

COMPUTING MODULAR POLYNOMIALS BY DEFORMATION

SABRINA KUNZWEILER, DAMIEN ROBERT

ABSTRACT. We present an unconditional CRT algorithm to compute the modular polynomial $\Phi_\ell(X, Y)$ in quasi-linear time. The main ingredients of our algorithm are: the embedding of ℓ -isogenies in smooth-degree isogenies in higher dimension, and the computation of m -th order deformations of isogenies. We provide a proof-of-concept implementation of a heuristic version of the algorithm demonstrating the practicality of our approach.

1. INTRODUCTION

The modular polynomial $\Phi_\ell(X, Y) \in \mathbb{Z}[X, Y]$ parametrizes ℓ -isogenies between elliptic curves. It is a fundamental object for building isogenies and exploring isogeny graphs, and is used in many applications, notably in the SEA point counting algorithm. We will assume ℓ to be prime for simplicity, the extension to the general case is not hard. It is well known that Φ_ℓ has degree $\ell + 1$ in each variable with height $O(\ell \log \ell)$, so the polynomial has size $O(\ell^3 \log \ell)$. More explicitly,

$$(1) \quad \log |a_{i,j}| \leq 6\ell \log(\ell) + 16\ell + \min(2\ell, 14\sqrt{\ell} \log \ell),$$

where $\Phi_\ell = \sum_{i,j} a_{i,j} X^i Y^j$, [6].

1.1. Algorithms for computing modular polynomials. The first quasi-linear algorithm to compute Φ_ℓ is due to Enge in [14] and uses analytic methods. Namely, it uses an evaluation-interpolation approach: fix $\ell + 1$ nice period matrices τ_i in the complex upper half plane, and evaluate $\Phi_\ell(j(\tau_i), Y) = \prod_{j=1}^{\ell+1} (Y - j(\tau_{i,j})) = Y^{\ell+1} + \sum_{j=0}^{\ell} c_j(j(\tau_i)) Y^j$, where the $\tau_{i,j}$ are period matrices of the $\ell + 1$ elliptic curves isogenous to E_{τ_i} . Each coefficient c_j , seen as polynomial in the variable X , can then be recovered by interpolation.

Since the modular polynomial has large height, the analytic method needs to work with large precision. To achieve quasi-linearity, the j -invariants thus need to be evaluated in quasi-linear time in the precision. Due to the use of large precision floating point arithmetic, Enge's algorithm is only heuristic: the assumption is that the loss of precision in the floating point arithmetic does not explode. This has been made rigorous in recent work by Kieffer [22, 20]. In [20], Kieffer focuses on the rigorous evaluation of dimension 2 theta functions; but the same arguments work (and are easier) to compute the modular polynomials in dimension 1.

In practice, the analytic algorithm works well; due to the quasi-linear complexity, the blocking factor to compute large modular polynomials are their size, hence the available memory, rather than the computing power. In [5], Bröker, Lauter and Sutherland introduce a quasi-linear algorithm based on the Chinese remainder theorem (CRT) to compute Φ_ℓ . Besides being purely algebraic, a great advantage is that thanks to the explicit CRT algorithm [2], the reduced modular polynomial $\Phi_\ell(X, Y) \bmod m$ (m a large integer) can be computed by using a memory of only $O(\ell^2 \log(m\ell))$ rather than $O(\ell^3 \log \ell)$.

In the CRT algorithm, the modular polynomial $\Phi_\ell(X, Y) \bmod p_i$ is computed for many $O(\ell)$ small primes p_i (with $\log p_i \approx \log \ell$). The Chinese remainder theorem is then used to reconstruct $\Phi_\ell(X, Y) \bmod \prod p_i$. To obtain a quasi-linear algorithm, $\Phi_\ell(X, Y) \bmod p_i$ needs to be computed in $O(\ell^2 \log^{O(1)} p_i)$. As in the analytic method, this is done by an evaluation-interpolation approach, but in a careful way to achieve the desired complexity.

Fix a small prime $p = p_i$. The naive evaluation-interpolation approach would be to select $\ell + 1$ elliptic curves E_i/\mathbb{F}_p , and compute all $\ell + 1$ -isogenous elliptic curves $E_{i,j}/\mathbb{F}_p$. But given an elliptic curve E_i/\mathbb{F}_p , it is already not obvious how to compute the $\ell + 1$ -isogenies efficiently: factoring the division polynomial (of degree ℓ^2) is too expensive. A solution is to require that $E_i[\ell]$ has rational torsion; we can then efficiently sample a basis of it, generate all kernels, and apply Vélu’s formula [38] to compute the isogenous curves. Vélu’s formula cost $O(\ell)$ arithmetic operations, so computing $\Phi_\ell(j(E_i), Y)$ in this way already costs $O(\ell^2)$ arithmetic operations. Applying the same algorithm to each E_i would therefore be too expensive (even using the recent square-root Vélu algorithm [3] instead of Vélu’s formula), since we would need to compute $O(\ell^2)$ ℓ -isogenies in total.

Instead, the solution proposed in [5] is to use isogeny volcanoes to find suitable ℓ -isogenous curves faster than in $O(\ell)$. They use the Hilbert class polynomial of a carefully crafted imaginary quadratic order to ensure that the isogeny volcano has the desired property. Their algorithm computes Φ_ℓ in $O(\ell^3 \log^3 \ell \log \log \ell)$ which is slightly faster than the analytic method, and as mentioned above requires less memory to compute $\Phi_\ell \bmod m$. However, the complexity proof relies on the generalised Riemann hypothesis (this is required to be sure that there exist small generators of the class group), hence is not unconditional.

1.2. New algorithmic tools. The field of isogeny based cryptography radically changed following the break of SIDH [8, 27, 33]. It was soon clear that the tools (Kani’s lemma, Zahrin’s trick) used for the break could also be used to achieve a very efficient representation of isogenies [31] by embedding them into higher dimensional isogenies, which in turn lead to new isogeny based cryptosystems (for instance SQISignHD [10] and FESTA [1]).

These new algorithmic tools are not restricted to isogeny based cryptography; in [32] the second author uses them to outline new algorithms to compute the endomorphism ring and canonical lifts of ordinary abelian varieties in polynomial time. A sketch of a quasi-linear algorithm to compute modular polynomials is described as well.

We summarize the efficient representation of isogenies of [31], which we will call an *HD representation* (for higher dimensional representation), as follows:

Theorem 1.1. *Let $f : E_1 \rightarrow E_2$ be an ℓ -isogeny between elliptic curves over a finite field \mathbb{F}_q . There exists an HD representation F of f which can be used to evaluate $f(P)$ on $P \in E(\mathbb{F}_q)$ in time $\tilde{O}(\log^\mathfrak{D} \ell)$ arithmetic operations in \mathbb{F}_q , for some constant \mathfrak{D} (not depending on q or ℓ)¹.*

If E_1 has full rational 2^n -torsion with $2^n > \ell$, there is a special HD-representation where the evaluation is in $O(\log \ell)$ arithmetic operations.

The key point of the SIDH attacks is that an elliptic curve isogeny f can be embedded into a higher dimensional N -isogeny as long as we know the evaluation of f on a basis of $E_1[N]$, $N > \ell$. The first representation in Theorem 1.1 follows by taking N powersmooth, and the second (which more generally could be used as long as the 2^n -torsion lies in a small extension of \mathbb{F}_q) by taking $N = 2^n$. Of course, we also have a special HD representation if the $3^n > \ell$ -torsion is rational and so on. We refer to [31, 32] for more details.

1.3. Our contributions. The goal of the current article is to give a fully fledged and rigorous analysis of the algorithm to compute modular polynomials sketched in [32, Section 6]. Notably, we prove in Section 6 the following theorem, which matches the complexity of [5].

Theorem 1.2. *There exists an unconditional CRT algorithm which computes the modular polynomial $\Phi_\ell(X, Y)$ in quasi-linear time $O(\ell^3 \log^3 \ell \log \log \ell)$.*

¹In practice we have $\mathfrak{D} = 12$ by [32, Proposition 2.9].

More precisely, we present a rigorous algorithm, along with a faster but heuristic version, which is the one we implemented. We remark that this can also be used to compute the reduction modulo m of Φ_ℓ in quasi-linear space $O(\ell^2 \log(m\ell))$. To simplify the exposition, we will always assume that $\ell > 2$ and the base characteristic $p > 2$, too.

We briefly explain the main ideas behind the algorithms. The key tool is the following theorem proved in Section 4, and stated here in an informal way:

Theorem 1.3. *Given an HD representation of an ℓ -isogeny $f : E_1 \rightarrow E_2$ over a finite field $\mathbb{F}_q = \mathbb{F}_{p^d}$, with ℓ prime to the characteristic p , and given an m -th order deformation of E_1 to an elliptic curve \mathcal{E}_1/R with $R = \mathbb{F}_q[\epsilon]/(\epsilon^{m+1})$, we can efficiently compute the deformation of f to \mathcal{E}_1 , that is $\tilde{f} : \mathcal{E}_1 \rightarrow \mathcal{E}_2$ in $\tilde{O}(\log^2 \ell)$ arithmetic operations in R .*

If we have a special HD representation of f , then we can compute \tilde{f} in $O(\log \ell)$ arithmetic operations in R .

In other words: once we have an efficient representation of an isogeny, we can efficiently deform it. Deformation techniques are not new in computer algebra, see for instance [34]. Here we exploit them in order to recover the modular polynomial Φ_ℓ modulo a prime p :

Corollary 1.4. *Let E_0/\mathbb{F}_q be an elliptic curve which has all its $\ell + 1$ ℓ -isogenies rational. Given an HD representation of these isogenies, we can compute $\Phi_\ell \bmod p$ in time $\tilde{O}(\ell^2 \log q)$. If we have a special HD representation of all isogenies, the cost is $O(\ell^2 \log^2 \ell)$ operations over \mathbb{F}_q .*

Proof. The idea is to use the HD representation of the isogenies to deform them to $\mathbb{F}_q[\epsilon]/(\epsilon^{m+1})$. When the precision m is high enough, we get enough information to reconstruct the full modular polynomial. We refer to Section 5 for more details. \square

We remark that if $q = p^d$ with d small, then Corollary 1.4 is quasi-linear in its output size $O(\ell^2 \log p)$. To apply Corollary 1.4, we need to find an elliptic curve over a small extension of \mathbb{F}_p on which we can efficiently compute the $\ell + 1$ ℓ -isogenies.

The easiest case is when we can find an elliptic curve E_0/\mathbb{F}_q with full rational ℓ -torsion. In this case we compute a basis, compute all kernels, and apply Vélú's formula. Given a basis of $E[\ell]$, computing all the ℓ -isogenies costs $O(\ell^2)$ arithmetic operations.

In particular, if $p \equiv 3 \pmod{4}$ and $\ell \mid p^2 - 1$, then we can take $E'_0 : y^2 = x^3 + x$, this is a supersingular curve over \mathbb{F}_p and $E'_0(\mathbb{F}_{p^4}) = (\mathbb{Z}/(p^2 - 1)\mathbb{Z})^2$, so we can efficiently sample a basis of the ℓ -torsion in $O(\log p)$ operations in \mathbb{F}_{p^4} (in practice we can work over E'_0/\mathbb{F}_{p^2} or its quadratic twist according to whether $\ell \mid p - 1$ or $\ell \mid p + 1$). For technical reasons (to kill the extra automorphisms of $y^2 = x^3 + x$), we will rather work with the 2-isogenous curve $E_0 : y^2 = x^3 + 6x^2 + x$.

If we add the congruence condition that $2^n \mid p^2 - 1$ for some $2^n > \ell$, so that the 2^n -torsion of E_0 is rational over \mathbb{F}_{p^4} , then we can even use a special HD representation of the ℓ -isogenies.

Proposition 1.5. *For $p \in \mathcal{P}_\ell^* = \{p > 11 \text{ prime} : 2^n \cdot \ell \mid p + 1, \text{ where } n = \lceil \log_2(\ell) \rceil\}$, we can compute $\Phi_\ell \bmod p$ in $O(\log p + \ell^2)$ arithmetic \mathbb{F}_p -operations.*

We prove this proposition in Section 5.4. We call a prime $p \in \mathcal{P}_\ell^*$ a suitable CRT prime. By the Dirichlet Arithmetic Progression Theorem, we know that there are sufficiently many such primes, with the appropriate density (see Lemma 5.3). Using these suitable CRT primes, we obtain Theorem 1.2.

We remark that the idea to use supersingular curves for modular polynomial computation is not new, and in [25] Leroux also gives a heuristic algorithm relying on supersingular isogenies to compute modular polynomials modulo a prime.

In practice, in the implementation, we use a subset $\mathcal{P}_\ell \subset \mathcal{P}_\ell^*$, defined in Section 5.1, where we replace the condition $n = \lceil \log_2(\ell) \rceil$ (which is enough to ensure that $2^n - \ell$ can be written as a

sum of four squares) by a condition $2^n - c_\ell \ell = a^2 + 4b^2$. When $p \in \mathcal{P}_\ell$, this stronger condition ensures that we can find a special HD representation in dimension 2 of our isogenies rather than in dimension 8. We prove this version of Proposition 1.5 for $p \in \mathcal{P}_\ell$ in Section 5.2. While this does not change the asymptotic complexity of the algorithm, it improves the constant and greatly simplifies the implementation. However, to achieve the quasi-linear CRT algorithm of Theorem 1.2 when using primes in \mathcal{P}_ℓ , we need to rely on a heuristic (Heuristic 5.1), which says that we can find a suitable n such that 2^n is not too large compared to ℓ . In this paper, we will give detailed algorithms for Theorems 1.2 and 1.3, Corollary 1.4, and Proposition 1.5 for the special case when $p \in \mathcal{P}_\ell$, because it is the one we implemented, and briefly explain how to generalize to the general cases.

1.4. Generalisations and perspectives. The techniques used in our algorithms, can also be applied to other settings. Here, we highlight some generalizations.

Computing Φ_ℓ modulo an arbitrary prime p . In this paper, we explain the computation for primes $p \in \mathcal{P}_\ell^*$. This may be extended to general primes p as follows. We also use supersingular curves, because all their ℓ -isogenies are rational over \mathbb{F}_{p^2} . We first need to sample a supersingular elliptic E_0/\mathbb{F}_{p^2} with explicit known endomorphism ring and endomorphism action. This can be seen as a precomputation depending only on p and not on ℓ . For instance, we can apply Bröker’s algorithm [6], which is in $O(\log^3 p)$ under GRH.

The explicit endomorphism ring action of E_0 allows us to compute the $\ell + 1$ ℓ -isogenies in polynomial time in $\log \ell + \log p$ by [39] under GRH, or unconditionally using the more recent CLAPOTIS method [29, Remark 2.10]. Namely, we compute the ideals in $\text{End}(E_0)$ corresponding to the $\ell + 1$ ℓ -isogenies, and then we use CLAPOTIS to compute an HD representation of the ℓ -isogenies associated to these ideals in $\tilde{O}((\log \ell + \log p)^\mathfrak{D})$ arithmetic operations by ideal.

This initialisation is only interesting when $\log p$ is sufficiently small compared to ℓ ; it is thus not suitable in applications like point counting where we usually have $\log p \approx \ell$. (Of course conversely if $\log p$ is sufficiently large compared to ℓ , it is faster to just compute Φ_ℓ directly.) An alternative initialisation approach is described in [32, 6.2] with better complexity in $\log p$, and another heuristic approach is described in [25].

From the discussion above, we get:

Corollary 1.6. *Let p be a fixed prime. Then there is a quasi-linear algorithm to compute $\Phi_\ell \bmod p$ in $\tilde{O}_p(\ell^2)$, where the notation \tilde{O}_p means that the constants depend on p .*

p -adic lifting. One can extend Theorem 1.3 to more general Artinian rings with residue field \mathbb{F}_q , notably rings of the form $R = \mathbb{Z}_q[\epsilon]/(\epsilon^{m_1}, p^{m_2})$. This allows us to combine our horizontal deformation with a p -adic lifting, and give a p -adic lifting version of Corollary 1.4 which results in an alternative way to compute the modular polynomial Φ_ℓ in quasi-linear time. However, the CRT algorithm is better because it can be used to compute $\Phi_\ell \bmod m$ by the explicit CRT.

Modular polynomials in higher dimension. In this paper, for simplicity we confine ourselves to computing deformations of isogenies of elliptic curves. Using theta models for abelian varieties, our methods to compute these deformations efficiently (Theorem 1.3) extend naturally to higher dimension. In subsequent work, we will explain how this paves the way to computing higher dimensional Siegel modular polynomials in quasi-linear time. Such a complexity was conjectured in [30, Conjecture 5.3.14]. The reason we focus on the dimension 1 case in this current paper is by lack of space, for the simplicity of the exposition², and also because we have only implemented this case so far.

²The main difficulty in higher dimension is that we now need to reconstruct multivariate rational functions from power series rather than simply polynomials; and also take into account the algebraic equations between the modular invariants.

For modular polynomials in higher dimension, there is currently no known quasi-linear analytic method. An important step in this direction is achieved in a recent work by Elkies and Kieffer, who developed a quasi-linear algorithm to evaluate theta functions in quasi-linear time in the precision in all dimensions. However, to compute the modular polynomial efficiently, one would also need a quasi-linear algorithm for computing period matrices from theta constants, because the evaluation-interpolation approach requires to select suitable points of interpolation to obtain a fast multivariate interpolation algorithm. Such an algorithm is only known in dimensions 1 and 2. The dimension 2 case follows from heuristic work of Dupont [13], and was then rigorously proven by Kieffer in [21]. Kieffer thus describes in [20] a rigorous algorithm to evaluate modular polynomials in dimension 2, which combined with [13, 28] gives a rigorous analytic algorithm in dimension 2 to compute the full modular polynomial.

Likewise, the CRT approach of [5] crucially uses the volcano structure; but isogeny graphs of principally polarised abelian varieties over a finite field are more complex, so it is not obvious how to extend this approach to higher dimensions.

By contrast, as we have mentioned, our deformation method generalises in a natural way to higher dimensions. We expect that similar deformation techniques could also be used to compute more general modular correspondences or actions of Hecke operators on modular forms.

1.5. Outline. In Section 3, we provide the theoretical background on deformations of abelian varieties and isogenies. The results of that section are made explicit in Section 4, where we describe different methods for computing with deformations. In Section 5, we present an explicit algorithm to compute modular polynomials over finite fields. Building on this, we describe a CRT algorithm for computing the modular polynomial in Section 6.

Some background on elliptic curves, abelian varieties and isogenies is provided in Appendix 2.

2. ELLIPTIC CURVES AND ABELIAN VARIETIES OVER FINITE FIELDS

2.1. Elliptic curves and isogenies. An *elliptic curve* over a field k is a smooth projective curve defined over k of genus 1 with a k -rational point. More concretely, elliptic curves can be described as planar curves by an (affine) equation of the form

$$E : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6 \quad \text{with } a_1, a_2, a_3, a_4, a_6 \in k.$$

Such an equation is called Weierstrass equation, and the distinguished k -rational point is the point at infinity $\mathcal{O} = [0 : 1 : 0]$.

For any field extension L/k , the set of L -rational points of an elliptic curve, $E(L)$, forms a group with identity element \mathcal{O} . For an integer N , we denote by $E[N] \subset E(\bar{k})$ the set of N -torsion points of E . Now, let k be a finite field with characteristic $\text{char}(k) = p$. Then

$$E[N] \cong \mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z} \quad \text{if } p \nmid N.$$

For the structure of $E[p]$, there are two possibilities. We say that E is *ordinary* if $E[p] \cong \mathbb{Z}/p\mathbb{Z}$, and E is called *supersingular* if $E[p] = \{\mathcal{O}\}$. We remark that for any supersingular curve E over a finite field k with $\text{char}(k) = p$, the trace of Frobenius is $0 \pmod{p}$, and moreover E is necessarily defined over \mathbb{F}_{p^2} . In particular, if E is also defined over \mathbb{F}_p , this implies

$$E(\mathbb{F}_{p^2}) \cong \mathbb{Z}/(p+1)\mathbb{Z} \times \mathbb{Z}/(p+1)\mathbb{Z}.$$

An *isogeny* $f : E \rightarrow E'$ is a non-constant morphism of elliptic curves satisfying $f(\mathcal{O}) = \mathcal{O}'$. This also defines a group homomorphism and in particular the kernel of an isogeny is a finite subgroup of $E(\bar{k})$ (we restrict to étale isogenies). Conversely for any finite subgroup $G \subset E(\bar{k})$, one can construct an isogeny $f : E \rightarrow E'$ with $\ker(f) = G$, where the codomain E' is up to isomorphism uniquely determined by the kernel group G .

We say that an isogeny f is *cyclic* if $\ker(f) \cong \mathbb{Z}/N\mathbb{Z}$ for some integer N . In that case, f is called an N -isogeny. If N is coprime to the characteristic p , then an N -isogeny is necessarily separable and has degree N . We note that ℓ -isogenies of prime degree ℓ can be computed in time $\tilde{O}(\sqrt{\ell})$ using the square-root Vélu algorithm from [3]. For small values of ℓ it is however faster to use the standard formulae from Vélu with complexity $O(\ell)$. Now N -isogenies of composite degree $N = \prod_i \ell_i$ can be computed in time $\tilde{O}(\sum_i \sqrt{\ell_i})$ or $\tilde{O}(\sum_i \ell_i)$ as the composition of prime degree isogenies. In particular, an N -isogeny with $N = \ell^n$ for some small prime ℓ can be computed in time $O(n \log(n))$ using the optimal strategies developed in [12]. We remark that the complexities given here, all count the number of operations over the field of definition of the kernel G .

2.2. The j -invariant and modular polynomials. To each elliptic curve, one can associate its j -invariant which can be computed from the Weierstrass equation defining the curve. However, the j -invariant does not depend on the equation, but it is an invariant of the isomorphism class of the elliptic curve.

The modular polynomial $\Phi_\ell(X, Y)$ parametrises pairs of j -invariants that are connected by an ℓ -isogeny. That is for two elliptic curves E, E' defined over some field k with characteristic $p \nmid \ell$, we have that $\Phi_\ell(j(E), j(E')) = 0$ if and only if there exists a cyclic ℓ -isogeny $f : E \rightarrow E'$. The polynomial Φ_ℓ has integer coefficients and the logarithmic height of these coefficients is in $O(\ell \log(\ell))$. More explicitly, it is shown³ in [6] that

$$(2) \quad \log |a_{i,j}| \leq 6\ell \log(\ell) + 16\ell + \min(2\ell, 14\sqrt{\ell} \log \ell),$$

where $\Phi_\ell = \sum_{i,j} a_{i,j} X^i Y^j$. Moreover Φ_ℓ is symmetric and has degree $\psi(\ell)$ in each variable, where ψ is the Dedekind Psi Function. In particular $\deg_X(\Phi_\ell) = \deg_Y(\Phi_\ell) = \ell + 1$ if ℓ is a prime.

2.3. Moduli space and invariants in dimension 2. Elliptic curves are abelian varieties of dimension 1 equipped with a principal polarisation. In a nutshell, an *abelian variety* A is a projective algebraic variety that is also an abelian group. And a *principal polarisation* is a symmetric positive isomorphism from A to its dual.⁴ We denote by \mathbf{A}_g the moduli space of principally polarised g -dimensional abelian varieties. This space has dimension $\frac{g(g+1)}{2}$.

Apart from elliptic curves, we are also mainly interested in the case $g = 2$ in this work. Elements in \mathbf{A}_2 are called *principally polarised abelian surfaces*. Any indecomposable element of \mathbf{A}_2 is the Jacobian of a genus-2 curve. Recall that a curve C of genus 2 is hyperelliptic, and if $\text{char}(k) \neq 2$ it can be described by an affine equation of the form

$$C : y^2 = f(x) = c \cdot \prod_{i=1}^6 (x - \alpha_i) \quad \text{with } \alpha_i \in \bar{k}.$$

The remaining elements of \mathbf{A}_2 are products of elliptic curves equipped with the product polarisation. Points in \mathbf{A}_2 are typically described by (Siegel) modular forms. In our applications, we will only use the Siegel modular cusp form χ_{10} . For an indecomposable abelian surface, this is related to the discriminant of the polynomial f from the curve equation. More precisely, it holds that

$$\chi_{10}(\text{Jac}(C)) = -2^{-12} \cdot \text{disc}(f), \quad \chi_{10}(E \times E') = 0,$$

where C is a genus-2 curve, and E, E' are elliptic curves.

³See [4] for refined bounds.

⁴Not all abelian varieties admit a principal polarisation.

2.4. Abelian surfaces and isogenies. For a principally polarised abelian surface defined over a field k with $\text{char}(k) = p$, the N -torsion is a free $\mathbb{Z}/N\mathbb{Z}$ -module, more precisely

$$A[N] \cong (\mathbb{Z}/N\mathbb{Z})^4 \quad \text{if } p \nmid N.$$

There exists an alternating, bilinear pairing $e_N : A[N] \times A[N] \rightarrow \mu_N$ called the *Weil pairing*, where μ_N denotes the group of N -th roots of unity.

A subgroup $G \subset A[N]$ is called *maximal N -isotropic* if the Weil pairing restricts trivially onto G , and there is no isotropic subgroup $H \subset A[N]$ properly containing G . Let $G \subset A[N]$ be a maximal N -isotropic subgroup, then up to isomorphism there exists a unique principally polarised abelian surface A' together with an isogeny $F : A \rightarrow A'$ with kernel $\ker(F) = G$. For kernel groups of rank 2, i.e. $G \cong \mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z}$, we say that F is an (N, N) -isogeny.

In this paper, we are mostly interested in the case $N = 2$ or $N = 2^n$. We will treat the corresponding isogeny algorithm as a black box. That is, we assume there exists an algorithm `2.2.isogeny` which on input $(A, \langle P_1, P_2 \rangle)$ with A an abelian surface and $G = \langle P_1, P_2 \rangle$ a maximal 2-isotropic group, returns the codomain curve A' .⁵ We only briefly describe the structure of 2-torsion points, since this will be important for the analysis of our algorithms. On an indecomposable abelian surface $\text{Jac}(C)$ with

$$C : y^2 = f(x) = c \cdot \prod_{i=1}^6 (x - \alpha_i) \quad \text{with } \alpha_i \in \bar{k},$$

there is a bijection between the non-trivial 2-torsion points of $\text{Jac}(C)$ and pairs $\{\alpha_i, \alpha_j\}$ with $i \neq j$. To be a bit more precise, the pair $\{\alpha_i, \alpha_j\}$ corresponds to the equivalence class of the divisor $(\alpha_i, 0) + (\alpha_j, 0) - \infty_- - \infty_+$ on C . We remark that α_i, α_j need not be k -rational, even when the corresponding 2-torsion point is rational. In a similar way, we present 2-torsion points on a product of elliptic curves $P \in E \times E'$ by a pair $\{\alpha_i, \alpha_j\}$, where in this case $(\alpha_i, \mathcal{O}') \in E[2] \times \mathcal{O}'$ and $(\mathcal{O}, \alpha_j) \in \mathcal{O} \times E'[2]$.

2.5. Kani's Lemma. Let E, E_a, E_b, E_{ab} be elliptic curves and f, g, f', g' isogenies of degree $d_a = \deg(f) = \deg(f')$ and $d_b = \deg(g) = \deg(g')$ that fit into the following commutative diagram.

$$(3) \quad \begin{array}{ccc} E & \xrightarrow{f} & E_a \\ \downarrow g & & \downarrow g' \\ E_b & \xrightarrow{f'} & E_{ab} \end{array}$$

In this case, we say that $g \circ f' = f \circ g'$ is a (d_a, d_b) -isogeny diamond. Any isogeny diamond gives rise to a product isogeny in dimension 2 as explained in the following lemma.

Lemma 2.1 (Kani's Lemma [19]). *Let $g \circ f' = f \circ g'$ be a (d_a, d_b) -isogeny diamond and denote $N = d_a + d_b$. Then $F = \begin{pmatrix} f & \hat{g}' \\ -g & \hat{f}' \end{pmatrix}$ is an N -isogeny $F : E \times E_{ab} \rightarrow E_a \times E_b$. If $\gcd(d_a, d_b) = 1$, then*

$$\ker(F) = \langle (-\hat{g}(P_N), f'(P_N)), (-\hat{g}(Q_N), f'(Q_N)) \rangle,$$

where $E_a[N] = \langle P_N, Q_N \rangle$.

Note that Kani's Lemma can be generalised to abelian varieties of arbitrary dimension, see [33], and is the key to the HD representation of isogenies defined in [31].

⁵In practice, our implementation is based on [23].

3. DEFORMATIONS AND ISOGENIES

In this section, we provide an overview on the topic of deformations of principally polarised abelian varieties, and in particular we discuss deformations of isogenies. For more detailed explanations, the reader is referred to [15, 35]. Throughout, k denotes a field with characteristic $\text{char}(k) = p > 2$, and R is an Artinian ring of the form $R = k[\epsilon]/(\epsilon^{m+1})$ for some integer $m \geq 0$.⁶

3.1. Deformations of (principally polarised) abelian varieties. Recall that an abelian variety is a group scheme over a field k which is also a proper, geometrically integral variety over k . This fits into the following more general framework of abelian schemes.

Definition 3.1. An *abelian scheme* over a scheme S is a group scheme $\mathcal{A} \rightarrow S$ which is smooth, proper and has geometrically connected fibres. If $S = \text{Spec}(k)$, then we call \mathcal{A} an abelian variety.

In our applications, we always have $S = \text{Spec}(R)$ with $R = k[\epsilon]/(\epsilon^{m+1})$ as above, but everything extends to a general local Artinian ring R with maximal ideal m_R and residue field $R/m_R = k$. Note that in this setting, along with the map $\text{Spec}(k) \rightarrow \text{Spec}(R)$ induced by the projection $R \rightarrow R/(\epsilon)$, there is also a map $\text{Spec}(R) \rightarrow \text{Spec}(k)$ induced by the canonical inclusion $k \hookrightarrow R$.

Example 3.2. Let $\mathcal{E} \rightarrow \text{Spec}(R)$ be an abelian scheme of dimension 1, in other words an elliptic curve over R . As is the case for elliptic curves over fields, we can make Definition 3.1 more explicit in this situation. Since R is local, it can be represented as a subscheme of \mathbb{P}_R^2 defined by an equation of the form

$$\mathcal{E} : Y^2Z + a_1XYZ + a_3YZ^2 = X^3 + a_2X^2Z + a_4XZ^2 + a_6Z^3$$

with $a_1, a_2, a_3, a_4, a_6 \in R$. Note that $\text{Spec}(R) = \{(\epsilon)\}$. Therefore, the scheme \mathcal{E} has only one fibre $E = \mathcal{E}_{(\epsilon)}$. It is obtained by reduction modulo (ϵ) .

Note that points of \mathcal{E} are sections $s : \text{Spec}(R) \rightarrow \mathcal{E}$. More explicitly, we may also view a point $P \in \mathcal{E}(R)$ as usual by three projective coordinates, i.e. $P = (x : y : z) \in \mathcal{E} \subset \mathbb{P}_R^2$.

Definition 3.3. Let A be an abelian variety over k . We say that an abelian scheme $\mathcal{A} \rightarrow \text{Spec}(R)$ is an *m -th order deformation* of A if its special fibre $\mathcal{A} \times_R k \rightarrow \text{Spec}(k)$ is isomorphic to A .

For any abelian variety A , there is a trivial deformation given by $\mathcal{A} = A \times_k \text{Spec}(R)$. In general, the group structure of a deformation \mathcal{A} of A is not obvious. However, we can describe the étale part of the torsion group explicitly which is explained in the next remark.

Remark 3.4. Let N be an integer coprime to p , and consider a deformation $\mathcal{A} \rightarrow \text{Spec}(R)$ of an abelian variety A . Since $\mathcal{A}[N]/\text{Spec}(R)$ is étale and R is Henselian, we have a canonical isomorphism $\mathcal{A}[n](R) \cong A[n](k)$. This isomorphism will be made explicit in Algorithm 2 for elliptic curves. In fact, because R is Henselian, the functor which associates to a finite étale cover $\mathcal{X}/\text{Spec}(R)$ its special fibre $X/\text{Spec}(k)$ is an equivalence of categories with finite étale covers of $\text{Spec}(k)$. Thus $A[n]$ deforms uniquely to $\mathcal{A}[n]/\text{Spec}(R)$, in our setting this is simply given by the trivial deformation.

We remark that in our applications, we always work with points of $\mathcal{A}[N]$, and only compute separable isogenies. In these cases, we can use the standard addition laws and formulae for isogeny computation known for abelian varieties over fields.

Although we won't need it, we can make explicit the full structure of $\mathcal{A}(R)$ as follows. The kernel of the reduction map $\mathcal{A}(R) \rightarrow A(k)$ is given by $\mathcal{A}(p)^0(R)$, where $\mathcal{A}(p)^0$ is the connected component of the p -divisible group of \mathcal{A} . We have $\mathcal{A}(p)^0(R) = \tilde{\Gamma}(m_R)$ where $\tilde{\Gamma}$ is the formal Lie group associated to \mathcal{A} . Furthermore, since k is perfect, the connected-étale sequence splits over k , hence we have an exact sequence $0 \rightarrow \mathcal{A}(p)^0(R) = \tilde{\Gamma}(m_R) \rightarrow \mathcal{A}(p)(R) \rightarrow \mathcal{A}(p)_{\text{étale}}(R) = A(p)_{\text{étale}}(k) \rightarrow 0$. We refer to [36] for more details.

⁶The exponent $m + 1$ is chosen, since we work with m -th order deformations (Definition 3.3).

Note that the deformation space of a g -dimensional abelian varieties has dimension g^2 which is equal to the dimension of the tangent space. We now specialize to the case of principally polarised abelian varieties. Recall that a *polarisation* L on an abelian variety A is given by an ample line bundle (defined up to translation and potentially over a separable field extension of k) which defines a rational morphism to the dual of A , $\phi_L : A \rightarrow \hat{A}$. And it is called *principal* if ϕ_L is an isomorphism.

Definition 3.5. Let (A, L) a principally polarised abelian variety over k . We say that $(\mathcal{A}, \mathcal{L})$ is an m -th order deformation, if \mathcal{A} is a deformation of A , and \mathcal{L} is an ample line bundle on \mathcal{A} with $\mathcal{L} \times_R \text{Spec}(k) \cong L^7$.

If the polarisation is clear from the context, we denote A (resp. \mathcal{A}) instead of (A, L) (resp. $(\mathcal{A}, \mathcal{L})$). The deformation space of principally polarised abelian varieties has dimension $g(g+1)/2$. More precisely, for a principally polarised abelian variety A/k , the functor of deformations of A is pro-representable by $k[[t_1, \dots, t_{g(g+1)/2}]]$. Concretely, it is given by the completion at A of the moduli space \mathbf{A}_g/k of principally polarised abelian varieties⁸. Essentially, this means that any m -th order deformation \mathcal{A} of A corresponds to a unique ring homomorphism

$$(4) \quad \psi_{\mathcal{A}} : k[[t_1, \dots, t_{g(g+1)/2}]] \rightarrow R.$$

This motivates the following definition.

Definition 3.6. Let A be a principally polarised abelian variety over k , and \mathcal{A} an m -th order deformation. Then we say that $\lambda_1 = \psi_{\mathcal{A}}(t_1), \dots, \lambda_{g(g+1)/2} = \psi_{\mathcal{A}}(t_{g(g+1)/2})$ with $\psi_{\mathcal{A}}$ as in Eq. 4 are *deformation parameters* of \mathcal{A} .

Note that the definition of the deformation parameters depends on the local representation $k[[t_1, \dots, t_{g(g+1)/2}]]$ of the moduli space at A . Below, we describe a canonical choice for the case $g = 1$.

Example 3.7. In the case of elliptic curves, the deformation space is one-dimensional, hence a deformation is defined by a single deformation parameter λ .

Let E be an elliptic curve and suppose that $j(E) \neq 0, 1728$. For an m -th order deformation \mathcal{E} of E , we can define the deformation parameter

$$\lambda_{\mathcal{E}} = j(\mathcal{E}) - j(E) \in (\epsilon) \triangleleft k[\epsilon]/(\epsilon^{m+1})$$

associated to the j -invariant.

Note that as long as $j(E) \neq 0, 1728$, we can construct the universal elliptic curve over the universal deformation ring. More explicitly, for a given deformation parameter $\lambda \in (\epsilon)$, we can consider

$$\mathcal{E} : y^2 = x^3 + ax + b, \quad \text{where } a = b = \frac{27(j(E) + \lambda)}{4(1728 - (j(E) + \lambda))},$$

which defines an elliptic curve over R with j -invariant $j(\mathcal{E}) = j(E) + \lambda$. Over $j(E) = 0, 1728$ the universal elliptic curve only exists as an algebraic stack, which is why we are going to use $y^2 = x^3 + 6x^2 + x$ rather than $y^2 = x^3 + x$ as the initial point in our algorithms.

⁷A slight technicality: deforming the polarisation means deforming the isogeny ϕ_L . However, if ϕ_L is induced by a line bundle over k (rather than over its separable closure), and ϕ_L deforms to R , then L does too. This is because the obstruction to descending L is given by a smooth torsor, which has a point over k by assumption, so over R too since R is Henselian.

⁸To handle deformations to a general Artinian ring, we would need to use the completion $\hat{\mathbf{A}}_{g,A}/\mathbb{Z} \simeq \mathbb{Z}_q[[t_1, \dots, t_{g(g+1)/2}]]$ at A of the moduli space over \mathbb{Z} rather than over k .

3.2. Deformations of isogenies. Let $f : A \rightarrow A'$ be an isogeny of abelian varieties over a field k . Here, we discuss lifts of such an isogeny to R . In particular, we are interested in lifts of product isogenies and the behaviour of deformation parameters under isogenies.

Recall that a morphism of group schemes $f : \mathcal{A} \rightarrow \mathcal{A}'$ over $\text{Spec}(R)$ is an isogeny if it is surjective and its kernel is a flat finite $\text{Spec}(R)$ -group scheme.

Proposition 3.8. *Let $f : A \rightarrow A'$ be a separable isogeny of principally polarised abelian varieties defined over k . Consider an m -th order deformation \mathcal{A} of A . Then up to isomorphism there exists a unique deformation \mathcal{A}' of A' so that f lifts to an isogeny $\tilde{f} : \mathcal{A} \rightarrow \mathcal{A}'$ over R .*

This is a classical result, which stems from the fact that the forgetting map $\mathbf{A}_g(\Gamma_0(\ell)) \rightarrow \mathbf{A}_g$ is étale. Recall that $\mathbf{A}_g(\Gamma_0(\ell))$ denotes the moduli space of principally polarized abelian varieties with $\Gamma_0(\ell)$ -structure, i.e. the elements are principally polarized abelian varieties together with the kernel of an ℓ -isogeny. Since we need to lift isogenies explicitly in our applications, we provide a short constructive proof below.

Proof of Proposition 3.8. Let G be the kernel of the isogeny $f : A \rightarrow A'$. By assumption, f is separable, hence $G \subset A[N]$ where N is coprime to p . Now consider the deformation \mathcal{A} of A to R . Since R is Henselian, it has the same finite étale covers as k . It follows that there is a unique lift $\tilde{G} \subset \mathcal{A}[N]$ of G . We can set $\mathcal{A}' = \mathcal{A}/\tilde{G}$ which is again a group scheme. Necessarily \mathcal{A}' is a deformation of A' , and the morphism $\tilde{f} : \mathcal{A} \rightarrow \mathcal{A}'$ is an isogeny lying above f . \square

The following corollary is an easy consequence of the unique lifting of isogenies and Kani's lemma (Lemma 2.1). It describes when product isogenies lift to product isogenies.

Corollary 3.9. *Let E, E_{ab} be elliptic curves over k and $F : E \times E_{ab} \rightarrow E_a \times E_b$ a product isogeny. Consider an m -th order deformation \mathcal{E} of E . Then up to isomorphism there exist unique deformations $\mathcal{E}_{ab}, \mathcal{E}_a, \mathcal{E}_b$ of E_{ab}, E_a, E_b , so that F lifts to a product isogeny*

$$\tilde{F} : \mathcal{E} \times \mathcal{E}_{ab} \rightarrow \mathcal{E}_a \times \mathcal{E}_b.$$

Remark 3.10. Let $F : E \times E_{ab} \rightarrow E_a \times E_b$ as in Corollary 3.9. Further let \mathcal{E} and \mathcal{E}_{ab} be arbitrary deformations of E and E_{ab} . Then the isogeny F lifts uniquely to an isogeny $\tilde{F} : \mathcal{E} \times \mathcal{E}_{ab} \rightarrow \mathcal{A}$, where \mathcal{A} is a deformation of $E_a \times E_b$. To understand, when \tilde{F} is a product isogeny, we recall from Kani's lemma that any such isogeny is induced by an isogeny diamond as in Eq. 3. This means that the deformation \mathcal{A} is a product if and only if there is an isogeny $\tilde{h} : \mathcal{E} \rightarrow \mathcal{E}_{ab}$ lifting the isogeny $f' \circ g = g' \circ f$ underlying the product isogeny F .

The following proposition tells us how deformations behave under isogenies. To obtain a description in terms of deformation parameters, we assume that some canonical choice on the local representation of the moduli space has been made. We simply refer to *the* deformation parameters, when we mean the deformation parameters induced by this choice.

Proposition 3.11. *Let $f : A \rightarrow A'$ be a separable isogeny of principally polarized abelian varieties over k and $m \geq 1$, then there exist polynomials*

$$h_1, \dots, h_{g(g+1)/2} \in k[x_1, \dots, x_{g(g+1)/2}]$$

of degree at most m with the following property:

For any m -th order deformation \mathcal{A} of A with parameters $(\lambda_1, \dots, \lambda_{g(g+1)/2})$, the deformation \mathcal{A}' with parameters $(\lambda'_1, \dots, \lambda'_{g(g+1)/2})$, where

$$\lambda'_i = h_i(\lambda_1, \dots, \lambda_{g(g+1)/2})$$

is the unique deformation of A' for which f lifts to an isogeny $\tilde{f} : \mathcal{A} \rightarrow \mathcal{A}'$.

Proof. Let $\hat{\mathbf{A}}_{g,A} = k[[t_1, \dots, t_{g(g+1)/2}]]$ and $\hat{\mathbf{A}}_{g,A'} = k[[t'_1, \dots, t'_{g(g+1)/2}]]$ be the local completions of the moduli space \mathbf{A}_g at A and A' , that induce our canonical choice of deformation parameters. The isogeny f corresponds to a point $(A, \text{Ker } f) \in \mathbf{A}_g(\Gamma_0(\ell))$. The modular correspondence $\overline{\Phi}_\ell : \mathbf{A}_g(\Gamma_0(\ell)) \rightarrow \mathbf{A}_g \times \mathbf{A}_g$ is given by two étale projection maps. This induces an (analytic, i.e. continuous) isomorphism of the completion of $\mathbf{A}_g(\Gamma_0(\ell))$ at $(A, \text{Ker } f)$ with $\hat{\mathbf{A}}_{g,A}$ and $\hat{\mathbf{A}}_{g,A'}$ respectively. In particular, f induces an isomorphism $\phi_f : \hat{\mathbf{A}}_{g,A} \rightarrow \hat{\mathbf{A}}_{g,A'}$. Thus, we can write

$$t'_i = \tilde{h}_i \in k[[\phi_f(t_1), \dots, \phi_f(t_{g(g+1)/2})]]$$

for each $i \in \{1, \dots, g(g+1)/2\}$. Now the polynomials $h_1, \dots, h_{g(g+1)/2}$ from the statement of the proposition are obtained by truncating these formal power series at degree m . \square

Remark 3.12. The modular polynomial Φ_ℓ describes the image of the modular correspondence $\overline{\Phi}_\ell : \mathbf{A}_1(\Gamma_0(\ell))/k \rightarrow \mathbf{A}_1/k \times \mathbf{A}_1/k$, it can also be seen as $\overline{\Phi}_\ell$ evaluated on the generic point of \mathbf{A}_1/k . From this point of view Corollary 1.4 is natural: we start with the modular correspondence $\overline{\Phi}_\ell$ evaluated at one point $E \in \mathbf{A}_1/k$, i.e. with all ℓ -isogenies starting from E . More precisely, the fibre $(\pi_1 \circ \overline{\Phi}_\ell)^{-1}(E)$ gives the isogenies and evaluating $(\pi_2 \circ \overline{\Phi}_\ell)$ on the fibre gives the codomains. We then evaluate $\overline{\Phi}_\ell$ at the universal deformation of E at a high enough precision to recover $\overline{\Phi}_\ell$ on the generic point.

A consequence of this proposition applied to product isogenies is described in the corollary below. It provides us with a convenient tool for lifting product isogenies. To make a precise statement, we require that isogenies are *normalised*. Essentially, this means that the isogeny acts as the identity on the basis of differentials.

Corollary 3.13. *Let E, E_{ab} be elliptic curves over a field k and $F : E \times E_{ab} \rightarrow E_a \times E_b$ a product isogeny. Then there exists a polynomial $h \in k[x_1, x_2]$ of degree at most m with the following property:*

Let \mathcal{E} and \mathcal{E}_{ab} be arbitrary m -th order deformations of E and E_{ab} with deformation parameters $\lambda = j(\mathcal{E}) - j(E)$ and $\lambda_{ab} = j(\mathcal{E}_{ab}) - j(E_{ab})$, respectively. Further let $\tilde{F} : \mathcal{E} \times \mathcal{E}_{ab} \rightarrow \mathcal{A}$ be a lift of F and assume the isogeny is normalised. Then

$$\chi_{10}(\mathcal{A}) = h(\lambda, \lambda_{ab}),$$

where χ_{10} is the Siegel modular cusp form (see Subsection 2.3).

Proof. On a decomposable abelian surface $A = E \times E_{ab}$, a canonical choice of deformation parameters is given by choosing two deformation parameters λ and λ_{ab} for E and E_{ab} , respectively. Note that altering only these two deformation parameters preserves the product structure on the resulting deformation. Let $S \subset \hat{\mathbf{A}}_{2,E \times E_{ab}}$ be the subring of the deformation ring given by product deformations.

A difference with Proposition 3.11 is that χ_{10} is a modular form rather than a modular function. Let $\chi_{10,\ell}$ be the modular form which associates to a normalised ℓ -isogeny $F : (A, \omega_A) \rightarrow (B, \omega_B)$ the value $\chi_{10}(B, \omega_B)$. (Normalised means that $F^*\omega_B = \omega_A$.) This is a modular form of level $\Gamma_0(\ell)$, hence a section of a power of the Hodge line bundle \mathfrak{h}^{10} on $\mathbf{A}_2(\Gamma_0(\ell))$.⁹

Now given curve equations for $\mathcal{E} \times \mathcal{E}_{ab}/R$, we get a canonical basis of differentials $((dx/y, 0), (0, dx_{ab}/y_{ab}))$, hence a trivialisation of the Hodge bundle \mathfrak{h} above S , i.e. an isomorphism $S \rightarrow S \otimes_{\mathbf{A}_2} \mathfrak{h}$. We can use the modular correspondence $\overline{\Phi}_\ell$ as in Proposition 3.11 to see S as a deformation ring of the isogeny F , and pulling back the trivialisation of the Hodge line bundle defined above, we can interpret $\chi_{10,\ell} \in \Gamma(S \otimes_{\mathbf{A}_2} \mathfrak{h}^{10})$ as an element of S , given by a polynomial h . \square

⁹Note that over \mathbb{C} , $\chi_{10,\ell}$ corresponds to the modular forms of level $\Gamma_0(\ell)$: $\tau \mapsto \chi_{10}(\tau/\ell)$.

Remark 3.14. In Section 4, we will apply Corollary 3.13 to a $(2^n, 2^n)$ -isogeny F induced by an isogeny diamond. The isogeny computation consists of a gluing isogeny, followed by a chain of Richelot isogenies, followed by a splitting isogeny, hence we need to ensure that each step is normalised. The differentials are implicitly kept track off by the curve equation: to an elliptic curve $y^2 = x^3 + ax + b$ we associate the differential dx/y , and to an hyperelliptic curve of genus 2 the differentials $(dx/y, xdx/y)$. Richelot formulas work at the level of curve equations and give normalised formulas.

4. COMPUTING WITH DEFORMATIONS

In this section we explain different algorithms related to the computation of deformations. This includes the lifting of torsion points to deformed elliptic curves (Algorithm 2), as well as the lifting of $(2, 2)$ -isogeny chains (Algorithm 5). Further, we present a method for lifting product isogenies (Algorithm 4) which lies at the heart of our algorithms for the computation of modular polynomials. The section concludes with the deformation of general isogenies and proves Theorem 1.3.

As in the previous section, k is a field with characteristic $p \neq 0$, and $R = k[\epsilon]/(\epsilon^{m+1})$.

4.1. Arithmetic in $R = \mathbb{F}_q[\epsilon]/(\epsilon^{m+1})$. In practice, we work with a finite field $k = \mathbb{F}_q$ in all algorithms. Note that computing in the Artin ring R is equivalent to computing in the formal power series ring $\mathbb{F}_q[[\epsilon]]$ with precision $m + 1$.

Throughout, we use

$$M(n) = O(n \log n \log \log n)$$

to describe the complexity for the multiplication of two n -bit integers. Further, we set

$$M(\mathbb{F}_q) = M(\log q), \quad M(R) = M(m \log q)$$

which correspond to the complexities for the multiplications of two elements in \mathbb{F}_q and R , respectively. These bounds can be obtained from Kronecker substitution. Similarly, the multiplication of two polynomials of degree $O(d)$ over R costs

$$M(R, d) = M(m e d (\log p + \log(m e d))), \quad \text{where } q = p^e.$$

For more details on these complexities, we refer to [16] and [17]. Moreover, we already note that later we will have $m \approx d \approx \ell$, $e \in \{1, 2\}$ and $\log(p) \in O(\log(\ell))$ for some integer ℓ , hence

$$M(R) = O(\ell \log^2 \ell \log \log \ell), \quad M(R, \ell) = O(\ell^2 \log^2 \ell \log \log \ell).$$

Inversions in R can be computed with the same asymptotic complexity, for instance by using *Newton iterations*. Remarkably, this approach also allows us to lift roots of polynomials over k to roots of polynomials over R . Since Newton iterations play an important role in our algorithms, this standard method is presented in Algorithm 1.

Algorithm 1 `Newton_lift`

Input: An element $\alpha \in k$ and a polynomial $f \in R[x]$ such that $f(\alpha) \equiv 0$ and $f'(\alpha) \not\equiv 0 \pmod{(\epsilon)}$.

Output: An element $\tilde{\alpha} \in R$ with $\tilde{\alpha} \equiv \alpha \pmod{(\epsilon)}$ and $f(\tilde{\alpha}) = 0$.

- 1: **for** $k \leftarrow 1, \dots, \lceil \log_2(m + 1) \rceil$ **do**
 - 2: $\alpha \leftarrow \alpha + O(\epsilon^{\min(2^k, m+1)})$
 - 3: $\alpha \leftarrow \alpha + \frac{f(\alpha)}{f'(\alpha)}$
 - 4: **end for**
 - 5: **return** α
-

4.2. Lifting torsion points. Let A be an abelian variety over k . Given a point $P \in A(k)$, and an m -th order deformation \mathcal{A} of A , there exist multiple points $\tilde{P} \in \mathcal{A}(R)$ reducing to P . On the other hand, there is an isomorphism $\mathcal{A}[N] \cong A[N]$ if $p \nmid N$. In order to be able to lift isogenies, we need to make this isomorphism explicit. The case of elliptic curves is covered by Algorithm 2 and the special case of abelian surfaces and $N = 2$ is covered by Algorithm 3.

Lemma 4.1. *On input a point $P \in E[N]$ with N coprime to $\text{char}(k)$, E an elliptic curve over k and an m -th order deformation \mathcal{E} , Algorithm 2 returns the unique lift $\tilde{P} \in \mathcal{E}[N]$ of P . For fixed N , the algorithm runs in time $O(1)$ over $R = k[\epsilon]/(\epsilon^{m+1})$.*

Proof. Since $\text{char}(k) \nmid N$, there is a unique lift $\tilde{P} = (\tilde{x}, \tilde{y})$ of $P = (x_0, y_0) \in E[N]$ which is also an N -torsion point. Let $f_{\mathcal{E}, N}$ denote the N -th division polynomial of \mathcal{E} . Then $f_{\mathcal{E}, N}(x_0) \equiv 0 \pmod{(\epsilon)}$, and we can lift x_0 to a root \tilde{x}_0 of $f_{\mathcal{E}, N}$ using the method `Newton_lift` (Algorithm 1). The corresponding y -coordinate \tilde{y}_0 lying above $y_0 \in k$ is computed by another call to `Newton_lift`.

It is clear that $\tilde{P} = (\tilde{x}_0, \tilde{y}_0)$ is in $\mathcal{E}[N]$. And the running time of the algorithm is determined by the running time of `Newton_lift`. \square

Algorithm 2 `lift_point`

Input: Elliptic curve E over k , a point $P \in E(k)[N]$ with $p \nmid N$, and an m -th order deformation $\mathcal{E} : y^2 = f(x)$.

Output: The lift $\tilde{P} \in \mathcal{E}[N]$ of P .

- 1: $(x_0, y_0) \leftarrow P$
 - 2: $f_{\mathcal{E}, N}$ the N -th division polynomial of \mathcal{E} .
 - 3: $\tilde{x}_0 \leftarrow \text{Newton_lift}(x_0, f_{\mathcal{E}, N})$
 - 4: $\tilde{y}_0 \leftarrow \text{Newton_lift}(y_0, y^2 - f(\tilde{x}_0))$
 - 5: **return** $(\tilde{x}_0, \tilde{y}_0) \in \mathcal{E}[N]$
-

There are different ways to generalise Algorithm 2 to higher dimensions. In our applications, we only need to lift 2-torsion points of abelian surfaces. On decomposable abelian surfaces, we may just apply Algorithm 2 on every component. Algorithm 3 describes a general method for lifting 2-torsion points on abelian surfaces.

Lemma 4.2. *On input a point $P \in A[2]$, where A is an abelian surface over k , and an m -th order deformation \mathcal{A} , Algorithm 2 returns the unique lift $\tilde{P} \in \mathcal{A}[2]$ of P . The algorithm runs in time $O(1)$ over $R = k[\epsilon]/(\epsilon^{m+1})$.*

Proof. Recall from Subsection 2.4 that we present 2-torsion points as pairs $\{\alpha_1, \alpha_2\}$.

First assume that A is indecomposable, i.e. $A = \text{Jac}(C)$ with $C : y^2 = f(x)$. Then $\mathcal{A} = \text{Jac}(\mathcal{C})$ and $\mathcal{C} : y^2 = \tilde{f}(x)$ is a deformation of C . In particular, lifting the torsion point P to $\text{Jac}(\mathcal{C})$ consists in lifting the roots α_1, α_2 of f to roots of the polynomial $\tilde{f} \in R[x]$. Note that we need to choose an equation $y^2 = \tilde{f}(x)$ so that $\tilde{f} \equiv f \pmod{(\epsilon)}$.¹⁰

On the other hand, if $A = E \times E'$, there are two cases to consider. If $\mathcal{A} = \mathcal{E} \times \mathcal{E}'$ is also a product of elliptic curves, then the situation is similar as above. We simply lift roots α_1, α_2 to roots of the defining polynomials for \mathcal{E} and \mathcal{E}' respectively.

The case where $\mathcal{A} = \text{Jac}(\mathcal{C})$ is indecomposable, but $A = E_1 \times E_2$ is a product of elliptic curves is more subtle. For $i \in \{1, 2\}$, we write $E_i : y_i^2 = f_i(x_i)$. Note that there necessarily exists a relation between the coordinates (x_1, y_1) and (x_2, y_2) . In particular, there exist polynomials \tilde{f}_i (of degree 5 or 6) in $k[\epsilon]/(\epsilon^{m+1})[x]$ such that the reduction modulo (ϵ) is equal to f_i . Consequently,

¹⁰In our applications (e.g. Algorithm 5), this is automatically the case.

Algorithm 3 lift_2_torsion

Input: An abelian surface A , a point $P \in A(k)[2]$, and an m -th order deformation \mathcal{A} .

Output: The lift $\tilde{P} \in \mathcal{A}[2]$ of P .

```
1:  $\{\alpha_1, \alpha_2\} \leftarrow P$ 
2: if  $A = \text{Jac}(C)$  with  $C : y^2 = f(x)$  then
3:    $\tilde{f} \leftarrow$  a lift of  $f$  with  $\mathcal{C} : y^2 = \tilde{f}(x)$  and  $\mathcal{A} = \text{Jac}(\mathcal{C})$ 
4:    $\tilde{f}_i \leftarrow \tilde{f}$  for  $i = 1, 2$ 
5: else
6:    $A = E_1 \times E_2$  with  $E_1 : y_1^2 = f_1(x_1)$  and  $E_2 : y_2^2 = f_2(x_2)$ 
7:   if  $\mathcal{A} = \mathcal{E}_1 \times \mathcal{E}_2$  then
8:      $\tilde{f}_i \leftarrow$  lift of  $f_i$  with  $\mathcal{E}_i : y_i^2 = \tilde{f}_i(x_i)$  for  $i = 1, 2$ 
9:   else
10:     $\mathcal{A} = \text{Jac}(\mathcal{C})$  with  $\mathcal{C} : y^2 = \tilde{f}(x)$ 
11:     $\tilde{f}_i \leftarrow$  lift of  $f_i$  with  $\mathcal{C} : y_i^2 = \tilde{f}_i(x_i)$  for  $i = 1, 2$ 
12:   end if
13: end if
14:  $\tilde{\alpha}_i \leftarrow \text{Newton\_lift}(\alpha_i, \tilde{f}_i)$  for  $i = 1, 2$ 
15: return  $\{\tilde{\alpha}_1, \tilde{\alpha}_2\} \in \mathcal{A}[2]$ 
```

the root α_i can be lifted to a root of \tilde{f}_i using `Newton_lift`. Note that in order to interpret the resulting representation $\{\tilde{\alpha}_1, \tilde{\alpha}_2\} \in \mathcal{A}[2]$ of the 2-torsion point correctly, one needs to consider the polynomials \tilde{f}_1 and \tilde{f}_2 . \square

4.3. Lifting an isogeny diamond. Given an isogeny diamond over a field k as in Subsection 2.5, there are two ways for lifting it to R . The straight-forward method is to lift the individual elliptic curve isogenies to isogenies over R . Here, we describe a different method (Algorithm 4), where the product isogeny induced by Kani's Lemma is directly lifted to R . The key ingredient for this algorithm is Corollary 3.13.

Proposition 4.3. *Let E, E_a, E_b, E_{ab} be elliptic curves which are the vertices of an isogeny diamond, and let $K \subset E \times E_{ab}$ the kernel of the product isogeny induced by Kani's Lemma. Further assume that none of the elliptic curves have extra automorphisms, that is we assume $j(E), j(E_a), j(E_b), j(E_{ab}) \notin \{0, 1728\}$. On input this data together with an m -th order deformation \mathcal{E} of E , Algorithm 4 outputs the deformations $(\mathcal{E}_a, \mathcal{E}_b, \mathcal{E}_{ab})$ such that F lifts to an isogeny $\tilde{F} : \mathcal{E} \times \mathcal{E}_{ab} \rightarrow \mathcal{E}_a \times \mathcal{E}_b$.*

Proof. The main part of the algorithm consists in computing the correct deformation of E_{ab} . Once this is done, we can lift the kernel K to a subgroup of $\mathcal{E} \times \mathcal{E}_{ab}$, and then $\mathcal{E}_a \times \mathcal{E}_b$ is computed as the codomain of the product isogeny with this kernel (Line 16).

To understand the main part, recall from Corollary 3.13 that there exists a polynomial $h \in k[x_1, x_2]$ with the property that

$$h(j(\mathcal{E}) - j(E), j(\mathcal{E}') - j(E_{ab})) = \chi_{10}(\mathcal{A}),$$

for any arbitrary deformation \mathcal{E}' of E_{ab} and where \mathcal{A} is the corresponding codomain of the lifted isogeny $F : \mathcal{E} \times \mathcal{E}' \rightarrow \mathcal{A}$. Our goal is to find the correct deformation $\mathcal{E}' = \mathcal{E}_{ab}$ so that \mathcal{A} is again decomposable, i.e. $\chi_{10} = 0$. If we knew the polynomial h , we could simply apply the Newton method to find the correct value for $j(\mathcal{E}_{ab})$. Since this is not the case, we use an interpolation based approach, where we evaluate h at two different values of $j(\mathcal{E}')$. Similar as in Newton's method, this approach lets us double the precision at each step.

Algorithm 4 lift_isogeny_diamond

Input: Elliptic curves E, E_a, E_b, E_{ab} which are the vertices of an isogeny diamond, and a group $K \subset E \times E_{ab}$ which is the kernel of the corresponding product isogeny. An m -th order deformation \mathcal{E} of E .

Assertion: $j(E), j(E_a), j(E_b), j(E_{ab}) \notin \{0, 1728\}$.

Output: The lifted isogeny diamond $(\mathcal{E}, \mathcal{E}_a, \mathcal{E}_b, \mathcal{E}_{ab})$ over R .

```
1:  $\mathcal{E}_{ab} \leftarrow E_{ab}, \tilde{K} \leftarrow K$ 
   /*The precision is doubled at each step.*/
2: for  $k \leftarrow 1, \dots, \lceil \log_2(m+1) \rceil$  do
3:    $j_{ab} = j(\mathcal{E}_{ab})$ 
4:    $R \leftarrow k[\epsilon] / \left( \epsilon^{\min(2^k, m+1)} \right)$ 
   /*We compute  $\chi_{10}$  for two different lifts of  $E_{ab}$ .*/
5:   for  $i = 0, 1$  do
6:      $j_i \leftarrow R(j_{ab}) + i \cdot \epsilon^{2^{k-1}}$ 
7:      $\mathcal{E}_i \leftarrow \text{deformation}(E_{ab})$  with  $j(E_i) = j_i$  over  $R$ .
8:      $K_i \leftarrow \text{lift}(\tilde{K}) \subset (\mathcal{E} \times \mathcal{E}_i)[N]$ .
9:      $A_i \leftarrow (\mathcal{E} \times \mathcal{E}_i) / K_i$ .
10:     $\delta_i \leftarrow \chi_{10}(A_i)$ 
11:   end for
   /*The correct lift of  $\mathcal{E}_{ab}$  is deduced from  $\delta_0, \delta_1$ .*/
12:    $j_{ab} \leftarrow j_{ab} - \frac{\delta_0}{\delta_1 - \delta_0}$ 
13:    $\mathcal{E}_{ab} \leftarrow \text{deformation}(E_{ab})$  with  $j(\mathcal{E}_{ab}) = j_{ab}$ 
14:    $\tilde{K} \leftarrow \text{lift}(\tilde{K}) \subset (E \times E')[N]$ 
15: end for
   /*We compute the lifted product isogeny.*/
16:  $\mathcal{E}_a \times \mathcal{E}_b \leftarrow (\mathcal{E} \times \mathcal{E}_{ab}) / \tilde{K}$ .
17: return  $(\mathcal{E}_a, \mathcal{E}_b, \mathcal{E}_{ab})$ 
```

To describe the idea in more detail, assume we are at Step k of the algorithm. Let $E_{ab, k-1}$ be correct deformation of order $2^{k-1} - 1$, that is

$$h(j(\mathcal{E}) - j(E), j(E_{ab, k-1}) - j(E_{ab})) \equiv 0 \pmod{(\epsilon^{2^{k-1}})}.$$

Now, we compute two different deformations of E_{ab} of order $2^k - 1$, and denote them by E_0 and E_1 . More precisely, we choose the lifts with j -invariants

$$j(E_0) = j(E_{ab, k-1}), \quad j(E_1) = j(E_{ab, k-1}) + \epsilon^{2^{k-1}}.$$

Note that we can compute arbitrary deformations, since $j(E_{ab}) \neq 0, 1728$ by assumption (cf. Example 3.7). For both cases $i = 0, 1$, we lift the kernel K to a subgroup $K_i \subset \mathcal{E} \times E_i$ using the method `lift_point` described in Algorithm 2. Then we compute the codomains $A_i = (\mathcal{E} \times E_i) / K_i$ of the resulting isogenies, and $\delta_i = \chi_{10}(A_i)$. With $j(E_{ab, k}) = j(E_{ab, k-1}) - \delta_0 / (\delta_1 - \delta_0)$, we find that

$$h(j(\mathcal{E}) - j(E), j(E_{ab, k}) - j(E_{ab})) \equiv 0 \pmod{(\epsilon^{2^k})}.$$

After computing the data corresponding to this deformation, one can proceed to Step $k + 1$. \square

4.4. Lifting (smooth-degree) isogenies. Given an isogeny $f : A \rightarrow A'$ and a deformation \mathcal{A} of A , Proposition 3.8 shows that there exists a unique lift $\tilde{f} : \mathcal{A} \rightarrow \mathcal{A}'$. Let $G \subset A[N]$ be the kernel of f . The straight-forward approach for lifting f consists in lifting the kernel generators, for instance

using Algorithms 2, in order to obtain a lifted kernel $\tilde{G} \subset \mathcal{A}[N]$. Then the isogeny can be computed using standard algorithms. This is the approach used in Algorithm 4.

However, if N is composite the above approach is not optimal. For instance, in the case $N = 2^n$, an isogeny can be lifted in time $O(n)$ (provided some auxiliary data from the original isogeny), while the naive method has complexity $O(n \log(n))$. To keep everything explicit, we restrict to the case of $(2^n, 2^n)$ -isogenies.

Algorithm 5 lift_2_2_chain

Input: A $(2^n, 2^n)$ -isogeny $f = f_n \circ \dots \circ f_1 : A_0 \rightarrow A_n$ over k , the data $\ker(f_i) = \langle P_{1,i}, P_{2,i} \rangle$ for each i , and an m -th order deformation \mathcal{A}_0 of A_0 .

Output: The deformation \mathcal{A}_n so that f lifts to an isogeny $\tilde{f} : \mathcal{A}_0 \rightarrow \mathcal{A}_n$.

```

1: for  $i = 1, \dots, n$  do
2:    $\tilde{P}_{j,i} \leftarrow \text{lift\_2\_torsion}(A_{i-1}, P_{j,i}, \mathcal{A}_{i-1})$  for  $j = 1, 2$ 
3:    $\mathcal{A}_i \leftarrow \text{2\_2\_isogeny}(\mathcal{A}_{i-1}, \langle \tilde{P}_{1,i}, \tilde{P}_{2,i} \rangle)$ 
4: end for

```

Lemma 4.4. *Given a $(2^n, 2^n)$ -isogeny $f = f_n \circ \dots \circ f_1 : \mathcal{A}_0 \rightarrow \mathcal{A}_n$ over a field k together with the data $\ker(f_i) = \langle P_{1,i}, P_{2,i} \rangle$ for each $i \in \{1, \dots, n\}$, and an m -th order deformation \mathcal{A}_0 of A_0 , Algorithm 5 outputs a deformation \mathcal{A}_n of A_n so that f lifts to an isogeny $\tilde{f} : \mathcal{A}_0 \rightarrow \mathcal{A}_n$. The algorithm runs in time $O(n)$ over $R = k[\epsilon]/(\epsilon^{m+1})$.*

Proof. The isogeny chain computed in Algorithm 5 is correct, since at each step the kernel generators are lifted correctly as per Lemma 4.2. The method `2_2_isogeny` is used as a black box. It is only important that it runs in time $O(1)$ over R , see Subsection 2.4 for more details. \square

Remark 4.5. An even faster method to lift a smooth degree isogeny is to use modular correspondences. For instance, a 2^n -isogeny over k is given by a sequence of modular points which are solutions of the modular correspondance $\overline{\Phi}_2$, and lifting the isogeny corresponds to lifting this sequence of points. In practice, since we want to be able to evaluate modular forms, we need to work with normalised isogenies, hence use a version of $\overline{\Phi}_2$ which keeps track of our differentials.

For instance, Richelot isogenies are well suited to write explicit modular equations. Explicit formulas are given by “multiradical isogeny” formulas [7, § 4.2] which describe the modular correspondence $\overline{\Phi}_2 : A_2(\Gamma_1(4)) \rightarrow A_2(\Gamma'_1(4)) \rightarrow A_2(\Gamma_1(2)) \times A_2(\Gamma_1(2))$ (see [7, § 3.1] for the notations). The theta duplication formula is also naturally 2-radical in dimension 2. Using one of these radical 2-isogeny formulas, lifting each such 2-isogeny in dimension 2 amounts to lifting (via Newton) three square roots. In our implementation, we follow precisely this approach. This has the advantage that it bypasses the need to lift torsion points and evaluating isogeny formulas on these lifted points. Another potential (constant-time) speed-up could be obtained by using 4-radical isogeny formulas instead.

4.5. Deforming a general isogeny. In this section we prove the general statement from Theorem 1.3. We are given an efficient representation of an isogeny $f : E_1 \rightarrow E_2$, and we want to lift it to $\tilde{f} : \mathcal{E}_1 \rightarrow \mathcal{E}_2$ for some given m -th order deformation \mathcal{E}_1/R .

Initialisation: We fix $N > \ell$ with N prime to ℓ and p , and such that the N -torsion is accessible (which means that its Sylow subgroups lie in small extensions of the base field), for instance N is powersmooth or $N = 2^n$ if the 2^n -torsion of E is rational. We write $N - \ell = a^2 + b^2 + c^2 + d^2$, and consider α , the 4×4 quaternion matrix with norm $N - \ell$. Then α induces an endomorphism on E_1^4 and E_2^4 . We can construct the isogeny diamond

$$\begin{array}{ccc}
E_1^4 & \xrightarrow{f} & E_2^4 \\
\downarrow \alpha & & \downarrow \alpha \\
E_1^4 & \xrightarrow{f} & E_2^4.
\end{array}$$

which allows us to embed f into an 8-dimensional N -endomorphism $F : E_1^4 \times E_2^4 \rightarrow E_1^4 \times E_2^4$. By assumption, we have an efficient representation of f , hence we can evaluate it on the N -torsion in order to compute the HD representation F of f . See [33] for more details.

Since N is taken to be smooth, we can decompose F as a product of small degree isogenies. In practice, F can be computed in the theta model, using the isogeny algorithm of [26], working with level 4 theta functions, and representing our intermediate abelian varieties by their theta constants. The codomain of F is equal to $E_1^4 \times E_2^4$ but it may have different theta constants than the domain (because the level $\Gamma(4, 8)$ -theta structure needs not be preserved), so we apply a symplectic change of basis at the end to get matching theta constants. A way to compute this matrix is described in [10, Appendix F].

Newton iterations: We then proceed by a Newton iteration, doubling the precision m at each stage. For simplicity we explain how to go from $m = 1$ to $m = 2$ here, the general case being similar.

Let $\tilde{f} : \mathcal{E}_1 \rightarrow \mathcal{E}_2$ be the deformation of f to \mathcal{E}_1 . The isogeny diamond above certainly deforms (with the same matrix α), and so F deforms to an endomorphism \tilde{F} of $\mathcal{E}_1^4 \times \mathcal{E}_2^4$. On the other hand, if we take an arbitrary lift \mathcal{E}'_2 of E_2 to R , the corresponding deformation \tilde{F}' is only an isogeny. In fact, if \tilde{F}' is an endomorphism, then it is given by a matrix, which reduces to the matrix giving F (by unicity of deformations), hence contains the deformation \tilde{f} .

In summary, \mathcal{E}'_2 is the correct codomain of the deformation of f if and only if the codomain of \tilde{F} is equal to the domain. Using modular invariants J (e.g. the theta constants), this can be tested via the equality $J(\mathcal{E}_1^4 \times \mathcal{E}'_2^4) = J(\text{codomain}(\tilde{F}))$. By the same arguments as in Section 3, the left and right member of this equality are analytic in terms of the deformation parameter λ of E_2 (i.e. are given by power series). Since we are doing a Newton iteration, everything becomes linear, so we just need to interpolate between two different evaluations of the deformation parameter λ to solve the equation. For each of these two choices of λ , we need to compute the deformation of \tilde{F} to recover its codomain.

Computing \tilde{F} : We proceed as in Section 4.4. Since F is decomposed into a product of small isogenies over \mathbb{F}_q , we deform these isogenies step by step. The easiest way is to deform the points in the kernel and apply the theta isogeny algorithm.

Deforming a point P of n -torsion: We also proceed by a Newton iteration. Division polynomials are not easy to compute in higher dimension. Instead, we just rely on the fact that there is a unique deformation \tilde{P} of P which is still of n -torsion, and that taking an arbitrary deformation \tilde{P}' we can efficiently compute $n \cdot \tilde{P}'$. In other words, it is easy to evaluate the n -division polynomials, which is enough for the Newton iteration.

Complexity: Since we double the precision m at each step, if we make the standard assumption (verified in practice) that the arithmetic of R is supra-linear in terms of m , the last Newton step is dominating.

The dominant step is in the computation of the deformation \tilde{F} ; it consists in deforming $O(\log \ell)$ isogenies of degree $O(\log \ell)$ in the worst case (e.g. when N is powersmooth). The algorithm cost is thus polynomial in $\log \ell$ arithmetic operations (for deforming the generators of the kernels and then computing the isogeny) in R at precision m .

In the best case, we can take $N = 2^n$ and the N -torsion is rational in E_1 , so we just need to deform $O(\log \ell)$ 2-isogenies.

5. COMPUTING MODULAR POLYNOMIALS OVER FINITE FIELDS

In this section, we present algorithms for computing the modular polynomial $\varphi_\ell(X, Y)$ over \mathbb{F}_p . We describe very explicit versions of the algorithm for a family of primes p depending on ℓ . These methods will be used in our implementation for computing the modular polynomial over \mathbb{Q} . In addition, Section 5.4 discusses the generalization to arbitrary primes.

In all cases, the main ingredient for our procedure is an algorithm to lift isogeny diamonds over \mathbb{F}_p to isogeny diamonds over $\mathbb{F}_p[[\epsilon]]$ with precision at least $\ell + 2$. The case distinction is necessary, since we construct (ℓ, d) -isogeny diamonds with $\ell + d$ smooth, which depends on the value of ℓ .

5.1. Suitable primes. Depending on $\ell \pmod{4}$, we define a set of primes \mathcal{P}_ℓ for which our algorithm can compute the modular polynomial over \mathbb{F}_p . For an odd prime ℓ , we set

$$\mathcal{P}_\ell = \{p > 11 \text{ prime} : \exists n, a, b \text{ with } 2^n - c_\ell \cdot \ell = a^2 + 4b^2 \text{ and } 2^n \cdot c_\ell \cdot \ell \mid p + 1\},$$

where

$$c_\ell = \begin{cases} 1 & \text{if } \ell \equiv 3 \pmod{4}, \\ 3 & \text{if } \ell \equiv 1 \pmod{4}. \end{cases}$$

Heuristic 5.1. Let ℓ be a positive integer and c_ℓ as defined above. Then we expect

$$\#\{n \in \mathbb{N} \mid 2^n - c_\ell \cdot \ell = a^2 + 4b^2, n \leq x\} \approx \frac{x}{\sqrt{x}}.$$

Under Heuristic 5.1, we expect to find a value $n \in \mathcal{O}(\log(\ell))$. The intuition behind this heuristic and our choice for c_ℓ is explained in the remark below.

Remark 5.2. Recall that an element $z \in \mathbb{N}$ can be written as a sum of two squares, $z = a^2 + b^2$, if and only if its prime decomposition $z = \prod_i p_i^{k_i}$ contains no prime factor $p_i = 3 \pmod{4}$ with k_i odd. In particular, an element $z \equiv 3 \pmod{4}$ cannot be written as a sum of two squares. Now $c_\ell \in \{1, 3\}$ is chosen so that this necessary congruence condition is satisfied for all $z = 2^n - c_\ell \cdot \ell$.

Note that in our setting, we need to write $z = a^2 + 4b^2$. However, we are always in the case that z is odd, hence this is equivalent to z being a sum of two squares. The fraction of elements smaller than some x that can be written as a sum of two squares is known to be in $\mathcal{O}\left(\frac{1}{\sqrt{\log(x)}}\right)$, hence

$$\#\{z = a^2 + 4b^2 \mid z \leq x\} \approx \frac{x}{\sqrt{\log(x)}}.$$

Essentially, Heuristic 5.1 means that we expect the elements of the form $a^2 + 4b^2$ to be uniformly distributed among the numbers of the form $2^n - c_\ell \cdot \ell$ for varying $n \in \mathbb{N}$.

Under Heuristic 5.1, we obtain that there are $\mathcal{O}(\ell)$ primes $p \in \mathcal{P}_\ell$ that have about the same logarithmic size as ℓ , i.e. $\log(p)$ in $\mathcal{O}(\log(\ell))$. A more precise estimate is provided in the lemma below.

Lemma 5.3. *Let ℓ be a positive integer and \mathcal{P}_ℓ as defined above. Then*

$$\#\{p \in \mathcal{P}_\ell \mid p \leq x\} \gtrsim \frac{x}{\ell 2^n \log(x)}.$$

More precisely, the m -th prime in \mathcal{P}_ℓ is bounded by $(2^n \ell)^L m \log(m)^2$, where $L \leq 5$ is Linnik's constant. If Heuristic 5.1 holds, the bound is $(\ell)^{L'} m \log(m)^2$, for some L' .

Proof. We look at the smallest integer so that there exist integers a, b with $2^n - c_\ell \cdot \ell = a^2 + 4b^2$. For this triple (n, a, b) , we can sieve through the primes to find those which satisfy

$$p \equiv -1 \pmod{2^n \cdot c_\ell \cdot \ell}.$$

It follows from the *prime number theorem for arithmetic progressions* that the proportion of primes of this form is $1/\varphi(2^n \cdot c_\ell \cdot \ell) \approx 1/\ell^2$ (with φ denoting Euler's totient function).

For the more precise statement, we need an upper bound on the smallest suitable prime. We need to find primes congruent to -1 modulo $A = 2^n \cdot \ell$. By Linnik's theorem [18, Corollary 18.8], the m -th such prime is bounded by $A^L m \log(m)^2$ where L is the Linnik constant. Under GRH it can be proved that $L \leq 2$, and unconditionally $L \leq 5$. Under Heuristic 5.1, we have $2^n \leq \ell^{\mathfrak{D}}$ for some constant \mathfrak{D} , which concludes the proof. \square

5.2. The case $\ell \equiv 3 \pmod{4}$. Here, we present an algorithm for computing the modular polynomial ϕ_ℓ over \mathbb{F}_p when $\ell \equiv 3 \pmod{4}$ for primes p in the set \mathcal{P}_ℓ .

Algorithm 6 modular_polynomial_modp $\ell \equiv 3 \pmod{4}$

Input: A prime $\ell \equiv 3 \pmod{4}$, and a prime $p \in \mathcal{P}_\ell$.

Output: The modular polynomial $\varphi_\ell(X, Y) \in \mathbb{F}_p[X, Y]$.

```

/*Setup*/
1: Set  $E : y^2 = x^3 + 6x^2 + x$  over  $\mathbb{F}_{p^2}$ 
2:  $\iota \leftarrow \iota \in \text{End}(E)$  with  $\iota \circ \iota = [-4]$ 
3:  $\gamma \leftarrow [a] + [b]\iota$ , where  $2^n - \ell = a^2 + 4b^2$  for some  $n$ .
4: Compute  $P_{2^n}, Q_{2^n}$  with  $E[2^n] = \langle P_{2^n}, Q_{2^n} \rangle$ 
5: Compute  $P_\ell, Q_\ell$  with  $E[\ell] = \langle P_\ell, Q_\ell \rangle$ 
6:  $\tilde{j} \leftarrow j(E) + \epsilon \in \mathbb{F}_{p^2}[\epsilon]/(\epsilon^{\ell+2})$ 
7:  $\mathcal{E} \leftarrow$  elliptic curve with  $j(\mathcal{E}) = \tilde{j}$ 
/*computing and lifting all  $\ell$ -isogenies*/
8: for  $k \leftarrow 0, \dots, \ell$  do
9:   if  $k = \ell$  then
10:     $P_k \leftarrow Q_\ell$ 
11:   else
12:     $P_k \leftarrow P_\ell + k \cdot Q_\ell$ 
13:   end if
14:    $E_k \leftarrow E/\langle P_k \rangle$  with  $f_k : E \rightarrow E_k$ 
/*constructing the  $(\ell, 2^n - \ell)$ -isogeny diamond*/
15:    $E'_k \leftarrow E/\langle \gamma(P_k) \rangle$  with  $f'_k : E \rightarrow E'_k$ 
16:    $K \leftarrow \langle (\hat{\gamma}(P_{2^n}), f'_k(P_{2^n})), (\hat{\gamma}(Q_{2^n}), f'_k(Q_{2^n})) \rangle$ 
17:    $(\mathcal{E}, \mathcal{E}_k, \mathcal{E}', \mathcal{E}'_k) \leftarrow \text{lift\_isogeny\_diamond}(E, E_k, E, E'_k, K, \mathcal{E})$ 
18:    $\tilde{j}_k \leftarrow j(\mathcal{E}_k)$ 
19: end for
/*final step*/
20:  $\varphi \leftarrow \prod_{k=0}^{\ell} (Y - \tilde{j}_k)(\epsilon = X - j(E)) \in \mathbb{F}_p[X, Y]$ 
21: return  $\varphi$ 

```

Theorem 5.4. *On input a prime $\ell \equiv 3 \pmod{4}$ and a prime $p \in \mathcal{P}_\ell$, Algorithm 6 returns the modular polynomial ϕ_ℓ over \mathbb{F}_p . The algorithm runs in*

$$O(\log p + \ell^2) \cdot M(\mathbb{F}_{p^2}) + O(n\ell) \cdot M(R),$$

with $R = \mathbb{F}_{p^2}(\epsilon)/(\epsilon^{\ell+2})$, and $M(\cdot)$ as defined in Subsection 4.1.

Proof. First, the algorithm sets $E : y^2 = x^3 + 6x^2 + x$ over \mathbb{F}_{p^2} . This curve is 2-isogenous to the elliptic curve with j -invariant 0, hence there exists an endomorphism $\iota : E \rightarrow E$ with the property $\iota^2 = [-4]$. Since $p \in \mathcal{P}_\ell$, there exist values n, a, b so that $2^n - \ell = a^2 + 4b^2$, and the algorithm sets

$\gamma = [a] + [b]\iota$ which is an endomorphism of degree $2^n - \ell$. Further, a basis (P_{2^n}, Q_{2^n}) for $E[2^n]$ and a basis (P_ℓ, Q_ℓ) for $E[\ell]$ are computed. Note that by assumption $p \equiv 3 \pmod{4}$, hence the elliptic curve E is supersingular and has cardinality $(p+1)^2$. As a consequence, the $\ell \cdot 2^n$ -torsion is \mathbb{F}_{p^2} -rational and it is not necessary to work in field extensions. The cost of the setup is dominated by the cost for computing the N -torsion bases for $N \in \{\ell, 2^n\}$. This can be done by sampling two points $P, Q \in E(\mathbb{F}_{p^2})$, computing $P_N = (p+1)/N \cdot P$, $Q_N = (p+1)/N \cdot Q$ and the Weil pairing (P_N, Q_N) . This is repeated until $e_N(P_N, Q_N)$ has order N .¹¹ This step costs

$$(5) \quad (\log(p) + \log \ell) \cdot M(\mathbb{F}_{p^2})$$

In the last line of the setup, we compute the deformation \mathcal{E} of E with j -invariant $j(\mathcal{E}) = j(E) + \epsilon$. Note that $j(E) = (2 \cdot 3 \cdot 11)^3 \notin \{0, 1728\}$ for $p > 11$, hence such a deformation exists, and it can be computed in time $O(1)$ in $R = \mathbb{F}_{p^2}[\epsilon]/(\epsilon^{\ell+2})$. This means the complexity of this step is simply

$$(6) \quad O(1) \cdot M(R).$$

The main part of the algorithm consists of computing the $\ell + 1$ different ℓ -isogenies emanating from the elliptic curve E ; and lifting these to isogenies emanating from the elliptic curve \mathcal{E} with j -invariant $\tilde{j} = j + \epsilon$. Using the Square-Root Vélú Algorithm, the isogeny computations over the ground field can be done in $\tilde{O}(\ell\sqrt{\ell})$ \mathbb{F}_{p^2} -multiplications. For the overall analysis it is sufficient to use the standard Vélú formulas, which results in

$$(7) \quad O(\ell^2) \cdot M(\mathbb{F}_{p^2})$$

To explain the lifting step in more detail, assume that we are at step k and want to lift the ℓ -isogeny $f_k : E \rightarrow E_k$ to an isogeny with domain \mathcal{E} . To this end, we construct an $(\ell, 2^n - \ell)$ -isogeny diamond as depicted below.

$$\begin{array}{ccc} E & \xrightarrow{f_k} & E_k \\ \downarrow \gamma & & \downarrow \gamma' \\ E & \xrightarrow{f'_k} & E'_k \end{array}$$

This gives rise to a product isogeny $F : E \times E'_k \rightarrow E \times E_k$ with kernel

$$K = \langle (\hat{\gamma}(P_{2^n}), f'_k(P_{2^n})), (\hat{\gamma}(Q_{2^n}), f'_k(Q_{2^n})) \rangle.$$

On input the vertices of the isogeny diamond E, E_k, E, E'_k and the lift \mathcal{E} together with the kernel K , the method `lift_isogeny_diamond` (Algorithm 4) outputs the vertices of the lifted isogeny diamond $(\mathcal{E}, \mathcal{E}_k, \mathcal{E}', \mathcal{E}'_k)$. In particular, this contains the lift \mathcal{E}_k of E_k such that f_k lifts to an isogeny $\tilde{f}_k : \mathcal{E} \rightarrow \mathcal{E}'_k$ over $\mathbb{F}_{p^2}[\epsilon]/(\epsilon^{\ell+2})$. We save its j -invariant \tilde{j}_k . Note that since we are lifting a $(2^n, 2^n)$ -chain here, we can use an improved lifting strategy in `lift_isogeny_diamond` as described in Subsection 4.4. For each ℓ -isogeny this gives us complexity $O(n)$ in R , hence overall the lifting step for all ℓ -isogenies costs

$$(8) \quad O(n\ell) \cdot M(R).$$

In the final step, the product

$$\varphi_\ell(j(E) + \epsilon, Y) = \prod_{k=0}^{\ell} (Y - \tilde{j}_k) \in R[Y].$$

is computed and evaluated at $\epsilon = X - j(E)$. This yields the modular polynomial $\varphi_\ell(X, Y)$ (modulo $(X - j(E))^{\ell+2}$). Given that $\varphi_\ell(X, Y)$ has degree $\ell + 1$, we may ignore the modulus and we conclude

¹¹There are better methods to compute a basis in practice. In particular, in the case of $N = 2^n$, one should use the methods developed in the framework of SIDH key compression [9].

that the output is the modular polynomial $\varphi_\ell(X, Y)$ in \mathbb{F}_p . The evaluation of the product can be done via a product-tree strategy. Here, the last step consisting of the multiplication of two polynomials of degree $O(\ell)$ is dominating. In our setting, the complexity for this multiplication is simply given by

$$(9) \quad O(\ell) \cdot M(R).$$

Taking the sum over the complexities from Equations 5 - 9, we obtain a complexity of

$$O(\log p + \ell^2) \cdot M(\mathbb{F}_{p^2}) + O(n\ell) \cdot M(R).$$

□

5.3. Generalization to $\ell \equiv 1 \pmod{4}$. In Algorithm 6 we described a method to compute the modular polynomial φ_ℓ in $\mathbb{F}_p[X, Y]$ for primes in \mathcal{P}_ℓ , when $\ell \equiv 3 \pmod{4}$. In this part, we sketch a modification of the algorithm which allows us to compute the modular polynomial for $\ell \equiv 1 \pmod{4}$.

The different treatment stems from the fact that in this case, there do not exist integers a, b, n so that $2^n - \ell = a^2 + 4b^2$. Instead we consider the equation $2^n - 3 \cdot \ell = a^2 + 4b^2$, and extend the ℓ -isogenies by an auxiliary 3-isogeny.

Theorem 5.5. *There exists a modification of Algorithm 6 which on input a prime $\ell \equiv 1 \pmod{4}$ and a prime $p \in \mathcal{P}_\ell$, returns the modular polynomial ϕ_ℓ over \mathbb{F}_p . The algorithm runs in*

$$O(\log p + \ell^2) \cdot M(\mathbb{F}_{p^2}) + O(n\ell) \cdot M(R),$$

with $R = \mathbb{F}_{p^2}(\epsilon)/(\epsilon^{\ell+2})$, and $M(\cdot)$ as defined in Subsection 4.1.

Proof. We outline the modifications of Algorithm 6. It will be clear that these changes have no effect on the asymptotic complexity of the algorithm.

In the setup, we choose an endomorphism γ as $\gamma = [a] + [b]\iota$ with $2^n - 3 \cdot \ell = a^2 + 4b^2$. And we also compute an auxiliary degree-3 isogeny $g : E \rightarrow C_0$.

The main change occurs in the construction of the isogeny diamond. In order to make γ and the ℓ -isogeny f_k fit into an isogeny diamond (with smooth induced product isogeny), it is necessary to expand f_k by the auxiliary isogeny g to obtain the following configuration.

$$(10) \quad \begin{array}{ccccc} C_0 & \xleftarrow{g} & E & \xrightarrow{f_k} & E_k \\ \downarrow \gamma_0 & & \downarrow \gamma & & \downarrow \gamma' \\ C_1 & \xleftarrow{g'} & E & \xrightarrow{f'_k} & E'_k \end{array}$$

The outer square is a $(3 \cdot \ell, 2^n - 3 \cdot \ell)$ -isogeny diamond and induces the product isogeny $F : C_0 \times E'_k \rightarrow C_1 \times E_k$ with kernel

$$K = \left\langle \left(\hat{\gamma}_0(P_{2^n}), f'_k \circ \hat{g}'(P_{2^n}) \right), \left(\hat{\gamma}_0(Q_{2^n}), f'_k \circ \hat{g}'(Q_{2^n}) \right) \right\rangle,$$

where (P_{2^n}, Q_{2^n}) is a basis for $C_1[2^n]$.

Recall that the goal is to lift the ℓ -isogeny f_k , for some fixed lift \mathcal{E} of E . For this purpose, one first computes the lift $\tilde{g} : \mathcal{E} \rightarrow C_0$ explicitly. Then we call `lift_isogeny_diamond` on input $(C_0, E_k, C_1, E'_k, K, C_0)$. The output contains the lift \mathcal{E}_k , and the algorithm proceeds as in the case $\ell \equiv 3 \pmod{4}$. □

The modified version of Algorithm 6 which is outlined in the proof above, is presented in Appendix A (Algorithm 8).

5.4. **The general case.** In this section, we give a quick overview of the proof of Corollary 1.4. If we start with an HD representation of the ℓ -isogenies starting from E_0 , we can deform them to $k[\epsilon]/(\epsilon^{m+1})$ using Theorem 1.3 as explained in Section 4.5. When $m > \ell$, this is enough to reconstruct $\Phi_\ell \pmod p$ in quasi-linear time as in the proof of Theorem 5.4.

In the previous two subsections, we provided algorithms to compute the modular polynomial ϕ_ℓ over a prime field \mathbb{F}_p , where p was an element in \mathcal{P}_ℓ . The number of primes of this form relies on Heuristic 5.1 which describes the size of n , for which $2^n - \ell$ can be written as a sum of two squares. Using a dimension 8 embedding instead of a dimension 2 embedding, as explained in Section 4.5, we generalize the techniques to a larger set of primes by expressing $2^n - \ell$ as a sum of four squares. More explicitly, this results in an algorithm for all primes p in

$$\mathcal{P}_\ell^* = \{p > 11 \text{ prime} : 2^n \cdot \ell \mid p + 1, \text{ where } n = \lceil \log_2(\ell) \rceil\}$$

of complexity $O(\log^2 p + \ell^2 \log p)$.

6. COMPUTING THE MODULAR POLYNOMIAL

Here, we present our algorithm for computing the modular polynomial $\Phi_\ell \in \mathbb{Z}[X, Y]$ of elliptic curves. The procedure is summarized in Algorithm 7. It is a CRT based approach, where the modular polynomial is computed modulo many small primes using Algorithm 6 or its modification described in Subsection 5.3. The runtime of the algorithm relies on Heuristic 5.1, but we prove that there also exists an unconditional version with the same asymptotic runtime.

Algorithm 7 modular_polynomial

Input: An odd prime ℓ .

Output: The modular polynomial $\Phi_\ell(X, Y) \in \mathbb{Z}[X, Y]$.

```

1:  $B \leftarrow 6\ell \log(\ell) + 16\ell + \min(2\ell, 14\sqrt{\ell} \log \ell) + \log(2)$ 
2:  $P \leftarrow 1, \Phi_\ell \leftarrow 0.$ 
3:  $c_\ell \leftarrow -\ell \pmod{4}.$ 
4:  $n \leftarrow \min\{n \mid \exists a, b : 2^n - c_\ell \cdot \ell = a^2 + 4b^2\}$ 
5: while  $P < \exp(B)$  do
6:    $p \leftarrow$  next prime with  $2^n \cdot c_\ell \cdot \ell \mid p + 1$ 
7:    $\varphi_\ell \leftarrow$  modular_polynomial_modp( $\ell, p$ )  $\in \mathbb{F}_p[X, Y]$ 
8:    $P \leftarrow P \cdot p$ 
9:    $\Phi_\ell \leftarrow$  CRT( $\Phi_\ell, \varphi$ )  $\in \mathbb{Z}/P\mathbb{Z}[X, Y]$ 
10: end while
11: return  $\Phi$ 
```

Theorem 6.1. *On input an odd prime ℓ , Algorithm 7 computes the modular polynomial $\Phi_\ell \in \mathbb{Z}[X, Y]$. Under Heuristic 5.1, the algorithm runs in time*

$$O(\ell^3 \log^3(\ell) \log \log(\ell)).$$

Proof. The bound

$$B = 6\ell \log(\ell) + 16\ell + \min(2\ell, 14\sqrt{\ell} \log \ell) + \log(2)$$

in the first line of the algorithm is equal to the bound on the logarithmic height of the coefficients of Φ_ℓ (see Eq. 2) plus $\log(2)$. Consequently, the coefficients of Φ_ℓ are uniquely determined by their residues modulo $\exp(B)$.

In Line 4, a minimal integer n is chosen for which there exists a pair (a, b) so that $2^n - c_\ell \cdot \ell = a^2 + 4b^2$. Under Heuristic 5.1, we may assume that $n = O(\log(\ell))$.

In the main part of the algorithm, the modular polynomial $\phi_\ell \in \mathbb{F}_p[X, Y]$ is computed for various suitable primes $p \in P_\ell$, until the modulus P is at least $\exp(B)$. At each step, we update $P = P \cdot p$ and compute $\Phi_\ell \in \mathbb{Z}/P\mathbb{Z}$ using an explicit version of the Chinese Remainder Theorem.

As per Theorems 5.4 and 5.5, the computation of $\phi_\ell \in \mathbb{F}_p[X, Y]$ is done in time

$$O(\log p + \ell^2) \cdot M(\mathbb{F}_{p^2}) + O(n\ell) \cdot M(R),$$

with $R = \mathbb{F}_{p^2}(\epsilon)/(\epsilon^{\ell+2})$. Further, we may assume that $n = O(\log(\ell))$ under Heuristic 5.1. Note that $B = O(\ell \log(\ell))$, hence it suffices to compute ϕ_ℓ for $O(p)$ many primes of size $\log(p) \in O(\log(\ell))$. Lemma 5.3 assures that there are enough primes of the desired form in \mathcal{P}_ℓ . In this setting,

$$M(\mathbb{F}_{p^2}) = O(\log \ell \log \log \ell \log \log \log \ell), \quad M(R) = O(\ell \log^2 \ell \log \log \ell \log \log \log \ell).$$

In conclusion, we obtain

$$O(\ell^3 \log^3(\ell) \log \log(\ell))$$

for the overall runtime of Algorithm 7. □

Theorem 6.2. *There exists an algorithm for computing the modular polynomial Φ_ℓ for any prime ℓ , with unconditional runtime*

$$O(\ell^3 \log^3(\ell) \log \log(\ell)).$$

Proof. We use the same proof as in Theorem 6.1, with the exception that P_ℓ is replaced by the set of primes \mathcal{P}_ℓ^* defined in Section 5.4. This requires to use dimension 8 embeddings instead of dimension 2 embeddings. However, the advantage of \mathcal{P}_ℓ^* is that we can set $n = \lceil \log_2(\ell) \rceil$ in the statement of Lemma 5.3 without heuristic assumptions. □

7. IMPLEMENTATION

A proof-of-concept implementation of Algorithm 7 (with $\ell \equiv 3 \pmod{4}$) in SageMath [37] is available in our GitHub repository [24]. The repository further contains the different subroutines for computing with deformations of elliptic curves and isogenies presented in this paper.

Our implementation works with elliptic curves in Montgomery form, and we use the available functions in SageMath to compute elliptic curve isogenies. The computation of $(2^n, 2^n)$ -isogenies is based on the formulas from [23], enhanced by explicit formulas for splitting and gluing isogenies. The individual $(2, 2)$ -isogenies in this framework are naturally represented by radical formulas which facilitates an efficient computation of the deformation of the isogeny chains.

In the future, we plan to switch to the faster theta formulas from [11] to compute the $(2^n, 2^n)$ -isogenies and then use radical $(2, 2)$ -isogenies in the theta model for computing the deformation of these isogenies.

REFERENCES

- [1] Andrea Basso, Luciano Maino, and Giacomo Pope. FESTA: fast encryption from supersingular torsion attacks. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology - ASIACRYPT 2023 - 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4-8, 2023, Proceedings, Part VII*, volume 14444 of *Lecture Notes in Computer Science*, pages 98–126. Springer, 2023.
- [2] Daniel Bernstein and Jonathan Sorenson. Modular exponentiation via the explicit chinese remainder theorem. *Mathematics of Computation*, 76(257):443–454, 2007.
- [3] Daniel J Bernstein, Luca De Feo, Antonin Leroux, and Benjamin Smith. Faster computation of isogenies of large prime degree. *Open Book Series*, 4(1):39–55, 2020.
- [4] Florian Breuer, Desirée Gijón Gómez, and Fabien Pazuki. Explicit bounds on the coefficients of the modular polynomials and the size of $x_0(n)$. *arXiv preprint arXiv:2310.14428*, 2023.
- [5] Reinier Bröker, Kristin Lauter, and Andrew Sutherland. Modular polynomials via isogeny volcanoes. *Mathematics of Computation*, 81(278):1201–1231, 2012.
- [6] Reinier Bröker and Andrew V Sutherland. An explicit height bound for the classical modular polynomial. *The Ramanujan Journal*, 22:293–313, 2010.

- [7] Wouter Castryck and Thomas Decru. Multiradical isogenies. *Arithmetic, Geometry, Cryptography, and Coding Theory*, 779:57–89, 2021.
- [8] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 423–447. Springer, 2023.
- [9] Craig Costello, David Jao, Patrick Longa, Michael Naehrig, Joost Renes, and David Urbanik. Efficient compression of sidh public keys. In *Advances in Cryptology–EUROCRYPT 2017: 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30–May 4, 2017, Proceedings, Part I 36*, pages 679–706. Springer, 2017.
- [10] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQISignHD: New dimensions in cryptography. *IACR Cryptol. ePrint Arch.*, page 436, 2023.
- [11] Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien Robert. An algorithmic approach to $(2, 2)$ -isogenies in the theta model and applications to isogeny-based cryptography. 11 2023.
- [12] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *Journal of Mathematical Cryptology*, 8(3):209–247, 2014.
- [13] R. Dupont. Moyenne arithmetico-geometrique, suites de Borchardt et applications. *These de doctorat, Ecole polytechnique, Palaiseau*, 2006.
- [14] Andreas Enge. Computing modular polynomials in quasi-linear time. *Mathematics of Computation*, 78(267):1809–1824, 2009.
- [15] Robin Hartshorne. *Deformation theory*, volume 257. Springer, 2010.
- [16] David Harvey, Joris Van Der Hoeven, and Grégoire Lecerf. Faster polynomial multiplication over finite fields. *Journal of the ACM (JACM)*, 63(6):1–23, 2017.
- [17] David Harvey and Joris Van Der Hoeven. Polynomial multiplication over finite fields in time. *Journal of the ACM (JACM)*, 69(2):1–40, 2022.
- [18] Henryk Iwaniec and Emmanuel Kowalski. *Analytic number theory*, volume 53. American Mathematical Soc., 2021.
- [19] Ernst Kani. The number of curves of genus two with elliptic differentials. 1997.
- [20] Jean Kieffer. Evaluating modular polynomials in genus 2. 2020.
- [21] Jean Kieffer. Sign choices in the AGM for genus two theta constants. Accepted à Publications Mathématiques de Besançon, 2020.
- [22] Jean Kieffer. Certified newton schemes for the evaluation of low-genus theta functions. 2022.
- [23] Sabrina Kunzweiler. Efficient computation of $(2^n, 2^n)$ -isogenies. *Cryptology ePrint Archive*, 2022.
- [24] Sabrina Kunzweiler and Damien Robert. Computing modular polynomials by deformation. <https://github.com/sabrinakunzweiler/modular-polynomials>, 2024.
- [25] Antonin Leroux. Computation of hilbert class polynomials and modular polynomials from supersingular elliptic curves. *arXiv preprint arXiv:2301.08531*, 2023.
- [26] David Lubicz and Damien Robert. Fast change of level and applications to isogenies. *Research in Number Theory (ANTS XV Conference)*, 9(1), 12 2022.
- [27] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery on sidh. In *Eurocrypt*. Springer-Verlag, 2023.
- [28] Enea Milio. A quasi-linear time algorithm for computing modular polynomials in dimension 2. *LMS Journal of Computation and Mathematics*, 18(1):603–632, 2015.
- [29] Aurel Page and Damien Robert. Introducing clapoti (s): Evaluating the isogeny class group action in polynomial time. 2023.
- [30] Damien Robert. Efficient algorithms for abelian varieties and their moduli spaces, 2021.
- [31] Damien Robert. Evaluating isogenies in polylogarithmic time. 2022.
- [32] Damien Robert. Some applications of higher dimensional isogenies to elliptic curves (overview of results). *Cryptology ePrint Archive*, 2022.
- [33] Damien Robert. Breaking SIDH in polynomial time. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 472–503. Springer, 2023.
- [34] Éric Schost. Computing parametric geometric resolutions. *Applicable Algebra in Engineering, Communication and Computing*, 13(5):349–393, 2003.
- [35] Edoardo Sernesi. *Deformations of algebraic schemes*, volume 334. Springer Science & Business Media, 2007.
- [36] John T Tate. p -divisible groups. In *Proceedings of a Conference on Local Fields: NUFFIC Summer School held at Driebergen (The Netherlands) in 1966*, pages 158–183. Springer, 1967.

- [37] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 10.0)*, 2024. <https://www.sagemath.org>.
- [38] Jacques Vélou. Isogénies entre courbes elliptiques. *Comptes-Rendus de l'Académie des Sciences*, 273:238–241, 1971.
- [39] Benjamin Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1100–1111. IEEE, 2022.

APPENDIX A. COMPUTING $\phi_\ell \in \mathbb{F}_p[x, y]$, WHEN $\ell \equiv 1 \pmod{4}$

Here, we present Algorithm 8, a variant of Algorithm 6, for computing the modular polynomial ϕ_ℓ over \mathbb{F}_p for $p \in \mathcal{P}_\ell$ and $\ell \equiv 1 \pmod{4}$. These modifications were outlined in Subsection 5.3. In particular, we obtain the following explicit version of Theorem 5.5.

Theorem A.1. *On input a prime $\ell \equiv 1 \pmod{4}$ and a prime $p \in \mathcal{P}_\ell$, Algorithm 8 returns the modular polynomial ϕ_ℓ over \mathbb{F}_p . The algorithm runs in time $O(\log(p) + \ell^2 \log^2(\ell) + n \ell^2 \log(\ell))$ over \mathbb{F}_{p^2} .*

Proof. This coincides with the proof of Theorem 5.5. □

INRIA BORDEAUX, INSTITUT DE MATHÉMATIQUES DE BORDEAUX

Algorithm 8 modular_polynomial_modp $\ell \equiv 1 \pmod{4}$

Input: A prime $\ell \equiv 1 \pmod{4}$, and a prime $p \in \mathcal{P}_\ell$.

Output: The modular polynomial $\varphi_\ell(X, Y) \in \mathbb{F}_p[X, Y]$.

```
  /*Setup*/
1: Set  $E : y^2 = x^3 + 6x^2 + x$  over  $\mathbb{F}_{p^2}$ 
2:  $\iota \leftarrow \iota \in \text{End}(E)$  with  $\iota \circ \iota = [-4]$ 
3:  $\gamma \leftarrow [a] + [b]\iota$ , where  $2^n - 3 \cdot \ell = a^2 + 4b^2$  for some  $n$ .
4: Sample  $P_3 \in E[3] \setminus \{\mathcal{O}\}$ 
5:  $C_0 \leftarrow E/\langle P_3 \rangle$ ,  $C_1 \leftarrow E/\langle \gamma(P_3) \rangle$ 
6:  $\gamma_0, g, g' \leftarrow$  isogenies with  $\gamma_0 \circ g = g' \circ \gamma$  as in Eq. 10
7: Compute  $P_{2^n}, Q_{2^n}$  with  $C_1[2^n] = \langle P_{2^n}, Q_{2^n} \rangle$ 
8: Compute  $P_\ell, Q_\ell$  with  $E[\ell] = \langle P_\ell, Q_\ell \rangle$ 
9:  $\tilde{j} \leftarrow j(E) + \epsilon \in \mathbb{F}_{p^2}[\epsilon]/(\epsilon^{\ell+2})$ 
10:  $\mathcal{E} \leftarrow$  elliptic curve with  $j(\mathcal{E}) = \tilde{j}$ 
11:  $\mathcal{C}_0 \leftarrow$  codomain of the lift  $\tilde{g} : \mathcal{E} \rightarrow \mathcal{C}_0$ 
  /*computing and lifting all  $\ell$ -isogenies*/
12: for  $k \leftarrow 0, \dots, \ell$  do
13:   if  $k = \ell$  then
14:      $P_k \leftarrow Q_\ell$ 
15:   else
16:      $P_k \leftarrow P_\ell + k \cdot Q_\ell$ 
17:   end if
18:    $E_k \leftarrow E/\langle P_k \rangle$  with  $f_k : E \rightarrow E_k$ 
  /*constructing the  $(3\ell, 2^n - 3\ell)$ -isogeny diamond*/
19:    $E'_k \leftarrow E/\langle \gamma(P_k) \rangle$  with  $f'_k : E \rightarrow E'_k$ 
20:    $K \leftarrow \langle (\hat{\gamma}_0(P_{2^n}), f'_k \circ \hat{g}'(P_{2^n})), (\hat{\gamma}_0(Q_{2^n}), f'_k \circ \hat{g}'(Q_{2^n})) \rangle$ 
21:    $(\mathcal{C}_0, \mathcal{E}_k, \mathcal{C}_1, \mathcal{E}'_k) \leftarrow \text{lift\_isogeny\_diamond}(C_0, E_k, C_1, E'_k, K, \mathcal{C}_0)$ 
22:    $\tilde{j}_k \leftarrow j(\mathcal{E}_k)$ 
23: end for
  /*final step*/
24:  $\varphi \leftarrow \prod_{k=0}^{\ell} (Y - \tilde{j}_k)(\epsilon = X - j(E)) \in \mathbb{F}_p[X, Y]$ 
25: return  $\varphi$ 
```
