

Computing newforms using supersingular isogeny graphs

Alex Cowan
cowan@math.harvard.edu

Abstract

We describe an algorithm that we used to compute the q -expansions of all weight 2 cusp forms of prime level at most 2,000,000 and dimension at most 6. We also present an algorithm that we used to verify that there was only one cusp form of dimension 7 or more per Atkin-Lehner eigenspace for prime levels between 10,000 and 1,000,000. Our algorithm is based on Mestre's Méthode des Graphes, and involves supersingular isogeny graphs and Wiedemann's algorithm for finding the minimal polynomial of sparse matrices over finite fields.

Contents

1	Introduction	2
2	Acknowledgments	3
3	Background	3
3.1	Wiedemann's algorithm	3
3.2	Newforms	4
3.3	La Méthode des Graphes	5
4	Overview of the algorithm	5
5	Computing the action of T_ℓ on $S_2(p)$	6
5.1	Finding a starting vertex	6
5.2	Exploring the graph	6
6	Computing the characteristic polynomial of T_2 mod ν	7
6.1	Shifting eigenvalues	7
6.2	Varying parameters	7
6.3	Getting the characteristic polynomial from the minimal polynomial	8
6.3.1	Comparing with known top coefficients	8
6.3.2	Checking eigenspace dimensions	9
6.4	Storing iterates	9
7	Finding \mathbb{Z}-eigenbases	9
7.1	Finding eigenvalues	10
7.2	Lifting 1-dimensional eigenspaces	10
7.3	Lifting higher dimensional eigenspaces	11
7.3.1	Finding an \mathbb{F}_ν -eigenbasis	11
7.3.2	Finding a \mathbb{Z} -eigenbasis for the full ρ -eigenspace	12
7.3.3	Finding a \mathbb{Z} -eigenbasis for each Galois orbit	12

8	Computing q-expansions	12
8.1	\mathbb{Z} -eigenbasis version of Mestre’s identity	13
8.2	Computing $j(q) \bmod p$	13
8.3	Fast power series algorithms	14
8.3.1	Atkin-Lehner eigenspace to Mestre’s formula	14
8.3.2	Evaluating Mestre’s formula	14
9	Checking for high degree factors of the characteristic polynomial	15
9.1	Factoring the characteristic polynomial modulo many small primes	15
9.2	Sieving possible degrees	15
9.3	Checking the Weil bound	16
10	Results	17

1 Introduction

Let $S_2(p)$ be the space of weight 2 cusp forms of prime level p . We say that the *dimension* of a newform of $S_2(p)$ is the degree of the number field its Hecke eigenvalues generate, or, equivalently, the size of its Galois orbit. There have been many efforts from computational number theorists to create databases with information about the newforms of $S_2(p)$; we highlight the Antwerp tables [6], Cremona’s database of elliptic curves [15], and the LMFDB [24]. The LMFDB builds on the previous two, and currently lists the q -expansions of every newform of $S_2(p)$ of dimension g at most 20 and level at most 10,000 [3].

In this paper, we describe an algorithm that we used to compute the q -expansions of all newforms with $g \leq 6$ and $p < 2,000,000$ up to the Sturm Bound, and also used to verify that there were no eigenforms with $g \geq 7$ and $10,000 < p < 1,000,000$ besides one factor of high dimension per Atkin-Lehner eigenspace. For each level p , our algorithm runs in time $\mathcal{O}(p^{2+\epsilon})$ and space $\mathcal{O}(p^{1+\epsilon})$.

In the field of arithmetic statistics there is a lot of interest in understanding how various properties of these eigenforms are distributed. This is partly because modular forms are interesting in their own right, but also partly because standard modularity conjectures [18] predict that, for each genus g factor of the modular jacobian $J_0(p)$, there is an associated weight 2 newform of level p and dimension g . The genus 1 case of elliptic curves has been studied extensively, and is one of the most important topics in modern number theory. The association between elliptic curves and 1-dimensional modular forms is a theorem [39, 36, 8], and the literature contains conjectures and theorems for how many related invariants are distributed, notably ranks [29, 4], Selmer groups [5], torsion subgroups [21], and other numbers which appear in the Birch and Swinnerton-Dyer formula [30]. Generalizations of these theorems to genus 2 or more are far out of reach for the most part, and in many situations it is poorly understood what the correct generalizations would even be. In particular, merely predicting the number of genus g factors of $J_0(p)$ has not been done whenever $g \geq 2$, whereas there are well established conjectures for the number of elliptic curves with bounded conductor [10, 37]. In light of this gap in understanding, databases of newforms of $S_2(p)$ are very useful: they give a way to observe generalizations of phenomena which occur in the genus 1 case, and they also allow one to formulate conjectures about the statistics of these objects. In [26], for instance, Martin computes the dimensions of the eigenforms of $S_2(p)$ for $p < 60000$ and uses this data to formulate conjectures related to counts of eigenforms of fixed dimension.

The main idea in our algorithm is from Mestre’s *Méthode des Graphes* [27]. In Mestre’s work, he relates the q -expansion of weight 2 newforms of prime level to “supersingular isogeny graphs”. The *supersingular ℓ -isogeny graph over $\overline{\mathbb{F}}_p$* is the graph whose vertices are supersingular j -invariants over $\overline{\mathbb{F}}_p$, and whose edges are ℓ -isogenies. These graphs have recently been of independent interest because of their applications to cryptography [11, 22, 19]. The relationship Mestre highlights depends on a trace formula: the action of the Hecke operator T_ℓ on the space $S_2(p)$ can be represented as the adjacency matrix of the supersingular ℓ -isogeny graph. We find simultaneous eigenvectors of these matrices, and then use a formula from Mestre’s work to compute the associated q -expansions.

Some of the building blocks of our algorithm come from more general-purpose techniques. In particular, we

use Wiedemann’s algorithm for finding the minimal polynomial of a sparse matrix over a finite field [38], Brent and Kung’s algorithm for fast power series composition [7], and a well known dynamic programming algorithm for solving the subset sum problem [14].

We implemented most of the algorithm in Sage. We used cython to multiply vectors by sparse matrices, and we used c++’s NTL package when doing power series manipulations. We ran our code on the Oklahoma University Supercomputing Center for Education & Research.

2 Acknowledgments

I’m extremely grateful to Noam Elkies for both suggesting I work on this problem, and for telling me about many of the ideas and techniques used in this paper. I thank Kimball Martin for many helpful discussions regarding ways to use the data, and for running and managing the computations. The computing for this project was performed at the OU Supercomputing Center for Education & Research (OSCER) at the University of Oklahoma (OU). I thank Drew Sutherland for help in adding the data to the LMFDB. I thank David Roe for helping me write some of the code. This work was done under Simons Collaboration grant number 550031.

3 Background

3.1 Wiedemann’s algorithm

We use [38] as a reference for what we call “Wiedemann’s algorithm”. Given an $n \times n$ nonsingular matrix M over a finite field F , Wiedemann gives a probabilistic algorithm for finding the minimal polynomial μ of M . Wiedemann’s algorithm is one of the major building blocks of the algorithm presented in this paper. See section 6 for details.

Let u be a vector in F^n , and let $i \in [1, n] \cap \mathbb{Z}$ be one of the indices of the coordinates of u . The key idea in Wiedemann’s algorithm is that the sequence $u_i, (Mu)_i, (M^2u)_i, \dots$ will satisfy a recursion relation, and that recursion relation will, for most choices of u and i , give the minimal polynomial μ .

For a given choice of u and i , computing the sequence $u_i, (Mu)_i, (M^2u)_i, \dots, (M^ru)_i$ takes time $\mathcal{O}(r\omega)$, where ω is the number of nonzero entries in M . Wiedemann then uses the following:

Proposition 3.1. *If $\mu(t) = \sum_{k=0}^n \mu_n t^k$ is the minimal polynomial of M , then, for any vector u , we have*

$$\sum_{k=0}^n \mu_n M^k u = 0.$$

As a consequence of proposition 3.1, we know that the sequence $u_i, (Mu)_i, (M^2u)_i, \dots, (M^ru)_i$ satisfies a recursion relation of length at most n . To determine what the recursion relation is, we’ll need a number of terms at least double the recursion length. Thus, we can take $r = 2n$. In our application, M will have $\mathcal{O}(n)$ nonzero entries, so we’ll compute the sequence $u_i, (Mu)_i, (M^2u)_i, \dots, (M^{2n}u)_i$ in time $\mathcal{O}(n^2)$.

Wiedemann then uses an algorithm of Berlekamp-Massey to find the coefficients of the recursion relation of the sequence $u_i, (Mu)_i, (M^2u)_i, \dots, (M^{2n}u)_i$. This takes time $\mathcal{O}(n^2)$. See [28] for a description of the Berlekamp-Massey algorithm in terms of continued fractions. One step in this algorithm involves writing $\frac{1}{\mu(t)}$ as a power series. If M is not invertible, then $\mu(t)$ will be divisible by t , making it impossible to do this. It’s possible to circumvent this problem in a number of ways. We chose to modify our matrices so that they would be invertible (see section 6.1), but Wiedemann gives a modification of his algorithm for this case in [38], and it’s also possible to modify the Berlekamp-Massey algorithm directly.

3.2 Newforms

We give [16] as a reference for this section. Define $\Gamma_0(p)$ to be the group

$$\Gamma_0(p) := \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}) : c = 0 \pmod{p} \right\}.$$

The group $\Gamma_0(p)$ acts on the upper half-plane $\mathbb{H} := \{x + iy : x \in \mathbb{R}, y \in \mathbb{R}_{>0}\}$ via Möbius transformations:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} z := \frac{az + b}{cz + d}.$$

A *weight 2 modular form on $\Gamma_0(p)$* is a holomorphic function $f : \mathbb{H} \rightarrow \mathbb{C}$ which satisfies the relation

$$f\left(\begin{pmatrix} a & b \\ c & d \end{pmatrix} z\right) = (cz + d)^2 f(z)$$

for every $z \in \mathbb{H}$ and every $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \Gamma_0(p)$. Weight 2 modular forms on $\Gamma_0(p)$ form a finite-dimensional complex vector space.

Because $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ is in $\Gamma_0(p)$, we have $f(z + 1) = f(z)$. It follows that modular forms have Fourier expansions, i.e. there exist complex numbers $a_n(f)$ such that

$$f(z) = \sum_{n=0}^{\infty} a_n(f) e^{2\pi i n z}.$$

The space $S_2(p)$ of cusp forms is the subspace of these modular forms which have $a_0 = 0$. Throughout this paper we use the shorthand $q := e^{2\pi i z}$, and we'll call these Fourier expansions *q-expansions*. We'll write a_n instead of $a_n(f)$ when the modular form f is clear from context.

The *Hecke operators* are linear operators T_n indexed by positive integers n which act on the space of modular forms as

$$a_m(T_n f) = \sum_{d|\mathrm{gcd}(m,n)} d \cdot a_{mn/d^2}(f).$$

The Hecke operators commute with one another, and hence they are simultaneously diagonalizable. Thus, there exist modular forms f with $a_1 = 1$ which, for all n simultaneously, satisfy the relations

$$T_n f = a_n f.$$

Modular forms with these properties are called *newforms*. They form a basis for the space $S_2(p)$.

In [34], Sturm proves the following:

Theorem 3.2 (Sturm bound). *If f and g are weight 2 newforms of level p and $a_n(f) = a_n(g)$ for all $n \leq \lfloor \frac{p+1}{6} \rfloor$, then $f = g$.*

There's are analogous results for other weights and levels as well.

There is a linear operator w_p called the *Atkin-Lehner involution* which acts on $S_2(p)$. This operator commutes with all of the Hecke operators, so newforms are also eigenforms of the Atkin-Lehner involution. As suggested by the name, the Atkin-Lehner involution is an involution. Thus, if f is a newform, then $w_p f = \pm f$, and $S_2(p)$ decomposes into two Atkin-Lehner eigenspaces (of roughly equal size; see [25]).

3.3 La Méthode des Graphes

The *supersingular j -invariants over $\overline{\mathbb{F}}_p$* are the j -invariants of the elliptic curves defined over $\overline{\mathbb{F}}_p$ which are supersingular (i.e. their endomorphism ring is an order in a quaternion algebra). It's known that there are $\lfloor \frac{p}{12} \rfloor + 0, 1, \text{ or } 2$ supersingular j -invariants over $\overline{\mathbb{F}}_p$, and that they're all defined over \mathbb{F}_{p^2} [32]. The *supersingular ℓ -isogeny graph over $\overline{\mathbb{F}}_p$* is the directed multigraph whose vertices are the supersingular j -invariants over $\overline{\mathbb{F}}_p$, and whose edges correspond to ℓ -isogenies over $\overline{\mathbb{F}}_p$ between the associated elliptic curves. These graphs have been and continue to be studied extensively [23, 35, 11, 22, 19], in part because of their potential applications to post-quantum cryptography.

As described in [20, §2,5], [27], and [17], one description of the action of the Hecke operator T_ℓ on $S_2(p)$ is as the adjacency matrix of the supersingular ℓ -isogeny graph over $\overline{\mathbb{F}}_p$. This connection involves a trace formula and an equivalence of categories between supersingular elliptic curves and orders in quaternion algebras.

As a consequence of the representation of T_ℓ as the adjacency matrix of the supersingular ℓ -isogeny graph over $\overline{\mathbb{F}}_p$, there is a bijection between newforms f of $S_2(p)$ and vectors $v = (v_j)_j$ with coordinates indexed by the supersingular j -invariants over $\overline{\mathbb{F}}_p$ which are simultaneous eigenvectors of all the Hecke operators T_ℓ . Mestre [27] uses this bijection to produce the identity of power series

$$\left(\sum_j v_j j \right) f(q) \frac{dq}{q} = \sum_j v_j \frac{dj(q)}{j(q) - j} \pmod{\mathfrak{p}},$$

where

- $j(q)$ is the modular j function,
- \mathfrak{p} is any prime above p in the number field generated by the Hecke eigenvalues a_ℓ , and
- the sums are over the supersingular j -invariants over $\overline{\mathbb{F}}_p$.

The Weil bound for f states that $|a_n| < 2\sqrt{n}$, so for $n \ll p^2$ this equality of power series is enough to know the values of a_n exactly. Since the Sturm bound is $\mathcal{O}(p)$, this identity is enough to distinguish newforms.

4 Overview of the algorithm

The algorithm is broken into five sections:

- Section 5: Computing the action of T_2 on $S_2(p)$.
- Section 6: Computing the characteristic polynomial of T_2 modulo some small auxiliary prime.
- Section 7: Determining the eigenvalues of degree 6 or less T_2 , and finding eigenbases over \mathbb{Z} whenever they'll correspond to newforms of dimension 6 or less.
- Section 8: Computing q -expansions using Mestre's formula.
- Section 9: Verifying that only one high genus factor per Atkin-Lehner eigenspace exists.

In section 5, we compute a representation of the action of T_2 on $S_2(p)$, or, more precisely, its action on each Atkin-Lehner eigenspace, by constructing a supersingular isogeny graph. We use "modular polynomials" to find edges, and explore the graph using a breath first search.

In section 6, we compute the characteristic polynomial χ_ν of T_ℓ modulo some small auxiliary prime ν . The main ingredient in this step is Wiedemann's algorithm [38] for computing minimal polynomials of sparse matrices defined over finite fields. This part of the algorithm contributes to the leading term in the overall asymptotic time complexity.

In section 7, we use the characteristic polynomial χ_ν to find the eigenvalues and the eigenvectors of T_ℓ which correspond to low degree factors of its characteristic polynomial $\chi_{\mathbb{Z}}$ over \mathbb{Z} . Our method for “lifting eigenspaces” in this way is based on the heuristic that each low degree factor of $\chi_{\mathbb{Z}}$ has an associated eigenbasis made up of vectors whose coordinates are small.

In section 8, we use a formula from Mestre’s Méthode des Graphes to get the q -expansion of the newforms in terms of a power series involving the previously computed eigenvectors. In evaluating this formula, we used an algorithm from [7] for composing power series.

Finally, in section 9, we find the degrees of the irreducible factors of $\chi_{\mathbb{Z}}$. This allows us to know how $J_0(p)$ decomposes as a product of Abelian varieties. As we mentioned before, we verified that, for $10^4 < p < 10^6$, there was only one factor per Atkin-Lehner eigenspace which had dimension 7 or more. This part of the algorithm uses a modified version of a well known dynamic programming algorithm for solving the subset-sum problem [14]. It is also a leading term in the asymptotic time complexity, and in practice is the most computationally expensive part of our algorithm, but it is also optional: the q -expansions of all newforms associated to factors of $J_0(p)$ of dimension at most 6 can be computed independently of this verification that only two other irreducible factors per level exist.

5 Computing the action of T_ℓ on $S_2(p)$

For a given prime ℓ , we generate two directed weighted multigraphs, G_ℓ^+ and G_ℓ^- , whose adjacency matrices are representations of the action of T_ℓ on the $+$ and $-$ Atkin-Lehner eigenspaces of $S_2(p)$ respectively.

The supersingular j -invariants over $\overline{\mathbb{F}}_p$ are all defined over \mathbb{F}_{p^2} , and it’s convenient to pick a generator of \mathbb{F}_{p^2} which has trace 0 because it simplifies the part of the algorithm described in section 8. Let σ be the nontrivial element of $\text{Gal}(\mathbb{F}_{p^2}/\mathbb{F}_p)$. Pick some arbitrary ordering $<$ of the supersingular j -invariants over $\overline{\mathbb{F}}_p$. The vertices of G_ℓ^+ are the pairs (j, j^σ) with $j < j^\sigma$, and the vertices of G_ℓ^- are the pairs (j, j^σ) with $j \leq j^\sigma$.

The graph G_ℓ^+ has a weight 1 edge from (j_1, j_1^σ) to (j_2, j_2^σ) for each ℓ -isogeny from j_1 to j_2 and for each ℓ -isogeny from j_1^σ to j_2^σ , and has a weight -1 edge from (j_1, j_1^σ) to (j_2, j_2^σ) for each ℓ -isogeny from j_1 to j_2^σ and for each ℓ -isogeny from j_1^σ to j_2 .

The graph G_ℓ^- has a weight 1 edge from (j_1, j_1^σ) to (j_2, j_2^σ) for each ℓ -isogeny from j_1 to j_2 , j_1 to j_2^σ , j_1^σ to j_2 , and j_1^σ to j_2^σ .

Constructing these graphs is done in two steps: finding a starting vertex (section 5.1), and exploring the graph (section 5.2).

5.1 Finding a starting vertex

Let $j \in \mathbb{Z}$ be the j -invariant of an elliptic curve over \mathbb{Q} with complex multiplication, and let D be the discriminant of the associated imaginary quadratic field. The reduction $j \bmod p$ is a supersingular j -invariant if D is not a square mod p [9]. Thus, over 99% of the time, the reduction of one of the 13 supersingular j -invariants over \mathbb{Q} will be a supersingular j -invariant over $\overline{\mathbb{F}}_p$. We use this as our starting vertex in these cases.

If p is such that every D is a square mod p , then we use code from Arpin, Camacho-Navarro, Lauter, Lim, Nelson, Scholl, and Sotáková to get the starting vertex [1].

5.2 Exploring the graph

For each prime ℓ there is a *modular polynomial* $\phi_\ell(x, y) \in \mathbb{Z}[x, y]$ with the property that $\phi_\ell(j, y)$ has a zero at $y = j'$ of order equal to the number of ℓ -isogenies from j to j' [35]. To generate the graphs G_ℓ^+ and G_ℓ^- , we do a breadth first search, finding the roots of $\phi_\ell(j, y)$ at the vertex (j, j^σ) at each step. Because ϕ_ℓ has coefficients in \mathbb{Z} , the roots of $\phi_\ell(j^\sigma, y)$ are the Galois conjugates of the roots of $\phi_\ell(j, y)$. We also make use of the fact that,

beyond the first vertex, we know at least one of the roots of $\phi_\ell(j, y)$, reducing the degree of the polynomial we have to solve by 1. Thus generating G_ℓ^\pm requires finding the roots in \mathbb{F}_{p^2} of one polynomial of degree $\ell + 1$ and $\mathcal{O}(p)$ polynomials of degree ℓ .

In our implementation we take $\ell = 2$, but later in the algorithm it is sometimes necessary to compute the action of T_ℓ for $\ell \geq 3$. We do this without making use of any information gained while computing the action of T_2 , because in practice this ran the fastest. We tested a different algorithm which made use of the fact that if there's an ℓ_1 -isogeny from j to j_1 , and an ℓ_2 -isogeny from j to j_2 , then there necessarily exists a j' which is both ℓ_1 -isogenous to j_2 and ℓ_2 -isogenous to j_1 , but our implementation took longer to evaluate the corresponding modular polynomials at all the neighbours of j_1 and j_2 than to find the roots of the modular polynomials directly.

6 Computing the characteristic polynomial of $T_2 \bmod \nu$

The next step of our algorithm for computing q -expansions is computing the characteristic polynomials of the adjacency matrices of the graphs G_ℓ^\pm (which we'll denote T_ℓ in a slight abuse of notation). Computing the characteristic polynomial over \mathbb{Z} directly appears to be infeasible. Instead we compute modulo some small arbitrary auxiliary prime ν , and use an algorithm from Wiedemann [38] with some small modifications and additions. The changes to Wiedemann's algorithm that we make serve two purposes: some of them result in speedups for our problem specifically, and others are needed to guarantee that we find all q -expansions. In this section we'll outline these changes.

6.1 Shifting eigenvalues

To compute the characteristic polynomial of T_2 , we first compute the characteristic polynomial of $T_2 + kI$ for some integer k , and then make a change of variables. We do this for two reasons: to try and avoid singular matrices, and to give another parameter to modify in the random algorithm.

Wiedemann's algorithm as described in [38] is significantly more involved for singular matrices. In our case, if there's a newform in $S_2(p)$ with $a_2 = 0$, then the matrix representation of T_2 will be singular over \mathbb{Z} . To avoid these we instead work with the matrix $T_2 + kI$ for $k > 3$. The matrix T_2 over \mathbb{Z} has one eigenvalue of 3, and the others are guaranteed to be real and at most $2\sqrt{2}$ in absolute value. Thus $T_2 + kI$ is guaranteed to be nonsingular over \mathbb{Z} whenever $k > 3$. The reduction of $T_2 + kI$ modulo ν might end up being singular anyway. We discuss this in section 6.2.

6.2 Varying parameters

Wiedemann's algorithm is a random algorithm, and, for any given random input, fails a non-negligible amount of the time. Our purposes give us the freedom to vary two parameters which would be fixed in other situations: the shift k (see section 6.1) and the modulus ν . We got significant speedups by tweaking our algorithm to vary these parameters while also varying the random inputs.

Overall, our implementation of Wiedemann's algorithm takes 4 inputs:

- a random starting vector u ,
- a random coordinate i ,
- a shift k , and
- a modulus ν .

To choose a random starting vector, we take the zero vector and set 50 random entries to 1. We found that setting only one entry to 1 causes the algorithm to fail more often.

The random coordinate i which we use for the Berlekamp-Massey part of the algorithm is chosen uniformly at random.

Even though the matrix $T_2 + kI$ is guaranteed to be nonsingular over \mathbb{Z} whenever $k > 3$, the reduction mod ν might happen to be singular, essentially “by chance”. In this case, the Berlekamp-Massey part of the algorithm will be given a power series which is not invertible. When this happens, we increment the shift k in addition to choosing new values of v and i .

The cases which are the most computationally intensive, by a wide margin, are cases where the characteristic polynomial of T_2 has repeated factors. The runtime of our implementation of Wiedemann’s algorithm is proportional to the multiplicity of the most frequently occurring factor (see section 6.3), and, if the repeated factors are ones which might be reductions of factors of the characteristic polynomial over \mathbb{Z} , then we’ll have to run the very expensive part of the algorithm which attempts to lift eigenspaces of dimension larger than 1 (see section 7). Genuine repeated factors of the characteristic polynomial \mathbb{Z} are quite rare. Thus, we’ve found that the fastest approach is to run Wiedemann’s algorithm for at most 2 random u ’s, and, if the algorithm fails for both choices, to change our small prime ν in case the failure was caused by a spurious repeated factor. After changing ν , we increase the maximum number of choices of u , so that the algorithm does eventually find genuine high dimensional eigenspaces. Every time we change ν , we have to recompute all of our iterates, so this results in a significant slowdown for levels which do have high dimensional eigenspaces, but there are very few of these.

6.3 Getting the characteristic polynomial from the minimal polynomial

For our purposes, it’s important to find not just the minimal polynomial of T_2 (which is what Wiedemann’s algorithm as described in [38] yields), but the full characteristic polynomial. We need to do this for two reasons.

First, we know the degree of the characteristic polynomial (because we know the dimension of $S_2(p)$), but not the degree of the minimal polynomial. Wiedemann’s algorithm has some chance to fail to find the minimal polynomial, and when it does so, it outputs a polynomial which properly divides the minimal polynomial, and does not detect that it has failed to find the full minimal polynomial. We want to provably find all Galois orbits of $S_2(p)$, so it’s necessary for us to be able to detect when the algorithm fails. By instead computing the characteristic polynomial, we can guarantee that the algorithm has succeeded by checking the degree.

Second, a repeated factor of the characteristic polynomial of T_2 over \mathbb{Z} indicates the presence of multiple Galois orbits of newforms with the same a_2 . If we were to only compute the minimal polynomial of T_2 , then there is some chance we would compute the q -expansion of only some of these newforms. In practice, this would be unlikely but not impossible; we elaborate on this in more detail in section 7.

Given the minimal polynomial of T_2 , we use two techniques to produce the characteristic polynomial: comparing with known top coefficients (6.3.1), and checking eigenspace dimensions (6.3.2).

6.3.1 Comparing with known top coefficients

The coefficient of the second-highest degree term of the characteristic polynomial of a matrix M is equal to $-\text{tr}(M)$. Similarly, the coefficient of the third-highest degree term of the characteristic polynomial is given by the expression [33]

$$\sum_{1 \leq i < j \leq \dim(M)} M_{ii}M_{jj} - M_{ij}M_{ji}. \quad (1)$$

Computing the trace of any matrix M takes a time of only $\mathcal{O}(\dim(M)^{1+\epsilon})$, and evaluating the expression (1) can be done in time $\mathcal{O}(\dim(M)^{2+\epsilon})$. Furthermore, because the matrix T_2 is sparse and nearly symmetric, these quantities can be computed more quickly. Expressions similar to the trace and (1) exist for other coefficients as well, and, while these would take too long to compute for general matrices, may be efficiently computable for sparse symmetric matrices like T_2 . We haven’t investigated this, but it would likely lead to a small improvement in the running time of our algorithm.

If the characteristic polynomial is of degree at most 2 more than than the degree of the minimal polynomial, and if the coefficients of the 3 leading terms of both polynomials are known, then the ratio of these two polynomials

can be found from an elementary calculation. We do this to find the characteristic polynomial whenever possible. Moreover, note that this trick still works even if the missing factors don't divide the minimal polynomial, so it also helps in some cases where Wiedemann's algorithm fails to find the minimal polynomial.

6.3.2 Checking eigenspace dimensions

This section of the overall algorithm is used only after Wiedemann's algorithm has been tried for 3 or more initial random starting vectors u , which we'll label u_1, \dots, u_n . We stored the values of $T_2^k u_i$ for $0 \leq k \leq \dim(T_2)$ and $1 \leq i \leq n$ (or, more precisely, the first 1000 entries of these vectors; see section 6.4). Let $\mu(t)$ denote the highest degree polynomial that's been returned by Wiedemann's algorithm thus far (so μ is a candidate for the minimal polynomial of T_2). Finding all the roots of μ takes time $\mathcal{O}(p^{1+\epsilon})$ using the Fast Fourier Transform (henceforth "FFT") [13]. Then, given a root λ of μ , the vector

$$v_{i,\lambda} := \frac{\mu(T_2)}{T_2 - \lambda} u_i$$

is an eigenvector of T_2 with eigenvalue λ . Given the iterates of u_i , computing the first 1000 entries of $v_{i,\lambda}$ takes time $\mathcal{O}(p^{1+\epsilon})$. The dimension of the span of the vectors $v_{i,\lambda}$ is at most the dimension of the λ -eigenspace (and it's very likely that these dimensions will be equal, provided the number of iterates n is at least the dimension of the eigenspace). This allows us to give lower bounds for the multiplicity with which linear factors of the minimal polynomial occur in the characteristic polynomial. In practice, the additional linear factors found with this trick, in combination with the trick from section 6.3.1, are usually enough to determine the characteristic polynomial of T_2 .

6.4 Storing iterates

To implement our version of Wiedemann's algorithm, we needed information about the iterates $T_2^k u$ of the random starting vector u for three different purposes.

1. We computed these iterates as $T_2^k u = T_2(T_2^{k-1} u)$, so every coordinate of the previous iterate is stored temporarily.
2. To run the Berlekamp-Massey part of Wiedemann's algorithm, we used the i^{th} coordinate of $T_2^k u$ for all $k < 2\dim(T_2) + 10$, for an arbitrary i (which we chose uniformly at random).
3. To implement the trick from section 6.3.2, we used the first 1000 coordinates of $T_2^k u$ for $k < \dim(T_2)$. Here the choice of 1000 is largely arbitrary; we just need to take enough coordinates to avoid spurious linear dependencies.

The runtime of the entire algorithm depended heavily on computing these iterates quickly. Asymptotically, computing characteristic polynomials is the dominant term of the overall time complexity, and computing the iterates $T_2^k u$ was the step in our implementation of Wiedemann's algorithm that took the longest. Thus, we wanted to compute and store as few of the iterates as possible. Moreover, we structured our code so that python and Sage governed the algorithm at the top level, while computationally intensive sections were run in faster languages; these iterates were computed using cython directly. Because our code interfaced different programming languages, we needed to use the cython data to generate corresponding python data. If we store too much data, then this translation can take a significant amount of time. In the extreme case of storing the entire iterates $T_2^k u$, the translation took much longer than the actual computation of the data in cython.

The other important reason to be judicious in how much of the iterates are stored is concern for the space complexity of the algorithm. We wanted our algorithm never use more than $\mathcal{O}(p^{1+\epsilon})$ memory, and to accomplish this we were only able to store a constant number of coordinates per iterate.

7 Finding \mathbb{Z} -eigenbases

Let $\chi_{\mathbb{Z}}$ be the characteristic polynomial of T_2 over \mathbb{Z} , and χ_{ν} its reduction modulo ν . At this point, the algorithm has computed χ_{ν} for one choice of ν . The objective of this section is to determine all irreducible factors of degree

6 or less of $\chi_{\mathbb{Z}}$, and, for each factor, an eigenbasis defined over \mathbb{Z} and a simultaneous eigenvector (defined over a number field) for all Hecke operators. Here we use the term “eigenspace” and related terms to mean $\ker(\rho(T_2))$ for some irreducible factor ρ of $\chi_{\mathbb{Z}}$. We’ll need \mathbb{Z} -bases of these spaces to use a formula from Mestre’s *Méthode des Graphes*, which we’ll discuss in section 8.

If an irreducible factor ρ of $\chi_{\mathbb{Z}}$ divides $\chi_{\mathbb{Z}}$ exactly once, then ρ corresponds to a single Galois orbit of a newform of $S_2(p)^{\pm}$, where \pm denotes the Atkin-Lehner eigenvalue. If an irreducible factor ρ divides $\chi_{\mathbb{Z}}$ twice or more then the situation is more complicated. Multiple Galois orbits of newforms with the same a_2 lead to repeated factors of $\chi_{\mathbb{Z}}$, as do newforms whose Hecke field is of strictly larger degree than the degree of a_2 . In these cases, Galois orbits of newforms will correspond to the minimal nontrivial subspaces of $\ker(\rho(T_2))$ which are invariant under the action of every T_ℓ simultaneously.

The key idea behind the lifting algorithm described in this section is taking advantage of the heuristic that the coordinates in the vectors making up the \mathbb{Z} -eigenbasis are very likely to be small in absolute value. We pick some number of “candidate lifts” of an \mathbb{F}_ν -eigenbasis in which the most common entries are small integers, and then check directly whether or not our candidate lifts are eigenvectors of T_2 over \mathbb{Z} . In a large majority of cases, if ρ does indeed divide $\chi_{\mathbb{Z}}$, then checking only a handful of candidate lifts is enough to find a \mathbb{Z} -eigenbasis.

7.1 Finding eigenvalues

Let $\{\rho_i\}$ denote the set of polynomials which occur as the minimal polynomial of a_2 of a simple Abelian variety over \mathbb{F}_2 of dimension 6 or less. There are 96795 such polynomials, and they can be found in the LMFDB. Every irreducible factor of $\chi_{\mathbb{Z}}$ of degree 6 or less is necessarily one of these.

To find the irreducible factors of $\chi_{\mathbb{Z}}$, we first determine which ρ_i ’s divide χ_ν . We do this by iterating over the set $\{\rho_i\}$ and checking for divisibility one by one, but one could use FFT [13] to do this more quickly if necessary.

It is possible for ρ_i to divide χ_ν but not $\chi_{\mathbb{Z}}$. When we find a ρ_i which divides χ_ν , we attempt to produce an eigenbasis defined over \mathbb{Z} for this factor using a method which we describe in the rest of this section. If we succeed, this proves that ρ_i does divide $\chi_{\mathbb{Z}}$. Our method for producing a \mathbb{Z} -eigenspace is not guaranteed to work, though it almost always does in practice. In the cases where it doesn’t, we compute $\chi_{\nu'}$ for some other small prime $\nu' \neq \nu$, and see if ρ_i divides $\chi_{\nu'}$. If $\rho_i \nmid \chi_{\nu'}$, then we’ve proven that $\rho_i \nmid \chi_{\mathbb{Z}}$, and, conversely, if $\rho_i \nmid \chi_{\mathbb{Z}}$, then there is guaranteed to be some prime ν' for which $\rho_i \nmid \chi_{\nu'}$ (and usually only one additional prime ν' needs to be checked).

It would be difficult to determine the Galois orbits of size 7 or greater using this approach, because the number of simple Abelian varieties over \mathbb{F}_2 of given dimension grows very quickly. In section 9, we describe the algorithm we used to determine the size of every Galois orbit, including the ones of size 7 or more. For every prime level between 10,000 and 1,000,000 and each Atkin-Lehner eigenspace, the characteristic polynomial $\chi_{\mathbb{Z}}$ had only one irreducible factor of degree 7 or more. For prime levels between 1,000,000 and 2,000,000 we computed the q -expansions of the newforms of degree 6 or less but did not investigate the decomposition of the rest of the space.

7.2 Lifting 1-dimensional eigenspaces

When an eigenspace is 1-dimensional (which is the case exactly when $\rho(t) = t - \lambda$ is linear and divides χ_ν exactly once), our algorithm is fairly straightforward. First, we find an eigenvector v of T_2 by computing

$$v := \frac{\mu(T_2)}{T_2 - \lambda} u$$

for some random starting vector u , where μ is the minimal polynomial of T_2 on the given Atkin-Lehner eigenspace. Computing v takes time $\mathcal{O}(p^{2+\varepsilon})$, so, for the levels for which require that we attempt to lift eigenspaces, this is a leading term of the over asymptotic time complexity of our algorithm, and these levels presumably make up a small but strictly positive proportion of all levels. The computation of v in this way involves successively

computing the iterates $T_2^k u = T_2(T_2^{k-1}u)$. This computation was done previously in the section of our algorithm which used Wiedemann’s algorithm, but it was impossible to store these iterates while meeting our goal of having a space complexity of $\mathcal{O}(p^{1+\epsilon})$. Implementations of our algorithm which can afford to store the iterates from the Wiedemann section can use those iterates here, saving some time.

With our eigenvector v defined over \mathbb{F}_ν , we generate our list of “candidate lifts” of v to an eigenvector \hat{v} defined over \mathbb{Z} by guessing that the most likely scenario is that the most common nonzero entry of \hat{v} is 1, the second most likely scenario is that the most common entry is 2, and so on. Note that it suffices to check for strictly positive most common entries, since both \hat{v} and $-\hat{v}$ span the \mathbb{Z} -eigenbasis.

Let $\alpha \in \mathbb{F}_\nu$ denote the most common nonzero entry of v . Then, our candidate lifts are $\frac{1}{\alpha}v, \frac{2}{\alpha}v, \dots$, where each \mathbb{F}_ν coordinate of these vectors is lifted to the integer of smallest absolute value in that residue class. As these candidate lifts are generated, we multiply them by T_2 to check directly whether or not they’re eigenvectors, and return the first candidate lift that is an eigenvector. This lift is guaranteed to have entries with a gcd of 1, so it spans the \mathbb{Z} -eigenspace (a condition which is necessary for the part of the algorithm described in section 8 to succeed). It is also guaranteed to be a simultaneous eigenvector of all the Hecke operators.

Our implementation of this algorithm generates 50 candidate lifts, and then, if no lift to \mathbb{Z} has been found, gives up and declares it has failed to lift the \mathbb{F}_ν eigenvector. This prompts the algorithm to try a different small prime $\nu' \neq \nu$, as discussed in section 7.1. We encountered no cases in which a lift existed but was not found.

7.3 Lifting higher dimensional eigenspaces

7.3.1 Finding an \mathbb{F}_ν -eigenbasis

We begin like we did in the 1-dimensional case, by computing

$$v := \frac{\mu(T_2)}{\rho(T_2)}u$$

for some random starting vector u , where μ is the minimal polynomial of T_2 on the given Atkin-Lehner eigenspace. The vector v is in the \mathbb{F}_ν -kernel of $\rho(T_2)$, and is what we’re calling an eigenvector (because it’s an eigenvector up to Galois conjugacy for the field generated by ρ). The same discussion as that in section 7.2 applies here: computing this eigenvector is a small but nonzero part of the leading term in asymptotic time complexity, and the iterates from the Wiedemann component of our algorithm can’t be used to save time in the computation in our implementation because of space limitations.

Let r denote the multiplicity with which ρ divides χ_ν . If $r = 1$, then we compute the dimension of the span of the vectors

$$v, T_2v, T_2^2v, \dots, T_2^{\deg(\rho)-1}v.$$

If the dimension of this span is equal to $\deg(\rho)$, then these vectors form an \mathbb{F}_ν -eigenbasis of the eigenspace. If this dimension is strictly less than $\deg(\rho)$, or if $r > 1$, then, starting from $\ell = 3$, we compute matrix which corresponds to the action of T_ℓ on the given Atkin-Lehner eigenspace by exploring the supersingular ℓ -isogeny graph in the way described in section 5, and compute the dimension of the span of the vectors

$$v, T_\ell v, T_\ell^2v, \dots, T_\ell^{r\deg(\rho)-1}v,$$

stopping our iteration over ℓ when this dimension is equal to $r \deg(\rho)$. We also periodically recompute v with a new random choice of u , because it’s possible (but unlikely) that v happens to be a simultaneous eigenvector of all Hecke operators, which would cause our iteration over ℓ to never terminate (in principle one could add a clause in our algorithm to take advantage of this whenever it happened, but in practice this never happened). In doing this, we find an ℓ such that all newforms of $S_2(p)^\pm$ have distinct a_ℓ ’s. We need these a_ℓ ’s to be distinct later, essentially so that we can separate the \mathbb{Z} -eigenbases. Let M denote the $\dim(S_2(p)^\pm) \times r \deg(\rho)$ matrix whose columns are $v, T_\ell v, T_\ell^2v, \dots, T_\ell^{r\deg(\rho)-1}v$.

7.3.2 Finding a \mathbb{Z} -eigenbasis for the full ρ -eigenspace

We now make use of our key heuristic that our \mathbb{Z} -eigenbases are very likely to be made up of vectors which have small entries. To do this, we start by making two lists (which we generate as needed, as opposed to storing in memory):

- \mathcal{M} : a list of the most common rows of M with the condition that any rows we add to this list are not in the span of the rows we've added previously, and
- \mathcal{C} : a list of “candidate columns”, which are $r\deg(\rho)$ -tuples with small integer entries, ordered in such a way such that, generally speaking, tuples with smaller entries are listed first.

Our implementation requires us to pick an ordering of the set $\mathcal{M}^{r\deg(\rho)} \times \mathcal{C}$ so that we may iterate over it. We do this more or less arbitrarily. Experimentally, we found that assigning a “size” to elements of \mathcal{C} which was proportional to the sum of squares of the entries, assigning a “size” to elements of $\mathcal{M}^{r\deg(\rho)}$ which was proportional to the sum of the squares of the inverse appearance frequencies, and then ordering $\mathcal{M}^{r\deg(\rho)} \times \mathcal{C}$ by the product of these sizes lead to having to check fewer candidate lifts than lexicographic orderings.

We iterate over $\mathcal{M}^{r\deg(\rho)} \times \mathcal{C}$. For each element (m, c) , let L be the linear combination of columns of M which yields c when restricted to the rows in m . We then produce a “candidate lift” \hat{v} , which is the vector with integer entries produced by taking the combination L of the columns of M , and then lifting each coordinate of the resulting \mathbb{F}_ν vector to the integer in the appropriate residue class with smallest absolute value. We then check directly whether or not \hat{v} is in $\ker(\rho(T_2))$. We continue this process until we've found $r\deg(\rho)$ linearly independent candidate lifts. This set of candidate lifts forms a \mathbb{Z} -basis of $\ker(\rho(T_2))$, but our algorithm requires a \mathbb{Z} -basis for each Galois orbit, so it remains to decompose our \mathbb{Z} -basis in this way. We do this next.

7.3.3 Finding a \mathbb{Z} -eigenbasis for each Galois orbit

The method we describe in this section for finding a \mathbb{Z} -eigenbasis for each Galois orbit is built around the observation that the action of T_ℓ fixes $\ker(\rho(T_2))$. This means that each column of $T_\ell M$ can be expressed as a linear combination of the columns of M . Let S denote the $r\deg(\rho) \times r\deg(\rho)$ integer matrix whose entries are the coefficients of these linear combinations (which are all small, since T_ℓ is given by a matrix with at most $2(\ell + 1)$ nonzero entries per row, each of which is ± 1).

Let χ_S denote the characteristic polynomial of S , and let h_1, h_2, \dots denote the minimal polynomials of the a_ℓ 's of the newforms we're considering. As discussed in section 7.3.1, we chose ℓ in a way that guarantees that each of these minimal polynomials is distinct. By construction, we have

$$\chi_S = \prod_i h_i.$$

We can then obtain, with some elementary linear algebra, a \mathbb{Z} -basis for $\ker(h_i(S))$ and a set of eigenvectors of S (which will be defined over number fields). Taking the linear combinations of the columns of M whose coefficients are given by the elements of these \mathbb{Z} -bases and eigenvectors then yields the \mathbb{Z} -eigenbases and simultaneous eigenvectors we need. Our implementation requires only one simultaneous eigenvector per Galois orbit, which we pick in an arbitrary way.

8 Computing q -expansions

Let v be a simultaneous eigenvector of all the Hecke operators whose coordinates are indexed by the supersingular j -invariants over $\overline{\mathbb{F}}_p$. In [27], Mestre gives the q -expansion of the associated newform f modulo any prime \mathfrak{p} above p in $K := \mathbb{Q}(a_2, a_3, \dots)$ via an equality of power series involving the q -expansion of the modular j -function $j(q)$:

$$\left(\sum_j v_j j \right) f(q) \frac{dq}{q} = \sum_j v_j \frac{dj(q)}{j(q) - j} \pmod{\mathfrak{p}}. \quad (2)$$

The Weil bound then allows one to determine a_ℓ for all $\ell \ll p^2$. Our algorithm uses this formula to compute a_n for n up to the Sturm bound.

Section 8.1 gives a version of equation (2) that does computations over \mathbb{F}_p using the \mathbb{Z} -eigenbases computed earlier. Sections 8.2 and 8.3 detail how we evaluate the right hand side of (2) efficiently.

8.1 \mathbb{Z} -eigenbasis version of Mestre's identity

We use the simultaneous eigenvector v to compute the values $a_{\ell_1}, a_{\ell_2}, \dots, a_{\ell_{\deg K}}$ (which might require computing new Hecke matrices using supersingular isogeny graphs). With each vector u_k in the \mathbb{Z} -eigenbasis, we evaluate the right hand side of (2), and let $\psi_k(q)$ denote the resulting power series. As we'll discuss in section 8.3, the power series $\psi_k(q)$ is guaranteed to be defined over \mathbb{F}_p . We write $\psi_k[\ell]$ to refer to the coefficient of q^ℓ in $\psi_k(q)$. Pick a basis $\{r_i\}$ of the ring of integers \mathcal{O}_K . Write each a_ℓ in terms of this basis:

$$a_\ell =: \sum_i \alpha_i[\ell] r_i.$$

Then, from (2), it follows that there exist coefficients $\beta_{i,k} \in \mathbb{F}_p$ such that, for every ℓ , we have the a linear combination

$$a_\ell = \sum_{i,k} \beta_{i,k} \psi_k[\ell] r_i \pmod{\mathfrak{p}},$$

so

$$\begin{pmatrix} \alpha_1[\ell_1] & \alpha_1[\ell_2] & \dots \\ \alpha_2[\ell_1] & \alpha_2[\ell_2] & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} = \begin{pmatrix} \beta_{1,1} & \beta_{1,2} & \dots \\ \beta_{2,1} & \beta_{2,2} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} \psi_1[\ell_1] & \psi_1[\ell_2] & \dots \\ \psi_2[\ell_1] & \psi_2[\ell_2] & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} \pmod{p}. \quad (3)$$

We solve the matrix equation (3) for the coefficients $\beta_{i,k}$ (as long as all the matrices involved are invertible; compute a_ℓ for more ℓ 's if there happens to be a linear dependence). We can then find the q -expansion of our newform as

$$f(q) \frac{dq}{q} = \sum_i \left(\sum_k \beta_{i,k} \psi_k[\ell] \right) r_i q^\ell.$$

We lift the coefficient $\sum_k \beta_{i,k} \psi_k[\ell]$ to \mathbb{Z} by ensuring that the Hecke eigenvalues satisfy the Weil bound. Usually the coefficient lifts to the integer in the residue class of smallest absolute value if one chooses $\{r_i\}$ to be a reduced basis of \mathcal{O}_K .

8.2 Computing $j(q) \pmod{p}$

To evaluate the right hand side of (2) up to the Sturm bound, we first need to compute $\mathcal{O}(p)$ coefficients of $j(q)$ modulo p . We compute these coefficients using the identity

$$j(q) = \frac{E_{12}(q)}{(\eta(q)^3)^8} - \frac{82104}{691} + 744, \quad (4)$$

where $E_{12}(q)$ is the weight 12 classical Eisenstein series (normalized to have constant term 1) and $\eta(q)$ is the Dedekind η -function. We chose equation (4) because the function $\eta(q)^3$ has a "sparse" q -expansion:

$$\eta(q)^3 = \sum_{k=0}^{\infty} (-1)^k (2k+1) q^{\frac{k(k+1)}{2}}.$$

Computing the first p terms of the quotient of two power series takes time $\mathcal{O}(p^{1+\varepsilon})$ using FFT, and empirically dividing 8 times by $\eta(q)^3$ was faster than alternatives, such as dividing by $\Delta(q)$ or 24 times by $\eta(q)$.

In section 8.3 we'll also need the power series $j'(q)$. Differentiating $j(q)$ takes time $\mathcal{O}(p^{1+\varepsilon})$.

8.3 Fast power series algorithms

8.3.1 Atkin-Lehner eigenspace to Mestre's formula

The eigenvectors $\{u_k\}$ we computed are indexed by pairs (j, j^σ) with $j < j^\sigma$ (in the $+$ space) or $j \leq j^\sigma$ (in the $-$ space), where $<$ is the arbitrary but fixed ordering from section 5. To use Mestre's formula (2), we need to construct eigenvectors $\{v_k\}$ that are indexed by supersingular j -invariants. Recall that in section 5 we chose a generator ξ of \mathbb{F}_{p^2} which satisfied $\xi^\sigma = -\xi$. This allows us to construct each v_k easily: if $j \notin \mathbb{F}_p$ we set $v_k(j) = u_k((j, j^\sigma))$ and $v_k(j^\sigma) = \mp u_k((j, j^\sigma))$, and if $j \in \mathbb{F}_p$ we set $v_k(j) = (1 \mp 1)u_k((j, j^\sigma))$. Then v_k is defined over \mathbb{Z} (because u_k is), and the sum in (2) is either Galois-invariant or Galois anti-invariant. In the latter case we then divide by ξ .

8.3.2 Evaluating Mestre's formula

Evaluating the sum of rational functions of q -expansions in (2) by evaluating each term separately and then then adding would take time $\gg p^2$. However, it is possible to evaluate this sum in time $\mathcal{O}(p^{\frac{3}{2}+\varepsilon})$ by doing a "binary tree decomposition" and using a power series composition algorithm from Brent and Kung [7].

Given constants $\gamma_1, \dots, \gamma_M$ and j_1, \dots, j_M , let P and Q be polynomials such that

$$\frac{P(x)}{Q(x)} := \sum_{i=1}^M \frac{\gamma_i}{x - j_i}.$$

Define

$$S(x, y) := \sum_{i=x}^y \frac{\gamma_i}{x - j_i}.$$

We compute P and Q recursively as

$$\begin{aligned} \sum_{i=1}^M \frac{\gamma_i}{x - j_i} &= S(1, M) \\ &= S(1, \frac{M}{2}) + S(\frac{M}{2} + 1, M) \\ &= [S(1, \frac{M}{4}) + S(\frac{M}{4} + 1, \frac{M}{2})] + [S(\frac{M}{2} + 1, \frac{3M}{4}) + S(\frac{3M}{4} + 1, M)] \\ &= \dots \\ &= \sum_{i=1}^M S(i, i). \end{aligned}$$

We start from the bottom expression and work upwards. At the k^{th} step we compute $\mathcal{O}(\frac{M}{2^k})$ sums of two terms. Each term at the k^{th} step is a rational function whose numerator and denominator have degree $\mathcal{O}(2^k)$. Multiplying two polynomials of degree d takes time $\mathcal{O}(d^{1+\varepsilon})$ using FFT. Thus, each step takes time $\mathcal{O}((2^k)^{1+\varepsilon} \frac{M}{2^k}) = \mathcal{O}(M^{1+\varepsilon})$. There are $\mathcal{O}(\log M)$ steps in this procedure. Thus, computing P and Q takes time $\mathcal{O}(M^{1+\varepsilon} \log M) = \mathcal{O}(M^{1+\varepsilon})$.

In the context of our problem, this means that we can, in time $\mathcal{O}(p^{1+\varepsilon})$, compute polynomials P, Q of degree $\mathcal{O}(p)$ such that

$$\sum_j v_j \frac{dj(q)}{j(q) - j} = \frac{P(j(q))}{Q(j(q))} j'(q) dq.$$

Writing $\frac{1}{Q(j(q))}$ as a power series in $j(q)$ with $\mathcal{O}(p)$ terms of precision takes time $\mathcal{O}(p^{1+\varepsilon})$, and computing the product $R(j(q)) := P(j(q)) \frac{1}{Q(j(q))}$ as a power series in $j(q)$ with $\mathcal{O}(p)$ terms of precision also takes time $\mathcal{O}(p^{1+\varepsilon})$.

Using [7], we compute the composition $R(j(q))$ to $\mathcal{O}(p)$ terms of precision in time $\mathcal{O}(p^{\frac{3}{2}+\varepsilon})$. Finally, multiplying by $j'(q) dq$ takes time $\mathcal{O}(p^{1+\varepsilon})$.

9 Checking for high degree factors of the characteristic polynomial

As mentioned in section 7.1, our algorithm doesn't find newforms whose a_2 has minimal polynomial of degree 7 or more. It's of value to at least determine the dimensions of all of the newforms of $S_2(p)$, even if we can't compute all of their q -expansions. This part of the algorithm is by far the most expensive time and space wise, and is the most technically involved. This part of the algorithm isn't required for computing q -expansions and can be omitted.

Our algorithm is designed to try and efficiently determine that there is exactly one irreducible factor of $\chi_{\mathbb{Z}}$ of degree 7 or more, since this was the case for every prime level between 10,000 and 1,000,000. The approach we take is based on the straightforward observation that if, for some small prime ν , the characteristic polynomial χ_{ν} has no factor of degree d , then $\chi_{\mathbb{Z}}$ cannot have a factor of degree d either. Thus, for each possible degree d between 7 and $\frac{1}{2}\deg\chi_{\mathbb{Z}}$ we aim to find some small prime ν such that χ_{ν} has no factor of degree d .

9.1 Factoring the characteristic polynomial modulo many small primes

The first part of the algorithm will require finding and factoring characteristic polynomials $\chi_{\nu_1}, \chi_{\nu_2}, \dots$ of T_2 on a given Atkin-Lehner eigenspace modulo many small primes, which we'll do as needed in what follows. Our algorithm is slightly more efficient if we replace χ_{ν_i} with χ_{ν_i} divided by all of the irreducible factors of degree 6 or less which we know divide $\chi_{\mathbb{Z}}$. Factoring polynomials over finite fields has high time complexity in theory, but experimentally we found that factoring χ_{ν_i} appeared to take time $\mathcal{O}(p^{2+\varepsilon})$ in our specific case. This is still the most significant chunk of the runtime, but it's not problematically expensive.

The reason factoring is feasible can be explained heuristically. There exists an algorithm based on FFT which factors polynomials over finite fields in three steps [12]:

1. Eliminating square factors of χ_{ν_i}
2. Factoring χ_{ν_i} into products of irreducible polynomials of equal degree
3. Factoring each of these products of irreducibles of equal degree

The first step is not problematic because we obtain the squarefree part of χ_{ν_i} (which is the minimal polynomial) directly through Wiedemann's algorithm. Even ignoring this, eliminating the square factors of χ_{ν_i} can be done efficiently by computing $\gcd(\chi_{\nu_i}, \chi'_{\nu_i})$.

The second step can be done in time $\mathcal{O}(p^{2+\varepsilon})$ using FFT.

The third step is the one which is theoretically challenging, but, in practice, it's very rare that χ_{ν_i} has multiple factors of the same degree if that degree is large. Thus, we can use Rabin's irreducibility test [31], which runs in time $\mathcal{O}(p^{2+\varepsilon})$, to determine which of the products from the second step require factoring. The products that do require factoring end up being products of a handful of small degree factors which can be factored quickly. There is also never a requirement to use any specific ν , so at worst one could abandon computations with ν 's that were stuck on this step.

With knowledge of this factorization we can use the algorithm described in section 9.2. We've found that, in our implementation, our algorithm ran more quickly if we only continued using this factorization if it didn't have "too many" irreducible factors, with a threshold determined experimentally and in a fairly arbitrary way.

9.2 Sieving possible degrees

Let E_i denote the set subset of $[7, \frac{1}{2}\deg(\chi_{\mathbb{Z}})] \cap \mathbb{Z}$ which we have yet to show cannot be degrees of factors of $\chi_{\mathbb{Z}}$ after running this part of the algorithm for ν_i . For each i , we determine

1. which elements of E_i occur as the degrees of (not necessarily irreducible) factors of χ_{ν_i} , and,

2. for each d in E_i , if there are only “a few” factors of χ_{ν_i} with this degree, we record them to later use in the part of the algorithm described in section 9.3. We specify what we mean by “a few” later in this section.

Define

$$D_i := \{d \in E_i : \text{there exists a factor of } \chi_{\nu_i} \text{ with degree } d\}.$$

Given a factorization of χ_{ν_i} into irreducibles, computing the set D_i is a manifestation of the well-studied subset sum problem. Direct enumeration of all possible combinations of irreducible factors of χ_{ν_i} takes time and space exponential in the number of irreducible factors and was infeasible in practice. There’s a dynamic programming algorithm [14] which computes D_i in time $\mathcal{O}(p^2)$ and space $\mathcal{O}(p)$, but says nothing about what the corresponding factors are. We’ll modify this dynamic programming algorithm so that it records which products of irreducible factors have degree d whenever there are only “a few” such products. Doing this allows us to use the the method described in section 9.3, which leads to having to compute fewer characteristic polynomials χ_{ν_i} and their factorizations. If one chooses, one can only compute D_i and never anything about the factors themselves, and then use the dynamic programming algorithm directly. We’ve found that omitting the method from 9.3 usually takes longer but requires less space for our implementation.

To compute D_i , we create a variable Δ_i which will ultimately be the function with domain D_i for which $\Delta_i(d)$ is either a set of combinations of irreducible factors whose product has degree d , or “null” if this set would have more than “a few” elements in it. We initialize Δ_i as the function with $\Delta_i(0) = \emptyset$ and no other elements in its domain.

For $k = 1, 2, \dots$, let h_k denote the irreducible factors of χ_{ν_i} . We sort these factors in descending order according to their degrees (breaking ties arbitrarily), and compute the partial sums $P_k := \sum_{\kappa \geq k} h_\kappa$. For each k , let R_k denote the set of elements of E_i which are not yet in Δ_i . We iterate over either the domain of Δ_i , or the set

$$\{d \in \mathbb{Z} : \min(R_k) - P_k \leq d \leq \max(R_k)\},$$

whichever is smaller. For each d , if $d \in \text{dom}(\Delta_i)$, then we check whether or not $d + \deg(h_k)$ is in $\text{dom}(\Delta_i)$. If it isn’t, then we set

$$\Delta_i(d + \deg(h_k)) = \{H \cup \{h_k\} : H \in \Delta_i(d), h_k \notin H\}.$$

If it is, then, if $\Delta_i(d + \deg(h_k)) \neq \text{“null”}$, we compute the set

$$\mathcal{H}' := \Delta_i(d + \deg(h_k)) \cup \{H \cup \{h_k\} : H \in \Delta_i(d), h_k \notin H\}.$$

If \mathcal{H}' is larger than some fixed parameter η given as input to the algorithm, we then replace \mathcal{H}' with “null”. We took $\eta = 5$ in our implementation. Then we update Δ_i by setting $\Delta_i(d + \deg(h_k)) = \mathcal{H}'$.

Previously, when we said we wouldn’t find which products of irreducible factors multiplied together to give factors of the given degree if there were more than “a few” such products, this process of replacing values of Δ_i by “null” whenever the values would have cardinality more than η is the condition we were referring to: at no intermediate step were there more than η ways of obtaining that intermediate degree. It was necessary to have some bound of this sort, since otherwise we are directly enumerating all possible products of irreducible factors, which is infeasible. We chose $\eta = 5$ because empirically it was the best balance we found between yielding non-null values of Δ_i and not using excessive memory.

After the iteration over k finishes, we set $E_{i+1} = E_i \cap \text{dom}(\Delta_i)$. Moreover, for each $d \in E_{i+1}$, we store the values of $\Delta_i(d)$ whenever they’re not “null”, and then run the part of our algorithm described in section 9.3 before continuing our iteration over i .

9.3 Checking the Weil bound

If $\theta(t) = t^d + \theta_1 t^{d-1} + \dots + \theta_d$ is a polynomial which divides $\chi_{\mathcal{Z}}$, then the coefficients of θ satisfy the Weil bound:

$$|\theta_j| \leq \binom{d}{j} (2\sqrt{2})^j.$$

This bound gives us a way to rule out the existence of a lift of a polynomial $\bar{\theta}(t) \in (\mathbb{Z}/m)[t]$ to a factor of $\chi_{\mathbb{Z}}$: if $\bar{\theta}$ has a coefficient which has no integer lifts that satisfy the Weil bound, then we know that $\bar{\theta}$ cannot lift to a factor of $\chi_{\mathbb{Z}}$. For our range of levels and our choice of small primes ν , this Weil bound condition ends up being trivial for $j \geq 3$, but for $j = 1, 2$ we use this bound to avoid computing and factoring more χ_{ν_i} 's than we would have to otherwise.

For each degree $d \in E_i$, if there are primes $\nu_{i_1}, \nu_{i_2}, \dots$ for which

1. all ways to obtain a factor of degree d from a product of irreducible factors of $\chi_{\nu_{i_k}}$ are known, and
2. $\prod_k \nu_{i_k}$ is large enough for the Weil bound strategy outlined above to not be trivial,

then we use the Chinese Remainder Theorem on all possible combinations of products of irreducible factors to produce a list of “candidate factors” modulo $m = \prod_k \nu_{i_k}$. We then use the Weil bound as outlined above on these candidate factors one at a time to try and rule them out. Even if not all candidate factors can be ruled out, usually some can, and we remove these candidates so that this method is more likely to succeed as we continue to iterate over i .

10 Results

We used the algorithm described in this paper to compute the q -expansions of all weight 2 eigenforms of dimension $g \leq 6$ and prime level between 10^4 and $2 \cdot 10^6$. The q -expansions of forms with level less than 10^4 were computed by Best, Bober, Booker, Costa, Cremona, Derickx, Lowry-Duda, Lee, Roe, Sutherland, and Voight [3]. Forms with $g = 1$ correspond to elliptic curves, and Weierstrass equations for all elliptic curves of prime conductor less than $2 \cdot 10^9$ were computed by Bennett, Gherga, and Reznitzner [2]. Both of these datasets are in the LMFDB [24].

Below, we tabulate the number of forms of prime level in the ranges $[1, 10^4]$, $[10^4, 10^6]$, and $[10^6, 2 \cdot 10^6]$, grouped by the discriminant Δ of their Hecke fields and omitting discriminants which don't appear in our dataset.

Number with prime level in...

g	Δ	$[1, 10^4]$	$[10^4, 10^6]$	$[10^6, 2 \cdot 10^6]$	Levels
1	1	329	8843	6406	11 ... 1999957
2	5	158	1900	986	23 ... 9973, 10103, 10267 ... 1997773, 1999867
	8	37	242	100	29 ... 9613, 12619, 13537 ... 1986043, 1991489
	13	13	40	6	73 ... 9967, 15193, 23473 ... 1773361, 1863347
	12	1	14	3	113, 13763, 15083, 15919, 22481 ... 1493441, 1894043
	21	1	3	1	1283, 112289, 329671, 577807, 1670563
3	17		1		75653
	49	34	90	30	97 ... 9857, 12569, 13121 ... 1929943, 1972423
	229	8	20	1	211 ... 6997, 11197, 14563 ... 614659, 1972651
	81	3	13		127, 3581, 8513, 10753, 35591 ... 277793, 336551
	169	2	6	3	1481, 6569, 22943, 42209 ... 1221239, 1856201
	257	9	6	1	71 ... 26713, 39089, 224057, 255877, 463249, 581657, 1120969
4	148	12	6		41 ... 2341, 14929, 31039, 133117, 319489, 429397, 707801
	321	2	1		113, 7057, 86813
	725	16	5	1	137 ... 9011, 13681, 14759, 35977, 264919, 794111, 1716109
	1957	4	2		47, 223, 863, 2593, 28789, 185599
5	2777	3	2		197, 359, 1301, 28057, 63607
	8768		1		10169
6	70601	2	1		193, 719, 26777
	11^4		1		86161
6	13^5		1		171713

References

- [1] Sarah Arpin, Catalina Camacho-Navarro, Kristin Lauter, Joelle Lim, Kristina Nelson, Travis Scholl, and Jana Sotáková. Adventures in Supersingularland. *arXiv e-prints*, page arXiv:1909.07779, September 2019.
- [2] Michael A. Bennett, Adela Gherga, and Andrew Rechnitzer. Computing elliptic curves over \mathbb{Q} . *Math. Comp.*, 88(317):1341–1390, 2019.
- [3] Alex J. Best, Jonathan Bober, Andrew R. Booker, Edgar Costa, John Cremona, Maarten Derickx, David Lowry-Duda, Min Lee, David Roe, Andrew V. Sutherland, and John Voight. Computing classical modular forms. *arXiv e-prints*, page arXiv:2002.04717, February 2020.
- [4] Manjul Bhargava, Daniel M. Kane, Hendrik W. Lenstra, Jr., Bjorn Poonen, and Eric Rains. Modeling the distribution of ranks, Selmer groups, and Shafarevich-Tate groups of elliptic curves. *Camb. J. Math.*, 3(3):275–321, 2015.
- [5] Manjul Bhargava and Arul Shankar. Binary quartic forms having bounded invariants, and the boundedness of the average rank of elliptic curves. *Ann. of Math. (2)*, 181(1):191–242, 2015.
- [6] B. J. Birch and W. Kuyk, editors. *Modular functions of one variable. IV*. Lecture Notes in Mathematics, Vol. 476. Springer-Verlag, Berlin-New York, 1975.
- [7] R. P. Brent and H. T. Kung. Fast algorithms for manipulating formal power series. *J. Assoc. Comput. Mach.*, 25(4):581–595, 1978.
- [8] Christophe Breuil, Brian Conrad, Fred Diamond, and Richard Taylor. On the modularity of elliptic curves over \mathbb{Q} : wild 3-adic exercises. *J. Amer. Math. Soc.*, 14(4):843–939, 2001.
- [9] Reinier Bröker. Constructing supersingular elliptic curves. *J. Comb. Number Theory*, 1(3):269–273, 2009.
- [10] Armand Brumer and Oisín McGuinness. The behavior of the Mordell-Weil group of elliptic curves. *Bull. Amer. Math. Soc. (N.S.)*, 23(2):375–382, 1990.
- [11] Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *J. Cryptology*, 22(1):93–113, 2009.
- [12] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993.
- [13] James W. Cooley and John W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Comp.*, 19:297–301, 1965.
- [14] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to algorithms*. MIT Press, Cambridge, MA, third edition, 2009.
- [15] John Cremona. ecddata: 2016-10-17, October 2016.
- [16] Henri Darmon. *Rational points on modular elliptic curves*, volume 101 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2004.
- [17] Isabelle Déchène. Quaternion algebras and the graph method for elliptic curves, 1998. Thesis (MSc.)—McGill University (Canada), <https://www.math.mcgill.ca/darmon/theses/dechene-master/thesis.pdf>.
- [18] Fred Diamond and Jerry Shurman. *A first course in modular forms*, volume 228 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2005.
- [19] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: reductions and solutions. In *Advances in cryptography—EUROCRYPT 2018. Part III*, volume 10822 of *Lecture Notes in Comput. Sci.*, pages 329–368. Springer, Cham, 2018.

- [20] Benedict H. Gross. Heights and the special values of L -series. In *Number theory (Montreal, Que., 1985)*, volume 7 of *CMS Conf. Proc.*, pages 115–187. Amer. Math. Soc., Providence, RI, 1987.
- [21] Robert Harron and Andrew Snowden. Counting elliptic curves with prescribed torsion. *J. Reine Angew. Math.*, 729:151–170, 2017.
- [22] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In *Post-quantum cryptography*, volume 7071 of *Lecture Notes in Comput. Sci.*, pages 19–34. Springer, Heidelberg, 2011.
- [23] David Russell Kohel. *Endomorphism rings of elliptic curves over finite fields*. ProQuest LLC, Ann Arbor, MI, 1996. Thesis (Ph.D.)—University of California, Berkeley.
- [24] The LMFDB Collaboration. The L-functions and modular forms database. <http://www.lmfdb.org>, 2020. [Online; accessed 20 September 2020].
- [25] Kimball Martin. Refined dimensions of cusp forms, and equidistribution and bias of signs. *J. Number Theory*, 188:1–17, 2018.
- [26] Kimball Martin. An on-average maeda-type conjecture in the level aspect. *Proceedings of the American Mathematical Society*, page 1, Sep 2020.
- [27] J.-F. Mestre. La méthode des graphes. Exemples et applications. In *Proceedings of the international conference on class numbers and fundamental units of algebraic number fields (Katata, 1986)*, pages 217–242. Nagoya Univ., Nagoya, 1986.
- [28] W. H. Mills. Continued fractions and linear recurrences. *Math. Comp.*, 29:173–180, 1975.
- [29] Jennifer Park, Bjorn Poonen, John Voight, and Melanie Matchett Wood. A heuristic for boundedness of ranks of elliptic curves. *J. Eur. Math. Soc. (JEMS)*, 21(9):2859–2903, 2019.
- [30] Bjorn Poonen. Heuristics for the arithmetic of elliptic curves. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures*, pages 399–414. World Sci. Publ., Hackensack, NJ, 2018.
- [31] Michael O. Rabin. Probabilistic algorithms in finite fields. *SIAM J. Comput.*, 9(2):273–280, 1980.
- [32] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, second edition, 2009.
- [33] Math Stackexchange. Formulas for the (top) coefficients of the characteristic polynomial of a matrix. <https://math.stackexchange.com/questions/23899/>. [Online; accessed 20 September 2020].
- [34] Jacob Sturm. On the congruence of modular forms. In *Number theory (New York, 1984–1985)*, volume 1240 of *Lecture Notes in Math.*, pages 275–280. Springer, Berlin, 1987.
- [35] Andrew V. Sutherland. Isogeny volcanoes. In *ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium*, volume 1 of *Open Book Ser.*, pages 507–530. Math. Sci. Publ., Berkeley, CA, 2013.
- [36] Richard Taylor and Andrew Wiles. Ring-theoretic properties of certain Hecke algebras. *Ann. of Math. (2)*, 141(3):553–572, 1995.
- [37] Mark Watkins. Some heuristics about elliptic curves. *Experiment. Math.*, 17(1):105–125, 2008.
- [38] Douglas H. Wiedemann. Solving sparse linear equations over finite fields. *IEEE Trans. Inform. Theory*, 32(1):54–62, 1986.
- [39] Andrew Wiles. Modular elliptic curves and Fermat’s last theorem. *Ann. of Math. (2)*, 141(3):443–551, 1995.