elliptic curve over a number field L with CM by a maximal order in an imaginary quadratic field K of discriminant D and let \mathfrak{p} be a prime ideal of \mathcal{O}_L with prime norm p such that E has good reduction at \mathfrak{p} . Prove that if $\left(\frac{D}{p}\right) = -1$ then the reduction of E modulo \mathfrak{p} is a supersingular elliptic curve over \mathbb{F}_p .

Let $n = 10^{14} + 2367$ (which is prime) and let $p = 2n - 1$ (also prime).

(e) Construct an elliptic curve E/\mathbb{F}_p such that $E[n] \subseteq E(\mathbb{F}_{p^2})$.

Representing \mathbb{F}_p by integers in $[0, p-1]$, choose $P_0 \in E(\mathbb{F}_p)$ so that $x(P_0)$ is the least integer greater than your student ID and $y(P_0)$ is minimal, and let $P = 2P_0$. In the unlikely event that $P = 0$, increment your student ID and repeat until $P \neq 0$. Then choose $Q_0 \in E(\mathbb{F}_p)$ so that $x(Q_0)$ is the least integer greater than twice your student ID and $y(Q_0)$ is minimal, and let $Q = 2Q_0$. Then both P and Q lie in $E[n]$ and you can use the above algorithm to compute $\log_P Q$, with $m = 2$. Here are a few tips to help you:

- To create the field \mathbb{F}_{p^2} in Sage use `F2 = GF(p**2)`.
- To generate R use `E.change_ring(F2).random_element()`.
- To compute $e_n(S, P)$ use `S.weil_pairing(P.change_ring(F2), N)`.
- To compute $\log_a b$ use `b.log(a)`.

(f) Compute $\log_P Q$. In your answer list the points P_0, P, Q_0, Q , the values of a and b , the integer $n = \log_a b$, and the running time of your algorithm (which should be well under a minute). Be sure to check your answer by verifying that $nP = Q$.

Remark. You should not be particularly impressed by this running time, since you can easily beat it with a careful implementation of the baby-steps giant-steps or Pollard rho algorithms. But for larger values of p and N this algorithm will easily outperform any generic method. Unfortunately Sage does not have a particularly fast implementation for discrete logarithms in non-prime finite fields, so I intentionally chose a small example.

Problem 5. A fast Las Vegas algorithm to compute $E(\mathbb{F}_p)$ (49 points)

Problem 2 of Problem Set 5 gave a Las Vegas algorithm to compute the structure of the group $E(\mathbb{F}_p) \simeq \mathbb{Z}/m\mathbb{Z} \oplus \mathbb{Z}/n\mathbb{Z}$, but its running time was $\exp(\frac{1}{4} \log p + o(1))$, exponential in $\log p$. In this problem, following Miller [5], you will develop a much faster algorithm to compute $E(\mathbb{F}_p)$. Strictly speaking it is not polynomial-time because it requires the factoring the integer $d = \gcd(\#E(\mathbb{F}_p), p-1)$, but typically d will either be small (in which case the problem is easy), or it will have only one large prime factor ℓ , in which case factoring d is easy but computing the structure of $E(\mathbb{F}_p)$ with the algorithm from Problem Set 5 will be very difficult if ℓ^2 divides $\#E(\mathbb{F}_p)$. In any case, this does give a subexponential-time Las Vegas algorithm, since we can always factor d in subexponential time using a Las Vegas algorithm (the best proven bound is $L[1/2, c]$, but heuristically this can be done in time $L[1/3, c]$).

(a) Let $\ell \neq p$ be prime. Prove that if $E[\ell] \subseteq E(\mathbb{F}_p)$ then $\ell | (p-1)$ and $\ell^2 | \#E(\mathbb{F}_p)$.

Let $N = \#E(\mathbb{F}_p)$ and write N as $N = N_0 N_1$, where N_0 and N_1 are relatively prime and N_1 is divisible only by primes ℓ that divide $p-1$ and whose square divides N . By part (a), when computing the structure of $E(\mathbb{F}_p)$, we can restrict our attention to $E(\mathbb{F}_p)[N_1]$. Consider the following algorithm to compute the structure of $E(\mathbb{F}_p)$, which takes as input the elliptic curve E/\mathbb{F}_p , the integers N_0 and N_1 , and the prime factorizations of N_1 .

1. Generate random points $P_0, Q_0 \in E(\mathbb{F}_p)$ and put $P := N_0 P_0$ and $Q := N_0 Q_0$.
2. Using the prime factorization of N_1 , compute $s := |P|$ and $t := |Q|$.
3. Let $r = \text{lcm}(s, t)$ and compute $\zeta := e_r(P, Q)$.
4. Using the prime factorization of N_1 , compute $n := |\zeta|$.
5. If $rn = N_1$ then put $m = rN_0$ and output $\mathbb{Z}/m\mathbb{Z} \oplus \mathbb{Z}/n\mathbb{Z}$, otherwise go to step 1.

(b) Prove that when the algorithm terminates we have $E(\mathbb{F}_p) \simeq \mathbb{Z}/m\mathbb{Z} \oplus \mathbb{Z}/n\mathbb{Z}$.

(c) Prove that in step 3 we have $\zeta \in \mathbb{F}_p^\times$.

(d) Prove that the expected number of times the algorithm repeats step 1 is $O(\log \log p)$. You may use Mertens' bound $\sum_{\ell \leq x} \frac{1}{\ell} = \log \log x + O(1)$, where ℓ ranges over primes.⁴

(e) Let $g \in G$ be an element of a generic group with exponent λ (so $\lambda g = 0$). Prove that, given the prime factorization of λ you can compute the order of g using $(\log \lambda)^{1+o(1)}$ group operations. (You may wish to review Lecture 9 on generic algorithms).

(f) Prove that the running time of the algorithm above is $(\log p)^{2+o(1)}$.

Remark. As noted above, this only gives a subexponential-time algorithm to compute $E(\mathbb{F}_p)$ if we are not given the factorization of N_1 . But it is known that we can do this in *average polynomial time* [2], in the following sense: for any prime p , if we pick a random $A, B \in \mathbb{F}_p$ the expected time to compute $E(\mathbb{F}_p) \simeq \mathbb{Z}/m\mathbb{Z} \oplus \mathbb{Z}/n\mathbb{Z}$ for the curve $E: y^2 = x^2 + Ax + B$ is polynomial in $\log p$.

⁴One can modify the algorithm so that the expected number of repeats is actually $O(1)$.

Problem 6. The Birch and Swinnerton-Dyer Conjecture (49 points)

The goal for this problem is to lead you to a formulation of the (weak) Birch and Swinnerton-Dyer conjecture. The main difficulty here is not for you to prove a precisely stated question, but rather for you to formulate a precise question, starting from some data and some general suggestions. This means that parts of the problem are deliberately vague; making the questions more precise is part of the problem. Other than part (a), you are not expected to prove anything; heuristic arguments are fine.

One of the outstanding [Millennium Prize Problems](#) is the famous conjecture of Birch and Swinnerton-Dyer, which concerns the ranks of elliptic curves over \mathbb{Q} . Recall from Lecture 1 that the Mordell-Weil theorem implies

$$E(\mathbb{Q}) = E(\mathbb{Q})_{\text{tors}} \oplus \mathbb{Z}^r.$$

As opposed to the torsion subgroup, the rank r of the Mordell-Weil group $E(\mathbb{Q})$ is far less understood; we do not have a fully general algorithm to compute r , and mathematicians do not even agree on whether r should be bounded or not (so not only can't we prove a conjecture, we don't even know what the right conjecture is!). However, in the 1960s, Birch and Swinnerton-Dyer carried out computations on the EDSAC computer at Cambridge University that led them to conjecture a deep relationship between the L -series of E and the rank r of the Mordell-Weil group. In this problem you will develop a conjecture and investigate the evidence for it.

Recall from Lecture 24 that the L -series of E/\mathbb{Q} is defined by

$$L_E(s) = \prod_p L_p(p^{-s})^{-1} = \prod_p (1 - a_p p^{-s} + \chi(p) p^{-2s+1})^{-1},$$

where the Dirichlet character $\chi(p) = 0$ if E has bad reduction at p and 1 otherwise.⁵ This converges for $\Re(s) > 3/2$; however by the modularity theorem, it admits an analytic continuation to the entire complex plane.

- (a) Assume that E is given in affine coordinates by the equation $y^2 = x^3 + Ax + B$ with $A, B \in \mathbb{Z}$.⁶ A rational point $(x_0 : y_0 : z_0) \in E(\mathbb{Q})$ gives a solution to

$$y^2 z \equiv x^3 + Axz^2 + Bz^3 \pmod{p^n},$$

and hence a point on E modulo p^n for all $p \nmid \Delta(E)$ and $n \geq 1$. Let N_{p^n} denote the number of projective solutions $(x_0 : y_0 : z_0)$ to this congruence (over $\mathbb{Z}/p^n\mathbb{Z}$). Using Hensel's lemma (see problem 5 of Problem Set 2) prove that

$$N_{p^n} = p^{n-1} N_p$$

for all $p \nmid \Delta(E)$. Conclude that

$$\lim_{n \rightarrow \infty} \frac{N_{p^n}}{p^n} = \frac{N_p}{p},$$

except for finitely many p . This quantity $\lim_{n \rightarrow \infty} \frac{N_{p^n}}{p^n}$ represents the density of p -adic points on E . Give a plausible relation with r .

⁵Recall that this assumes a minimal Weierstrass model for E .

⁶This need not be a minimal Weierstrass model, but you may assume it is if you wish.

- (b) Let S be the set of primes of bad reduction of E . Define

$$f_E(X) = \prod_{\substack{p \notin S \\ p \leq X}} \frac{N_p}{p},$$

What is the relationship between $f_E(X)$ and $L_E(s)$?

- (c) Consider the elliptic curve⁷ E with rank 3 given by

$$y^2 = x^3 - 82x.$$

Compute and plot the values of

$$f_E(X) = \prod_{\substack{p \notin S \\ p \leq X}} \frac{N_p}{p}$$

for X up to at least 10^6 (or further if you like). You can use the Sage method `E.aplist` to efficiently compute a list of a_p values at all primes $p \leq X$ (including primes of bad reduction). For primes of good reduction a_p is the trace of Frobenius of $E \bmod p$, from which you can derive N_p . What appears to be the asymptotic growth of the function $f_E(X)$? Make a plot of your results with an appropriate choice of scale on the coordinate axes so your answer is apparent. Attach relevant plots in your solutions.

- (d) Repeat part (c) for the following curves of the form

$$E_i : y^2 = x^3 - d_i^2 x,$$

for the values:

	d_i	rank
E_1	1	0
E_2	5	1
E_3	34	2
E_4	1254	3
E_5	29274	4

Display your final plots (with the appropriate scaling of the axes) together.

- (e) Combining your results from parts (b), (c), and (d) above, make a precise conjecture on the relationship between $L_E(s)$ and the rank of E . Your conjecture should be precise enough that different ranks give rise to different behaviors of the L -function. (Hint: You might need to do more work in (b) in order to give a more precise relationship – there is a natural interplay between conjecture and computation.)
- (f) Choose 5 random elliptic curves with $|A|, |B| < 100$ and conjecturally assign their ranks. Note: you can use Sage to check your answer, but you must provide evidence based upon your work in previous parts to receive full credit.⁸

⁷In 1938, this was the highest rank known, due to Billings.

⁸While we do not have a general algorithm to compute the rank of E/\mathbb{Q} , for small $|A|$ and $|B|$, Sage can easily do so.

Problem 7. Survey (2 points)

Complete the following survey by rating each problem you attempted on a scale of 1 to 10 according to how interesting you found it (1 = “mind-numbing,” 10 = “mind-blowing”), and how difficult you found it (1 = “trivial,” 10 = “brutal”). Also estimate the amount of time you spent on each problem to the nearest half hour.

	Interest	Difficulty	Time Spent
Problem 1			
Problem 2			
Problem 3			
Problem 4			
Problem 5			
Problem 6			

Also, please rate each of the following lectures that you attended, according to the quality of the material (1=“useless”, 10=“fascinating”), the quality of the presentation (1=“epic fail”, 10=“perfection”), the pace (1=“way too slow”, 10=“way too fast”, 5=“just right”) and the novelty of the material (1=“old hat”, 10=“all new”).

Date	Lecture Topic	Material	Presentation	Pace	Novelty
5/12	Isogeny volcanoes				
5/15	Modular forms				
5/17	Fermat’s Last Theorem				

Finally, if you have any comments about the course as a whole, in particular, things that you think would improve it in the future, please let me know!

References

- [1] R. Bröker, K. Lauter, and A.V. Sutherland, *Modular polynomials via isogeny volcanoes*, Mathematics of Computation **81** (2012), 1201–1231.
- [2] J.B. Friedlander, C. Pomerance, I.E. Shparlinski, *Finding the group structure of elliptic curves over finite fields*, Bulletin of the Australian Mathematical Society **72** (2005), 251–263.
- [3] David Kohel, *Endomorphism rings of elliptic curves over finite fields*, PhD thesis, University of California at Berkeley, 1996.
- [4] A.J. Menezes, T. Okamoto, and S.A. Vanstone, *Reducing elliptic curve logarithms to logarithms in a finite field*, IEEE Transactions on Information Theory **39** (1993), 1639–1646.
- [5] V.S. Miller, *The Weil pairing and its efficient calculation*, J. Cryptology **17** (2004), 235–261.
- [6] L.C. Washington, *Elliptic curves: Number theory and cryptography*, second edition, CRC Press, 2008.