

Some applications of higher dimensional isogenies to elliptic curves

Overview of results

DAMIEN ROBERT

ABSTRACT. We give some applications of the “embedding Lemma”. The first one is a deterministic polynomial time (in $\log q$) algorithm to compute the endomorphism ring $\text{End}(E)$ of an *ordinary* elliptic curve E/\mathbb{F}_q , provided we are given the factorisation of Δ_π . In particular, the full endomorphism ring computation can be done in quantum polynomial time.

The second application is an algorithm to compute the canonical lift of E/\mathbb{F}_q , $q = p^n$, (still assuming that E is ordinary) to precision m in time $\tilde{O}(nm \log^{O(1)} p)$. We deduce a point counting algorithm of complexity $\tilde{O}(n^2 \log^{O(1)} p)$. In particular the complexity is polynomial in $\log p$, by contrast of what is usually expected of a p -adic cohomology computation.

The third application is a quasi-linear CRT algorithm to compute Siegel modular polynomials of elliptic curves, which does not rely on any heuristic or conditional result (like GRH).

We also outline how to generalize these algorithms to (ordinary) abelian varieties.

1. INTRODUCTION

A spectacular application of dimension 2 isogenies to attack SIDH was given in [CD22; MM22] (relying on a lemma by Kani), followed by the use of even higher dimensional isogenies in [Rob22a]. In [Rob22a] we asked the question of whether the tools used to break SIDH could be used constructively. This was soon afterwards answered affirmatively in [Rob22b] for the evaluation of isogenies. Namely, the embedding lemma (see Section 2.1) shows that once we have evaluated a given N -isogeny f on an N' -torsion basis with N' -smooth (and the points of its Sylow components living in extensions of small degrees), we may then embed f into a smooth higher dimensional N' -isogeny F , and then use F to evaluate f on any other points efficiently.

A key obstacle to the use of [Rob22b] is the requirement of evaluating f on a basis of $E[N']$ first. In this paper, we show how to exploit [Rob22b] further. First, if f is an endomorphism given by an explicit polynomial in the Frobenius, evaluating f is easy at points of order prime to the denominators. We will exploit this in Section 4 to compute the endomorphism ring of an ordinary elliptic curve, and in Section 5 to compute the cardinal modulo p of an ordinary elliptic curve. The second idea is that once we have embedded f into the smooth F , we can lift f by lifting F , see Section 3. We will use this idea in Section 5 to compute the canonical lift of an ordinary elliptic curve (and as an application recover its full cardinal rather than only its cardinal modulo p), and in Section 6 for computing modular polynomials (and various variants).

The outline is as follow. In Section 2 we recall the embedding lemma, and then explain in Section 2.3 how to efficiently evaluate endomorphisms. In Section 3 we explain how to lift isogenies using the embedding lemma. Our first application is the computation of the

endomorphism ring of an elliptic curve in Section 4. Our second concern canonical lift and is described in Section 5. Our third is about modular polynomials in Section 6. This paper is just an overview of our results, and its aim is to give a brief leisurely description of the main algorithms. It will be followed by more technical papers giving more details.

1.1. **Thanks.** I thank Andrew Sutherland who asked me if higher dimensional isogenies could help computing the endomorphism ring of an elliptic curve. This led to Section 4.

I thank Jean-Marc Couveignes and Pierrick Gaudry for various discussions about other applications of canonical lifts than point counting, and Aurel Page for brainstorming sessions about trying to apply the same techniques as Section 5 to compute the crystalline cohomology of a general ordinary scheme.

I thank Antonin Leroux for various discussions about the computation of modular polynomials of elliptic curves. In particular, the algorithms presented in Section 6 at different points use supersingular elliptic curves for convenience; this idea is due to him. In [Ler23], Leroux gives other quasi-linear algorithms for modular polynomials that exclusively rely on supersingular curves. We explain the differences between his algorithms and ours in Section 6.

2. EVALUATING ISOGENIES AND ENDOMORPHISMS

2.1. **The embedding lemma.** If α_1, α_2 are two endomorphisms of an elliptic curve E of degree a_1 and a_2 , then $\alpha_1 \circ \alpha_2$ is of degree $a_1 a_2$. However it is harder to control the degree of the sum; by Cauchy-Schwartz we can bound it as: $(a_1^{1/2} - a_2^{1/2})^2 \leq \deg(\alpha_1 + \alpha_2) \leq (a_1^{1/2} + a_2^{1/2})^2$ (unless $\alpha_1 = -\alpha_2$). And $\alpha_1 + \alpha_2$ is of degree $a_1 + a_2$ if and only if $\alpha_1 \tilde{\alpha}_2$ is of trace 0.

If α_1 commutes with α_2 , we can instead use Kani's lemma [Kan97, § 2] to build an endomorphism F in dimension 2 on E^2 which is an $(a_1 + a_2)$ -isogeny (so is of degree $(a_1 + a_2)^2$ since we are in dimension 2). So by going to higher dimension we can combine degrees additively. The proof of this lemma is very simple (a simple two by two matrix computation), but its powerful algorithmic potential went unnoticed until Castrick and Decru applied it in [CD22] to attack on SIDH.

Lemma 2.1 (Kani). *An isogeny diamond is a commutative diagram of isogenies (between polarised abelian varieties)*

$$\begin{array}{ccc} A & \xrightarrow{f_1} & A_1 \\ \downarrow f_2 & & \downarrow f'_1 \\ A_2 & \xrightarrow{f'_2} & B \end{array}$$

with f_1 a d_1 -isogeny, f_2 a d_2 -isogeny and $f = f'_1 \circ f_1 = f'_2 \circ f_2$ a $d = d_1 d_2$ -isogeny (equivalently f'_1 is a d_2 -isogeny or f'_2 is a d_1 -isogeny). Then $F = \begin{pmatrix} f_1 & \tilde{f}'_1 \\ -f_2 & \tilde{f}'_2 \end{pmatrix}$ is a d -isogeny $A \times B \rightarrow A_1 \times A_2$ where $d = d_1 + d_2$.

Furthermore, $\text{Ker } F = \{(\tilde{f}'_1 x + \tilde{f}'_2 y, f'_1 x + f'_2 y), x \in A_1[d], y \in A_2[d]\}$, $\text{Ker } F = \{(\tilde{f}'_1 x, f'_1 x), x \in A_1[d]\}$ if $\text{Ker } \tilde{f}'_1 \cap \text{Ker } f'_1 = 0$ (eg if d_1 is prime to d_2), and if d_1 is prime to d_2 , then $\text{Ker } F = \{(d_1 x, f x), x \in A[d]\}$.

Proof. Kani's lemma is stated for elliptic curves in [Kan97, § 2], its extension to abelian varieties and the statements about the kernels is immediate, see [Rob22a; Rob22b]. \square

Remark 2.2. Our situation above with endomorphisms α_1, α_2 is a special case of Kani's lemma where we take $f_1 = f_2' = \alpha_1, f_2 = f_1' = \alpha_1$. Given $f_1 : A \rightarrow A_1$ a d_1 -isogeny, $f_2 : A \rightarrow A_2$ a d_2 -isogeny, then provided that d_1 is prime to d_2 , the pushforward of f_1 by f_2 gives an isogeny diamond.

Remark 2.3. A recursive application of Kani's lemma shows that we can embed m isogenies $f_1 : A \rightarrow A_1, \dots, f_m : A \rightarrow A_m$ of type β_1, \dots, β_m with the β_i totally real positive numbers primes with each other into a big $2^{m-1} \times 2^{m-1}$ matrix such that the corresponding isogeny F is a $\beta_1 + \dots + \beta_m$ isogeny $A \times A_{12} \times A_{13} \times \dots \times A_{ij} \dots \times A_{1234} \times \dots \times A_{ijkl} \dots \times \dots \rightarrow A_1 \times A_2 \times \dots \times A_i \dots \times A_{123} \times \dots \times A_{ijk} \dots \times A_{12345} \times \dots \times A_{ijklm} \dots \times \dots$ where A_{ijk} denotes the pushforward of f_i, f_j, f_k . The relative primality condition on the β_i is to ensure that these pushforward stays β_i isogenies. If the f_i are commuting endomorphism this is not necessary; F will then be an endomorphism on $A^{2^{m-1}}$.

We can combine Kani's lemma (extended to higher dimension) with Zarhin's trick: for any $m \in \mathbb{N}$, it is possible to build an m -endomorphism α_A on A^u where $u = 1, 2$ or 4 depending on whether m is a sum of 1, 2 or 4 squares, see [Rob22a]. The endomorphism α_A will be given by a $u \times u$ matrix with integer coefficients. So if $f : A \rightarrow B$ is an N -isogeny, we can consider the diagonal matrix $f : A^u \rightarrow B^u$, it will commute with α_A and α_B , so we can apply Lemma 2.1.

Lemma 2.4 (Embedding lemma). *For any $m > 0$, an N -isogeny $f : A \rightarrow B$ in dimension g of principally polarised abelian varieties can always be efficiently embedded into an $N + m$ -isogeny F in dimension $2ug$, according to whether m is a sum of $u = 4, 2$, or 1 squares (so F will be in dimension $8g, 4g$ and $2g$ respectively).*

More precisely, we embed f into the endomorphism F of $A^u \times B^u$ given by $F = \begin{pmatrix} \alpha_A & -f \\ f & \alpha_B \end{pmatrix}$.

The kernel of F is given by $\{(\tilde{\alpha}_A x + \tilde{f} y, -f x + \alpha_B y), x \in A^u[N + m], y \in A^u[N + m]\}$. If m is prime to N , then given the image of f on a basis (P_1, \dots, P_g) of $A[N_1]$ and (Q_1, \dots, Q_g) of $A[N_2]$ with $N + m = N_1 N_2$, we can explicitly decompose F as $F = F_2 \circ F_1$, F_1 a N_1 -isogeny with a basis of its kernel given by $(\alpha_A(P_i), -f(P_i))$, and F_2 a N_2 -isogeny, with a basis of the kernel of \tilde{F}_2 given by $(\alpha_A(Q_i), f(Q_i))$.

Proof. This comes from Lemma 2.1, except the description of the kernels when m is prime to N . But in this case the kernel of F is cyclic of rank g , hence we can split it as stated, cf [Rob22a, § 6.4]. \square

Remark that F embeds both f and its dual \tilde{f} .

Remark 2.5. For efficiency, we'd like to take u as small as possible in Lemma 2.4. If A has efficient endomorphisms, we can try to use them to compute an appropriate m -isogeny α_A on A^v , with $v < u$. If m is prime to N , we can then build the isogeny diamond in dimension vg given by the pushforward of $f : A^v \rightarrow B^v$ and α_A , and apply Lemma 2.1 directly.

An example of this situation is given by the curve $E_0 : y^2 = x^3 + x$ which has the explicit endomorphism i . We can use it to construct α_A on E_0 if m is a sum of two squares, or on E_0^2 if m is a sum of four squares. This allows to gain a factor 2 on u compared to Lemma 2.4.

2.2. Evaluating isogenies.

Definition 2.6. Let us define the N -evaluation problem as follow: given an N -isogeny $f : A/k \rightarrow B/k$ and a point $Q \in A(k)$, evaluate $f(Q)$. Here we remain deliberately vague about how f is specified, usually it will be by its kernel K , which is a maximal isotropic subgroup in $A[N]$.

Definition 2.7. The converse problem may be defined as follow: given an N -isogeny f as above, $P \in A[N']$ and the tuple $(P, f(P))$ along with a point $Q \in A(k)$, the (N, N') -interpolation problem ask to evaluate $f(Q)$. Of course, N' needs to be large enough compared to N so that f is uniquely determined by the data $(P, f(P))$.

We will be interested in the following weaker variant: the (N, N') -weak interpolation problem ask to evaluate $f(Q)$ provided we are given the value of f on a basis of $A[N']$.

Note that if $N = N'$, given the value of f on a basis of $A[N]$ we can (up to DLP computations¹) recover the kernel of f , hence the weak evaluation problem reduces to the evaluation problem in this case.

We may apply the embedding lemma to reduce the weak interpolation problem to the evaluation problem in higher dimension:

Lemma 2.8. *If $N' \geq N$, or more generally if we can find two (not necessarily distincts) divisors N'_1, N'_2 of N' with $N'_1 N'_2 \geq N$ and $N'_1 N'_2$ prime to N , then the weak (N, N') -interpolation problem reduces to the N' evaluation problem in higher dimension.*

Proof. The embedding lemma (Lemma 2.4) gives us an N' -isogeny F that embeds f , so evaluating $f(Q)$ can be done by evaluating $F(Q)$. If $N' \geq N$ we have the kernel of F directly, and so we may use it to evaluate F . In particular, if N' is prime to N , $\text{Ker } F$ can be completely determined by the value of $f(A[N])$: $\text{ker } F = \{(\alpha_A x, -fx), x \in A^u[N]\}$. A fun fact is that in this case we do not even need to compute DLPs to recover $\text{Ker } F$.

The more general case follow from the statement about the decomposition of F in Lemma 2.4: we have the kernel of F_1 so we may evaluate it, and we have the kernel of \tilde{F}_2 , so we can evaluate it on $(A^u \times B^u)[N'_2]$ to recover the kernel of F_2 . \square

This reduction is interesting because if $k = \mathbb{F}_q$ is a finite field and N' is powersmooth (or if N' is smooth and $A[N']$ lives in a small extension), the N' -evaluation problem can be done in polynomial time in $\log q$ and the smoothness bound B of N' (here we assume the dimension g fixed). This has the following application to the N -evaluation problem: if we can evaluate f on the N' -torsion, the evaluation problem reduces trivially to the (N, N') -weak interpolation problem, and we have just seen that this reduces to the N' -evaluation problem in higher dimension. So assuming that we have an oracle giving us this evaluation of f on $A[N']$, we can reduce the N -evaluation problem into the N' -evaluation problem (in higher dimension), which can be computed in polynomial time if N' is powersmooth. In other words, we embed the N -isogeny f into a powersmooth N' -isogeny F . This application is described in more details in [Rob22b]. For our complexity analysis, we need to briefly review the complexity results obtained there.

Proposition 2.9. *Let $f : A \rightarrow B$ be an N -isogeny between polarised abelian varieties of dimension g defined over a finite field $k = \mathbb{F}_q$. Let $N' = \prod_{i=1}^s \ell_i^{e_i}$, and let $u = 1, 2, 4$ according to whether $m = N' - N$ is prime to N and a sum of 1, 2 or 4 squares. Let F be the embedding of f given by the embedding lemma.*

Assume that we are given the image of $f : A \rightarrow B$ on a basis of each $A[\ell_i^{e_i}]$. Let B_1 be a bound on the ℓ_i , D_1 a bound on the degrees of the extensions where the points of $A[\ell_i^{e_i}]$ are defined, and D_2 a bound on the degree where the points of the compositum $A[\ell_i^{e_i} \ell_j^{e_j}]$ are defined.

¹If $f : A \rightarrow B$, an alternative strategy that do not require DLP is to extract a basis of $\text{Ker } \tilde{f} = f(A[N])$ from the image of f on $A[N]$. This only requires to compute Weil pairings and find a $g \times 2g$ submatrix with determinant of order N . The order check can be done if we know the factorisation of N . Then we recover generators of $\text{Ker } f$ via $\text{Ker } f = \tilde{f}(A[N])$, from which we extract a basis too.

Then we can decompose F as a product of e_i ℓ_i -isogenies in time $\tilde{O}(s(\sum e_i)B_1^{2ug}D_2 \log q)$. Evaluating F then requires $\tilde{O}((\sum e_i)D_1B_1^{2ug} \log q)$ operations.

Proof. We can write $F = F_2 \circ F_1$ where F_1 is a $\ell_1^{e_1}$ isogeny. From the image of f on a basis of $A[\ell^e]$, we easily recover $\text{Ker } F_1 = \text{Ker } F[\ell^e]$ since α_A is easy to evaluate. We decompose F_1 as a product of e_1 ℓ_1 -isogenies, since an ℓ_1 -isogeny in dimension $G = 2ug$ costs $O(\ell_1^G)$ to evaluate, this can be done in $\tilde{O}(e_1\ell_1^{2ug})$ operations over an extension of degree less than D_1 . We then need to push the image of f on a basis of $A[\ell_i^{e_i}]$ through f_1 , we have $2gs$ points to push and we work over an extension of degree less than D_2 , so this costs $O(se_1\ell_1^{2ug}D_2 \log q)$. Then we iterate.

Once this decomposition is done, evaluating F amount to evaluating the e_i ℓ_i -isogenies we have decomposed it into, each costing $O(\ell_i^g)$ operations over an extension of degree less than D_1 . \square

Corollary 2.10. *With the notations of Proposition 2.9, let B be a powersmooth bound on N' , ie a bound on the $\ell_i^{e_i}$. Then we can decompose F with $\tilde{O}(B^{2g(2+u)} \log^2 N')$ arithmetic operations over \mathbb{F}_q , and then evaluate it on a point in $O(B^{2g(1+u)} \log N')$ arithmetic operations (over the field of definition of this point).*

If $A[N']$ is rational, we can decompose F with $\tilde{O}(B^{2gu} \log^2 N')$ arithmetic operations over \mathbb{F}_q , then evaluate it in $O(B^{2gu} \log N')$ operations.

We can always find $N' = O(N)$ with a powersmooth bound of $B = \log N$, so in the general case, we can decompose F in time $\tilde{O}(\log^{2+g(4+2u)} N)$ arithmetic operations over \mathbb{F}_q , and then do subsequent evaluations in $\tilde{O}(\log^{1+g(2+2u)} N)$ arithmetic operations. And in the rational case, we can decompose F in time $\tilde{O}(\log^{2+2gu} N)$ arithmetic operations over \mathbb{F}_q , and then do subsequent evaluations in $\tilde{O}(\log^{1+2gu} N)$ arithmetic operations.

Proof. For the first statement, apply Proposition 2.9 with $D_1 = B^{2g}$, $D_2 = B^{4g}$, $s = O(\log N')$, $\sum e_i = O(N')$. For the second one, we use $D_1 = D_2 = 1$. \square

So the smaller u , the better complexity, but the harder to find a suitable N' . The easiest case is $u = 4$, we just need to find a powersmooth $N' > N$ and prime to N . We simply take the product of the first $O(\log N)$ primes to N , and then decompose $N' - N$ as a sum of squares. This cost $O(\log^2 N)$. The hardest case is $u = 1$, we need to find N' such that $N' - N$ is a square. In general this will not be possible. This could still have some applications, eg as in Section 5 where $N = p$, if we take the base field to be of a special form. The middle case is $u = 2$. It is difficult to test if an integer $N' - N$ is a sum of two squares (this requires factorizing it), so a solution is to test if $N' - N$ is prime and a sum of squares. A probabilistic algorithm (missing a few primes) cost $O(\log^2(N' - N))$. There is a heuristically a probability of $\Omega(1/\log N)$ that $N' - N$ is both a square and a sum of two primes, so we need to test $O(\log N) N'$. So we can find a suitable N' in *heuristic time* $O(\log^3 N)$. Of course once N' and the decomposition of $N' - N$ as a sum of two squares is found, it is easy to check that N' work.

Remark 2.11. Assume that we have a β -isogeny f on an abelian variety A with RM by K_0 , and K_0/\mathbb{Q} is Galoisian. Let $\beta_1 = \beta, \beta_2, \dots, \beta_g$ be the Galois conjugates of β , and choose any β_i isogeny f_i , with $f_1 = f$. Then we can use Remark 2.3 to embed the f_i (hence f) into F_1 a $\text{Tr } \beta$ -isogeny (assuming β is a prime power), and then use Kani's lemma again to embed F_1 into F_2 a N' -isogeny, $N' > \text{Tr } \beta$. To determine the kernel of F_2 we need to compute the action of the f_i (and various pushforwards) on $A[N']$ (and pushforwards $A_{ijk}[N']$).

2.3. Evaluating endomorphisms. Now the main obstacle of this idea is the need to evaluate f on the N' -torsion first. The idea of this paper is that if A/\mathbb{F}_q is an ordinary abelian variety, then $\mathbb{Z}[\pi]$ is an order in $\text{End}(A)$ (recall that for an ordinary abelian variety the endomorphism ring is invariant by a field extension, so $\text{End}(A) = \text{End}_{\mathbb{F}_q}(A) = \text{End}_{\overline{\mathbb{F}_q}}(A)$). So any element $\alpha \in \text{End}(A)$ can be written as $P_\alpha(\pi)/D$ where P_α is a polynomial of degree $d < 2g$ with integer coefficients, and D an integer dividing the index $f_\pi = [O_K : \mathbb{Z}[\pi]]$ where O_K is the maximal order in $\text{End}^0(A) = \text{End}(A) \otimes_{\mathbb{Z}} \mathbb{Q}$.

Note that since A is principally polarised, it contains $\mathbb{Z}[\pi, \bar{\pi}]$ where $\bar{\pi} = q/\pi$ (the Verschiebung) is the image of π by the Rosatti involution. This allows to write α as a polynomial in $\pi, \bar{\pi}$ where this time the denominator D divides $[O_K : \mathbb{Z}[\pi, \bar{\pi}]]$, so can be smaller. We won't need this in the following.

Evaluating α on a point $P \in A$ can be done as follow: find any point P' such that $P = DP'$. Then $\alpha(P) = P_\alpha(\pi)(P')$. We remark that π is easy to evaluate: it requires $O(\log q)$ arithmetic operations, and of course integer multiplications $[m]$ can be evaluated in $O(\log m)$ operations on the abelian variety. But if D has a large prime factor, finding P' will be very expensive in general. Still, in the particular case when $P \in A[N']$, with N' prime to D , then finding P' amount to inverting D modulo N' and a scalar multiplication. So we can evaluate α on $A[N']$, provided that N' is prime to D , in time polynomial in $\log q$ and the height of the coefficients of P_α/D . This allow us to efficiently embed α into a higher dimensional endomorphism F_α .

Thus, if α is an N -isogeny, taking $N' > N$ powersmooth and prime to N and the index f_π , we can evaluate the endomorphism α represented abstractly as above on any point $Q \in A(\mathbb{F}_q)$ in time polynomial in $\log q$ and the height of α :

Proposition 2.12. *If α is an N -endomorphism of height H , N' B powersmooth, $B = O(\log N)$, as in Corollary 2.10, then we can evaluate α in time $\tilde{O}((H + \log q)B^{2g} \log q + B^{4g+2ug} \log^2 N' \log q)$.*

Proof. We can use Mahler's bound to bound linearly the height of P_α from the height of α and of the characteristic polynomial χ_π of π (we assume the dimension g fixed here). By Weil's theorem, the height of χ_π is linear in $\log q$. So the coefficients of P_α are of height $O(H + \log q)$, and we need to evaluate the multiplication of these coefficients on points defined over an extension of degree $O(B^{2g})$ of \mathbb{F}_q ; this costs $O((H + \log q)B^{2g} \log q)$ operations. The Frobenius evaluation costs $O(\log q B^{2g} \log q)$. The remaining complexity follows from Proposition 2.9. \square

3. LIFTING AN ISOGENY

Once we have embedded an isogeny $f : A \rightarrow B$ into a higher dimensional one F , we can use F to lift f .

We will consider two kind of lifting: from \mathbb{F}_q to \mathbb{Z}_q , ie a p -adic lift. This will be used to compute canonical lifts in Section 5. Another lift we will use is from \mathbb{F}_q to $\mathbb{F}_q[[\epsilon]]$, this will be used in Section 6. In other words, given a deformation of A , we will compute the corresponding deformation of f (lifting an N -isogeny is unique when N is prime to the base characteristic). We could state this section for an arbitrary deformation data, namely given a ring R and an ideal I with $I^2 = 0$, an isogeny f over R/I and a lift of A to R , compute the lift of f to R . But for simplicity we will stick to the two cases mentioned above.

We recall that the moduli stack $\mathcal{A}_g(N)$ of ppav with a level N structure is smooth over $\mathbb{Z}[1/N]$ and finite etale over $\mathcal{A}_g/\mathbb{Z}[1/N]$, so in particular N -isogenies lift uniquely when N is prime to p .

Proposition 3.1. *Let $f : A \rightarrow B$ be an N -isogeny between polarised abelian varieties of dimension g defined over a finite field $k = \mathbb{F}_q$. Assume that we are given the image $\text{off} : A \rightarrow B$ on a basis of $A[\ell_i^{e_i}]$. Let F, N', B_1, D_1, D_2, u be as in the notations of Proposition 2.9. Assume that N and N' is prime to the characteristic p .*

Let $O = \mathbb{Z}_q$ or $\mathbb{F}_q[[\epsilon]]$, and m a target precision. Let \tilde{A} be a lift of A to O at precision m , ie to \mathbb{Z}_q/p^m or $\mathbb{F}_q[[\epsilon]]/\epsilon^m$. Then we can lift f to O at precision m , in time $\tilde{O}(s(\sum e_i)B_1^{2ug}D_2 \log q + (\sum e_i)D_1B_1^{2ug}m \log q)$.

If B is a powersmooth bound, then as in Corollary 2.10 we get a complexity of $\tilde{O}(B^{g(4+2u)} \log^2 N' \log q + B^{g(2+2u)}m \log N' \log q)$, so if we specialize further to $B = O(\log N)$: $\tilde{O}(\log^{2+g(4+2u)} N \log q + m \log^{1+g(2+2u)} N \log q)$. If the N' torsion is rational, this reduces to $\tilde{O}(B^{2gu} \log^2 N' \log q + mB^{2gu} \log N' \log q)$, so with $B = O(\log N)$ to $\tilde{O}(\log^{2+2gu} N \log q + m \log^{1+2gu} N \log q)$,

Proof. We first decompose $F : A^u \times B^u \rightarrow A^u \times B^u$ into ℓ_i -isogenies as in Proposition 2.9, this gives the first term of the complexity analysis. We will then lift F by a Newton iteration, doubling the precision at each step. Let us explain how to go from precision 1 to precision $m = 2$. The isogeny $f : A \rightarrow B$ will lift to $\tilde{f} : \tilde{A} \rightarrow \tilde{B}$, and the endomorphisms α_A, α_B too, hence F will lift as a matrix on $\tilde{A}^u \times \tilde{B}^u$.

A difficulty is that we don't yet know \tilde{B} yet. So we take an arbitrary candidate \tilde{B}_1 . We lift F (or more precisely the kernel of its decompositions) to \tilde{F}_1 starting on $\tilde{A}^u \times \tilde{B}_1^u$. Since \tilde{B}_1 is arbitrary, the codomain \tilde{C}_1 of \tilde{F}_1 has no reason to be $\tilde{A}^u \times \tilde{B}_1^u$ (ie our lift \tilde{F}_1 may not be an endomorphism), or even a product. Still, the deformation space of B is of dimension $g(g+1)/2$, and the codomain \tilde{C} depends linearly on the $g(g+1)/2$ deformation parameters. So taking $g(g+1)/2 + 1$ arbitrary lifts \tilde{B}_i and computing the codomain \tilde{C}_i each time, we can by linear algebra express the deformation of \tilde{C} linearly in terms of the deformation parameters of \tilde{C}_i .

As an example, if $g = 1$, and $O = \mathbb{F}_p[[\epsilon]]$, a lift of B to precision $m = 2$ correspond to an elliptic curve \tilde{B} with j -invariant $j(\tilde{B}) = j(B) + \lambda\epsilon$. Given a modular invariant J in dimension g , the J -invariant of the codomain \tilde{C} will be linear in ϵ . It suffices to compute the \tilde{C} corresponding to (for instance) $\lambda = 0$ and $\lambda = 1$ to recover this linear equation expressing \tilde{C} in function of \tilde{B} .

We then solve the linear equation $J(\tilde{C}) \simeq J(\tilde{A}^u \times \tilde{B}^u)$ (in terms of the deformation parameters for \tilde{B}). For this \tilde{B} , our \tilde{F} lifts as an endomorphism $\tilde{A}^u \times \tilde{B}^u$. Since the polarisation uniquely determine the product decomposition of an abelian variety (up to permutation), our decomposition $\tilde{A}^u \times \tilde{B}^u$ reduces to our starting decomposition $A^u \times B^u$. Since \tilde{F} reduces to f , the coefficients of the matrix corresponding to \tilde{F} reduces to the coefficients of the matrix corresponding to F . In particular, we have lifted f .

Since the lift of f is unique, the above discussion shows that the linear algebra step is invertible. Hence the whole algorithm requires $1 + g(g+1)/2$ isogeny computations of F at precision m . Working in this algebra cost $\tilde{O}(m \log q)$. So by Proposition 2.9, we get the second term of the complexity analysis, for $m = 2$.

For a general m , we do a Newton iteration, doubling the precision at each step. The last step will be at least as costly as all the preceding terms using the standard sub-linearity assumptions on the complexity of the multiplication at precision m , hence the final complexity analysis. \square

Remark 3.2. In the context of Remark 2.5, the lifting of F (hence f) will be a bit more complex. Indeed if we use a full isogeny diamond $F : A^v \rightarrow B^v \rightarrow C = A_1 \times B_1$ as in

Lemma 2.1 rather than an endomorphism, then given a lift \tilde{A} of A , we now need to find the lift of B, A_1, B_1 simultaneously.

As in Proposition 3.1, taking $1 + g(g+1)/2$ lifts of B we can compute the full deformation data giving the codomain \tilde{C} of the lift \tilde{F} associated to \tilde{B} . Now we need to know the local/formal locus $T = 0$ around C of abelian varieties of dimension $2v$ splitting into two abelian varieties of dimension v . For instance in dimension 2, the locus of product of elliptic curves is given by $\chi_{10} = 0$.

In our Newton process, the locus T will be given by linear equations. If $\tilde{C} = \tilde{A}' \times \tilde{B}'$ is in this locus, then the converse of Kani's lemma shows that \tilde{F} is given by a matrix induced by an isogeny diamond at precision m . \tilde{A}' and \tilde{B}' reduces to A_1 and B_1 (using an appropriate permutation if needed), so the isogeny diamond at precision m reduces to our isogeny diamond giving F . Hence we have lifted our full isogeny diamond at once, and in particular \tilde{F} lifts f . In particular, this also means that \tilde{C} has to be unique, so our linear system is invertible.

4. COMPUTING THE ENDOMORPHISM RING OF AN ORDINARY ELLIPTIC CURVE

If E/\mathbb{F}_q is an ordinary elliptic curve, we can recover the characteristic polynomial $\chi_\pi = X^2 - tX + q$ of π in polynomial time in $\log q$ by a point counting algorithm. We can thus recover $\Delta_\pi = t^2 - 4q$. If we know the factorisation of this discriminant, we can compute its associated fundamental discriminant, hence the maximal order $O_K = \mathbb{Z}[\omega]$ of $K = \mathbb{Q}(\sqrt{\Delta_\pi}) = \text{End}^0(E)$, and the factorisation of the conductor $f_\pi = [O_K : \mathbb{Z}[\pi]]$. We can write $\pi = a + f_\pi \omega$ (where a will depend on the trace of π , so has height $O(\log q)$). We know that $\pi - a \in \text{End}(E)$. To determine $\text{End}(E)$ is equivalent to determining the index of $\text{End}(E)$ in O_K or the index of $\mathbb{Z}[\pi]$ in $\text{End}(E)$, and so is equivalent to determining the largest divisor f_E of f_π such that $\frac{\pi-a}{f_E} \in \text{End}(E)$.

Since we know the factorisation of f_π , we are reduced to the following problem: let g be a factor of f_π . Is $\frac{\pi-a}{g} \in \text{End}(E)$? This can be done by checking that $\pi - a$ is trivial on $E[g]$, but computing the g torsion will be expensive if g has a large prime power as a factor.

Remark 4.1. This approach to endomorphism ring computations is used in [ELo7; FLo8] in dimension 2. The standard approach to compute the endomorphism ring of an ordinary elliptic curve is to follow paths in the isogeny volcano and is due to Kohel [Koh96] (see also [FMo2]). These algorithms are exponential in the worst case. An heuristic subexponential algorithm is presented in [BSo9], and further improved in [Bis11] to only rely on the GRH. This later algorithm has subexponential complexity (when provided with a factorisation of the discriminant) of $L(1/2, 1/\sqrt{2} + o(1))(\Delta_\pi)$.

Instead we use the embedding lemma.

Theorem 4.2. *Given an ordinary elliptic curve E/\mathbb{F}_q and the factorisation of the discriminant of the Frobenius π , $\text{End}(E)$ can be determined in polynomial time $O(\log^{7+2u} q)$ arithmetic operations.*

Here we can take $u = 4$ to get a proven complexity, or $u = 2$ to get an heuristic one.

Proof. We know how $\alpha = \frac{\pi-a}{g}$ is supposed to act on $E[N']$ (taking $N' > N(\alpha)$ prime to g and $N(\alpha)$), if it exists as an endomorphism. If α exists, we get an endomorphism F of E^{2u} (where $u = 1, 2, 4$) that embeds α as one of its matrix coefficient. If $N = \deg(\pi - a)$, then $\deg(\alpha) = N(\alpha) = N/g^2$. If $m = N' - N(\alpha)$ and γ an m -endomorphism on E^u , then we can build $\text{Ker } F$ as $\text{Ker } F = \{(\gamma P, -\alpha P) \mid P \in E^u[N']\}$. Since g is prime to N' , the action

of α on $E[N']$ is well defined even if it is not a real endomorphism, and it is easy to check that $\text{Ker } F$ is always isotropic in $E^{2u}[N']$.

So we first compute $E^{2u}/\text{Ker } F$ and check that F is indeed an endomorphism. This can be done in polynomial time if N' is powersmooth. If not, we know that α cannot be an endomorphism.

It is instructive to look at what happens if F is an endomorphism of E^{2u} . Let us assume $u = 1$ here for simplicity. Then by the converse of Kani's lemma, we know that F must be of the form $F = \begin{pmatrix} f_1 & -\widetilde{g}_1 \\ f_2 & \widetilde{g}_2 \end{pmatrix}$ for endomorphisms f_1, f_2, g_1, g_2 such that $g_2 g_1 = f_2 f_1$ and $\deg g_1 = \deg f_2$, and $\deg f_1 + \deg f_2 = N'$, and of course its kernel has to be the one specified above. So there is no guarantee a priori, even if F is an endomorphism, that it embeds α and not other endomorphisms.

But, since we can evaluate F efficiently, we can check if one of the matrix coefficient β of F acts like α on $E[N'']$, where N'' is powersmooth (we just need to check it on a basis of the N'' -torsion).² Since F is an N' -isogeny (because we have specified its kernel to be maximal isotropic in the N' -torsion), the individual components are ($\leq N'$)-isogenies.

Now by Cauchy-Schwarz, if α and β are two endomorphisms of degree $\leq M$, then $\alpha + \beta$ is of degree $\leq 4M$. So if the endomorphisms α, β agree on $E[N'']$, they are equal as long as $N''^2 > 4M$.

So we check if we can find a matrix coefficient β that acts like α on $E[N'']$. Then $g\beta$ acts like $\pi - a$ on $E[N'']$, so by the above result we have that $g\beta = \pi - a$ as long as $N''^2 > 4 \max(g^2 N', \deg(\pi - a)) = 4g^2 N'$ (since we take $N' > \deg((\pi - a)/g)$). In this case, $(\pi - a)/g$ is indeed an endomorphism, and the converse is immediate.

Of course we will follow this approach step by step, so we already know that say $(\pi - a)\ell/g$ (with $\ell \mid g$) is an endomorphism and we just need to check that $\ell\beta$ acts like $(\pi - a)\ell/g$, which allows to take a smaller N'' .

We do at most $\log|\Delta_\pi|$ steps, and the index f_π , hence its divisors, are at most $|\Delta_\pi|$. The full computation is thus polynomial in $\log q$ and $\log|\Delta_\pi|$. Since $\log|\Delta_\pi| = \log(q^2 - 4t) = O(\log q)$, we obtain the final complexity result by Corollary 2.10. \square

Remark 4.3. The dominating step of the endomorphism ring computation is thus the factorisation of the discriminant. The (unconditional randomised) proven complexity of the factorisation is $L(1/2, 1 + o(1))(\Delta_\pi)$ by [LP92], and the heuristic complexity of the NFS algorithm is of $L(1/3, (64/9)^{1/3} + o(1))(\Delta_\pi)$ by [BLP93]. Since factorisation can be done in polynomial time on a quantum computer by Schor's algorithm [Sho94], the endomorphism ring computation is in quantum polynomial time. Surprisingly it seems that no such quantum polynomial time algorithm was known before this article.

Remark 4.4. In the supersingular setting, then given an explicit suborder $O \subset \text{End}(E)$, generated by endomorphisms that we can efficiently evaluate on torsion points, then we can use the same ideas as above to check if a given order O' such that $O \subset O'$ is of finite index is still in $\text{End}(E)$.

Remark 4.5. Using Remarks 2.3 and 2.11, the same framework allows to compute the endomorphism ring of an ordinary abelian variety, provided that the real multiplication is Galoisian over \mathbb{Q} and we can compute the real Galois action efficiently on an abstract

²To be more precise, we need to test $\gamma\beta$ for all automorphisms γ of E . But E has no automorphisms apart from $[-1]$, unless $j(E) = 0$ or 1728. And we know the endomorphism ring of these curves.

representation of $\mathbb{Q}(\pi + \bar{\pi}) \subset \mathbb{Q}(\pi)$. Indeed, we can use Remark 2.11 to embed a β -endomorphism into a $\text{Tr}(\beta)$ -endomorphism, and then embed this one into a smooth endomorphism. Note that seeing π as a quadratic element over the real subfield, and σ an element of the Galois group of the real subfield, we can “conjugate” π by conjugating the coefficient of its minimal quadratic polynomial, and then conjugate an endomorphism $\alpha = a + b\pi$ via $\sigma(\alpha) = \sigma(a) + \sigma(b)\sigma(\pi)$.

It would be very interesting to be able to move in the ℓ -isogeny volcano in time polynomial in $\log \ell$.

5. POINT COUNTING AND CANONICAL LIFTS

5.1. The action of the Verschiebung on the tangent space. Let E/\mathbb{F}_q , $q = p^n$, be an ordinary elliptic curve. The Frobenius π_q has two eigenvalues, one λ which is invertible modulo p , and the other is q/λ . Since π_q is easy to evaluate, we can evaluate its action on the tangent space T_0E , but this gives us 0 since it is inseparable. The action of the Verschiebung $\bar{\pi}_q$ on T_0E allows us to recover $\lambda \pmod{p}$, hence the trace of π modulo p . Since $[q] = \bar{\pi}_q \circ \pi_q^3$, it is easy to evaluate the Verschiebung on a point P which is in the image of π_q . Unfortunately this does not help us to evaluate it on the tangent space, since the image of the Frobenius there is trivial. An alternative is to compute the kernel of the Verschiebung and apply Vélú’s formula, but since the degree of the Verschiebung is q , this is too expensive. (At this point we would actually compute the small Verschiebung instead which is of degree p).

Instead, since the Verschiebung is easy to compute on the N' -torsion ($N' > q$ powersmooth), we can embed it into a higher dimensional endomorphism F of E^{2u} ; this also embeds its dual π_q . We can then evaluate F on the tangent space at 0, this recover the action of $\bar{\pi}_q$ and π_q on T_0E . We thus get a polynomial time algorithm to recover $\lambda \pmod{p}$. Like above, it is more efficient to only embed π_p and $\bar{\pi}_p$ and recover λ via a norm, see [Rob21, § 6]. Using Corollary 2.10, this algorithm to recover $\lambda \pmod{p}$ costs $O(\log^{6+2u} p)$ arithmetic operations.

Notice the similarity with Schoof algorithm: in Schoof we compute the action of π_q on small ℓ_i -torsions groups $E[\ell_i]$, recover $\chi_\pi \pmod{\ell_i}$ via some DLP computations in $E[\ell_i]$, then reconstruct $\chi_\pi \pmod{\prod \ell_i}$ by the CRT. In our approach, we also compute π_q (or π_p) on these $E[\ell_i]$, but we instead use the action to reconstruct F a $\prod \ell_i$ isogeny embedding π_q and $\bar{\pi}_q$ (or π_p and $\bar{\pi}_p$).

5.2. Canonical lift. The above approach seems like a lot of trouble for less information than Schoof’s algorithm. But the nice thing about having the isogeny F is that lifting F gives a lift of the Frobenius. We can thus use F to see how π_p acts on the deformation space of E , and recover the canonical lift to precision m as in [MR22].

Usually, the action of π_p on the deformation space was computed using the modular polynomial ϕ_p . The modular polynomial ϕ_p is of size $O(p^3)$, and then evaluating to p -adic precision m cost $\tilde{O}(nmp^2)$. In [MR22], we explained how to compute the action via lifting the kernel of the Verschiebung $\bar{\pi}_p$ instead; since it is of degree p this allows to compute canonical lift in time $\tilde{O}(nmp)$. (A slight annoyance is that by using the Verschiebung rather than the Frobenius, we lose one bit in the p -adic precision at each step. In particular we need another method to bootstrap to precision $m = 2$: we use the fact that the étale p -torsion only lifts to \tilde{E} if $\tilde{E} = \hat{E}$ modulo p^2). Here we are going to use F instead, this way we can recover

³We can also write $\bar{\pi}_q = t - \pi_q$, this is closer in spirit to the description of Section 1, but of course at this point we do not know the trace t yet.

the action of π_p rather than $\bar{\pi}_p$ so there is no loss of precision, but more importantly F (and its lift) can be evaluated in time polynomial in $\log p$:

Theorem 5.1. *Given E/\mathbb{F}_q an ordinary elliptic curve, $q = p^n$, the canonical lift \hat{E} of E can be computed to precision m in time $\tilde{O}(nm \log^{4+2u} p + n \log^{7+2u} p)$, and the cardinal of E in time $\tilde{O}(n^2 \log^{4+2u} p + n \log^{7+2u} p)$.*

Here $u = 1, 2$ or 4 . We can only take $u = 1$ when p is a special form. We can always take $u = 4$. We can also take $u = 2$, the cost of finding N' described in Section 2.2 is heuristic, but once it is found it is easy to check that N' works. Furthermore this can be seen as a precomputation depending only on p .

Proof. This almost follow from Section 3, but there is a priori one technical difficulty towards applying the results of Section 3 to our situation: the Frobenius and Verschiebung have degree p not prime to the characteristic. In fact, the Verschiebung does not even lift uniquely. The key point will be that we lift both of them together.

Let us describe this in more details. Assume for now for simplicity that our F is in dimension 2. Let σ be the lift of the Frobenius to \mathbb{Q}_q , and \hat{E} denote the canonical lift of E , $\sigma(\hat{E})$ is then the canonical lift of $\sigma(E)$. F is an endomorphism of $E \times \sigma(E)$. The canonical lift \hat{E} is the unique lift \tilde{E} of E such that π_p lifts to $\tilde{\pi}_p : \tilde{E} \rightarrow \sigma(\tilde{E})$. We thus look for \tilde{E} such that the unique lift of F (as an isogeny) to $\tilde{E} \times \sigma(\tilde{E})$ is still an endomorphism (the lift is unique since F is étale). We remark that lifting F amount to lifting its kernel, which can be done by lifting generators of this kernel to points of N' torsion in \tilde{E} via a Newton iteration.

Let us look at how to lift from precision $m = 1$ to precision $m = 2$, then $m = 4$, and so on. We proceed as in Section 3: let $E_1 = E, E_2 = \sigma(E)$, we fix an arbitrary lift \tilde{E}'_1 of E and another \tilde{E}'_2 of $\sigma(E)$. We lift F to compute its action on $\tilde{E}'_1 \times \tilde{E}'_2$. We can then deform \tilde{E}'_1 to another lift \tilde{E}''_1 , compute the action of F again, and then deform \tilde{E}'_2 to \tilde{E}''_2 and compute the action of F . This is enough, via linear algebra, to be able to compute the action of F on arbitrary lifts of E_1 and E_2 , namely if $j(\tilde{E}_1) = j(\tilde{E}'_1) + \varepsilon_1 p, j(\tilde{E}_2) = j(\tilde{E}'_2) + \varepsilon_2 p$, we can compute $J(\tilde{E}_1 \times \tilde{E}_2 / \text{Ker } \tilde{F}) = J(\tilde{E}'_1 \times \tilde{E}'_2 / \text{Ker } \tilde{F}) + U\varepsilon_1 + V\varepsilon_2$, where J is a set of modular invariants in dimension 2. Note that we only care about the deformation of $E_1 \times E_2$ to a product abelian surface, that is why we only have two parameters $\varepsilon_1, \varepsilon_2$ rather than three.

If \tilde{E} is a lift of E , the Frobenius $\pi_p : E \rightarrow \sigma(E)$ lifts uniquely to $\tilde{E} \rightarrow \tilde{E}_2$. However, as mentioned above, in general the Verschiebung $\sigma(E) \rightarrow E$ does not lift to an arbitrary lift \tilde{E}_2 , and if it does the lift is not unique. In other words, the stack of elliptic curves with a degree p isogeny is étale at (E, π_p) when E is ordinary, but not at $(E, \bar{\pi}_p)$. In fact, by looking at the Serre-Tate formal moduli, it is classical that if $\tilde{E} = \hat{E}$ to precision m , and $\tilde{\pi}_p : \tilde{E} \rightarrow \tilde{E}_2$ is a lift of π_p , then $\tilde{E}_2 = \sigma(\hat{E})$ to precision $m + 1$. Hence the Verschiebung $\bar{\pi}_p$ can be lifted to \tilde{E}_2 if $\tilde{E}_2 = \hat{E}$ to precision at least 2, and in this case, among the multiple possible lifts, there is a canonical one which is the dual of the lift of the Frobenius $\tilde{E}_1 \rightarrow \tilde{E}_2$. It is characterised by being the unique lift whose kernel lies in the maximal unramified extension of \mathbb{Q}_q .

Anyway going back to our situation, when taking an arbitrary lift \tilde{E}_1 and \tilde{E}_2 of E and $\sigma(E)$, the lift of π_p to \tilde{E}_1 has codomain another elliptic curve $\tilde{E}_{2,can}$, and so the codomain of the lift \tilde{F} of F will not be a product abelian surface unless $\tilde{E}_2 = \tilde{E}_{2,can}$. On the moduli of abelian surfaces, the modular form χ_{10} has for locus the split surfaces, so plugging up χ_{10} in the expression of $J(\tilde{E}_1 \times \tilde{E}_2 / \text{Ker } \tilde{F})$ above we get a linear equation between ε_1 and ε_2 giving the locus where $\tilde{E}_2 = \tilde{E}_{2,can}$. On this locus, the Verschiebung lifts from \tilde{E}_2 to \tilde{E}_1

by the above discussion, hence F lifts as a matrix. Alternatively, we could plug the equation $J(\tilde{E}_1 \times \tilde{E}_2 / \text{Ker } \tilde{F}) = J(\tilde{E}_1 \times \tilde{E}_2)$.

The canonical lift \hat{E} at precision 2 can then be recovered by plugging the further equation $j(\tilde{E}_{2,can}) = \sigma(j(\tilde{E}_1))$. This way we obtain an Artin-Schreier equation $A\sigma(\varepsilon_1) + B\varepsilon_1 + C = 0$. Since the lifting solution is unique, A and B are not both 0, so they are uniquely determined (up to normalising C) from $j(\hat{E})$ and $\sigma(j(\hat{E}))$. In the general case where we are in dimension $2u$, we also use the equations $j(\tilde{E}_{2,can}) = \sigma(j(\tilde{E}_1))$ and $J(\tilde{E}_1 \times \tilde{E}_2 / \text{Ker } \tilde{F}) = J(\tilde{E}_1 \times \tilde{E}_2)$ where J is a set of modular equations to recover this Artin-Schreier equation.

From the Serre-Tate formal moduli, we then know that A is of valuation 0 and B of valuation 1. We can thus solve the equation to precision $m' = 1$ and then lift it via Newton iterations to the precision $m' = 2m$ that we need. This allows us to compute our canonical lift from precision 1 to 2, and we iterate.

Of course, we can also use the lift \tilde{F} to compute the action of $\hat{\pi}_p$ on $T_0\sigma\hat{E}$ to precision m . By Section 2.1, the dominating cost is the initial decomposition of \tilde{F} as a product of small isogenies which cost $O(\log^{6+2u} p)$ arithmetic operations, then the evaluations of \tilde{F} at precision m which cost $O(nm \log^{3+2u} p)$ arithmetic operations. \square

Remark 5.2. In Theorem 5.1, the \tilde{O} notation is actually hiding some ϵ dependency on the exponent rather than a usual quasi-linear dependency. We will call this pseudo-linear time. The trouble is the evaluation of σ on \mathbb{Z}_q , unless we have a normal basis (lifted to \mathbb{Z}_q), the best method I know is to evaluate it via modular composition. Using [KU11], modular composition over \mathbb{Z}_q to precision m can be done in $\tilde{O}(n^{1+\epsilon} m \log p)$, but is impractical.

Note that using Teichmüller representatives, the evaluation σ can be done in $O(p)$ multiplications over \mathbb{Z}_q , which is too much when p is large but can be used if p is small compared to n . A baby step giant step approach shows that one can also evaluate σ in $O(\sqrt{n})$ multiplications in \mathbb{Z}_q ; this can be used when n is small compared to p .

In the statement of Theorem 5.1, we implicitly assume that the cost of the algorithms will be dominated by the isogeny evaluations in higher dimension, and the arithmetic evaluation of σ on \mathbb{Z}_q is not dominant. We have seen this will be the case if we have a Gaussian normal basis, if n is small with respect to p , or if p is small with respect to n . If that is not the case, one should add a pseudo-linear $\tilde{O}(n^{1+\epsilon} m \log p)$ to the complexity estimates.

We can thus list the complexity of the different point counting algorithm, according to the underlying cohomology theory they use, as follow:

- Étale cohomology: Schoof's algorithm [Sch85] is in $O(\log^5 q) = O(n^5 \log^5 p)$, and SEA's algorithm [Sch95] in $\tilde{O}(\log^4 q) = \tilde{O}(n^4 \log^4 p)$.
- Rigid (Monsky-Washnitzer) cohomology: Kedlaya's algorithm [Kedo1] is in $\tilde{O}(n^3 p)$ and Harvey's variant [Har07] in $\tilde{O}(n^{3.5} p^{1/2} + n^5 \log p)$.
- Crystalline cohomology: Satoh's algorithm [Sat00] (after improvements by Harley) is in $\tilde{O}(n^2 p^2)$, and it has been improved to $\tilde{O}(n^2 p)$ in [MR22]. The (proven version of the) current algorithm is in $\tilde{O}(n^2 \log^{15} p)$ and the heuristic version in $\tilde{O}(n^2 \log^{11} p)$.

Remark 5.3. Over an ordinary abelian variety, the same method allows to recover the tangent matrix of $\hat{\pi}_p$ and $\tilde{\pi}_p$ to precision m in time $O(nm \log^{O(1)} p)$ (where the $O(1)$ hides a dependency at least linear in g).

On a supersingular elliptic curve E/\mathbb{F}_{p^2} , if α is a non trivial endomorphism, there exist a unique lift \tilde{E} of E such that α lift as an endomorphism $\tilde{\alpha}$ on \tilde{E} . If α can be efficiently evaluated

on torsion points, we can embed it into a higher dimensional endomorphism, and compute $(\tilde{E}, \tilde{\alpha})$ with the same approach as for an ordinary curve.

Remark 5.4. Another way to compute a canonical lift with a complexity sublinear in p is to compute the endomorphism ring and its class group, and then find a decomposition of the Frobenius as a product of small ideals. In other word, to find a cycle of small isogenies from E to E . (To forgo having to compute $\text{End}(E)$, one can also work with the class group of $\mathbb{Z}[\pi_E]$.) This gives an algorithm which is subexponential (under GRH) in p , see [CHo2, Theorem 2]. (A similar approach is also implicit in [Koh08, § 4.2], where Kohel tries to find a path of small isogenies from E to $\sigma(E)$.) Our present algorithm improves this complexity from subexponential to polynomial.

Remark 5.5. Another standard application of canonical lifts is the computation of class polynomials. If we start with an abelian variety A/\mathbb{F}_q , we can compute its endomorphism ring O , compute the action of the type norm to get the conjugate abelian varieties, lift them to \mathbb{Z}_q with enough precision, and then reconstruct the class polynomial, see eg [Rob21, Chapter 7]. If the class polynomial is of degree N and of height H (we assume here that $\Theta(\log N) = \Theta(\log H) = \Theta(\log \Delta_O)$), we need to lift to precision $m = H/\log p$ the N conjugate varieties. Under GRH we can span the group generated by the type norms with small generators, hence isogenies of small degree ($O(\log^2 \Delta_O)$). So we can find the N conjugate abelian varieties in time $\tilde{O}(N \log^{O(1)} q)$, and using the fast lifting algorithm lift them to precision m in time $\tilde{O}(N m n \log^{O(1)} p)$, hence compute the class polynomial in $\tilde{O}(N H n \log^{O(1)} p)$. This yields a quasi-linear algorithm when $\log q = O(\log^{O(1)}(NH))$.

Unfortunately, this does not suffices to compute the class polynomial of O in quasi-linear time. To bootstrap the algorithm we need to start with an abelian variety A/\mathbb{F}_q with CM by O . Under GRH the smallest prime p that totally splits in the class field will be in $O(N^2 \log^2 \Delta_O)$, and since there is $O(p^{g(g+1)/2})$ over \mathbb{F}_p , finding an abelian variety with CM by O will take too long (even if we only try to find the correct isogeny class).

5.3. Canonical lifts, crystalline cohomology, and impact on isogeny based cryptography.

We study the impact of canonical lifts on isogeny based cryptography. We have two ordinary elliptic curves E_1, E_2 with CM by O_K , and an ideal I inducing an isogeny $f : E_1 \rightarrow E_2$. The goal is to recover I . The case of CSIDH would be similar.

The canonical lifts \tilde{E}_1, \tilde{E}_2 still have CM by O_K and f lift to $\tilde{f} : \tilde{E}_1 \rightarrow \tilde{E}_2$. If we had the complex embedding of \tilde{E}_1, \tilde{E}_2 , to sufficiently high precision, we could recover I as follow. First, recover the period lattices $\Lambda_1 = H^1(\tilde{E}_{1,\mathbb{C}}, \mathbb{Z})$, $\Lambda_2 = H^1(\tilde{E}_{2,\mathbb{C}}, \mathbb{Z})$. This can be done in time quasi-linear in the precision thanks to the AGM. In fact, the AGM gives (an approximation of) the elements τ_1, τ_2 in the fundamental domain associated to the lattices Λ_1, Λ_2 . This allows to recover the lattices not only as \mathbb{Z} -modules, but also as O_K -modules. Here the action of O_K comes from the canonical embedding of O_K in \mathbb{C} and the action of \mathbb{C} on τ_1, τ_2 . More generally, whenever we have an explicit description of these lattices as \mathbb{Z} -module, along with the explicit action of O_K so that we also obtain the O_K -module structure (eg via normalizing this action), then (an approximation of) the ideal $I = [\Lambda_1 : \Lambda_2]$ can be recovered. Of course, to recover the complex embedding of $j(\tilde{E})$ from its p -adic embedding, the only way I know is via its minimal polynomial (which will be the class polynomial of O_K , cf Remark 5.5), which will be too big in cryptographic situations.

Still, one way of computing canonical lifts in characteristic p is via the degree p version of the AGM, so the similarity with recovering the period lattice in the complex case is striking. In this section we explain why canonical lifts allows to recover the crystalline cohomology

group, which behaves like the p -adic version of the period lattice, and also why this does not help *a priori* to attack (commutative) isogeny based cryptosystems.

First, recall that Serre-Tate theory gives an equivalence of category between lifts of (polarised) abelian varieties A/\mathbb{F}_q to \mathbb{Z}_q and the lift of their p -divisible group $A[p^\infty]$ (along with a polarisation). Now if A/\mathbb{F}_q is ordinary, the connected-étale exact sequence splits: $A[p^\infty] = A_{\text{étale}}[p^\infty] \oplus A_{\text{mult}}[p^\infty]$, where $A_{\text{étale}}$ is the étale component of rank g , and A_{mult} its Cartier dual (thanks to the Weil pairing), hence of multiplicative (more precisely toric) type. The étale component lifts uniquely to $\tilde{A}_{\text{étale}}[p^\infty]/\mathbb{Z}_q$, hence the multiplicative component lifts by duality. So lifting A amount to choosing an extension

$$(1) \quad 0 \rightarrow \tilde{A}_{\text{mult}}[p^\infty] \rightarrow G \rightarrow \tilde{A}_{\text{étale}}[p^\infty] \rightarrow 0.$$

The canonical lift corresponds to the unique split extension. Since an isogeny sends the toric and étale part to their counterparts, it lifts uniquely along these components. As canonical lifts correspond to the split extension (which is unique), this isogeny lifts canonically to the canonical lifts by the universal property of split extensions. As a particular case, all endomorphisms lift. Conversely if the Frobenius lifts, then it is easy to see that it induces a splitting of Equation (1), hence the lift is canonical.

More generally, an Ext group computation associated to the extension given by Equation (1) shows that the formal moduli (of ppav) is in bijection with $\widehat{\mathbb{G}_m}^{g(g+1)/2}$; these are the Serre-Tate formal coordinates. More canonically, the formal moduli is given by $\text{Hom}_{\mathbb{Z}_p}(T_p A \otimes T_p A^\vee, \widehat{\mathbb{G}_m})$, principal polarisations corresponds to symmetric elements, and an isogeny $f : A \rightarrow B$ lift to $\tilde{f} : \tilde{A} \rightarrow \tilde{B}$ where \tilde{A}, \tilde{B} are lifts (non necessarily canonical) associated to given Serre-Tate coordinates if and only f satisfy some symmetric conditions with respect to the Serre-Tate local coordinates associated to \tilde{A} and \tilde{B} [Kat81, Theorem 2.1]. Since the canonical lifts correspond to the trivial coordinates, this symmetry condition is automatic, and we recover the fact that all isogenies lift to the canonical lifts. It is worth pointing out that if we have two isogenies $f, g : A \rightarrow B$, and take an arbitrary lift \tilde{A} of A , then both isogenies f, g lift to \tilde{A} (provided that they are N_1, N_2 -isogenies with N_1, N_2 prime to p), but the codomain of the lifted isogenies need not be the same. However, if \tilde{A} is the canonical lift, then in this case the codomain is the same, namely the canonical lift \tilde{B} of B ! We refer to [Mes72; Kat81] for more details, along with the explicit form that the Kodaira-Spencer isomorphism and Gauss-Manin connection takes on the Serre-Tate formal moduli.

We will now make the link between the crystalline cohomology of A/\mathbb{F}_q and its canonical lift, following [DO86]. We assume that $p > 2$, so that the canonical divided power basis $x^n/n!$ on \mathbb{Z}_q is nilpotent at p on the Artin rings $\mathbb{Z}_q/p^m\mathbb{Z}_q$. Messing associates in [Mes72] to a p -divisible group \mathbb{G} over $k = \mathbb{F}_q$ a crystal $\mathbb{D}(\mathbb{G})$ (on these canonical divided powers). Its value at \mathbb{Z}_q (with the usual abuse of notation of patching the Artin algebras $\mathbb{Z}_q/p^m\mathbb{Z}_q$ together) is the usual ⁴ Dieudonné module over $\mathbb{Z}_q\{F, V\}$ associated to the p -divisible group \mathbb{G} . We can classify lifts of \mathbb{G} in terms of its associated crystal: the value of the crystal $\mathbb{D}(\mathbb{G})$ at $R = \mathbb{Z}_q/p^n\mathbb{Z}_q$ is given by the Lie algebra of the universal vectorial extension of any lift $\tilde{\mathbb{G}}$ of \mathbb{G} to R . This Lie algebra comes with a natural Hodge filtration, which reduces to the Hodge filtration on $\mathbb{D}(\mathbb{G}) | \mathbb{F}_q$ (this is called an admissible filtration), and while $\mathbb{D}(\mathbb{G}) | R$ does not depends on $\tilde{\mathbb{G}}$, the corresponding Hodge filtration does. Grothendieck-Messing theory states that this is an equivalence of category: lifts corresponds to admissible filtrations on $\mathbb{D}(\mathbb{G})$, and morphisms lifts when they respect the filtrations.

⁴The point of the Berthelot-Grothendieck-Messing-Mazur theory of crystals associated to Barsotti-Tate groups is that it extends to schemes, at least when the divided power structure is locally nilpotent.

When A/\mathbb{F}_q is an abelian variety, its crystalline cohomology group $H_{\text{crys}}^1(A/\mathbb{Z}_q)$ is the crystal associate to its p -divisible group $A[p^\infty]$. By the description of the crystal $\mathbb{D}(A[p^\infty])$ above, and since the formal Lie group associated to $A[p^\infty]$ is the formal group of A , we get that the \mathbb{Z}_q -module structure of $H_{\text{crys}}^1(A/\mathbb{Z}_q)$ is given by the De Rham cohomology of any lift \tilde{A}/\mathbb{Z}_q . Moreover the Hodge to De Rham spectral sequence always degenerates on an abelian variety X/R (even in characteristic p), so we have the exact sequence:

$$(2) \quad 0 \rightarrow \text{Lie}(X/R) \rightarrow H_{\text{DR}}^1(X/R) \rightarrow \text{Lie}(X^\vee/R)^\vee \rightarrow 0$$

since $\text{Lie}(X^\vee/R)^\vee \simeq H^1(X, \mathcal{O}_X)$. So given a lift \tilde{A}/\mathbb{Z}_q , the Hodge filtration on $H_{\text{DR}}^1(\tilde{A}/\mathbb{Z}_q)$ induces a Hodge filtration on the crystalline group $H_{\text{crys}}^1(A/\mathbb{Z}_q) = \mathbb{D}(A[p^\infty])$. This is exactly the Hodge filtration on the crystal $\mathbb{D}(A[p^\infty])$ corresponding via Grothendieck-Messing theory to the lifting $\tilde{A}[p^\infty]$ of the p -divisible group $A[p^\infty]$.

So in particular, given lifts \tilde{A}, \tilde{B} , an isogeny $f : A \rightarrow B$ lifts to $\tilde{A} \rightarrow \tilde{B}$ if and only if it respects the corresponding Hodge filtrations. By linear algebra, we recover Deuring's theory on the existence of a lift of an endomorphism of an elliptic curve (possibly supersingular).

Now if A/\mathbb{F}_q is ordinary, $H_{\text{crys}}^1(A/\mathbb{Z}_q) = U \oplus T$ splits canonically into subspaces stable by π_A , where U corresponds to the toric and T to the étale decomposition of $A[p^\infty]$: π_A is an isomorphism on U and p times an isomorphism on T . Indeed, the Hodge filtration over \mathbb{F}_q is given by the kernel of the Frobenius π_A on A acting on $H_{\text{DR}}^1(A/\mathbb{F}_q)$; and A/\mathbb{F}_q is ordinary iff the image of π_A gives the quotient of the Hodge filtration, ie iff $H_{\text{DR}}^1(A/\mathbb{F}_q)$ is split by the kernel and image of π_A . In particular, the components of the Hodge filtration are then stable by the Frobenius. The decomposition above is then the unique lifting of the Hodge filtration over \mathbb{F}_q to \mathbb{Z}_q that is stable by the Frobenius, with T a lift of the kernel and U a lift of the image. So the lift \tilde{A}/\mathbb{Z}_q associated to this decomposition by Grothendieck-Messing and Serre-Tate is the unique lift which admits a lift of the Frobenius, this is the canonical lift! For this canonical lift \tilde{A} , $T_0\tilde{A} \simeq T$, so on the tangent space we recover the action of F, V on $T \subset H_{\text{crys}}^1(A/\mathbb{Z}_q)$, and by duality also of U , ie we recover the full $\mathbb{Z}_q\{V, F\}$ -module structure. In other words, while any lift gives the \mathbb{Z}_q -module structure, only from the canonical lift \tilde{A}/\mathbb{Z}_q can we read of the action of F, V from the action of a lift of the Verschiebung or Frobenius acting on $T_0\tilde{A}$. This explains the link between crystalline cohomology and canonical lifts.

Now if A and B are abelian varieties over \mathbb{F}_q , Tate's ℓ -adic and p -adic isogeny theorems state that

$$\text{Hom}(A, B) \otimes \mathbb{Z}_\ell \simeq \text{Hom}_\pi(H_{\text{etale}}^1(B, \mathbb{Z}_\ell), H_{\text{etale}}^1(A, \mathbb{Z}_\ell))$$

(for $\ell \neq p$, and where we take the étale cohomology over $\overline{\mathbb{F}}_q$ and the isomorphism comes from pullback) and

$$\text{Hom}(A, B) \otimes \mathbb{Z}_p \simeq \text{Hom}_{\mathbb{Z}_q\{F, V\}}(H_{\text{crys}}^1(B/\mathbb{Z}_q), H_{\text{crys}}^1(A/\mathbb{Z}_q)).$$

We recall that $H_{\text{etale}}^1(A, \mathbb{Z}_\ell) = T_\ell(A)^\vee$, so Tate's ℓ -adic isogeny theorem also reads as

$$\text{Hom}(A, B) \otimes \mathbb{Z}_\ell \simeq \text{Hom}_\pi(T_\ell(A), T_\ell(B)).$$

So étale and crystalline cohomology behave like local (ℓ -adic and p -adic respectively) versions of the lattices we had (via singular cohomology) when working over \mathbb{C} . This is even more striking if we consider the canonical lifts \tilde{A}, \tilde{B} of A, B embedded into \mathbb{C} : $H_{\text{etale}}^1(A, \mathbb{Z}_\ell) \simeq H_{\text{etale}}^1(\tilde{A}, \mathbb{Z}_\ell) \simeq \Lambda_{\tilde{A}} \otimes \mathbb{Z}_\ell$ by the étale-singular comparison theorem, so we really recover the \mathbb{Z}_ℓ -local part of the lattice from the étale cohomology of A/\mathbb{F}_q ! We also have $H_{\text{etale}}^1(\tilde{A}, \mathbb{Z}_p) \simeq \Lambda_{\tilde{A}} \otimes \mathbb{Z}_p$ and by the étale-crystalline comparison theorem (since \tilde{A} has good reduction

modulo p) we have $\mathbb{D}_{\text{crys}}(H_{\text{étale}}^1(\tilde{A}, \mathbb{Q}_p)) = H_{\text{crys}}^1(A/\mathbb{Z}_p) \otimes_{\mathbb{Z}_p} \mathbb{Q}_p$ where \mathbb{D}_{crys} is Fontaine's functor associated to his period ring B_{crys} . We even have an integral comparison theorem in this case by Fontaine and Messing because $e = 1, i = 1$ so $ei < p - 1$. Since \mathbb{D}_{crys} has a quasi-inverse, we also recover the \mathbb{Z}_p -local part of the lattice from the crystalline cohomology of A/\mathbb{F}_q .

From all this discussion, it is now clear that canonical lift provides the \mathbb{Z}_p -local structure of the isogeny module: $\text{Hom}(E_1, E_2) \otimes \mathbb{Z}_p$. Fortunately for isogeny based cryptography, all ideals I of O_K are isomorphic over \mathbb{Z}_p . In fact in this setting, since E_1 and E_2 are linked by isogenies of degree prime to p , their crystalline cohomology are isomorphic. So we cannot read off the ideal connecting them uniquely from the crystalline cohomology, hence from their canonical lifts. The key difference with the situation over \mathbb{C} is that singular cohomology gave the full integral structure, ie the lattice over \mathbb{Z} rather than over \mathbb{Z}_p . The key difficulty here would be to know when a morphism given on the crystalline cohomology is integral, ie is induced by a real isogeny $f : A \rightarrow B$, rather than by an element $f \otimes \alpha, \alpha \in \mathbb{Z}_p$. A solution to check for integrality would be to check the action of f on the tangent space of the lifts (from the description above, a morphism on cohomology induces a morphism on tangent space which commutes with the Frobenius) and solve a differential equation, but this is not practical when f has cryptographic degree. In fact, if $f : A \rightarrow B$ is an isogeny of degree $< p$, then we could already reconstruct it (in $\tilde{O}(p)$, ie exponential time!) from its action on the tangent space T_0A, T_0B over \mathbb{F}_q . Canonical lifts would help only to reconstruct isogenies of degrees $> p$. (We will see this situation again in Section 6.)

The same difficulty appears when considering the étale cohomology: namely given a morphism on the cohomology group, ie an action on the ℓ^∞ torsion (commuting with the Frobenius), it is not obvious how to check if this action is induced by a real isogeny, without eg doing an interpolation. And without a way of gluing all these cohomology groups together in a coherent way, it seems difficult to extract an integral structure from these local structures.

We remark that even an efficient way to test an individual candidate would be not enough. For instance, say we had a large group of rational 2^m torsion in E_1/\mathbb{F}_q , and we tried to reconstruct an isogeny $f : E_1 \rightarrow E_2$ of degree $d < 2^m$. Then we could use the ideas of Section 2 to test efficiently if a candidate for f given by an application $E_1[2^m] \rightarrow E_2[2^m]$ is really induced by an isogeny. But even in this case, where integrality testing is somewhat easy, the difficulty remains that there are too many possibilities. We would really need a way to extract in batch the integral family $\text{Hom}(E_1, E_2)$ from the family $\text{Hom}(E_1, E_2) \otimes \mathbb{Z}_\ell$ seen via étale (or crystalline if $\ell = p$) cohomology.

6. MODULAR POLYNOMIALS

We explain how lifting isogenies can be used to compute modular polynomials. We first restrict to elliptic curves for simplicity, but the algorithms we describe can be generalised to abelian varieties, see [Rob21, § 5.3].

There are several related problems:

- (1) Compute $\phi_\ell(X, Y)$, it takes $\tilde{O}(\ell^3)$ space, and we will compute it in quasi-linear time.
- (2) Given a prime number $p \neq \ell$, compute $\phi_\ell(X, Y)$ modulo p ; it takes $\tilde{O}(\ell^2 \log p)$ space, and we will compute it in quasi-linear time.
- (3) Given an elliptic curve E/\mathbb{F}_p , compute $\phi_\ell(j(E), Y)$. It takes $\tilde{O}(\ell \log p)$ space, and I don't know how to compute it in quasi-linear time. Instead we will compute it

in $\tilde{O}(\ell^2 \log p)$ by lifting E to \mathbb{Z} and invoking Problem 4. We could also invoke Problem 2 to get a similar complexity.

- (4) Given an elliptic curve E/K of height H over a number field, compute $\phi_\ell(j(E), Y)$. It takes space $\tilde{O}(H\ell^2)$ since it is given by a polynomial of degree $O(\ell)$ with coefficients of height $\tilde{O}(H\ell)^5$. We will compute it in quasi-linear time.

We survey some known results on these problems. In [Engo9], Enge gives a $\tilde{O}(\ell^3 \log^4 \ell)$ (the notation \tilde{O} involving log means that we ignore log log factors) analytic algorithm for Problem 1, under the heuristic assumption that the loss of precision when evaluating the modular polynomials is not too large. In [BLS12], the authors give a $\tilde{O}(\ell^3 \log^3 \ell)$ algorithm under GRH. The same bound is achieved by [Ler23], also under some heuristics. In Section 6.3, we will get a similar bound without any heuristic.

For Problem 2, Leroux gives in [Ler23] two algorithms, one in $\tilde{O}(\ell^2(\log l \log p + \log^2 p) + p \log p)$ and another in $\tilde{O}(\ell^2(\log^4 p + \log^2 \ell \log p))$, again with some heuristics. In Section 6.2, we give a quasi-linear algorithm (without heuristics) of $\tilde{O}(\ell^2 \log^{3+2u} \ell \log p)$. Depending on the relative size of ℓ and $\log p$, our asymptotic may or may not be better than [Ler23]. Of course, if $\log p$ is too large, it is faster to compute ϕ_ℓ directly then reduce it modulo p .

For Problem 3, in [Sut13] Sutherland gives (under the GRH) algorithms in $\tilde{O}(\ell^2(\ell \log \ell + \log q) \log^2(\ell + \log q)) = \tilde{O}(\ell^3 + \ell^2 \log p)$, and $\tilde{O}(\ell^3(\log q + \log \ell) \log \ell)$ if $\log q = O(\ell^{O(1)})$, and with excellent space complexity (the second one has quasi-linear space complexity).

For Problem 4, in [Kie20b], Kieffer gives a quasi-linear algorithm to evaluate modular polynomials over a number field in dimension 2. The same analytic method works (and is easier) in dimension 1, one day the details will be written in [KR22]. Meanwhile the reader can consult [Rob21, Remarks 5.3.8 and 5.3.9]. We will give an algebraic (based on lifting) algorithm in Theorem 6.3, which also has quasi-linear complexity $\tilde{O}(H\ell^2)$. This allows to solve Problem 3 in $\tilde{O}(\ell^2 \log p)$. However the space complexity is also in $\tilde{O}(\ell^2 \log p)$, so in particular not quasi-linear as in Sutherland's approach, because via a p -adic lifting method we cannot do the same space saving trick as in the explicit CRT.

6.1. Evaluating $\phi_\ell(j(E), Y)$. Let us start with Problem 3. As explained above, one way to compute $\phi_\ell(j(E), Y)$ is to take an arbitrary lift \tilde{E} of E to \mathbb{Z} , evaluate $\phi_\ell(j(\tilde{E}), Y)$ via an analytic, p_0 -adic or CRT method, and reduce modulo p . We refer to [Rob21, § 5.3.8] for more details, and how to adapt these to the case where E is defined over a finite field \mathbb{F}_q .

The evaluation on \tilde{E} , $\phi_\ell(j(\tilde{E}), Y)$, takes space $\tilde{O}(\ell^2 \log p)$, so the best we can hope is an algorithm quasi-linear in this. The analytic method for evaluation in dimension 1 and 2 can be made quasilinear [KR22; Kie20b], but the p_0 -adic or CRT methods described in [Rob21, § 5.3.8] were not quasi-linear.

We now describe how to use the same ideas as in Section 5 to get a quasi-linear method, answering [Rob21, Conjecture 5.3.14]. We will have an initialization step, then a lifting step.

Initialisation: Find some suitable p_0 of good reduction, and compute all $\ell + 1$ isogenies f_i on $\tilde{E} \otimes \mathbb{F}_{p_0}$;

Lifting: Lift these isogenies to \mathbb{Z}_{p_0} to sufficiently high precision to recover $\phi_\ell(j(\tilde{E}), Y)$.

We need to do the lifting step for each isogeny and go to precision $m = O(\ell)$, so to get a quasi-linear algorithm we need the lifting step to be quasi-linear in the precision and more importantly in polylogarithmic time with respect to ℓ . We will use Section 3 for this.

⁵This bound is not uniform with respect to K , in particular the constants will involve the degree of K . See [Rob21, Remark 5.3.8] for a more refined discussion stating the explicit dependency on K .

A tricky part is the initialisation. The naive method is to work over the field of definition of the points of ℓ -torsion of $\tilde{E} \otimes \mathbb{F}_{p_0}$, compute all kernels and then the corresponding isogenies via Vélú's formula. In the worst case, the ℓ torsion on $\tilde{E} \otimes \mathbb{F}_{p_0}$ will live in an extension of degree ℓ^2 , so each isogeny will take $O(\ell)$ operations in a field of extension ℓ^2 and we have $O(\ell)$ isogenies to compute, so the complexity will be $\tilde{O}(\ell^4 \log p_0)$, which is too much.

The situation is better if $\tilde{E} \otimes \mathbb{F}_{p_0}$ has all its ℓ -torsion rational: we can find a basis in $\tilde{O}(\ell^2 \log^2 p_0)$, and then compute the $\ell + 1$ isogenies in $\tilde{O}(\ell^2 \log p_0)$. Heuristically, if we take a random p_0 this will happen with probability $1/\ell^2$, which is large enough (since rationality can be tested quickly) to obtain a quasi-linear algorithm. We will describe our algorithm with this heuristic (which we will state for a general number field rather than just \mathbb{Q}), then we will explain how to get rid of it.

Heuristic 6.1. *Let K be a number field. Given an elliptic curve \tilde{E}/K , and a prime \mathfrak{p}_0 of good reduction, if $E_0 = \tilde{E} \otimes \mathbb{F}_{\mathfrak{p}_0}$, the probability that $E_0[\ell]$ has rational points is $\Omega(1/\ell^2)$ (where the constants may depend on K).*

Remark 6.2. For elliptic curves over \mathbb{Q} , Heuristic 6.1 follows from the Sato-Tate theorem, at least if our lift \tilde{E} does not have CM. We state it as an Heuristic because the Sato-Tate conjecture is not proved for all number fields or for abelian varieties. We will see below an alternative method that does not rely on this Heuristic.

Theorem 6.3. *If \tilde{E} is a curve defined over a number field K of height H , $\phi_\ell(j(\tilde{E}), Y)$ can be evaluated in quasi-linear time $\tilde{O}(H\ell^2)$ ⁶.*

In particular, if E/\mathbb{F}_p is an elliptic curve, $\phi_\ell(j(E), Y)$ can be evaluated in time $\tilde{O}(\ell^2 \log p)$.

Remark 6.4. We can also compute the derivate $\partial\phi_\ell/\partial X(j(E), Y)$ using the derivative trick of [Rob21, Remark 5.3.10].

Proof. We prove Theorem 6.3 for $K = \mathbb{Q}$ for simplicity. For the general case, the only difficulty is to make the constants depending on K explicit, and we refer to [Rob21, § 5.3.8] for this.

Given a p_0 with good reduction, we have several strategies to test if $E_0 = E \otimes \mathbb{F}_{p_0}$ has rational ℓ -torsion and if so find a basis. The first strategy is to compute Ψ_{l, E_0} and find its roots, this costs $\tilde{O}(\ell^2 \log^2 p_0)$. The second strategy is to do point counting on E_0 , this takes polynomial time in $\log p_0$, and then we can find a basis in time $O(\ell \log p_0 + \log^2 p_0)$.

Under Heuristic 6.1, one can then find in $O(\ell^2)$ tries a p_0 with $\log p_0 = O(\log \ell)$ with good reduction such that $E_0[\ell]$ is rational. Then, via Vélú's formula we can evaluate our $\ell + 1$ isogenies $\phi_1, \dots, \phi_{\ell+1}$ in time in $\tilde{O}(\ell^2)$.

Now we can lift these isogenies to \tilde{E} ; since the evaluated modular polynomial has height $\tilde{O}(\ell \log p)$, we need to lift them to \mathbb{Z}_{p_0} at precision $m = \tilde{O}(\ell \log p)$. This can be done for instance by lifting a generator of the kernel then doing Vélú's formula at precision m , this cost $\tilde{O}(\ell m \log p)$ by isogeny, hence we do not get a quasi-linear algorithm, even if we use the sqrtVélú's algorithm [BDL+20] instead of Vélú.

Instead, we invoke Section 3, which allows us to lift our isogenies in time polynomial in $\log \ell$ (and the arithmetic operations on \mathbb{Z}_q at precision m) by isogeny. Lifting our $\ell + 1$ isogenies then cost $\tilde{O}(\ell m \log p)$, and we recover $\phi_\ell(j(\tilde{E}), Y)$ in quasi-linear time by a product tree.

⁶Again, the constants depends on K

The quasi-linear complexity under Heuristic 6.1 follows from the above discussion. We now explain how to get rid of it.

Fix any p_0 of good reduction. In Section 6.2, we explain how to evaluate ϕ_ℓ modulo p_0 in quasi-linear time, this allows to compute $\phi_\ell(j(E_0), Y)$ in $\tilde{O}(\ell^2 \log p_0)$.

First, let's assume that this evaluation splits (we could do a variant of Heuristic 6.1 for this, we expect it to happen with probability roughly $1/(\ell + 1)$). We can find the roots in $\tilde{O}(\ell \log^2 p_0)$, then recover the isogenies by solving a differential equation [BMS+08], assuming that the derivatives $\partial\phi_\ell/\partial X$ do not vanish. Each isogeny can be recovered in quasi-linear time by using a Newton iteration to solve the differential equations. So in this case the initialisation step can be done in quasi-linear time.

But in fact, we don't need $\phi_\ell(j(E_0), Y)$ to split. We simply work over the degree $\ell + 1$ algebra $A[T] = \mathbb{F}_{p_0}[T]/\phi_\ell(j(E_0), Y)$; T encode the (j -invariant of the codomain of the) universal isogeny on E_0 . We solve the differential equation over this algebra⁷ to recover the universal isogeny from E_0 over A . Then we lift to $\tilde{A} = \mathbb{Z}_{p_0}[T]/\phi_\ell(j(E_0), Y)$ at precision m , we obtain the j -invariant of the universal codomain \tilde{E}' over \tilde{A} . The modular polynomial is then the characteristic polynomial of $j(\tilde{E}')$ over $\tilde{A}/\mathbb{Z}_{p_0}$. This characteristic polynomial can be computed using power projection in pseudo-linear time $\tilde{O}(\ell^{1+\epsilon} m \log p)$ by [KU11] (see Remark 5.2 for the terminology).

Here we can do better: the characteristic polynomial $\phi_\ell(j(\tilde{E}), Y)$ is a lift/deformation of $\phi_\ell(j(E_0), Y)$ to precision m corresponding to the lift/deformation $\tilde{E}/\mathbb{Z}_{p_0}$ of E_0/\mathbb{F}_{p_0} . During our Newton iteration for lifting the universal isogeny, we can compute the corresponding deformation of the characteristic polynomial at each step, indeed everything become linear (since we double the precision, the deformation data has square 0, so the correction is given by linear data). So if we compute $\phi_\ell(j(\tilde{E}), Y)$ at precision m in parallel with lifting the isogeny, tweaking the algebra \tilde{A} each time⁸, we achieve a quasi-linear algorithm. \square

Remark 6.5. In the proof of Theorem 6.3, when using $\phi_\ell(j(E_0), Y)$ to reconstruct the isogenies, we need that the evaluation at $j(E_0)$ of the derivative modular polynomial $\partial\phi_\ell/\partial X(j(E_0), Y)$ is prime to the evaluated modular polynomial $\phi_\ell(j(E_0), Y)$. This failure can happen because while the moduli stack of elliptic curves with an ℓ -isogeny is smooth over $\mathbb{Z}[1/\ell]$, the modular polynomial ϕ_ℓ only describe a scheme birational to the coarse moduli space. In particular, we can have two different ℓ -isogenies from E_0 to the same elliptic curve E_1 ; these two isogenies then induces a cyclic ℓ^2 -endomorphism on E_0 . So this situation happens with low probability. A solution when this happens is, if $\text{Aut } E_0 = \pm 1$, to compute the normalisation of the modular polynomial. Another solution is to rigidify the data by imposing a small level $n \geq 3$ structure, this will allow us to distinguish between the two possible isogenies $E_0 \rightarrow E_1$ as long as they differ on the n -torsion, and we can always find a small n such that it is so. We will gloss over these details in the following.

Another thing that can go wrong is solving the differential equation: it involves division by numbers less than ℓ , which can cause loss of precision when working p_0 -adically with $p_0 \leq \ell$. Thankfully, this loss of precision can be well controlled [LV16; CEL20]. And for Problem 4, we might as well choose our starting p_0 large enough.

⁷Again assuming that the derivative of ϕ_ℓ at $j(E_0)$ is prime to ϕ_ℓ .

⁸This is key to make the algorithm work in quasi-linear time, otherwise we would need to compute some modular compositions. I owe this idea to Xavier Caruso.

Note also that Theorem 6.3 adapts immediately to the evaluation of ϕ_ℓ on $E/(\mathbb{Z}_p/p^m\mathbb{Z}_p)$ (or $E/(\mathbb{Z}_q/p^m\mathbb{Z}_q)$, we lift E to \mathbb{Q} (resp. a number field), and apply the Theorem. The lift has height $H = m \log p$, so the evaluation takes time $\tilde{O}(\ell^2 m \log p)$.

6.2. Computing ϕ_ℓ modulo p . We can now describe Problem 2. We will also use an initialisation followed by a lifting step.

Initialisation: Find some suitable E_0/\mathbb{F}_p , and compute all $\ell + 1$ isogenies f_i on E_0 .

Lifting: Lift these isogenies to $\mathbb{F}_p[[\epsilon]]$ to sufficiently high precision to recover $\phi_\ell(j(\tilde{E}), Y)$.

Like in Section 6.1, for simplicity, we will first rely on the following heuristic to get a quasi-linear algorithm for the initialisation step, which we will show later how to remove.

Heuristic 6.6. *Given a prime number p , the probability that a random elliptic curve E/\mathbb{F}_p has its ℓ -torsion rational is $\Omega(1/\ell^2)$.*

Heuristic 6.6 is the “horizontal” pendant of the “vertical” Heuristic 6.1.

Theorem 6.7. *The modular polynomial ϕ_ℓ can be evaluated modulo p in quasi-linear time $\tilde{O}(\ell^2 \log p)$.*

Proof. Using Heuristic 6.6, we can find E_0 modulo p with rational ℓ -torsion in around ℓ^2 tries. Like in Theorem 6.3, we test if E_0 has rational ℓ -torsion either by looking at the roots of ψ_ℓ , or by doing a point counting. Once we have found a suitable E_0 , we compute the $\ell + 1$ isogenies $f_i : E_0 \rightarrow E_i$ in time $\tilde{O}(\ell^2 \log p)$ using Vélú’s formula. The trick now, is that rather than lifting them to \mathbb{Z}_p , in this case we will lift them to $\mathbb{F}_p[[\epsilon]]$ to some precision m (ie working modulo ϵ^m). We let \tilde{E}_0 be the elliptic curve with j -invariant $j(E_0) + \epsilon$.

Again, we need to invoke Section 3 in order to lift these isogenies in time polynomial in $\log \ell$ and the arithmetic operations in $\mathbb{F}_p[[\epsilon]]$ at precision m . Using a product tree, we can then compute $\phi_\ell(j(E_0) + \epsilon, Y) \in \mathbb{F}_p[[\epsilon]]$ at precision m , which is enough to recover $\phi_\ell(X, Y)$ as long as $m > \ell + 1$.

Now we’d like to get rid of Heuristic 6.6. We use Theorem 6.3: given E_0/\mathbb{F}_p , we can evaluate $\phi_\ell(j(E_0), Y)$ in time $\tilde{O}(\ell^2 \log p)$. Then we work over the algebras of degree $\ell + 1$ $A[T] = \mathbb{F}_q[T]/\phi_\ell(j(E_0), T)$, $\tilde{A}[[\epsilon][T] = \mathbb{F}_q[[\epsilon][T]]/\phi_\ell(j(E_0), T)$, as in the proof of Theorem 6.3. Another solution is to pick up a curve E_0 such that $\phi_\ell(j(E_0), Y)$ splits, an easy way to choose such a curve is to take E_0 supersingular. Then we compute the $\ell + 1$ roots in $\tilde{O}(\ell \log^2 p)$, work over each $\ell + 1$ isogeny separably (lifting them to precision m), and compute a product tree at the end. According to the relative size of ℓ and $\log p$, this second approach can be faster than the first.

The astute reader will remark that we now have a recursive dependency between Theorems 6.3 and 6.7. In Theorem 6.3, to evaluate $\phi_\ell(j(E), Y)$ modulo p we require ϕ_ℓ modulo some small p_0 , but in Theorem 6.7 to compute ϕ_ℓ modulo p_0 we start with some evaluation $\phi_\ell(j(E_0), Y)$ modulo p_0 . At some point, we need to bootstrap this process. This is actually easy: we start with the curve $\tilde{E}_0 : y^2 = x^3 + x$ of j -invariant 1728. If $p \equiv 3 \pmod{4}$ it is supersingular and in this case $E(\mathbb{F}_{p^2}) = \mathbb{Z}/(p+1)\mathbb{Z} \oplus \mathbb{Z}/(p+1)\mathbb{Z}$. Hence if both $p \equiv 3 \pmod{4}$ and $\ell \mid p+1$, E_0 has its ℓ -torsion rational over \mathbb{F}_{p^2} . And the distribution of these p is easy to control thanks to the Dirichlet arithmetic density theorem (see Section 6.3); in particular we can find a small p_0 satisfying these conditions, with $\log(p_0) = O(\log \ell)$.

Starting with this E_0 and p_0 , we can evaluate $\phi_\ell(j(E_0), Y)$ modulo p_0 in time $\tilde{O}(\ell^2)$. We can then use this evaluation to vertical lift to \mathbb{Z}_{p_0} hence evaluate $\phi_\ell(j(E_0), Y)$ over

\mathbb{Z} in $\tilde{O}(\ell^2)$, then evaluate $\phi_\ell(j(E_0))$ modulo any p in $\tilde{O}(\ell \log p)$, then evaluate $\phi_\ell(X, Y)$ modulo p in $\tilde{O}(\ell^2 \log p)$ by horizontal lifting, for a total cost of $\tilde{O}(\ell^2 \log p)$.

As a special case, if we then specialize to $X = j(E)$, we obtain $\phi_\ell(j(E), Y)$ in $\tilde{O}(\ell^2 \log p)$, giving an alternative to Theorem 6.3. In this approach, we start from (E_0, p_0) , lift vertically to (E_0, p) (going through \mathbb{Z}), then horizontally to (E, p) . We could go the other way: evaluate $\phi_\ell(X, Y)$ modulo p_0 by horizontal lift in $\tilde{O}(\ell^2)$ (remember that $\log p_0 = O(\log \ell)$), then evaluate $\phi_\ell(j(\tilde{E}), Y)$ modulo p_0 in $\tilde{O}(\ell)$ (assuming p_0 is of good reduction for \tilde{E}), then evaluate $\phi_\ell(j(E), Y)$ modulo p by vertical lift in $\tilde{O}(\ell^2 \log p)$. \square

Remark 6.8. The proof of Theorem 6.7 shows that from $\phi_\ell(j(E_0), Y)$ modulo p_0 , one can get $\phi_\ell(j(E), Y)$ either by horizontal lifting from $(j(E_0), p_0)$ to $(j(E), p_0)$ and then by vertical lifting from $(j(E), p_0)$ to $(j(E), p)$, or by vertical lifting from $(j(E_0), p_0)$ to $(j(E_0), p)$ and then by horizontal lifting from $(j(E_0), p)$ to $(j(E), p)$!

Keeping track of the $\log \ell$ factors via Proposition 3.1, the complexity of computing $\phi_\ell \bmod p$ or $\phi_\ell(j(E), Y)$ can be done in $\tilde{O}(\ell^2 \log^{3+2u} \ell \log p)$, with the notations of Section 2.2. This neglects the complexity of $\tilde{O}(\log^4 \ell \log p + \log^3 \ell \log^2 p)$ needed to compute a basis of $E[\ell_i]$ for the small primes $\ell_i = O(\log \ell)$ of Section 2.2 (computation that can be done by factorisation of the division polynomials), if $\log p$ is too big it is better to evaluate ϕ_ℓ directly then reduce it modulo p anyway. This is also assuming that the derivatives of the modular polynomial are non zero. Otherwise, as explained in Remark 6.5, we need to work with level n modular invariants, for n a small prime up to $O(\log \ell)$.

Remark 6.9. When $p < \ell$, as mentioned in Remark 6.5, there will be division problems when solving the differential equation. In this case the solution is to do a p -adic lifting to precision m_2 with p^{m_2} large enough, using the control of the loss of precision given by [LV16; CEL20]. So in that case in the complexity of Remark 6.8, the $\log p$ factor should be replaced by a $\log \ell$. We remark that the algorithm is still quasi-linear.

6.3. A CRT algorithm to evaluate ϕ_ℓ . From Theorem 6.7, it is easy to see that we have a CRT algorithm to compute ϕ_ℓ in quasi-linear time. We recall that ϕ_ℓ has coefficients of heights $O(\ell \log \ell)$. Such an algorithm is not new, under GRH [BLS12] gives an algorithm in $O(\ell^3 \log^3 \ell \log \log \ell)$ time to compute ϕ_ℓ .

Here we show how in a CRT algorithm, we can sieve our primes p_i that we use to compute $\phi_\ell \bmod p_i$ in time $O(\ell \log \ell M(\ell \log p) + M(\ell^2 \log p) \log \ell)$ where $M(n)$ denotes the complexity of multiplication in a finite algebra with $O(1)$ separated variables and of degree n (in practice our algebra will be $\mathbb{F}_q[\epsilon]/\epsilon^m$ for Theorem 6.10). Taking $M(n) = O(n \log n \log \log n)$, and since we will have $O(\ell)$ primes p_i of size $\log p_i = O(\log \ell)$, the CRT reconstruction will cost $O(\ell^3 \log^3 \ell \log \log \ell)$ too, except that we do not rely on GRH.

The idea is to use the method of Section 6.2 with our curve $E_0 : y^2 = x^3 - x$. Fix $v = \log_2 \ell$ rounded up so that $2^v > \ell$. We sieve for a CRT prime p such that $p \equiv 3 \pmod{4}$, $p \equiv -1 \pmod{\ell}$, $p \equiv -1 \pmod{2^u}$.

For such a p , over \mathbb{F}_{p^2} the curve E_0 will be supersingular and have its full ℓ -torsion and 2^u -torsion rational. The algorithm to evaluate ϕ_ℓ modulo such a p is as follow:

- (1) Find a basis of $E_0[\ell]$ and $E_0[2^v]$; this costs $O((\log p + \log \ell)M(\log p))$ because we already know the cardinal of E_0 . The $\log p$ is the multiplication by the cofactor, and the $\log \ell$ corresponds to a Weil pairing computation.
- (2) Compute all $(\ell + 1)$ isogenies f_i and their image on the basis of the 2^v torsion via Vélu's formula, this costs $O(\ell M(\log p))$.

- (3) Let $a = 2^v - \ell$. According to whether a is a sum of at most 2 squares or of at most 4 squares, we can lift the isogenies f_i into 2^v -isogenies F_i in dimension 2 or 4. Here we gain a factor 2 with respect to the dimension because on E_0 we can make use of the endomorphism i , see also Remarks 3.2 and 2.5. In practice we will try to change v , so that $2^v - \ell$ is a sum of two squares. Given the density of sum of two squares, this should be possible while keeping v small (ie $v = O(\log \ell)$), and allows to work in dimension 2.

Anyway we decompose each F_i as a product of v 2-isogenies in higher dimension; this costs $O(v^2 M(\log p))$ by isogeny via the naive decomposition algorithm, or $\tilde{O}(vM(\log p))$ via the fast decomposition algorithm.

- (4) Now we lift E_0 to \tilde{E}_0 of j -invariant $1728 + \varepsilon$ to precision m . We lift the decomposition of 2-isogenies to $\mathbb{F}_p[[\varepsilon]]/(e^m)$ via a Newton iteration. This costs $O(vM(m \log p))$ by isogeny. We go to precision $m = \ell + 2$, so the total cost is $O(\ell \log \ell M(\ell \log p))$.
- (5) We take the product tree $\prod (Y - j(\tilde{E}_i))$ where \tilde{E}_i is the lifted codomain of the isogeny $f_i : E_0 \rightarrow E_i$. The product tree costs $O(M(\ell m \log p) \log \ell) = O(M(\ell^2 \log p) \log \ell)$. This gives us $\phi_\ell(1728 + \varepsilon, Y)$ modulo p .

By the Dirichlet density theorem, we can find $O(\ell)$ such primes p with size $\log p = O(\log \ell)$. Notice that unlike effective bounds for the Chebotarev density theorem, which are exponentially worse without GRH (polynomials in the discriminant Δ , vs in $\log^2 \Delta$ with GRH), these effective bounds on Dirichlet's theorem on arithmetic progression are only polynomially worse without GRH. Since the complexity involve $\log p_0$, this only affect the constants, so this is enough for our application.

Hence the final complexity result:

Theorem 6.10. *There exists a CRT algorithm to compute ϕ_ℓ in time $O(\ell^3 \log^3 \ell \log \log \ell)$.*

An alternative strategy to compute ϕ_ℓ would be to fix only one suitable p , proceed as in the first steps of the algorithm above to find the $\ell + 1$ isogenies from E_0 modulo p , and then lift them to $\mathbb{Z}_p[[\varepsilon]]/(p^{m_1}, \varepsilon^{m_2})$. In other words, combine the algorithm of Theorem 6.10 with a p -adic evaluation algorithm like Theorem 6.3.

Remark 6.11. In the CRT algorithms [BLS12; Ler23] to compute ϕ_ℓ , the authors also choose suitable ‘‘CRT primes’’ p_i , and reconstruct each $\phi_\ell \bmod p_i$ in quasi-linear time. They do that by finding a clever way ([BLS12] use the volcano structure and the class group action, [Ler23] use the supersingular graph structure) to select $\ell + 1$ different j -invariants j_0, \dots, j_ℓ and from each of them construct the $\ell + 1$ isogenies starting from them. This gives the polynomials $\phi_\ell(j_i, Y)$, and then $\phi_\ell \bmod p$ is reconstructed by an interpolation in each coefficient. Our method only uses one j -invariant j_0 and lift the $\ell + 1$ isogeny to $\mathbb{F}_p[[\varepsilon]]$ to precision $\ell + 1$. We could combine both approaches: namely fix some j -invariants along with some precision $(j_0, m_0), (j_1, m_1), \dots, (j_r, m_r)$ with $\sum m_i \geq \ell + 1$, lifts the $\ell + 1$ isogenies from j_i to precision m_i , obtain the coefficients of $\phi_\ell(j_i + \varepsilon, Y)$ to precision m_i , and then do a Hermite-Padé interpolation to reconstruct ϕ_ℓ .

6.4. Modular polynomials for abelian varieties. It is easy to generalize Theorem 6.10 to computing Siegel modular polynomials for abelian varieties: we select the same CRT primes and work with E_0^g . Since we know how to compute isogenies in the theta model [LR22], we obtain:

Theorem 6.12. Fix a modular invariant J in dimension g which can be computed efficiently from the theta constants (say of level 4). Then we can compute the (rational or integral version⁹ of the) Siegel modular polynomials Φ_ℓ with respect to J in quasi-linear time.

We can adapt Theorem 6.12 to the case of Hilbert modular polynomials as follow. We are given a Galoisian totally real field K_0 of dimension g , and $\beta \in K_0$ a totally positive element, we want to compute the Hilbert modular polynomials Φ_β . We suppose that we have an algorithm to compute a β -isogeny efficiently (say quasi-linear in the norm of β); this is still somewhat a work in progress, see [DJR+22; LR22]. We suppose also that if p is a CRT prime as above, we have an efficient way to find A_0 a product of g supersingular curves and an embedding of K_0 into $\text{End}(A_0)$. For instance, when $g = 2$, any D -isogeny $E_0 \rightarrow E_1$, E_0 our supersingular curve $y^2 = x^3 + x$, gives an embedding $\mathbb{Q}(\sqrt{D}) \rightarrow \text{End}(E_0 \times E_1)$. Then using Remarks 2.3 and 2.11 combined with the algorithm of Theorem 6.12, we also get a quasi-linear algorithm to compute Φ_β .

Remark 6.13. We can also extend Problems 2 and 4 to abelian varieties as follow. If $N = g(g+1)/2$, the Siegel modular polynomials have height $\tilde{O}(\ell^N)$ by [Kie20a], degree in Y $O(\ell^N)$, and each coefficients have degree $O(\ell^N)$ in each of the N variables X_1, \dots, X_N . So their total size is $\tilde{O}(\ell^{N(N+2)})$. Their size modulo p is $\tilde{O}(\ell^{N(N+1)} \log p)$ and the size of the evaluation at $J(\tilde{A})$, where \tilde{A}/K is an abelian variety over a number field K with J -invariant of height H is $\tilde{O}(\ell^{N(N+1)}H)$.

For the initialisation step, the worst case to get the full ℓ -torsion is to work over an extension of degree ℓ^{2g} of our base finite field, then each of the $O(\ell^N)$ isogeny take $O(\ell^g)$ operations over this extension, for a total cost of $\tilde{O}(\ell^{N+3g})$. For $g = 2$, $N = 3$, $\ell^{N(N+1)} = \ell^{12}$, while $\ell^{N+3g} = \ell^9$, so even there the initialisation step is not dominant even if we don't optimize it. We can still optimize the initialisation as follow: for Problem 2, we can start with $E_0^g, E_0/\mathbb{F}_{p^2}$ any supersingular curve. Indeed the full ℓ -torsion is defined over an extension of degree at most $O(\ell)$, rather than $O(\ell^{2g})$.

The situation is different for Hilbert modular polynomial though. Here we need to be careful with our initialisation step, so to get a quasi-linear algorithm we need similar algorithms as in Sections 6.1 and 6.2, using a recursive computation of Φ_β along different abelian varieties and primes. For this, we need to be able to (efficiently!) recover the isogenies from the evaluated modular polynomials. For $g = 2$ this is done in [KPR20], and a very brief outline of an algorithm with modular polynomials given by theta constants is described in [Rob21, § 5.7].

REFERENCES

- [BDL+20] D. Bernstein, L. De Feo, A. Leroux, and B. Smith. “Faster computation of isogenies of large prime degree”. In: *Algorithmic Number Theory Symposium*. 2020. arXiv: [2003.10118](#).
- [Bis11] G. Bisson. “Computing endomorphism rings of elliptic curves under the GRH”. In: *Journal of Mathematical Cryptology* (2011). arXiv: [1101.4323](#).
- [BS09] G. Bisson and A. Sutherland. “Computing the endomorphism ring of an ordinary elliptic curve over a finite field”. In: *Journal of Number Theory* (2009).
- [BMS+08] A. Bostan, F. Morain, B. Salvy, and E. Schost. “Fast algorithms for computing isogenies between elliptic curves”. In: *Mathematics of Computation* 77.263 (2008), pp. 1755–1778.

⁹See [Rob21, § 5.3.6]

- [BLS12] R. Bröker, K. Lauter, and A. Sutherland. “Modular polynomials via isogeny volcanoes”. In: *Mathematics of Computation* 81.278 (2012), pp. 1201–1231. arXiv: [1001.0402](https://arxiv.org/abs/1001.0402).
- [BLP93] J. Buhler, H. Lenstra, and C. Pomerance. “Factoring integers with the number field sieve”. In: *The development of the number field sieve* (1993), pp. 50–94.
- [CEL20] X. Caruso, E. Eid, and R. Lercier. “Fast computation of elliptic curve isogenies in characteristic two”. 2020.
- [CD22] W. Castryck and T. Decru. *An efficient key recovery attack on SIDH (preliminary version)*. Cryptology ePrint Archive, Paper 2022/975. 2022. URL: <https://eprint.iacr.org/2022/975>.
- [CH02] J.-M. Couveignes and T. Henocq. “Action of modular correspondences around CM points”. In: *International Algorithmic Number Theory Symposium*. Springer, 2002, pp. 234–243.
- [DJR+22] A. Dudeanu, D. Jetchev, D. Robert, and M. Vuille. “Cyclic Isogenies for Abelian Varieties with Real Multiplication”. Accepted for publication at *Moscow Mathematical Journal*. Feb. 2022. URL: <http://www.normalesup.org/~robert/pro/publications/articles/cyclic.pdf>. HAL: [hal-01629829](https://hal.archives-ouvertes.fr/hal-01629829).
- [DO86] B. Dwork and A. Ogus. “Canonical liftings of Jacobians”. In: *Compositio Mathematica* 58.1 (1986), pp. 111–131.
- [EL07] K. Eisentrager and K. Lauter. “A CRT algorithm for constructing genus 2 curves over finite fields”. In: *AGCT-11* (2007).
- [Eng09] A. Enge. “Computing modular polynomials in quasi-linear time”. In: *Math. Comp* 78.267 (2009), pp. 1809–1824.
- [FM02] M. Fouquet and F. Morain. “Isogeny volcanoes and the SEA algorithm”. In: *Algorithmic number theory (Sydney, 2002)*. Vol. 2369. Lecture Notes in Comput. Sci. Berlin: Springer, 2002, pp. 276–291. DOI: [10.1007/3-540-45455-1_23](https://doi.org/10.1007/3-540-45455-1_23).
- [FL08] D. Freeman and K. Lauter. “Computing endomorphism rings of Jacobians of genus 2 curves over finite fields”. In: *Algebraic geometry and its applications* (2008), pp. 29–66.
- [Har07] D. Harvey. “Kedlaya’s algorithm in larger characteristic”. In: *Int. Math. Res. Notices* (2007).
- [Kan97] E. Kani. “The number of curves of genus two with elliptic differentials.” In: *Journal für die reine und angewandte Mathematik* 485 (1997), pp. 93–122.
- [Kat81] N. Katz. “Serre-Tate local moduli”. In: *Surfaces algébriques*. Springer, 1981, pp. 138–202.
- [Ked01] K. Kedlaya. “Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology”. 2001.
- [KU11] K. S. Kedlaya and C. Umans. “Fast polynomial factorization and modular composition”. In: *SIAM Journal on Computing* 40.6 (2011), pp. 1767–1802.
- [Kie20a] J. Kieffer. “Degree and height estimates for modular equations on PEL Shimura varieties”. Accepted at London Mathematical Society. 2020. HAL: [hal-02436057](https://hal.archives-ouvertes.fr/hal-02436057).
- [Kie20b] J. Kieffer. “Evaluating modular polynomials in genus 2”. 2020. HAL: [hal-02971326](https://hal.archives-ouvertes.fr/hal-02971326).
- [KPR20] J. Kieffer, A. Page, and D. Robert. “Computing isogenies from modular equations between Jacobians of genus 2 curves”. Oct. 2020. URL: http://www.normalesup.org/~robert/pro/publications/articles/modular_isogenies_g2.pdf. HAL: [hal-02436133](https://hal.archives-ouvertes.fr/hal-02436133).

- [KR22] J. Kieffer and D. Robert. “Fast evaluation of modular polynomials and compact representation of isogenies between elliptic curves”. Aug. 2022. In preparation.
- [Koh96] D. Kohel. “Endomorphism rings of elliptic curves over finite fields”. PhD thesis. University of California, 1996.
- [Koh08] D. R. Kohel. “Complex multiplication and canonical lifts”. In: *Algebraic Geometry And Its Applications: Dedicated to Gilles Lachaud on His 60th Birthday*. World Scientific, 2008, pp. 67–83.
- [LV16] P. Lairez and T. Vaccon. “On p-adic differential equations with separation of variables”. In: *Proceedings of the ACM on International Symposium on Symbolic and Algebraic Computation*. 2016, pp. 319–323.
- [LP92] H. W. Lenstra and C. Pomerance. “A rigorous time bound for factoring integers”. In: *Journal of the American Mathematical Society* 5.3 (1992), pp. 483–516.
- [Ler23] A. Leroux. “Computation of Hilbert class polynomials and modular polynomials from supersingular elliptic curves”. In: *Cryptology ePrint Archive* (2023).
- [LR22] D. Lubicz and D. Robert. “Fast change of level and applications to isogenies”. Accepted for publication at **ANTS XV Conference** — Proceedings. Aug. 2022. URL: http://www.normalesup.org/~robert/pro/publications/articles/change_level.pdf.
- [MR22] A. Maiga and D. Robert. “Towards computing canonical lifts of ordinary elliptic curves in medium characteristic”. Mar. 2022. URL: http://www.normalesup.org/~robert/pro/publications/articles/fast_canonical_lift_g1.pdf.
- [MM22] L. Maino and C. Martindale. *An attack on SIDH with arbitrary starting curve*. Cryptology ePrint Archive, Paper 2022/1026. 2022. URL: <https://eprint.iacr.org/2022/1026>.
- [Mes72] W. Messing. “The crystals associated to Barsotti-Tate groups”. In: *The crystals associated to Barsotti-Tate groups: with applications to abelian schemes*. Springer, 1972, pp. 112–149.
- [Rob21] D. Robert. “Efficient algorithms for abelian varieties and their moduli spaces”. HDR thesis. Université Bordeaux, June 2021. URL: <http://www.normalesup.org/~robert/pro/publications/academic/hdr.pdf>. Slides: <2021-06-HDR-Bordeaux.pdf> (1h, Bordeaux).
- [Rob22a] D. Robert. “Breaking SIDH in polynomial time”. Aug. 2022. URL: http://www.normalesup.org/~robert/pro/publications/articles/breaking_sidh.pdf. eprint: 2022/1038.
- [Rob22b] D. Robert. “Evaluating isogenies in polylogarithmic time”. Aug. 2022. URL: http://www.normalesup.org/~robert/pro/publications/articles/polylog_isogenies.pdf. eprint: 2022/1068.
- [Sat00] T. Satoh. “The canonical lift of an ordinary elliptic curve over a finite field and its point counting”. In: *J. Ramanujan Math. Soc.* 15.4 (2000), pp. 247–270.
- [Sch85] R. Schoof. “Elliptic curves over finite fields and the computation of square roots mod p ”. In: *Mathematics of computation* 44.170 (1985), pp. 483–494.
- [Sch95] R. Schoof. “Counting points on elliptic curves over finite fields”. In: *J. Théor. Nombres Bordeaux* 7.1 (1995), pp. 219–254.
- [Sho94] P. W. Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In: *Proceedings 35th annual symposium on foundations of computer science*. Ieee. 1994, pp. 124–134.

- [Sut13] A. Sutherland. “On the evaluation of modular polynomials”. In: *The Open Book Series* 1.1 (2013), pp. 531–555.

INRIA BORDEAUX-SUD-OUEST, 200 AVENUE DE LA VIEILLE TOUR, 33405 TALENCE CEDEX FRANCE

Email address: damien.robert@inria.fr

URL: <http://www.normalesup.org/~robert/>

INSTITUT DE MATHÉMATIQUES DE BORDEAUX, 351 COURS DE LA LIBÉRATION, 33405 TALENCE CEDEX FRANCE