Some applications of higher dimensional isogenies to elliptic curves
Preliminary version

DAMIEN ROBERT

Abstract. We give two applications of the "embedding Lemma". The first one is a deterministic polynomial time (in $\log q$) algorithm to compute the endomorphism ring $\mathrm{End}(E)$ of an ordinary elliptic curve $E/\mathbb{F}_q$, provided we are given the factorisation of Δ_π. In particular, this 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 $\widetilde{O}(nm \log^{O(1)} p)$. We deduce a point counting algorithm of complexity $\widetilde{O}(n^2 \log^{O(1)} p)$. In particular the complexity is polynomial in $\log p$, by contrast of what is usually expected of a p-adic cohomology computation. This algorithm generalizes to ordinary abelian varieties.

1. Introduction

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.

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-isogeny on E^u where $u = 1, 2$ or 4 depending on whether m is a sum of 1, 2 or 4 squares, see [Rob22a]. The same ideas hold for an abelian variety, which yield the following embedding lemma: for any $m > 0$, an N-isogeny $f : A \to B$ in dimension g of principally polarised abelian varieties can always be efficiently embedded into an $N + m$-isogeny F in dimension $8g$ (and sometimes $4g$ or $2g$). Indeed, if u is as above, we can build m-isogenies α_A, α_B on A^u, B^u such that $\alpha_B f = f \alpha_A$, and take F to be endomorphism of $A^u \times B^u$ given by $F = \begin{pmatrix} \alpha_A & -\tilde{f} \\ f & \widetilde{\alpha_B} \end{pmatrix}$. We remark that F embeds both f and its dual $\tilde{f}$. This has been applied to break SIDH in [CD22; MM22; Rob22a].

More generally, let us define the N-evaluation problem as follow: given an N-isogeny $f : A/k \to B/k$ 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 $A[N]$. 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

Date: December 13, 2022.

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 computations) 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 all case. Namely the embedding lemma gives us an N'-isogeny F that embeds f, so evaluating $f(Q)$ can be done by evaluating $F(Q)$. Furthermore, if N' is prime to N, $\ker F$ can be completely determined by the value of $f(A[N])$: $\ker F = \{(\alpha_A x, -fx), x \in A^u[N]\}$. A fun fact is that in this case we do not even need to compute DLPs to recover $\ker F$. So the weak (N, N')-interpolation problem can always be reduced to an N'-evaluation problem in higher dimension, provided that $N' > N$ is prime to N. In fact, by considering the contragredient isogeny of F, we only need $N'^2 > N$, see [Rob22a, § 6.4].

This 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].

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 $\mathrm{End}(A)$ (recall that for an ordinary abelian variety the endomorphism ring is invariant by a field extension, so $\mathrm{End}(A) = \mathrm{End}_{\mathbb{F}_q}(A) = \mathrm{End}_{\overline{\mathbb{F}}_q}(A)$). So any element $\alpha \in \mathrm{End}(A)$ 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 $\mathrm{End}^0(A) = \mathrm{End}(A) \otimes_{\mathbb{Z}} \mathbb{Q}$.

Note that since A is principally polarised, it contains $\mathbb{Z}[\pi, \overline{\pi}]$ where $\overline{\pi} = q/\pi$ (the Verschiebung) is the image of π by the Rosatti involution. This allows to write α as a polynomial in $\pi, \overline{\pi}$ where this time the denominator D divides $[O_K : \mathbb{Z}[\pi, \overline{\pi}]]$, 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. Indeed, 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). And by Weil's theorem, the height of χ_π is linear in $\log q$.

We will see how to apply these techniques to the computation of endomorphism rings and canonical lifts. This paper is just a preliminary version to give a brief leisurely description of the main algorithms, it will be followed by two technical papers giving more details and a finer complexity analysis.

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 3. I thank Jean-Marc Couveignes and Pierrick Gaudry for various discussion about other applications of canonical lifts than point counting, and Aurel Page for brainstorming sessions about trying to apply the same techniques as Section 4 to compute the crystalline cohomology of a general ordinary scheme.

2. Embedding an isogeny

For our complexity analysis, we need to briefly review the results of [Rob22b].

Given an N-isogeny: $f : E \to E'$ over $\mathbb{F}_q$, we try to find N' powersmooth (with powersmoothness bound B), such that $m = N' - N$ is a sum of $1, 2$ or 4 squares. This allows to embed f into a N'-isogeny in dimension $2u$.

To recover the kernel of F and decompose it as a product of $(\leq B)$-isogenies, we need to work with algebras of degree up to $O(B^4)$. We need to push up to $\log N'$ points at each step, and each $(\leq B)$-isogeny evaluation cost $O(B^{2g})$. Since there are up to $\log N'$ steps, the complexity of decomposing F is $O(B^4 B^{2u} \log^2 N')$ arithmetic operations. For subsequent isogeny evaluations, to evaluate $f(Q)$ if $Q \in E(\mathbb{F}_q)$, we work with algebra of degree up to $O(B^2)$, and follow $\log N'$ $(\leq B)$-isogenies, for a total cost of $O(B^2 B^{2u} \log N')$ arithmetic operations. In practice, we will take a bound $B = O(\log N)$ and try to find N' such that $\log N' = O(\log N)$, so the decomposition cost is $O(\log^{6+2u} B)$ arithmetic operations and further evaluations are in $O(\log^{3+2u} B)$ arithmetic operations.

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 4 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.

3. 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 $K = \mathbb{Q}(\sqrt{\Delta_\pi}) = \mathrm{End}^0(E)$, and the factorisation of the conductor $f_\pi = [O_K : \mathbb{Z}[\pi]]$. We can write $\pi = a + f_\pi\omega$ (where a will depends on the trace of π, so has height $O(\log q)$). We know that $\pi - a \in \mathrm{End}(E)$. To determine $\mathrm{End}(E)$ is equivalent to determining the index of $\mathrm{End}(E)$ in O_K or the index of $\mathbb{Z}[\pi]$ in $\mathrm{End}(E)$, and so is equivalent to determining the largest divisor f_E of f_π such that $\frac{\pi-a}{f_E} \in \mathrm{End}(E)$.

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 $\mathrm{End}(E)$? 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 3.1. This approach to endomorphism ring computations is used in [EL07; FL08] 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 [FM02]). These algorithms are exponential in the worst case. An heuristic subexponential algorithm is presented in [BS09], 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. 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 $\mathrm{Ker}\, F$ as $\mathrm{Ker}\, F = \{(\gamma P, -\alpha P) \mid P \in E^u[N']\}$. 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 $\mathrm{Ker}\, F$ is always isotropic in $E^{2u}[N']$.

So we first compute $E^{2u}/\mathrm{Ker}\, F$ 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 $F = \begin{pmatrix} f_1 & -\widehat{g_1} \\ f_2 & \widehat{g_2} \end{pmatrix}$ 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, 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).[1] 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 $N"^2 > 4\max(g^2N', \deg(\pi - a)) = 4g^2N'$ (since we take $N' > \deg((\pi - a)/g)$). In this case, $(\pi - a)/g$ is indeed an endomorphism, and the converse is immediate.

[1] To 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.

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 get using Section 2:

Theorem 3.2. *Given an ordinary elliptic curve $E/\mathbb{F}_q$ and the factorisation of the discriminant of the Frobenius π,* $\mathrm{End}(E)$ *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.

Remark 3.3. 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 3.4. The same framework should allow to compute the endomorphism ring of an ordinary abelian variety, provided that we can work with real multiplication isogenies (and embed them powersmoothly). We leave that for future work. It would also be very interesting to be able to move in the ℓ-isogeny volcano in time polynomial in $\log \ell$.

4. Point counting and canonical lifts

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 $\overline{\pi}_q$ on T_0E allows us to recover $\lambda \mod p$, hence the trace of π modulo p. Since $[q] = \overline{\pi}_q \circ \pi_q$[2], 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 Vélu'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 $\overline{\pi}_q$ and π_q on T_0E. We thus get a polynomial time algorithm to recover $\lambda \mod p$. Like above, it is more efficient to only embed π_p and $\overline{\pi}_p$ and recover λ via a norm, see [Rob21, § 6].

Using Section 2, this algorithm to recover $\lambda \mod p$ 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 $\chi_\pi \mod \ell_i$ via some DLP computations in $E[\ell_i]$,

[2]We can also write $\overline{\pi}_q = t - \pi_q$, 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.

then reconstruct $\chi_\pi \mod \prod \ell_i$ 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 $\overline{\pi}_q$ (or π_p and $\overline{\pi}_p$).

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 $\widetilde{O}(nmp^2)$. In [MR22], we explained how to compute the action via lifting the kernel of the Verschiebung $\overline{\pi}_p$ instead; since it is of degree p this allows co compute canonical lift in time $\widetilde{O}(nmp)$. (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 boostrap to precision $m = 2$: we use the fact that the étale p-torsion only lifts to $\widetilde{E}$ if $\widetilde{E} = \hat{E}$ modulo p^2). Here we are going to use F instead, this way we can recover the action of π_p rather than $\overline{\pi}_p$ so there is no loss of precision, but more importantly F (and its lift) can be evaluated in time polynomial in $\log p$.

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 $\widetilde{E}$ of E such that π_p lifts to $\hat{\pi}_p : \widetilde{E} \to \sigma(\widetilde{E})$. We thus look for $\widetilde{E}$ such that the unique lift of F (as an isogeny) to $\widetilde{E} \times \sigma(\widetilde{E})$ 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 $\widetilde{E}$ 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 fix an arbitrary lift $\widetilde{E}'_1$ of E and another $\widetilde{E}'_2$ of $\sigma(E)$. We lift F to compute its action on $\widetilde{E}'_1 \times \widetilde{E}'_2$. We can then deform $\widetilde{E}'_1$ to another lift $\widetilde{E}''_1$, compute the action of F again, and then deform $\widetilde{E}'_2$ to $\widetilde{E}''_2$ 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 $j(\widetilde{E}_1) = j(\widetilde{E}'_1) + \varepsilon_1 p$, $j(\widetilde{E}_2) = j(\widetilde{E}'_2) + \varepsilon_2 p$, we can compute $J(\widetilde{E}_1 \times \widetilde{E}_2 / \operatorname{Ker} \tilde{F}) = J(\widetilde{E}'_1 \times \widetilde{E}'_2 / \operatorname{Ker} \tilde{F}) + U\epsilon_1 + V\epsilon_2$, 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 $\widetilde{E}$ is a lift of E, the Frobenius $\pi_p : E \to \sigma(E)$ lifts uniquely to $\widetilde{E} \to \widetilde{E}_2$. However in general the Verschiebung $\sigma(E) \to E$ does not lift to an arbitrary lift $\widetilde{E}_2$, 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 $(E, \overline{\pi}_p)$. In fact, by looking at the Serre-Tate formal moduli, it is classical that if $\widetilde{E} = \hat{E}$ to precision m, and $\widetilde{\pi}_p : \widetilde{E} \to \widetilde{E}_2$ is a lift of π_p, then $\widetilde{E}_2 = \sigma(\hat{E})$ to precision $m + 1$. Hence the Verschiebung $\overline{\pi}_p$ can be lifted to $\widetilde{E}_2$ if $\widetilde{E}_2 = \hat{E}$ 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 $\widetilde{E}_1 \to \widetilde{E}_2$. 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 $\widetilde{E}_1$ and $\widetilde{E}_2$ of E and $\sigma(E)$, the lift of π_p to $\widetilde{E}_1$ has codomain another elliptic curve $\widetilde{E}_{2,can}$, and so the codomain of the lift $\tilde{F}$ of F will not be a product abelian surface unless $\widetilde{E}_2 = \widetilde{E}_{2,can}$. 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 $J(\widetilde{E}_1 \times \widetilde{E}_2/\operatorname{Ker}\tilde{F})$ above we get a linear equation between ϵ_1 and ϵ_2 giving the locus where $\widetilde{E}_2 = \widetilde{E}_{2,can}$. On this locus, the Verschiebung lifts from $\widetilde{E}_2$ to $\widetilde{E}_1$ by the above discussion, hence F lifts as a matrix. Alternatively, we could plug the equation $J(\widetilde{E}_1 \times \widetilde{E}_2/\operatorname{Ker}\tilde{F}) = J(\widetilde{E}_1 \times \widetilde{E}_2)$.

The canonical lift $\hat{E}$ at precision 2 can then be recovered by plugging the further equation $j(\widetilde{E}_{2,can}) = \sigma(j(\widetilde{E}_1))$. 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 $j(\widetilde{E}_{2,can}) = \sigma(j(\widetilde{E}_1))$ and $J(\widetilde{E}_1 \times \widetilde{E}_2/\operatorname{Ker}\tilde{F}) = J(\widetilde{E}_1 \times \widetilde{E}_2)$ 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 $\hat{\tilde{\pi}}_p$ on $T_0\sigma\hat{E}$ to precision m. By Section 2, the dominating cost is the initial decomposition of 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. In summary:

Theorem 4.1. *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 $\widetilde{O}(nm\log^{4+2u} p + n\log^{7+2u} p)$, and the cardinal of E in time $\widetilde{O}(n^2\log^{4+2u} p + n\log^{7+2u} p)$.*

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 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.

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 $\widetilde{O}(\log^4 q) = \widetilde{O}(n^4\log^4 p)$.
- Rigid (Monsky-Washnitzer) cohomology: Kedlaya's algorithm [Ked01] is in $\widetilde{O}(n^3 p)$ and Harvey's variant [Har07] in $\widetilde{O}(n^{3.5}p^{1/2} + n^5\log p)$.
- Crystalline cohomology: Satoh's algorithm [Sato0] (after improvements by Harley) is in $\widetilde{O}(n^2p^2)$, and it has been improved to $\widetilde{O}(n^2 p)$ in [MR22]. The (proven version of the) current algorithm is in $\widetilde{O}(n^2\log^{15} p)$ and the heuristic version in $\widetilde{O}(n^2\log^{11} p)$.

Remark 4.2. Over an ordinary abelian variety, the same method allows to recover the tangent matrix of $\hat{\tilde{\pi}}_p$ and $\hat{\pi}_p$ to precision m in time $O(nm\log^{O(1)} p)$ (where the $O(1)$ hides a dependency at least linear in g).

Remark 4.3. 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 $\operatorname{End}(E)$, 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 [CH02, 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.

References

[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).

[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.

[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.

[EL07] K. Eisentrager and K. Lauter. "A CRT algorithm for constructing genus 2 curves over finite fields". In: *AGCT-11* (2007).

[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.

[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.

[Ked01] K. Kedlaya. "Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology". 2001. arXiv: math/0105031.

[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.

[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.

[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.

[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.

[Satoo] 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.

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 liberation, 33405 Talence cedex FRANCE