HENSEL-LIFTING BLACK-BOX ALGORITHMS AND FAST TRACE COMPUTATION FOR ELLIPTIC-CURVE ENDOMORPHISMS

LORENZ PANNY[†], DAMIEN ROBERT[*], AND ALESSANDRO SFERLAZZA[†]

Abstract. We demonstrate a general and efficient technique to Hensel-lift a solution to a system of (p-adically analytic) equations which may be given *implicitly* in the form of an efficient evaluation algorithm. Contrary to textbook Hensel lifting, we do not require the equations to be represented explicitly; indeed, our main application uses the method for a system of equations that can be exponentially larger than its representation as an arithmetic circuit: We show how to compute traces of elliptic-curve endomorphisms over a finite field $\mathbb{F}_q$ by constructing an (approximate) lift to $\mathbb{Z}_q$. Our examples include endomorphisms represented as a chain of Vélu, $\sqrt{\text{élu}}$, modular, or radical isogenies, as well as HD-embedded endomorphisms. The resulting trace-computation algorithm outperforms the state of the art both asymptotically and concretely.

1. Introduction

Hensel lifting is an essential tool in computer algebra: In its most basic form, it lifts any simple root $\overline{\alpha} \in \mathbb{Z}/p$ of a polynomial $f \in \mathbb{Z}[x]$ to a root $\alpha \in \mathbb{Z}/p^k$ of f such that $\alpha \bmod p = \overline{\alpha}$. This immediately generalizes to solutions of sufficiently well-behaved *systems* of equations. The technique is most naturally viewed in the context of p-adic analysis, where it is equivalent to Newton's method, and "closer" means "congruent modulo higher powers of p".

Classical descriptions (and typical implementations) of Hensel's method assume that the system of equations under consideration is not just known, but also small enough to be represented explicitly (say, on a computer). In this work, we demonstrate that this is not needed: To apply Hensel lifting to a system $F(x) = 0$ with an analytic map $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$, all that is required is an explicit *evaluation algorithm* for F which need not admit a succinct formula at all. One particular case covering many practical situations is when F is given as an arithmetic circuit.

Our main application concerns the lifting of elliptic-curve endomorphisms from $\mathbb{F}_q$ to $\mathbb{Z}_q$ using Hensel's method. This gives rise to a new algorithm for computing the *trace* of an elliptic-curve endomorphism, very similar in spirit to Satoh's algorithm [Sat00] for point counting. We show that our approach can be applied to all commonly used representations of elliptic-curve endomorphisms (chains of Vélu, $\sqrt{\text{élu}}$, and radical isogenies, as well as HD-embedded endomorphisms) by expressing the properties required from a correct lift via a system of algebraic constraints on the data defining the isogeny.

In the process, we also determine the action on invariant differentials (the "isogeny scaling factor") induced by those various isogeny representations by providing a generic method to compute it, requiring only access to a sufficiently general (e.g.,

[†]Technische Universität München, Germany
[*]Centre Inria de l'Université de Bordeaux, France
E-mail addresses: lorenz@yx7.cc, damien.robert@inria.fr, alessandro.sferlazza@tum.de.
Authors listed alphabetically; see https://ams.org/profession/leaders/CultureStatement04.pdf.
Date of this document: 2026-02-25.

algebraic) evaluation algorithm for the isogeny. In addition, it is also usually possible to explicitly determine the scaling factors for a given isogeny representation theoretically; see Appendix A.

Our new algorithm for trace computation resulting from these techniques is significantly faster than the current state of the art [MPSW25] both in theory (with a quasiquadratic complexity in typical situations, rather than quasiquartic) and in practice. It is also more general, since contrary to [MPSW25] we do not assume supersingularity. (Note, however, that the standard SEA algorithm *heuristically* achieves the same complexity as [MPSW25] even for ordinary curves.)

Thanks. The original formulas for the 2-radical and modular 2-radical isogenies in the Montgomery model we used in Examples A.1 and A.6 were more complicated than the current version. Their simplification is due to Sabrina Kunzweiler.

2. Hensel lifting

2.1. Preliminaries: p-adic fields.

To introduce p-adic lifting, we start by recalling the basic construction of p-adic integers and fix some relevant notation.

Theoretical construction. Let p be a prime. We denote by $\mathbb{Z}_p = \varprojlim_n \mathbb{Z}/p^n\mathbb{Z}$ the ring of p-adic integers, and by $\mathbb{Q}_p = \operatorname{Frac}\mathbb{Z}_p$ its fraction field, the *field of p-adic numbers*. Elements α of $\mathbb{Q}_p$ can be identified with formal Laurent series $\sum_{n=n_0}^{\infty} a_n p^n$ with all coefficients $a_n \in \{0, ..., p-1\}$. The (p-adic) valuation $\nu_p(\alpha) \in \mathbb{Z} \cup \{\infty\}$ of α is the smallest $n \in \mathbb{Z}_{\geq n_0}$ such that $a_n \neq 0$, and ∞ if no such n exists.

The p-adic valuation extends uniquely to any degree-d extension $K/\mathbb{Q}_p$ via the norm map $N_{K/\mathbb{Q}_p} \colon K \to \mathbb{Q}_p$; concretely, we have $\nu_p(\alpha) = \frac{1}{d}\nu_p(N_{K/\mathbb{Q}_p}(\alpha))$ for $\alpha \in K$. The ring of integers of K is the subring $\mathcal{O}_K = \{\alpha \in K \mid \nu_p(\alpha) \geq 0\}$; it is a local ring with (unique) maximal ideal $\mathfrak{p}_K = \{\alpha \in K \mid \nu_p(\alpha) > 0\}$. The extension $K/\mathbb{Q}_p$ is *unramified* if $\mathcal{O}_K \,/\, \mathfrak{p}_K \cong \mathbb{F}_{p^m}$; in that case, we have $\mathfrak{p}_K = p\,\mathcal{O}_K$. There exists an (up to $\mathbb{Q}_p$-isomorphism) unique unramified extension of $\mathbb{Q}_p$ for every degree $m \in \mathbb{Z}_{\geq 1}$, which we will denote by $\mathbb{Q}_q$, where $q = p^m$. The ring of integers $\mathcal{O}_{\mathbb{Q}_q}$ is denoted by $\mathbb{Z}_q$. Letting $C \subseteq \mathbb{Z}_q$ with $0 \in C$ denote a complete set of representatives of the quotient ring $\mathbb{Z}_q/p\mathbb{Z}_q$, every element α of $\mathbb{Q}_q$ can be identified with a formal Laurent series $\sum_{n=n_0}^{\infty} a_n p^n$ with all $a_n \in C$. As before, $\nu_p(\alpha)$ equals the smallest n with $a_n \neq 0$. Concretely, thanks to the isomorphism $\mathbb{Z}_q/p\mathbb{Z}_q \cong \mathbb{F}_q$, the "digits" in C can be represented as elements of the residue field $\mathbb{F}_q$: We will employ this abuse of notation throughout and view $\alpha \in \mathbb{Q}_q$ as a formal Laurent series in p "with coefficients in $\mathbb{F}_q$". Under this viewpoint, the projection from $\mathbb{Z}_q$ to $\mathbb{Z}_q/p\mathbb{Z}_q = \mathbb{F}_q$ is given by simply extracting the constant coefficient.

For practical computations, elements α of $\mathbb{Z}_q$ are commonly truncated to some finite precision $k \in \mathbb{Z}_{\geq 1}$: Formally, this means working with the coset $\alpha + p^k\mathbb{Z}_q$ of α in the quotient $\mathbb{Z}_q/p^k\mathbb{Z}_q$. We also introduce the suggestive notation $\alpha + O(p^k)$ for $\alpha + p^k\mathbb{Z}_q$, and we will (by abuse of notation) write "$\alpha + O(p^k) \in \mathbb{Z}_q$" to indicate that α is only determined as a coset representative modulo $p^k\mathbb{Z}_q$. A *lift* of a truncated element $\overline{\alpha} + O(p^k) \in \mathbb{Z}_q$ from precision k to precision $k' \geq k$ refers to any $\alpha + O(p^{k'}) \in \mathbb{Z}_q$ such that $\alpha + O(p^k) = \overline{\alpha} + O(p^k)$. Conversely, in the same situation, $\overline{\alpha} + O(p^k)$ is a *truncation* of $\alpha + O(p^{k'})$ from precision k' to precision k.

The p-adic topology. The *p-adic absolute value* of an element $\alpha \in \mathbb{Q}_q$ is defined as $|\alpha|_p = p^{-\nu_p(\alpha)}$ for $\alpha \neq 0$ and $|0|_p = 0$. The field $\mathbb{Q}_q$ is complete with respect to this (non-Archimedean) absolute value, which enables the use of concepts from analysis: In particular, one can define analytic functions and Taylor expansions exactly like

in the Archimedean case. Throughout this paper, terminology from analysis will always be understood to refer to the topology induced by the p-adic absolute value.

Arithmetic complexity. Let $q = p^m$. As customary, $\mathbb{F}_q$ is constructed as a degree-m algebraic extension $\mathbb{F}_p(\omega)$, hence an element $a \in \mathbb{F}_q$ is represented as a polynomial over $\mathbb{F}_p$ of degree $< m$. Consequently, representing an element $\alpha + O(p^k) \in \mathbb{Z}_q$ as a sequence of k coefficients in $\mathbb{F}_q$ takes $O(k \cdot m \log p)$ bits of space. Using asymptotically fast multiplication, the cost of arithmetic in $\mathbb{Z}_q$ lies in $\widetilde{O}(k \cdot m \log p)$: Indeed, viewing $\mathbb{Z}_q$ as a degree-m algebraic extension of $\mathbb{Z}_p$ by choosing a suitable lift of the minimal polynomial of ω to $\mathbb{Z}_p$, each multiplication in $\mathbb{Z}_q/p^k\mathbb{Z}_q$ reduces to a multiplication of polynomials of degree $< m$ over the ring $\mathbb{Z}_p/p^k\mathbb{Z}_p$, followed by a Euclidean division, which also reduces (via Newton iteration [GG13, Theorem 9.6]) to a logarithmic number of multiplications of polynomials of degree $O(m)$ over $\mathbb{Z}_p/p^k\mathbb{Z}_p$. All of this can be done in quasi-linear time following [HHL17, § 9.2].

Notation 2.1. Suppose given $q = p^m$ a prime power. Let M_k denote the maximum bit-operation cost of an arithmetic operation (addition, multiplication, inversion) over $\mathbb{Z}_q/p^k\mathbb{Z}_q$. By the discussion above, we have $\mathsf{M}_k \in \widetilde{O}(k \log q)$.

2.2. **Hensel lifting.** We recall a standard formulation of multivariate Hensel lifting: It essentially consists of Newton's method applied in the p-adic topology, where values divisible by large powers of p are considered small.

Theorem 2.2. *Suppose* $|\mathsf{in}| \geq |\mathsf{out}|$. *Consider an analytic map* $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$ *and let* $\mathrm{D}F(x) \in \mathbb{Z}_q^{\mathsf{out}\times\mathsf{in}}$ *denote its Jacobian matrix at a point* $x \in \mathbb{Z}_q^{\mathsf{in}}$.

Let $\overline{x} + O(p^k) \in \mathbb{Z}_q^{\mathsf{in}}$ *be a zero of the map* F *in precision* k, *i.e.,* $F(\overline{x}) \in O(p^k)$, *and assume* $\mathrm{D}F(\overline{x}) \bmod p$ *has full rank* $|\mathsf{out}|$.

Then there exists a $\delta + O(p^k) \in \mathbb{Z}_q^{\mathsf{in}}$ *such that* $F(\overline{x} + p^k\delta) \in O(p^{2k})$. *Furthermore, the set of such* $\delta + O(p^k)$ *forms a coset of a free rank-*r *submodule of* $(\mathbb{Z}_q/p^k\mathbb{Z}_q)^{\mathsf{in}}$, *where* $r = |\mathsf{in}| - |\mathsf{out}|$.

Proof. A constructive proof is given in Proposition 2.3 below. □

2.3. **Implicit Hensel lifting.** We show how to generalize Hensel lifting to analytic maps that may not be given explicitly, but rather as an efficient algorithm that evaluates the map at inputs lying in a quotient $\mathbb{Z}_q/p^k\mathbb{Z}_q$.

The key idea is that for any representative $\overline{x} \in \mathbb{Z}_q^{\mathsf{in}}$ of $\overline{x} + O(p^k)$ and any $\delta \in \mathbb{Z}_q^{\mathsf{in}}$, we obtain the first-order Taylor expansion

$$(*) \qquad F(\overline{x} + p^k\delta) = F(\overline{x}) + \mathrm{D}F(\overline{x}) \cdot p^k\delta + O(p^{2k})\,.$$

Thus, in double precision $2k$, we can view $F(\overline{x} + p^k\delta)$ as an affine-linear map in δ. As a consequence, the Jacobian matrix $\mathrm{D}F(\overline{x})$ can be computed in precision k by evaluating F at the two elements $\overline{x}$ and $\overline{x} + p^k$ in precision $2k$.

Proposition 2.3 (Implicit Hensel lifting)**.** *Let* $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$ *and* $\overline{x} + O(p^k) \in \mathbb{Z}_q^{\mathsf{in}}$ *be as in Theorem 2.2. Suppose given a (black-box) algorithm* $\mathcal{A}$ *which, given some* $a + O(p^{2k}) \in \mathbb{Z}_q^{\mathsf{in}}$, *outputs* $F(a) + O(p^{2k}) \in \mathbb{Z}_q^{\mathsf{out}}$. *Then a lift* $x + O(p^{2k}) \in \mathbb{Z}_q^{\mathsf{in}}$ *of* $\overline{x} + O(p^k)$ *such that* $F(x) \in O(p^{2k})$ *can be computed using* $|\mathsf{in}| + 1$ *calls to* $\mathcal{A}$ *and by solving a linear system of* $|\mathsf{out}|$ *equations in* $|\mathsf{in}|$ *variables over* $\mathbb{Z}_q/p^k\mathbb{Z}_q$.

Proof. To ease notation, suppose $\mathsf{in} = \{1, ..., m\}$. For each $i \in \mathsf{in}$, let e_i denote the ith basis vector of $\mathbb{Z}_q^m$ over $\mathbb{Z}_q$.

Fix a lift $x_0 + O(p^{2k}) \in \mathbb{Z}_q$ of $\overline{x} + O(p^k)$ and define $x_i := x_0 + p^k e_i \in \mathbb{Z}_q/p^{2k}\mathbb{Z}_q$ for each $i \in \{1, ..., m\}$. Run $\mathcal{A}$ on the $m+1$ elements $x_0, x_1, ..., x_m \in \mathbb{Z}_q/p^{2k}\mathbb{Z}_q$,

yielding results $y_0, y_1, ..., y_m \in \mathbb{Z}_q/p^{2k}\mathbb{Z}_q$. By construction, for all $i \in \{1, ..., m\}$,

$$y_i = F(x_i) \equiv F(x_0) + \mathrm{D}F(x_0) \cdot p^k e_i \pmod{p^{2k}}.$$

The expression $\mathrm{D}F(x_0) \cdot e_i$ equals the ith partial derivative $\partial_i F(x_0)$ of F at x_0, hence by subtracting y_0 from each y_i for $i \in \{1, ..., m\}$, we get the m linear equations

$$y_i - y_0 \equiv p^k \cdot \partial_i F(x_0) \pmod{p^{2k}},$$

or equivalently

$$\partial_i F(x_0) \equiv (y_i - y_0)/p^k \pmod{p^k},$$

allowing us to recover $\mathrm{D}F(x_0)$ in precision k.

Finally, to compute the desired lift $x + O(p^{2k})$ of $\overline{x} + O(p^k)$, we may compute $\delta \in \mathbb{Z}_q/p^k\mathbb{Z}_q$ satisfying $F(\overline{x} + p^k\delta) \in O(p^{2k})$ by substituting the Taylor expansion from (∗), resulting in the final system

$$(\dagger) \qquad \mathrm{D}F(x_0) \cdot \delta \equiv -F(x_0)/p^k \pmod{p^k}.$$

Note that this system is always solvable thanks to the assumption that $\mathrm{D}F(x_0)$ has full rank modulo p (and therefore modulo p^k). The set of solutions is thus a coset of the kernel of $\mathrm{D}F(x_0)$ mod p^k, which is of rank $|\mathsf{in}| - |\mathsf{out}|$. □

2.4. **Hensel lifting for constraint systems.** Fundamentally, our Hensel lifting algorithm is based on an(y) algorithm to evaluate the analytic function $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$ that defines the conditions to be satisfied by the lifted solution. For efficiency, it is often beneficial to split F into multiple separate functions $F_j\colon \mathbb{Z}_q^{X_j} \to \mathbb{Z}_q^{Y_j}$, where the Y_j form a partition of out, and each X_j is a (typically minimal) subset of in such that F restricts to a well-defined map $\mathbb{Z}_q^{X_j} \to \mathbb{Z}_q^{Y_j}$. We call the collection of triplets (X_j, Y_j, F_j) a **constraint system** for F. One special case that occurs in practice is when F is given by rational functions computed by an *algebraic algorithm* (see for instance [BDLS20, § 1.1]), or equivalently by an *arithmetic circuit*; in the latter case, each Y_j can be chosen as a subset of the output nodes and the corresponding X_j as the set of input nodes v for which there exists a path from v to any node in Y_j.

For a given $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$, and $k \in \mathbb{Z}_{\geq 0}$, Proposition 2.3 gives the complexity of lifting an input $\overline{a}$ with $F(\overline{a}) \equiv 0 \pmod{p^k}$ to an input a with $F(a) \equiv 0 \pmod{p^{2k}}$.

Proposition 2.4. *Suppose given an algorithm that evaluates $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$ with $|\mathsf{in}| \geq |\mathsf{out}|$ in precision $2k$ with cost $C_{F,2k}$, as well as $\overline{a} + O(p^k)$ such that $F(\overline{a}) \in O(p^k)$ and $\mathrm{D}F(\overline{a})$ has full rank modulo p.*

Then we can compute $a + O(p^{2k})$ satisfying $a \equiv \overline{a} \pmod{p^k}$ and $F(a) \in O(p^{2k})$ via Proposition 2.3 using $O(|\mathsf{in}| \cdot C_{F,2k} + |\mathsf{in}|^\omega \mathsf{M}_k)$ bit operations, where ω is a linear-algebra exponent.[1]

In some cases, taking into account the structure of F can make lifting much cheaper. Suppose given a constraint system $((X_j, Y_j, F_j))_j$ for F, as defined above, such that $|X_j| \leq m$ for some constant $m \ll |\mathsf{in}|$). In this case, it can be beneficial to lift the outputs sequentially: Computing the system (†) applied to F_j only costs up to $m + 1$ evaluations of F_j (instead of $|\mathsf{in}| + 1$ evaluations), and the resulting linear system is sparse, hence easier to solve.

Using Algorithm 1 and the following proposition, we show that lifting by considering output constraints "one (or few) at a time" leads to a better strategy than Hensel-lifting the whole system at once.

[1] Certainly $\omega \leq 3$ using naïve matrix arithmetic, and there is a trivial lower bound of $2 \leq \omega$. The best currently known algorithms achieve $\omega \approx 2.37...$.

Algorithm 1: Lifting solutions to a constraint system to double precision.

Input: A constraint system $((X_j, Y_j, F_j))_{j\in\{1,\dots r\}}$, for an analytic function $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$ with $|\mathsf{in}| \geq |\mathsf{out}|$; an assignment $\overline{a} + O(p^k) \in \mathbb{Z}_q^{\mathsf{in}}$ with $F(\overline{a}) \in O(p^k)$, where $\mathrm{D}F(\overline{a})$ has full rank modulo p.

Output: A double-precision assignment $a + O(p^{2k})$ satisfying $F(a) \in O(p^{2k})$.

Initialize $H \in (\mathbb{Z}_q/p^k\mathbb{Z}_q)^{\mathsf{out}\times\mathsf{in}}$ and $b \in (\mathbb{Z}_q/p^k\mathbb{Z}_q)^{\mathsf{out}}$ with zeroes.

For j **from** 1 **to** r **do**

  $y_0 \leftarrow F_j(\overline{a}_{X_j}) \bmod p^{2k} \in (\mathbb{Z}_q/p^{2k}\mathbb{Z}_q)^{Y_j}$.

  For $i \in X_j$ **do**

    // e_i has a 1 at position i and 0 at all positions in $X_j\backslash\{i\}$

    $y_i \leftarrow F_j(\overline{a}_{X_j} + p^k e_i) \bmod p^{2k} \in (\mathbb{Z}_q/p^{2k}\mathbb{Z}_q)^{Y_j}$.

    Store $H_{Y_j,i} \leftarrow (y_i - y_0)/p^k \in (\mathbb{Z}_q/p^k\mathbb{Z}_q)^{Y_j}$.

  Store $b_{Y_j} \leftarrow -y_0/p^k \in (\mathbb{Z}_q/p^k\mathbb{Z}_q)^{Y_j}$.

// Now $H = \mathrm{D}F(\overline{a}) + O(p^k)$ and $b = -F(\overline{a})/p^k + O(p^k)$.

Compute an arbitrary solution $\delta \in (\mathbb{Z}_q/p^k\mathbb{Z}_q)^{\mathsf{in}}$ of $H\delta = b$. // Sparse.

Return $\overline{a} + p^k\delta$.

Proposition 2.5. *Suppose given a constraint system $(X_j, Y_j, F_j)_{j\in\{1,\dots,r\}}$ for an analytic function $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$ with $|\mathsf{in}| \geq |\mathsf{out}|$, as well as an input $\overline{a} + O(p^k) \in \mathbb{Z}_q^{\mathsf{in}}$ satisfying $F(\overline{a}) \in O(p^k)$ and such that $\mathrm{D}F(\overline{a})$ has full rank modulo p.*

Denote by $C_{j,2k}$ the cost of evaluating F_j in precision $2k$, and by $C_{\mathrm{sys},k}$ the bit-operation cost of solving an inhomogeneous linear system with matrix $\mathrm{D}F(\overline{a})$ in precision k (note that $\mathrm{D}F(\overline{a})$ has at most $m \cdot |\mathsf{out}|$ nonzero entries). Write $m = \max_{j\in\{1,\dots,r\}} |X_j|$. Then Algorithm 1 computes a lift $a + O(p^{2k})$ of $\overline{a}$, satisfying $F(a) \in O(p^{2k})$, using at most $(m+1)\cdot\sum_j C_{j,2k} + C_{\mathrm{sys},k}$ bit operations.

Proof. Correctness follows from Proposition 2.3, as for Proposition 2.4.

Regarding the complexity: For each $j \in \{1, \dots, r\}$, we perform $|X_j| + 1 \leq m + 1$ evaluations of F_j, and finally we solve the system $\mathrm{D}F(\overline{a}) \cdot \delta = b$ over $\mathbb{Z}_q/p^k\mathbb{Z}_q$. □

We've seen above how to lift a given solution to a constraint system to double p-adic precision. By repeatedly applying this procedure, starting from a solution modulo p, i.e., in precision 1, we can lift it to arbitrarily high precision k.

Corollary 2.6. *Suppose given an algorithm to evaluate $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$, and a tuple $\overline{a} + O(p) \in \mathbb{Z}_q^{\mathsf{in}}$ (or equivalently, a tuple in $\mathbb{F}_q^{\mathsf{in}}$) satisfying $F(\overline{a}) \in O(p)$ and such that $\mathrm{D}F(\overline{a})$ has full rank modulo p. Fix $k = 2^h \in \mathbb{Z}$, and let C_k be the cost of lifting solutions from precision $k/2$ to precision k, as given by Proposition 2.4 and Proposition 2.5. We can compute an arbitrary lift $a + O(p^{2k})$ of $\overline{a}$, satisfying $F(a) \in O(p^{2k})$, with a cost of $O(C_k)$ bit operations.*

Proof. We apply Proposition 2.4 or Proposition 2.5 to repeatedly double the precision from 1 to $2, 4, \dots, 2^h$. Given that the bit-operation cost of arithmetic modulo p^k grows superlinearly in k, each lifting step costs at most half as much as the next one. In particular, we have $C_k + C_{k/2} + C_{k/4} + \dots + C_2 \in O(C_k)$. □

Remark 2.7. In addition to constraints that are required to be zero for a solution, practical computations often also require that other constraints are *non*zero for a solution. (We call these "closed" resp. "open" constraints, since their solution sets have those respective properties in the Zariski topology.)

Since clearly $\not\equiv 0 \pmod p$ implies $\neq 0 \in \mathbb{Z}_q$, all "open" constraints are trivially satisfied for *any* lift of a solution modulo p. Therefore, we need not pay any attention to "open" constraints during lifting: A lift to $\mathbb{Z}_q$ of any variable that only appears in "open" constraints can be chosen completely arbitrarily and hardcoded into the system of "closed" constraints before the lifting of "closed" constraints commences. (Alternatively, such variables could still be viewed as inputs to F anyway, but they will always contribute a full dimension to the space of lifts.)

Remark 2.8. In the situation of proposition 2.5, if $\mathrm{D}F(\overline{a}) = p^e M$ where the matrix M has full rank modulo p and $e < k$, then we can apply a similar method but we lose a precision e at each Newton iteration. More precisely given $\overline{a} + O(p^k)$ such that $F(\overline{a}) \in O(p^{k+e})$ (note that due to our assumption on $\mathrm{D}F$, $F(\overline{a})$ is well defined modulo p^{k+e}), then a Newton iteration will give $\overline{a} + O(p^{2k-e})$ such that $F(\overline{a}) \in O(p^{2k})$.

Even more general type of Jacobian matrices $\mathrm{D}F(\overline{a})$ can be covered by switching to a lattice model of the p-adic precision [CRV14].

2.5. **A geometric description of Hensel lifting.** In this section, we reinterpret the algebraic Hensel lifting algorithm from a geometric point of view, recasting the language of constraint systems into the language of schemes. The reader only interested in the algorithmic applications can safely skip this section. We refer to Remark 2.14 for an outline of a more general version.

In the context of Hensel lifting, we're given a system of equations over $\mathbb{Z}_q$ of the form $F(x_1, \dots, x_n) = 0$, where $F = (F_1, \dots, F_s)$ is a polynomial map, and a solution $P = (a_1, \dots, a_n)$ of the polynomial system modulo p. Theorem 2.2 tells us how to lift the solution P to higher precision, say to a $\widetilde{P}$ such that $F(P) = 0 \pmod{p^k}$, but requires us to check that the Jacobian matrix has maximal rank. In this section, using the language of schemes, we get geometric tools to check this criterion — and we achieve a greater level of generality.

To fix notations, suppose we have a map of schemes $\varphi\colon X \to Y$, a point $\widetilde{Q} \in Y(\mathbb{Z}/p^k\mathbb{Z})$ which reduces modulo p to $Q \in Y(\mathbb{Z}/p\mathbb{Z})$, and a point $P \in X(\mathbb{Z}/p\mathbb{Z})$ such that $\varphi(P) = Q$. In the setting of Theorem 2.2, X plays the role of the affine scheme $\operatorname{Spec} \mathbb{Z}_q[x_1, \dots, x_n]/(F_1, \dots, F_m)$, and φ is the structure morphism onto $\operatorname{Spec} \mathbb{Z}_q$. Our goal is to describe all lifts $\widetilde{P} \in X(\mathbb{Z}/p^k\mathbb{Z})$ of P such that $\varphi(\widetilde{P}) = Q$.

The Jacobian rank condition: smoothness and étaleness. If the morphism φ is *formally smooth* then a lift as above always exists [Stacks, Tag 02GZ], and if φ is *formally étale* the lift is moreover unique [Stacks, Tag 02HF]. By [Stacks, Tag 02H6, Tag 02HM], $\varphi\colon X \to Y$ is smooth (resp. étale) if and only if it is formally smooth (resp. formally étale) and locally of finite presentation[2]. Note that φ is étale at a point P if and only if it is smooth at P of relative dimension 0 [Stacks, Tag 02GU].

Smoothness and étaleness are Zariski-local conditions; more precisely φ is smooth (resp. étale) at P if and only if we can find small affine opens $U = \operatorname{Spec} R_2 \subset X$ and $V = \operatorname{Spec} R_1 \subset Y$, with $P \in U$ and $\varphi(U) \subset V$, such that $\varphi|_U$ is standard smooth [Stacks, Tag 01V7, Tag 01V9] (resp. standard étale [Stacks, Tag 02GU]). Recall that an affine scheme morphism $\varphi\colon U \to V$ as above is *standard smooth at P of relative dimension r* if and only if there exist $F_1, \dots, F_m \in R_1[X_1, \dots, X_n]$ such that $R_2 \cong R_1[X_1, \dots, X_n]/(F_1, \dots, F_s)$ with $r = n - s$, and the matrix $(\frac{\partial F_j}{\partial X_i}(P))_{i,j \in \{1,\dots,s\}}$ is invertible over R_1 [Stacks, Tag 00T6]. So we see that, up to permutation of the variables, this condition is exactly the rank condition of

[2]Being locally of finite presentation means that φ behaves well with respect to filtered limits [Stacks, Tag 01ZC].

Theorem 2.2. In the case $r = 0$ (i.e., étaleness of φ at P) we can even reduce to a single variable: φ is standard étale if we can find $F, H \in R_1[X]$, with F monic and F' invertible in R_2, such that $R_2 = R_1[X]_H/F(X)$ [Stacks, Tag 02GI, Tag 00UB].

Remark 2.9.

- The description of a standard smooth morphism above shows that a smooth morphism is étale-locally of the form $\mathbb{A}^r_R \to \operatorname{Spec}(R)$ [Stacks, Tag 039P]. Indeed, $\varphi|_U$ decomposes as $U \simeq \operatorname{Spec} R_2 \to \operatorname{Spec} R_1[X_{s+1}, \dots, X_n] \simeq \mathbb{A}^r_{R_1} \to V \simeq \operatorname{Spec} R_1$ where the map $U \to \operatorname{Spec} R_1[X_{s+1}, \dots, X_n]$ is standard smooth of relative dimension 0, hence étale.
- Let $\varphi\colon X/S \to Y/S$ be a S-map of relative smooth schemes X, Y over S, $P \in X$, $Q = \varphi(P)$. Then φ is smooth (resp. étale, resp. unramified) at P if and only if $d\varphi\colon T_{X/S,P} \to T_{Y/S,Q}$ is surjective (resp. bijective, resp. injective), where $Q = \varphi(P)$ [GD67, IV Théorème 17.11.1, Corollaires 17.2.2 et 17.11.2].

Reformulating the lifting theorems. We now specialize the description above to our case of interest. We have X a $\mathbb{Z}_q$-scheme and $\varphi\colon X \to \operatorname{Spec} \mathbb{Z}_q$ the structure morphism, and assume that we are given a point $P \in X(\mathbb{F}_q)$. We will consider the problem of lifting this point to $\widetilde{P} \in X(\mathbb{Z}_q)$. Formally, a lift $\widetilde{P}$ correspond to a dotted morphism in the diagram below, which makes it commutative:

$$\begin{array}{ccc} \operatorname{Spec} \mathbb{F}_q & \xrightarrow{P} & X \\ \downarrow & \overset{\widetilde{P}}{\nearrow} & \downarrow{\scriptstyle \varphi} \\ \operatorname{Spec} \mathbb{Z}_q & = & \operatorname{Spec} \mathbb{Z}_q \end{array}$$

As argued above, the situation is well behaved when φ is smooth at P. The following proposition gives a refinement of Theorem 2.2 where the coset of lifts is given a geometric description.

Proposition 2.10. *Let X be a $\mathbb{Z}_q$-scheme and $\varphi\colon X \to \operatorname{Spec} \mathbb{Z}_q$ its structure morphism. Assume given a point $P \in X(\mathbb{F}_q)$ such that φ is smooth at P of relative dimension r. The point P can be lifted (non uniquely) to a point $\widetilde{P} \in X(\mathbb{Z}_q)$, and the set of such lifts $\widetilde{P}$ forms a torsor under the natural action of $\Gamma(P, T_{X/\mathbb{Z}_q} \otimes p\mathbb{Z}_q) \simeq p\mathbb{Z}_q^r$.*

Proof. Since $\operatorname{Spec} \mathbb{Z}_q$ is the inductive limit of the $\operatorname{Spec} \mathbb{Z}_q/p^k\mathbb{Z}_q$, we reduce to lifting P to a point $\widetilde{P} \in X(\mathbb{Z}_q/p^k\mathbb{Z}_q)$. The fact that we can lift P to $\widetilde{P}$ follows from the fact that φ is formally smooth. To obtain the more precise description of the set of lifts, we use the fact from remark 2.9 that there is an affine open U of X containing P such that $\varphi|_U$ factors as an étale map $\pi\colon U \to \mathbb{A}^r_{\mathbb{Z}_q}$ followed by the smooth projection map $\mathbb{A}^r_{\mathbb{Z}_q} \to \operatorname{Spec}(\mathbb{Z}_q)$.

Since π is étale, lifts of P to $X(\mathbb{Z}_q)$ correspond bijectively to lifts of $\pi(P)$ to $\mathbb{A}^n_{\mathbb{Z}_q}(\mathbb{Z}_q)$, which are parametrized by $\Gamma(\pi(P), T_{\mathbb{A}^r_{\mathbb{Z}_q}/\mathbb{Z}_q} \otimes p\mathbb{Z}_q)$. Since π induces an isomorphism on tangent spaces (by étaleness), we see that the lifts $\widetilde{P}$ are parametrized by $\Gamma(P, T_{X/\mathbb{Z}_q} \otimes p\mathbb{Z}_q)$ (non-canonically, as the parametrization depends on the choice of π). However, the action of $\Gamma(\pi(P), T_{\mathbb{A}^r_{\mathbb{Z}_q}/\mathbb{Z}_q} \otimes p\mathbb{Z}_q)$ on the lifts of $\pi(P)$ induces an action of $\Gamma(P, T_{X/\mathbb{Z}_q} \otimes p\mathbb{Z}_q) \cong p\mathbb{Z}_q^r$ on the lifts of P, and the latter does not depend on the choice of π. This action on the lifts of P that we've just obtained also has a more intrinsic description: See for example the proof of [Stacks, Tag 06JI], setting $\mathcal{F}$ to be the deformation functor associated to X, so that (by [Stacks, Tag 06SS]) $T\mathcal{F}$ is the relative tangent space $T_{X/\mathbb{Z}_q}$. □

Towards a clearer connection between the geometric point of view and the previous sections, we now explicitly describe *in coordinates* the lifts of P whose existence is given by Proposition 2.10. The following Proposition is a reformulation of Proposition 2.3 in geometric language:

Proposition 2.11. *Let $U \subseteq X$ be an open affine subscheme containing the point P, and let $i = (X_1, \dots, X_n) \colon U \to \mathbb{A}^n_{\mathbb{Z}_q}$ denote local affine coordinates around P as above, and $\pi_r = (X_1, \dots, X_r) \colon \mathbb{A}^n_{\mathbb{Z}_q} \to \mathbb{A}^r_{\mathbb{Z}_q}$ the projection onto the first r coordinates. If the structure morphism φ is smooth of relative dimension r at P, up to reindexing of the coordinates we may assume that the differential $d(\pi_r \circ i)$ is injective on the tangent space $T_{X/\mathbb{Z}_q,P}$, so $\pi_r \circ i$ is étale at P.*

Now let $I = (F_1, \dots, F_c)$ be an ideal of definition of $i(X) \cap U \subset U$ in an affine open U of $\mathbb{A}^n_{\mathbb{Z}_q}$ containing $i(P)$. Assume that we have access to a black-box algorithm $\mathfrak{B}$ that returns the value of the $F_i(Q)$ for all $Q \in U$. Given any lift $\widetilde{P'}$ of $P' = \pi_r \circ i(P)$ to p-adic precision m, we can compute the unique lift $i(\widetilde{P})$ of $i(P)$ such that $\pi_r \circ i(\widetilde{P}) = \widetilde{P'}$ to p-adic precision k in $O(\log k)$ call to the black-box algorithm $\mathfrak{B}$ on points living in $\mathbb{A}^n_{\mathbb{Z}_q}(\mathbb{Z}_q/p^k\mathbb{Z}_q)$ along with linear algebra on $(\mathbb{Z}_q/p^k\mathbb{Z}_q)^n$.

Example 2.12. Anticipating Section 3, we give two examples, related to elliptic curves, where the lifts are naturally parametrized by a space of dimension 1.

(1) Let $E \colon y^2 = x^3 + ax + b$ be an elliptic curve over $\mathbb{F}_q$, and $j(E)$ its j-invariant. The coarse moduli space of elliptic curves is isomorphic to $\mathbb{A}^1$, via the j-invariant, so any $\widetilde{j} \in \mathbb{Z}_q$ lifting $j(E)$ is the j-invariant of an elliptic curve $\widetilde{E}/\mathbb{Z}_q$ lifting E.

Let $\mathfrak{A}_1$ be the fine moduli space (over $\mathbb{F}_q$) of elliptic curves as defined in [Stacks, Tag 072K]. By definition, to give a map $\varphi \colon S \to \mathfrak{A}_1$ is the same as giving an elliptic curve E_φ/S. Via this definition, the map $\mathrm{Id} : \mathfrak{A}_1 \to \mathfrak{A}_1$ defines the *universal elliptic curve* $\mathfrak{X}_1/\mathfrak{A}_1$, where $\mathfrak{X}_1$ (like $\mathfrak{A}_1$) is an algebraic Deligne–Mumford stack; any elliptic curve E_φ as above is then obtained as the pullback of $\mathfrak{X}_1/\mathfrak{A}_1$ by φ. If $j(E) \neq 0, 1728$, then, étale-locally above $\widetilde{j}$, the universal elliptic curve is given by a scheme which is smooth of relative dimension 1. In other words, the Weierstraß coefficients $\widetilde{a}, \widetilde{b}$ of the elliptic curves with $\widetilde{E}$ such that $j(\widetilde{E}) = \widetilde{j}$ can be parametrized by a linear space of dimension 1: fixing any of the two coefficients uniquely determines the other.

For curves with j-invariant $j \in \{0, 1728\}$, lifting curve coefficients requires extra care, as these values of j correspond to stacky points in the fine moduli space $\mathfrak{A}_1$. While the universal elliptic curve is smooth of relative dimension 1 above $\mathfrak{A}_1$, it is only represented by a scheme if we specify enough level structure (i.e., twist information) to remove the ambiguity given by the extra automorphisms at these j-invariants.

(2) Let $\widetilde{E} \colon y^2 = x^3 + \widetilde{a}x + \widetilde{b}$ be an elliptic curve over $\mathbb{Z}_q$. Then $\widetilde{E}$ is smooth of relative dimension 1, and $\widetilde{E}$ is naturally embedded into $\mathbb{P}^2_{\mathbb{Z}_q}$. If $P \in \widetilde{E}(\mathbb{F}_q)$ is a point, the lifts of P are parametrized by a linear space of dimension 1. For instance, if $y(P) \neq 0$ (i.e., $P \notin \widetilde{E}(\mathbb{F}_q)[2]$), the projection map $(x, y) \mapsto x$ is étale at P. It follows that the lifts $\widetilde{P} \in E$ of P are parametrized by their x-coordinate: $x(\widetilde{P}) = x(P) + px_1 + p^2x_2 + O(p^3)$. The y-coordinate of $\widetilde{P}$ is then uniquely determined by $y(P)$ and the curve equation.

Remark 2.13. Using the notation of Proposition 2.11, we observe the following:

(1) We can use the methods of Section 2.3 to compute the Jacobian matrix of the generators of I at $i(P)$, provided we have the full set of polynomial constraints locally at P (i.e., a basis of our ideal of equations I). This allows to check if φ is smooth at P, and pick up our choice of coordinates for the projection.

However, in practice, often we know from geometry that φ is smooth of relative dimension r at P, but we do not know if we have a full set of constraints (or we have a very large set of constraints that works generically, but we want a smaller subset for our specific P for efficiency reasons.) In this case, it suffices to add enough constraints until the Jacobian matrix has the correct rank $n-r$, we then know that we have enough local equations around $i(P)$.

(2) Rather than directly embedding $U \to \mathbb{A}^n$, we can first look for an intermediate embedding of $\mathbb{Z}_q$-schemes $i\colon U \to X'$, such that the structure morphism of X' is smooth at $i(P)$, then embed $i'\colon X' \to \mathbb{A}^n$. Then if we have black-box algorithms for evaluating the equations of i' and the equations defining $i(U)$ inside X', then we can apply the tools of Proposition 2.11 to compute the lifts of $i(P)$ that belong to $i(U)$. This corresponds to the formulation of "lift constraints one after another" of Algorithm 1.

As an example, to lift a point of n-torsion $P \in E[n]$ of an elliptic curve, we could either embed $E[n]$ directly into affine space, or decompose this via the natural embedding of $E[n] \to E$ followed by the embedding of E into $\mathbb{P}^2$, and solve first the lifting problem to E, using the algebraic constraint that $\widetilde{P}$ should be on the curve (as in the example above), and then to $E[n]$ by adding a further algebraic constraint that $\widetilde{P}$ should be of n-torsion. The requirement of i being an embedding can be weakened to i being unramified at P, since by [Stacks, Tag 04HG] an unramified morphism is étale-locally an immersion.

(3) Our assumption that we have a map $\varphi\colon X \to \operatorname{Spec}\mathbb{Z}_q$ of schemes is more general that what we need. Indeed, the scheme X, seen as a presheaf on the category of affine schemes, has a well defined functor of points $X(R)$ for any ring R, but in our lifting problem we are only concerned with $X(R)$ for Artinian local $\mathbb{Z}_q$-algebras. The correct setting for this section would thus be the context of smooth morphisms between deformation functors or deformation categories, for which we refer to [Stacks, Tag 06G7]. A deformation category is a predeformation category that satisfies the Rim-Schlessinger conditions [Stacks, Tag 06J9]; this is enough to describe the lifts as a torsor under the tangent space of the deformation category, see [Stacks, Tag 06JI].)

Remark 2.14. In this section we looked at the problem of lifting a point $P \in X(\mathbb{F}_q)$ from characteristic p to a point $\widetilde{P} \in X(\mathbb{Z}_q)$ of characteristic zero. A similar approach would work to compute lifts/deformations of P to $X(\mathbb{F}_q[[\varepsilon]])$ instead. We refer to [KR24] for an algorithmic application.

3. Elliptic curves, endomorphisms, lifts

In this section, we fix notations and recall useful facts concerning elliptic curves and their morphisms, and we discuss the theory of p-adic canonical lifting.

3.1. Elliptic curves and their endomorphisms.

Let k be a field and $a, b \in k$ with $4a^3 + 27b^2 \neq 0$. We consider elliptic curves $E = E_{a,b}$ in short Weierstraß form over k, (we assume $\operatorname{char}(k) \notin \{2, 3\}$ throughout).

Two elliptic curves $E_0 = E_{a_0,b_0}$ and $E_1 = E_{a_1,b_1}$ are isomorphic over k if and only if there exists $u \in k^*$ with $a_0 = a_1u^4, b_0 = b_1u^6$, and in this case $(x, y) \mapsto (x/u^2, y/u^3)$ is an isomorphism. They are isomorphic over the algebraic closure of k if and only if they have the same j-invariant $j(E_0) = j(E_1)$, with $j(E_{a,b}) = 4a^3/(4a^3 + 27b^2)$.

For $N \in \mathbb{Z}_{\geq 1}$, the N-torsion subgroup $E[N]$ is the kernel of the scalar multiplication by N. We denote by $\psi_{a,b,N}(x) = \psi_N(a, b, x)$ the N-division polynomial

on $E = E_{a,b}$, which has integer coefficients and is such that for $P \in E$, we have $P \in E[N]$ is of N-torsion if and only if its x-coordinate $x(P)$ is a root of $\psi_{E,N}$.

Remark 3.1. In this paper we will also work with elliptic schemes over rings like $R = \mathbb{Z}_q$ or $\mathbb{Z}_q/p^m\mathbb{Z}_q$. But if $E \to \operatorname{Spec} \mathbb{Z}_q$ is an elliptic scheme, it is proper smooth (hence flat) over $\mathbb{Z}_q$, and E is the Néron model of its generic fiber $E_{\mathbb{Q}_q}$. Hence we can adapt the theory of elliptic curves over $\mathbb{Q}_q$ to these rings.

Isogenies and endomorphisms. An isogeny $\varphi\colon E_0 \to E_1$ is a rational map that sends ∞_{E_0} to ∞_{E_1}; every isogeny is a homomorphism of abelian groups. We say an isogeny is defined over a field k (resp. separable, resp. of degree $d \in \mathbb{Z}$) if it is defined over k (resp. separable, resp. of degree d) as a rational map. A separable isogeny is uniquely determined by its kernel, up to post-composition with an isomorphism on the codomain. Given a separable kernel $\ker \varphi = \langle P \rangle$ with $\operatorname{ord} P = N = \prod_i \ell_i$, the isogeny φ splits as a chain $\varphi = \varphi_n \circ \ldots \circ \varphi_1$ with $\deg \varphi_i = \ell_i$. If the degree is of the form $N = \ell^n$, we also refer to φ as an *isogeny path* in the ℓ-isogeny graph. Given two curves E_0, E_1 whose j-invariants are respectively j_0, j_1, there exists an ℓ-isogeny $E_0 \to E_1$ if and only if $\Phi_\ell(j_0, j_1) = 0$, where Φ_ℓ is the ℓ-modular polynomial.

Given a curve E over k, the isogenies $\varphi\colon E \to E$, together with the zero morphism, form the endomorphism ring $\operatorname{End}(E)$. If $\operatorname{char}(k) = 0$, then $\operatorname{End}(E)$ can be either $\mathbb{Z}$ or an order in a quadratic imaginary number field. If $\operatorname{char}(k) = p > 0$, then $\operatorname{End}(E)$ is either an order in a quadratic imaginary number field (in this case E is called *ordinary*) or a maximal order in the quaternion algebra $B_{p,\infty}$ ramified at p and ∞. In both cases, a nonscalar $\varphi \in \operatorname{End}(E)\backslash\mathbb{Z}$ is a quadratic imaginary algebraic integer, i.e., it is a root of a quadratic polynomial $x^2 - tx + d$ with negative discriminant, where $d = \deg \varphi$ and $t = \operatorname{tr} \varphi$ are respectively the degree and trace of the endomorphism.

Higher-dimensional isogenies. Principally polarized abelian varieties (PPAVs) are a useful computational tool to represent elliptic curves isogenies whose degree can't be factored into small primes. We refer the reader to [Mil86] for the underlying theory and [Rob24b] for the computational applications, and recall some relevant facts here. A polarized N-isogeny of PPAVs $\Phi\colon (A, \lambda_A) \to (B, \lambda_B)$ is an isogeny that respects the principal polarizations and such that $\Phi \circ \Phi^\dagger$ is the multiplication-by-N endomorphism on A, with $\Phi^\dagger = \lambda_A^\vee \circ \Phi^\vee \circ \lambda_B$.

Let, E_0, E_1 be two elliptic curves over k, and $P = (P_0, P_1), Q = (Q_0, Q_1)$ be points on $(E_0 \times E_1)[N]$, with $N \in \mathbb{Z}_{\geq 1}$ that generate a maximal Weil-isotropic subgroup, i.e., $\langle P, Q \rangle \cong (\mathbb{Z}/N\mathbb{Z})^2$ and $e_N(P, Q) = e_N(P_0, Q_0) \cdot e_N(P_1, Q_1) = 1$. Setting $\ker \Phi = \langle P, Q \rangle$ uniquely determines a polarized N-isogeny $\Phi\colon E_0 \times E_1 \to E_2 \times E_3$ up to post-composition with an isomorphism, provided the products of elliptic curves are equipped with the product polarization. If $N = \prod_i \ell_i$, the isogeny Φ splits as a chain $\Phi = \Psi_n \circ \ldots \circ \Psi_1$ with Ψ_i an ℓ_i-isogeny, as in the 1-dimensional case.

In dimension 2, let $U^A_{11,11}(z)$ be the level-$(2,2)$ *even theta function*, referred to as θ_{1111} in [Igu62]. The constant $\chi_{\text{split}}(A) := U^A_{11,11}(0)$ is a tool to recognize products of elliptic curves: given a PPAV A, the constant $\chi_{\text{split}}(A)$ vanishes if and only if A is a product $E_1 \times E_2$ (i.e., the polarization of A is a product polarization) with product theta structure.

3.2. **Scaling factors.** Let $\varphi\colon E_0 \to E_1$ be an isogeny over k. As a rational map, it can be written as $\varphi(x, y) = (f(x), c\, y f'(x))$ for some $f \in k(x), c \in k$ [Gal12, Theorem 9.7.5]. The factor c, which we also denote by c_φ, is the *differential scaling factor* of φ [Gal12, § 9.7]. An isogeny φ is said to be *normalized* if its scaling factor c_φ is 1. Given an elliptic curve E over k, the map $\operatorname{End}(E) \to k$ defined by $\varphi \mapsto c_\varphi$ is a ring homomorphism.

Retrieving the scaling factor from a given isogeny is immediate in some cases: if $\eta\colon E_0 \to E_1, \eta(x,y) = (x/u^2, y/u^3)$ is an isomorphism, its scaling factor c_η is u. Secondly, by multiplicativity of scaling factors $c_{\varphi_1\circ\varphi_0} = c_{\varphi_1}\cdot c_{\varphi_0}$, the scaling factor of an isogeny chain is the product of the scaling factor of its steps.

For a general isogeny φ, via a straightforward use of the formula $\varphi(x,y) = (f(x), cyf'(x))$, Algorithm 2 can efficiently compute the scaling factor c_φ. We use the well-known fact that for the "first-order thickening" $S = k[[\varepsilon]]/(\varepsilon^2)$ the Taylor expansion $f(x+\varepsilon) = f(x) + f'(x)\varepsilon$ holds for any f sufficiently regular, so f' can be retrieved from an evaluation of f over S. We assume given a (black-box) algebraic x-only evaluation algorithm eval_φ, that takes input the x-coordinate of a point $Q \in E(S')$ defined over any k-algebra S' and outputs the x-coordinate of $\varphi(P)$.

Algorithm 2: Computing the scaling factor of an isogeny over k.

Input: An x-only evaluation algorithm eval_φ of an k-isogeny $\varphi\colon E_0 \to E_1$;
a point $P = (x,y) \in E_0$ and its image $P_1 = \varphi(P) = (x_1, y_1)$ under φ.
Output: The differential scaling factor c_φ of φ.

$S \leftarrow k[\varepsilon]/(\varepsilon^2)$
$a + \varepsilon b \leftarrow \mathsf{eval}_\varphi(x+\varepsilon) \in S$ // $a = f(x),\ b = f'(x)$
Return $y_1/(y\cdot b)$ // $(x_1, y_1) = (f(x), cyf'(x)) = (a, cyb)$

Remark 3.2. An x-only evaluation algorithm for φ only determines φ up to a sign, i.e., up to post-composition with $[\pm 1]$, hence c_φ up to a sign.

3.3. **Lifts of endomorphisms.** In this section, we use Serre and Tate's theory of canonical lifts, as extended by Grothendieck and Messing, to describe the problem of p-adic lifting of isogenies and endomorphisms of elliptic curves.

We first show that the lift of an isogeny of degree prime to p is uniquely determined by the choice of lift of its domain. Denote by $\mathfrak{A}_1$ the moduli stack of elliptic curves and, if Γ is a congruence subgroup, denote by $\mathfrak{A}_1(\Gamma)$ the moduli stack of elliptic curves with Γ-level structure (cf. [DR72]). The stack $\mathfrak{A}_1(\Gamma^0(\ell))$ parametrizes pairs (E,K) where $K \subset E[\ell]$ is a cyclic subgroup of order ℓ, and the forgetting map $\mathfrak{A}_1(\Gamma^0(\ell)) \to \mathfrak{A}_1$ given by $(E,K) \mapsto E$ is étale over $\mathbb{Z}[1/\ell]$. As a consequence, if $\widetilde{E}/\mathbb{Z}_q$ is an elliptic curve with reduction E over $\mathbb{F}_q$, and $\varphi\colon E \to E'$ is an ℓ-isogeny over $\mathbb{F}_q$ with $p \nmid \ell$, then φ lifts uniquely to an isogeny $\widetilde{\varphi}\colon \widetilde{E} \to \widetilde{E'}$. We can compute $\widetilde{\varphi}$ by lifting the points of the kernel of φ: since $E[\ell]/\mathbb{Z}_q$ is étale, the points of ℓ-torsion of $E_{\mathbb{F}_q}$ lift uniquely.

Now let $\gamma\colon E \to E$ be an endomorphism over $\mathbb{F}_q$. If γ is of degree ℓ prime to p, for every $\widetilde{E}/\mathbb{Z}_q$ that lifts E, then γ lifts to an isogeny $\widetilde{\gamma}\colon \widetilde{E} \to \widetilde{E}'$, but in general $\widetilde{E}'$ is not isomorphic to $\widetilde{E}$. Only for an (essentially) unique choice of lift $\hat{E}$ over $\mathbb{Z}_q$ the domain and the codomain coincide, that is, γ lifts to an endomorphism $\widetilde{\gamma} \in \mathrm{End}(\hat{E})$ provided suitable conditions on γ. We call such a lift $\hat{E}$ a *γ-canonical lift* of E.

Proposition 3.3 (Lifting endomorphisms)**.** *Let $E/\mathbb{F}_q$ be an elliptic curve and γ be a non-scalar endomorphism over $\mathbb{F}_q$ of discriminant not divisible by p. There exists a γ-canonical lift $\widehat{E}_\gamma$ over $\mathbb{Z}_q$ (resp. over $\mathbb{Z}_q/p^k$ for $k \in \mathbb{Z}_{\geq 1}$) such that γ lifts as an endomorphism to $\widehat{E}_\gamma$. This lift is unique up to a $\mathbb{Z}_q$-isomorphism (resp. $\mathbb{Z}_q/p^k$-isomorphism) that reduces to the identity modulo p.*

If $E_1, E_2/\mathbb{F}_q$ are two γ-oriented curves [Onu21], then lifting gives a canonical bijection from the γ-oriented isogenies $\mathrm{Hom}_{\mathbb{F}_q,\mathbb{Z}[\gamma]}(E_1,E_2)$ between E_1, E_2 to the $\mathbb{Z}_q$-isogenies $\mathrm{Hom}_{\mathbb{Z}_q}(\widehat{E}_{1,\gamma}, \widehat{E}_{2,\gamma})$ between the respective γ-canonical lifts.

We devote the rest of the section to proving Proposition 3.3, by recalling important results in the theory of p-adic lifting of abelian schemes.

Theorem 3.4 (Serre–Tate, Grothendieck–Messing)**.** *Assume that $p > 2$. The functor $\mathcal{A}/\mathbb{Z}_q \mapsto (\mathcal{A}_{\mathbb{F}_q}, \Omega^1_{\mathcal{A}/\mathbb{Z}_q} \subset H^1_{\mathrm{dR}}(\mathcal{A}/\mathbb{Z}_q))$, which associates to an abelian scheme $\mathcal{A}/\mathbb{Z}_q$ its special fiber $A = \mathcal{A}_{\mathbb{F}_q}$ and the Hodge filtration on $H^1_{\mathrm{dR}}(\mathcal{A}/\mathbb{Z}_q) \simeq H^1_{\mathrm{crys}}(A/\mathbb{Z}_q)$, is an equivalence of category with tuples $(A/\mathbb{F}_q, \mathrm{Fil})$ where Fil is a direct summand lift of the Hodge-Filtration $\Omega^1_{A/\mathbb{F}_q} \subset H^1_{\mathrm{dR}}(A/\mathbb{F}_q) \simeq H^1_{\mathrm{crys}}(A/\mathbb{F}_q)$ to $H^1_{\mathrm{crys}}(A/\mathbb{Z}_q)$. The morphisms in the second category are given by the morphisms $\varphi\colon A/\mathbb{F}_q \to B/\mathbb{F}_q$ such that the induced map on cohomology $H^1_{\mathrm{crys}}(\varphi)\colon H^1_{\mathrm{crys}}(B/\mathbb{Z}_q) \to H^1_{\mathrm{crys}}(A/\mathbb{Z}_q)$ (a morphism of $\mathbb{Z}_q\{F,V\}$-modules, where $\mathbb{Z}_q\{F,V\}$ is the Dieudonné ring) furthermore respects the given filtrations.*

Proof. By standard inductive limit arguments and Grothendieck's algebraization (see also the proof of [Mes72, Theorem V.3.3]), it suffices to understand the deformations (i.e., the lifts) of $A/\mathbb{F}_q$ to $\mathbb{Z}_q/p^m\mathbb{Z}_q$. The theorem is obtained by combining the theorem by Serre and Tate that encodes the deformations of an abelian variety $A/\mathbb{F}_q$ in terms of the deformation of its p-divisible group $A(p)$ [Kat06], along with the theory by Grothendieck-Messing that encodes deformation of a p-divisible group G in terms of admissible liftings of the Hodge filtration of its crystal $\mathbb{D}(G)$ [Mes72, Theorem V.1.6]. We also use the fact that $H^1_{\mathrm{crys}}(A/\mathbb{Z}_q)$ induces the crystal $\mathbb{D}(A(p))$ associated to the p-divisible group $A(p)$ of A, and that given any lift $\mathcal{A}/(\mathbb{Z}_q/p^m\mathbb{Z}_q)$ of A, the Hodge filtration on $H^1_{\mathrm{dR}}(\mathcal{A}/(\mathbb{Z}_q/p^m\mathbb{Z}_q)) \simeq H^1_{\mathrm{crys}}(A/\mathbb{Z}_q) \otimes_{\mathbb{Z}_q} \mathbb{Z}_q/p^m\mathbb{Z}_q$ gives an admissible lift of the Hodge filtration of $\mathbb{D}(A(p))_{\mathbb{F}_q} \simeq H^1_{\mathrm{dR}}(A/\mathbb{F}_q)$ [BBM82, Théorème 2.5.6, Propositions 2.5.8 and 3.3.7]. The restriction to $p > 2$ is needed to ensure that the canonical divided power $p^n/n!$ be locally nilpotent. We refer to Appendix B for more details. □

Remark 3.5. If $\mathcal{A}, \mathcal{B}$ are abelian schemes over $\mathbb{Z}_q$, the reduction map gives an injection $\mathrm{Hom}_{\mathbb{Z}_q}(\mathcal{A}, \mathcal{B}) \subset \mathrm{Hom}_{\mathbb{F}_q}(A, B)$ where $A = \mathcal{A}_{\mathbb{F}_q}, B = \mathcal{B}_{\mathbb{F}_q}$. By Theorem 3.4, the image of the reduction map is given by the morphisms $\varphi\colon A \to B$ that respect the Hodge filtration induced by $\mathcal{A}, \mathcal{B}$ on $H^1_{\mathrm{crys}}(A/\mathbb{F}_q)$ and $H^1_{\mathrm{crys}}(B/\mathbb{F}_q)$ respectively. Using the isomorphism $\mathrm{Hom}_{\mathbb{F}_q}(A, B) \otimes \mathbb{Z}_p \simeq \mathrm{Hom}_{\mathbb{Z}_q\{F,V\}}(H^1_{\mathrm{crys}}(B/\mathbb{Z}_q), H^1_{\mathrm{crys}}(A/\mathbb{Z}_q))$ given by Tate's theorem, we conclude that a morphism $\varphi\colon \mathcal{A} \to \mathcal{B}$ is completely determined by its action on cohomology $H^1_{\mathrm{crys}}(\varphi)\colon H^1_{\mathrm{crys}}(B/\mathbb{Z}_q) \to H^1_{\mathrm{crys}}(A/\mathbb{Z}_q)$.

Proof of Proposition 3.3. As the discriminant $\Delta(\gamma) \neq 0 \pmod p$ is nonzero, γ is diagonalisable on $H^1_{\mathrm{crys}}(E/\mathbb{Z}_q)$ with distinct eigenvalues modulo p. The associated filtration Fil_γ lifts the one on $H^1_{\mathrm{dR}}(E/\mathbb{F}_q)$ since γ is an endomorphism over $\mathbb{F}_q$. Thus, $\widetilde{E}$ is the lift associated with this filtration as given by Theorem 3.4. Conversely, since the eigenvalues are distinct, any filtration stable by γ has to be the one given by the eigenvectors. From the proof of Theorem 3.4, we also get that $\hat{E}_{\mathbb{Z}_q/p^k\mathbb{Z}_q}$ is the unique γ-canonical lift modulo p^k.

By Theorem 3.4, the isogenies $\widehat{E}_{1,\gamma} \to \widehat{E}_{2,\gamma}$ are precisely the (unique) lifts of the isogenies $\varphi\colon E_1 \to E_2$ such that $H^1_{\mathrm{crys}}(\varphi)$ preserves the filtrations Fil_γ on E_i. Now $H^1_{\mathrm{crys}}(\varphi)$ preserves the filtrations if an only if it commutes with $H^1_{\mathrm{crys}}(\gamma)$ and, by Remark 3.5, this happens if and only if φ commutes with γ. □

Example 3.6 (Canonical lifts of an ordinary elliptic curve). If $E/\mathbb{F}_q$ is ordinary, $\Delta(\pi_q) \neq 0 \pmod p$ and there is a unique lift $\widehat{E}/\mathbb{Z}_q$ where the Frobenius lifts as an endomorphism, usually referred to as the canonical lift. Since the endomorphism ring is commutative, all other endomorphisms respect the filtration induced by the Frobenius on $H^1_{\mathrm{crys}}(E/\mathbb{Z}_q)$ hence they all lift to $\widehat{E}$.

We remark however that, though endomorphisms on E lifts uniquely to endomorphisms of $\widehat{E}$, they may not lift uniquely as isogenies if their degree is not coprime with p. For instance, the Verschiebung endomorphism admits several lifts as isogenies to $\widehat{E}$, only one of which is still an endomorphism (namely, the one whose kernel is generated by the unramified points of q-torsion of $\widehat{E}$). The Frobenius endomorphism π_q, on the other hand, does lift uniquely even as an isogeny.

Remark 3.7 (Canonical lifts of an ordinary abelian variety). If $A/\mathbb{F}_q$ is ordinary, there is a unique canonical lift $\widehat{A}/\mathbb{Z}_q$ such that the Frobenius π_q lifts as an endomorphism. The canonical lift $\widehat{A}/\mathbb{Z}_q$ can also be characterized as the unique lift such that the reduction map $\mathrm{End}_{\mathbb{Z}_q}(\widehat{A}) \to \mathrm{End}_{\mathbb{F}_q}(A)$ is bijective, or also as the unique lift such that the small Frobenius $\pi_p \colon A \to A^{(p)}$ lifts to an isogeny $\widehat{A} \to \widehat{A}^\sigma$ (or the same thing for an iterate of the small Frobenius or Verschiebung). If $\widehat{A}, \widehat{B}/\mathbb{Z}_q$ are canonical lifts of $A, B/\mathbb{F}_q$, the reduction map $\mathrm{Hom}_{\mathbb{Z}_q}(\widehat{A}, \widehat{B}) \to \mathrm{Hom}_{\mathbb{F}_q}(A, B)$ is bijective. We refer to [Mes72, Appendix] for more details.

4. Efficient representations of isogenies

In this section, we review some widely used forms of efficient representations of isogenies, and we describe how to lift a given isogeny over $\mathbb{F}_q$, with $q = p^m$, to the ring $\mathbb{Z}_q$ (we will always assume the degree of φ to be coprime to p). Throughout the section, we will use Notation 2.1 to describe the cost of arithmetic over $\mathbb{Z}_q/p^k\mathbb{Z}_q$.

Definition 4.1. Let $\varphi \colon E \to E'$ be an isogeny between elliptic curves over a field k. An **efficient representation** of φ is given by its "data" $\mathcal{D} \in \{0,1\}^s$, usually encoding a tuple of elements of k, together with three algorithms which are assumed to be polynomial-time (in the size of the input):

- domain_φ takes $\mathcal{D}$ and outputs the domain curve E.
- $\mathsf{codomain}_\varphi$ takes $\mathcal{D}$ and outputs the codomain curve E'.
- eval_φ takes $\mathcal{D}$ and some point $P \in E(R)$, where R is a k-algebra, and outputs the image point $\varphi(P) \in E(R)$.

4.1. **Lifting a path of j-invariants.** Suppose given an isogeny φ of composite degree $N = \prod_i \ell_i$. To describe it efficiently, it is convenient to split it as a *chain*

$$E_0 \xrightarrow{\varphi_1} E_1 \xrightarrow{\varphi_2} \dots \xrightarrow{\varphi_{r-1}} E_{r-1} \xrightarrow{\varphi_r} E_r \qquad (\text{with } E_0 \xrightarrow{\varphi} E_r)$$

where the isogeny step φ_i has degree ℓ_i. Via the ℓ-modular polynomial, the existence of an ℓ-isogeny between two curves can be encoded as a polynomial relation between the j-invariants of the two curves. Through this idea (already appeared in [VPV01] for Satoh's point counting algorithm) we can encode isogenies as solutions of polynomial systems, which allows us to compute their lifts via Hensel lifting.

Lemma 4.2. *Suppose given a tuple $\overline{\mathcal{D}} = (\overline{j}_0, \dots, \overline{j}_n) + O(p^k) \in \mathbb{Z}_q^{n+1}$ that satisfies the polynomial constraints:*

$$\Phi_{\ell_i}(\overline{j}_{i-1}, \overline{j}_i) \in O(p^k) \qquad i = 1, \dots, n$$

where Φ_{ℓ_i} is the ℓ_ith classical modular polynomial.

For every elliptic curve $E_0/\mathbb{F}_q$ with $j(E_0) \equiv \overline{j}_0 \pmod p$, the tuple $\overline{\mathcal{D}}$ defines an isogeny chain $\varphi \colon E_0 \xrightarrow{\varphi_1} E_1 \xrightarrow{\varphi_2} \dots \xrightarrow{\varphi_n} E_n$, with $j(E_i) \equiv \overline{j}_i \pmod p$ for all i.

If the tuple also satisfies the additional constraint

$$j_n - j_0 \in O(p^k)\,,$$

then $\overline{\mathcal{D}}$ defines an endomorphism $\varphi \in \mathrm{End}(E_0)$ over $\mathbb{F}_q$.

Suppose that for all i, the isogeny φ_i is the only ℓ_i-isogeny from E_{i-1} to E_i up to post-composition with automorphisms of E_i. Then, letting $\ell = \max_i \deg(\varphi_i) < p$, we can compute a lift $\mathcal{D} = (j_0, \dots, j_n)$ of $\overline{\mathcal{D}}$ in $\widetilde{O}(\ell^2 n)\mathsf{M}_{2k}$ bit operations.

Proof. The first part of the claim stems from the classical theory of modular curves over a field. Let $F = (F_1, \dots, F_n)$ with $F_i(j_0, \dots, j_n) = \Phi_{\ell_i}(j_{i-1}, j_i)$. The uniqueness condition on φ_i ensures $\partial_{j_i} F(\overline{\mathcal{D}}) \neq 0 \pmod p$, so the Jacobian matrix $\mathrm{D}F(\overline{\mathcal{D}})$ has full rank modulo p. We also observe that this Jacobian matrix is a sparse matrix (banded with only 2 nonzero diagonals, except possibly the last row).

To compute the lift, we apply Proposition 2.5. Note that that each polynomial constraint depends at most $m = 2$ variables, and the complexity of evaluation of a modular polynomial of degree ℓ is $\widetilde{O}(\ell^2)$; finally, because of the sparsity pattern of $\mathrm{D}F(\overline{\mathcal{D}})$, solving the relative linear system in precision k costs $O(n)\mathsf{M}_k$. □

A path of j-invariants as above is technically not enough to be an unambiguous efficient representation: we must also specify a domain curve E_0 with the given j-invariant j_0 (or an algorithm to produce E_0 from j_0), which is what we do next.

4.2. **Lifting an isogeny step.** Given a cyclic isogeny $\varphi\colon E_0 \to E_1$ of small (typically prime) degree ℓ, several algorithms in the literature allow us to get an efficient representation of φ from the domain E_0 and a single extra parameter x. All common such representations can be encoded as solutions of an algebraic (hence analytic) constraint system, hence they're suited for Hensel lifting.

Vélu, $\sqrt{élu}$. Vélu's formulas [Vél71] and the $\sqrt{\text{élu}}$ algorithm [BDLS20] can be used to represent a (separable) isogeny of (typically) prime degree ℓ by a single point of order ℓ that generates the kernel. Such representations can be chained: suppose given a chain $\varphi\colon E_0 \xrightarrow{\varphi_1} E_1 \xrightarrow{\varphi_2} \dots \xrightarrow{\varphi_n} E_n$ of Vélu or $\sqrt{\text{élu}}$ isogenies φ_i, with each individual prime-degree step φ_i having small degree ℓ_i, often assumed to be polynomial in $\log(\deg \varphi)$. Whenever $N = \deg \varphi$ is a smooth integer, we get an efficient representation for φ. The data $\mathcal{D}$ is given by a tuple of field elements $(a_0, b_0, x_1, ..., x_r) \in k^{r+2}$ such that a_0, b_0 are the short Weierstraß coefficients of E_0, and each $\ker \varphi_i$ is generated by a point on E_{i-1} with x-coordinate x_i.

Lemma 4.3 ([Vél71; BDLS20])**.** *Fix $\ell \in \mathbb{Z}_{\geq 1}$ with $p \nmid \ell$ and let $\varphi\colon E_0 \to E_1$ be a normalized cyclic isogeny over k with kernel $\ker \varphi = \langle P \rangle$, $P \in E_0[\ell]$. Vélu's formulas (resp. the $\sqrt{élu}$ algorithm) are algebraic algorithms that compute the Weierstraß coefficients of E_1 as a rational map $(a(E_1), b(E_1)) = \mathsf{Vélu}(a(E_0), b(E_0), x(P))$ in $O(\ell)$ (resp. $\widetilde{O}(\sqrt{\ell})$) arithmetic operations over k.*

Lemma 4.4. *Fix $\ell \in \mathbb{Z}_{\geq 1}$ with $p \nmid \ell$. Let $\overline{\mathcal{D}} = (\overline{a}_0, \overline{b}_0, \overline{x}_1, \overline{a}_1, \overline{b}_1) + O(p^k) \in \mathbb{Z}_q^5$ with $4\overline{a}_i^3 + 27\overline{b}_i^2 \neq 0 \pmod p$, $i = 0, 1$ satisfy the following rational constraints:*

$$\begin{cases} \psi_\ell(\overline{a}_0, \overline{b}_0, \overline{x}_1) \in O(p^k) \\ (\overline{a}_1, \overline{b}_1) - \mathsf{Vélu}(\overline{a}_0, \overline{b}_0, \overline{x}_1) \in O(p^k) \end{cases} \tag{1}$$

This tuple defines an ℓ-isogeny $\varphi\colon E_0 \to E_1$ defined over $\mathbb{Z}_q/p^k\mathbb{Z}_q$. For every lift $(a_0, b_0) + O(p^{2k})$, using Vélu's formulas (resp. the $\sqrt{élu}$ algorithm), there exists a lift $(a_0, b_0, x_1, a_1, b_1) + O(p^{2k})$ that satisfies the constraints in precision $2k$. We can compute this lift in $O(\ell)\mathsf{M}_{2k}$ (resp. $\widetilde{O}(\sqrt{\ell})\mathsf{M}_{2k}$) bit operations.

Proof. The Jacobian matrix of the above constraints is of the form

$$\begin{pmatrix} * & * & \partial_{x_1}\psi_\ell(\overline{a}_0, \overline{b}_0, \overline{x}_1) & & \\ * & * & 0 & 1 & \\ * & * & * & 0 & 1 \end{pmatrix}$$

where $\partial_{x_1}\psi_\ell(\overline{a}_0, \overline{b}_0, \overline{x}_1) = \psi'_{\ell,\overline{a}_0,\overline{b}_0}(\overline{x}_1) \neq 0 \pmod{p^k}$ is nonzero since a division polynomial with $p \nmid \ell$ is separable. We use Algorithm 1, noting that the pivots of $\mathrm{D}F(\overline{\mathcal{D}})$ correspond to variables x_1, a_1, b_1, and see that the solution set is parametrized by the lifts of a_0, b_0. □

The BMSS algorithm. While Vélu and $\sqrt{\text{élu}}$ compute an isogeny from the domain curve and a kernel point, an isogeny can also be computed from the data of the domain curve and the j-invariant of the codomain.

More precisely, given an elliptic curve $E_0 = E_{a_0,b_0}$ over $\mathbb{Z}_q/p^k\mathbb{Z}_q$ and the j-invariant of a curve ℓ-isogenous to E_0, the algebraic algorithm from [BMSS08] can compute the codomain E_1 of a normalized ℓ-isogeny $\varphi\colon E_0 \to E_1$ provided $\{j(E_0), j(E_1)\} \cap \{0, 1728\} = \emptyset$. We denote by $(a_1, b_1) = \mathsf{BMSS}(a_0, b_0, j_1)$ the resulting algorithm, which runs in $\widetilde{O}(\ell)\mathsf{M}_k$ bit operations.

The following lemma shows how to lift the representation of an ℓ-isogeny using this algorithm, and can be proved analogously to Lemma 4.4.

Lemma 4.5. *Fix $\ell \in \mathbb{Z}_{\geq 1}$ with $p \nmid \ell$. Let $(\overline{a}_0, \overline{b}_0, \overline{j}_1, \overline{a}_1, \overline{b}_1) + O(p^k) \in \mathbb{Z}_q^5$ with $j(\overline{a}_i, \overline{b}_i) \bmod p \notin \{0, 1728\}$, $i = 0, 1$ satisfy the following rational constraints:*

$$\begin{cases} \overline{j}_1 - j(\overline{a}_1, \overline{b}_1) \in O(p^k) \\ (\overline{a}_1, \overline{b}_1) - \mathsf{BMSS}(\overline{a}_0, \overline{b}_0, \overline{j}_1) \in O(p^k) \end{cases}$$

This tuple defines an ℓ-isogeny $\varphi\colon E_0 \to E_1$ defined over $\mathbb{Z}_q/p^k\mathbb{Z}_q$. For every lift $(a_0, b_0) + O(p^{2k})$, there exists a lift $(a_0, b_0, j_1, a_1, b_1) + O(p^{2k})$ that satisfies the constraints in precision $2k$. We can compute this lift in $\widetilde{O}(\ell)\mathsf{M}_{2k}$ bit operations.

Radical isogenies. Radical isogenies, introduced in [CDV20], provide a different algorithm to parametrize ℓ-isogenies outgoing from a given curve, not via their kernel point but via a choice of ℓth root. In this case, a curve E is described in Tate normal form $E\colon y^2 + (1-c)xy - by = x^3 - bx^2$, so that the point $P = (0, 0)$ is of order ℓ.[3] Suppose given coefficients (b_0, c_0) defining a curve E_0, and α an ℓth root of $\rho = f_{\ell,P}(-P)$, where $f_{\ell,P}$ is a Miller function for P. The tuple (b, c, α) defines an ℓ-isogeny $\varphi\colon E_0 \to E_1$ to another curve in Tate normal form, and the coefficients of E_1 are algebraic expressions $(b_1, c_1) = \mathsf{RadIsog}(b, c, \alpha)$. For $\ell \leq 13$, these expressions are more efficient to compute than the other representations mentioned above. We can use Hensel lifting once again to compute p-adic lifts of radical isogenies.

Lemma 4.6. *Fix a prime $\ell \leq 13$, $\ell \neq p$. Suppose given $\overline{\mathcal{D}} = (\overline{b}_0, \overline{c}_0, \overline{\alpha}_1, \overline{b}_1, \overline{c}_1) + O(p^k) \in \mathbb{Z}_q^5$, where (b_i, c_i) define a curve E_i in Tate normal form with $P_i = (0, 0) \in E_i[\ell]$ of order ℓ, for $i = 0, 1$. Assume $\overline{\mathcal{D}}$ satisfies the following rational constraints:*

$$\begin{cases} \overline{\alpha}_0^n - t_N(P_0, -P_0) \in O(p^k) \\ (\overline{b}_1, \overline{c}_1) - \mathsf{RadIsog}(\overline{b}_0, \overline{c}_0, \overline{\alpha}_1) \in O(p^k) \end{cases}$$

This tuple defines an ℓ-isogeny $\varphi\colon E_0 \to E_1$ defined over $\mathbb{Z}_q/p^k\mathbb{Z}_q$. For every lift $(b_0, c_0) + O(p^{2k})$, there exists a lift $(b_0, c_0, j_1, b_1, c_1) + O(p^{2k})$ that satisfies the constraints in precision $2k$, thus defines a radical isogeny lifting the given one.

[3] However, we present in Appendix A.1 formulas to compute radical 2-isogenies on curves in Montgomery form.

4.3. **Lifting isogeny and endomorphism chains.** Once we know how to lift a small-degree isogeny, we show how to lift isogenies and endomorphisms that can be factored into a chain of small-degree isogenies (possibly including isomorphisms). We show how to do that in the case of Vélu steps, but remark that the same algorithms can be used to lift chains of $\sqrt{\text{élu}}$, BMSS, radical-isogeny steps.

Lifting an isogeny chain is immediate from the lifting of a single steps. From the repeated application of Lemma 4.4, we get the following

Corollary 4.7. *Let φ be the composition $E_0 \xrightarrow{\varphi_1} E_1 \xrightarrow{\varphi_2} \dots \xrightarrow{\varphi_n} E_n$, let $\overline{a}_i, \overline{b}_i$ be the Weierstraß coefficients of E_i and $\overline{x}_i = x(P_i)$ where $P_i \in E_i[\ell_i]$ is a generator of $\ker \varphi_i$. Let $\ell = \max_i \ell_i$.*

Suppose given a tuple $\overline{\mathcal{D}} = (\overline{a}_0, \overline{b}_0, \overline{x}_1, \overline{a}_1, \overline{b}_1, \dots, \overline{x}_n, \overline{a}_n, \overline{b}_n) + O(p^k)$ such that the elements $(\overline{a}_{i-1}, \overline{b}_{i-1}, \overline{x}_i, \overline{a}_i, \overline{b}_i)$ are a solution of (1) for all $i = 1, \dots, n$. For all lifts $(a_0, b_0) + O(p^{2k})$ of $(\overline{a}_0, \overline{b}_0)$, we can compute a lift $D + O(p^{2k})$ of $\overline{\mathcal{D}}$ satisfying the system in precision $2k$, using $O(n\ell)\mathsf{M}_{2k}$ bit operations.

To lift an endomorphism, we need extra data: introducing a variable u representing an isomorphism factor, we add constraints to ensure that the last step of the chain is isomorphic to the domain.

Proposition 4.8. *We follow the notation of Corollary 4.7. Suppose given a tuple $\overline{\mathcal{D}} = (\overline{a}_0, \overline{b}_0, \overline{x}_1, \dots, \overline{a}_n, \overline{b}_n, \overline{u})$ satisfying*

$$\begin{cases} Y_{\text{divpoly},i}\colon & \psi_{\overline{a}_i, \overline{b}_i, \ell_i}(\overline{x}_i) \in O(p^k) \qquad i = 1, \dots, n \\ Y_{\mathsf{Vélu},i}\colon & (a_i, b_i) - \mathsf{Vélu}(a_{i-1}, b_{i-1}, x_i) \in O(p^k) \qquad i = 1, \dots, n \\ Y_{\text{endo}}\colon & (a_n - a_0 \cdot u^4, b_n - b_0 \cdot u^6) \in O(p^k). \end{cases}$$

This tuple defines an endomorphism $\varphi = \eta \circ \varphi_n \circ \dots \circ \varphi_1$ over $\mathbb{Z}_q/p^k\mathbb{Z}_q$ of degree $N = \prod_i \ell_i$, factored as a chain of normalized Vélu isogenies followed by a Weierstraß isomorphism $\eta\colon E_n \to E_0$ with isomorphism factor u. Assume that φ does not reduce to a scalar endomorphism modulo p.

Set $\ell = \max_i \ell_i$ assume $p \nmid \ell$. For all lifts $a_0 + O(p^{2k})$ (resp. lifts $b_0 + O(p^{2k})$, if $j(E_0) \equiv 0 \pmod p$) we can compute a tuple $(a_0, b_0, \dots, a_n, b_n, u)$ satisfying the system in precision $2k$, using $O(n\ell)\mathsf{M}_{2k}$ bit operations.

Proof. The division polynomial constraints ensure that x_i defines a point of order ℓ_i on the curve E_i, hence a separable ℓ_i isogeny. The last (2-dimensional) constraint $Y_{\text{endo}} = (Y_a, Y_b)$ ensures the codomain of this isogeny is isomorphic to the domain curve, hence the isogeny is actually an endomorphism.

We must prove that the Jacobian matrix $\mathrm{D}F(\overline{\mathcal{D}})$ of the evaluation function F at the given tuple is full-rank modulo p. First, if $\overline{a}_0 \neq 0 \pmod p$, the partial derivative $\partial_u Y_a(\overline{\mathcal{D}})$ is nonzero, since the isomorphism factor $\overline{u}$ is always nonzero. Otherwise $\overline{b}_0 \bmod p$ is nonzero, so $\partial_u Y_b(\overline{\mathcal{D}}) \neq 0 \pmod p$. We prove for simplicity the case $\overline{a} \notin O(p)$; for the other case, replace Y_a with Y_b. Proceeding as in the proof of Lemma 4.4, we see that the rows relatives to the constraints $(Y_{\text{divpoly},1}, Y_{\mathsf{Vélu},1}, \dots, Y_{\text{divpoly},n}, Y_{\mathsf{Vélu},n}, Y_a)$ are linearly independent, since the top-right $(3n{+}1)$-by-$(3n{+}1)$ submatrix of $\mathrm{D}F(\overline{\mathcal{D}})$ is triangular with nonzero diagonal. We're left to prove that the last constraint is independent of the others. Suppose not. Since a lift of φ must exist, the system (†) has solution, and the set of solutions is parametrized by the (arbitrary) lifts of $\overline{a}_0, \overline{b}_0$. If Y_b were lineary dependent with the previous constraints, having free choice on the lifts a_0, b_0 would yield at least two non-isomorphic endomorphisms that lift of (E_0, φ_0), which contradicts Proposition 3.3.

To evaluate the complexity of lifting, we use Proposition 2.5. In precision $2k$, the cost of evaluating $Y_{\mathrm{divpoly},i}$ is $O(\log \ell)\mathsf{M}_{2k}$, the cost of $Y_{\mathsf{Vélu},i}$ is $O(\ell)\mathsf{M}_{2k}$ and the endomorphism constraints take $O(1)\mathsf{M}_{2k}$, for a total evaluation cost of $O(n\ell)\mathsf{M}_{2k}$ bit operations. Since the Jacobian matrix $\mathrm{D}F(\overline{\mathcal{D}})$ is a sparse matrix (banded with only 6 nonzero diagonals, except the last row), the relative linear system (†) can be solved in $O(n)\mathsf{M}_k$. □

Remark 4.9. We remark a substantial difference in the above two statements. In the first, we see that for a general isogeny $\overline{\varphi}\colon \overline{E}_0 \to \overline{E}_1$, any arbitrary lift $\overline{E}_0$ of the domain defines a valid lift $\varphi\colon E_0 \to E_1$. In the case of endomorphisms, however, the lift is unique, as we know from Proposition 3.3. More precisely, given a separable $\overline{\varphi} \in \mathrm{End}(E_0)$, for every E_0 lifting $\overline{E}_0$ there is a lift $\varphi\colon E_0 \to E_0'$ as an isogeny, where E_0' lifts of $\overline{E}_0$ as well, but for only one E_0 (up to isomorphism) we have $E_0' \cong E_0$.

Remark 4.10. In this paper, we focus on $p \geq 5$, where it is always possible to consider all curves to be defined in short Weierstraß form. The algorithms we present generalize to curves of arbitrary shape, as long as one replaces (a_0, b_0) with the tuple of relevant coefficients.

Remark 4.11. This section represents Vélu isogeny chains of degree ℓ^n with the kernel points of each step of the chain. It is possible to represent chains with a single kernel point $P \in E_0[\ell^n]$ generating the kernel, the resulting algorithm is slightly less efficient, since $x_1, \ldots, x_n$ must be recovered from P via repeated doubling and isogeny evaluations [DJP14, § 4.2.2]. Therefore, instead of lifting a representation from a single kernel point (exactly as in Equation (1)), it is faster to first transform the representation (E_0, P) into a chain $(E_0, x_1, \ldots, x_n)$ in low precision, then perform the lifting.

4.4. **Lifting HD representations.** Suppose given a separable isogeny $\varphi\colon E \to E'$ over $\mathbb{F}_q$ of odd degree d, not necessarily polynomially smooth in $\log(qd)$. Though φ cannot in general be split into a chain of small-degree factors, [Rob22a] describe a way to "embed" φ as a component of a N-isogeny (of principally polarized abelian varieties) between products of elliptic curves, where N *is* polylog(qd)-smooth. The dimension of the abelian varieties involved is $r \in 2, 4, 8$. We consider here for simplicity the 2-dimensional case with $N = 2^n$.

Lemma 4.12. *Fix $N = 2^n$ for some $n \in \mathbb{Z}_{\geq 0}$. Suppose given elliptic curves E_0, E_1 over $R = \mathbb{Z}_q/p^k\mathbb{Z}_q$, let (a_i, b_i) denote the short Weierstraß coefficients of E_i for $i = 0, 1$, and let $(P_0, P_1), (Q_0, Q_1)$ be points on $(E_0 \times E_1)(R)[N]$ such that the subgroup $K = \langle [4](P_0, P_1), [4](Q_0, Q_1)\rangle$ is maximally isotropic. Let $\Phi\colon E_0 \times E_1 \to A$ be the polarized N-isogeny defined by $\ker \Phi = K$ and $R = (R_0, R_1) \in (E_0 \times E_1)(R)$ a point with $U_{11,11}(\Phi(R)) \neq 0 \pmod p$. Using $O(n \log n)\mathsf{M}_k$ bit operations, on input $((a_i, b_i, x_{P_i}, x_{Q_i}, x_{R_i})_{i=0,1})$, we can compute a theta structure A for the codomain of Φ, the splitting constant $\chi_{\mathrm{split}}(A)$ and, in case $\chi_{\mathrm{split}}(A) = 0$, compute E_2 (and E_3). All the algorithms are algebraic in the coordinates of the input.*

Proof. We refer to [DMPR24, § 4.1] for the details of computation. We remark that, if $A \cong E_2 \times E_3$ is isomorphic to a product, there is a linear change of coordinates with coefficients in $\mathbb{Z}$, computable from the reduction of A modulo p, that maps A to the product structure $E_2 \times E_3$. To compute the splitting constant, [DMPR24] computes $(\chi_{\mathrm{split}}(A))^2 = U_{11,11}(0)^2$, but we need the non-squared value for our purposes. To compute it, we can use the duplication formula [Rob21, Eq. (2.9)] with $\widetilde{x} = R, \widetilde{y} = 0$ to get $\theta^A_{11,11}(0)$, using the fact that $\theta^A_{11,11}(R)$ is invertible. □

In the following, we denote by $\mathsf{split}((a_i, b_i, x_{P_i}, x_{Q_i})_{i=0,1})$ the algorithm to compute the splitting constant as above, where we treat x_R as a constant, fixed at the

beginning depending on the input curves and kernel modulo p (see Remark 2.7). Similarly, we denote by $\mathsf{HDcodom}_1((a_i, b_i, x_{P_i}, x_{Q_i})_{i=0,1})$ the algorithm to compute the first factor of the codomain, as given by the lemma.

The lemma above states that a two-dimensional isogeny $\Phi\colon E_0 \times E_1 \to E_2 \times E_3$ between two elliptic products yields an efficient representation of the one-dimensional isogeny $\varphi\colon E_0 \to E_2$ given by the composition $E_0 \hookrightarrow E_0 \times E_1 \xrightarrow{\Phi} E_2 \times E_3 \twoheadrightarrow E_2$. We show how to lift this representation when $E_0 = E_2$, i.e., when $\varphi \in \mathrm{End}(E_0)$.

Proposition 4.13. *Let $\overline{\mathcal{D}}$ be a tuple $(\overline{a}_i, \overline{b}_i, \overline{x}_{P_i}, \overline{x}_{Q_i}, \overline{u})_{i=0,1} \in \mathbb{Z}_q^9$ satisfying*

$$\begin{cases} (\psi_{2^{n+2}}(\overline{a}_i, \overline{b}_i, \overline{x}_{P_i}), \psi_{2^{n+2}}(\overline{a}_i, \overline{b}_i, \overline{x}_{Q_i})) \in O(p^k) & i = 0, 1 \\ \mathsf{split}((\overline{a}_i, \overline{b}_i, \overline{x}_{P_i}, \overline{x}_{Q_i})_{i\in\{0,1\}}) \in O(p^k) \\ \mathsf{HDcodom}_1((\overline{a}_i, \overline{b}_i, \overline{x}_{P_i}, \overline{x}_{Q_i})_{i\in\{0,1\}}) - (\overline{a}_0\overline{u}^4, \overline{b}_0\overline{u}^6) \in O(p^k) \end{cases}$$

The tuple $\overline{\mathcal{D}}$ defines a two-dimensional 2^n-isogeny $\Phi\colon E_0 \times E_1 \to E_0 \times E_3$ over $\mathbb{Z}_q/p^k\mathbb{Z}_q$. For all lifts $a_0 + O(p^{2k})$ (resp. for all lifts $b_0 + O(p^{2k})$, if $\overline{a}_0 \in O(p)$) and of $a_1 + O(p^{2k})$ (resp. b_1, if $\overline{a}_1 \in O(p)$) we can compute a tuple D lifting $\overline{\mathcal{D}}$ satisfying the system in precision $2k$, using $O(n \log n)\mathsf{M}_{2k}$ bit operations.

We remark that the lifting approach is the p-adic analogue of what [KR24] does over $\mathbb{F}_q[\varepsilon]/(\varepsilon^2)$, see also Remark 2.14.

Proof. In the constraint system, the division polynomial constraints make sure (x_{P_i}, x_{Q_i}) define points P_i, Q_i on E_i of order 2^{n+2}. Write $P_i' = [4]P_i$, $Q_i' = [4]Q_i$. As by assumption the points $(P_0', P_1'), (Q_0', Q_1')$ generate an isotropic subgroup when reduced to lower precision, the lifted points themselves also generate an isotropic subgroup. Indeed, the pairing $e_{2^n}((P_0', P_1'), (Q_0', Q_1'))$ must be a lift of $\overline{e} = e_{2^n}([4](\overline{P}_0, \overline{P}_1), [4](\overline{Q}_0, \overline{Q}_1)) = 1$, but also a lift of $\overline{e}$ as a root of $f(x) = x^{2^n} - 1$ since Weil pairings are always 2^nth roots of unity. Since the polynomial $f(x)$ is separable, the root 1 has a unique lift, namely 1 itself. The splitting constraint makes sure that the isogeny Φ lands on a product. Finally, the last constraint ensures that the one-dimensional isogeny embedded as a component of Φ is an endomorphism.

We show that the Jacobian matrix $\mathrm{D}F(\overline{\mathcal{D}})$ of the constraint function F has full rank modulo p. First, by similar arguments as in the proof of Lemma 4.4 invoking separability of ψ_{2^n}, we observe the top-right 4-by-4 submatrix is triangular with nonzero diagonal elements. The splitting constraint must have nonzero partial derivative modulo p with respect to one between a_1 and b_1 (depending on which one is nonzero modulo p). If not, in precision $2k$ there would be an E_0 and at least two non-isomorphic choices of E_1, that admit a separable $d(2^n - d)$-isogeny $E_0 \to E_1$ lifting the isogeny diamond that we're given in precision p. This contradicts the fact that lifts of separable isogenies are unique once the domain and a kernel generator are fixed, as shown in Lemma 4.4. To show that the differential of the endomorphism constraint is linearly independent from the rest of the Jacobian matrix when evaluated at the given tuple $\overline{\mathcal{D}}$ modulo p, we invoke the uniqueness of lifts of endomorphisms as in the proof of Proposition 4.8.

The complexity of the constraint system amounts to four evaluations of division polynomials of degree 2^n, which costs $O(n)\mathsf{M}_{2k}$ bit operations, and the evaluation of the splitting character, as well as the codomain algorithm, both of which have complexity $O(n \log n)\mathsf{M}_{2k}$. Having a constant number of variables, the complexity of the linear algebra step is negligible here. □

Remark 4.14. The same observations as for 1-dimensional chains apply here: if Φ is a general isogeny that doesn't have to be an endomorphism, every choice of

lift of E_0, E_3 uniquely determines a lift of Φ. Via small edits, the theorem can be generalized to polylog-smooth polarized degree N (not of the form 2^n). Instead of two kernel generators on $E_0 \times E_3$, we can shave a factor $\log\log N$ off the complexity of lifting if we represent the 2D isogeny as a chain and save the kernel points of each step, although this adds complications since the intermediate steps here are non-split theta structures, therefore the classical algorithms we use on elliptic curves (e.g., evaluating division polynomials to test point orders) are nontrivial to generalize. Finally, the above lifting algorithm (presented on 2D representations for simplicity) can be generalized to 4D and 8D representations using the same techniques.

5. Computing the trace of an endomorphism

We now turn to the problem of computing the trace of an $\mathbb{F}_q$-endomorphism $\varphi\colon E \to E$ via gradual p-adic lifting. Our approach is inspired by [MPSW25]. It is based on the observation that the differential scaling factor c_φ of an endomorphism $\varphi \in \mathrm{End}(E)$ is itself a root of the characteristic polynomial of φ and the trace $\mathrm{tr}(\varphi)$ is easily computable from c_φ. Being a coefficient of this characteristic polynomial, the trace of an endomorphism over a ring R lies on the image of $\mathbb{Z}$ in R. If (E, φ) is defined over $\mathbb{F}_q$, then $\mathrm{tr}\,\varphi$ is an integer modulo the characteristic p of $\mathbb{F}_q$. If instead we first lift the endomorphism to characteristic 0, the same observation yields the full trace as an integer.

Proposition 5.1. *Consider an efficient representation of a nonscalar separable endomorphism $\overline{\varphi} \in \mathrm{End}(\overline{E})\backslash\mathbb{Z}$ defined over $\mathbb{F}_q$ with $q = p^m$, described as a zero of an efficiently evaluatable analytic function $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$. Let $k = \lceil \log_p(4\sqrt{\deg\overline{\varphi}})\rceil$. Then Algorithm 3 computes the trace $\mathrm{tr}(\overline{\varphi}) \in \mathbb{Z}$ using $O(\log_2(k))$ calls[4] to Algorithm 1 (to construct a lifted endomorphism φ in precision k), followed by one call to Algorithm 2 (to recover the scaling factor of φ).*

Proof. We first prove that Algorithm 3 is correct. Let (E, φ) be the lift over $\mathbb{Z}_q/p^k\mathbb{Z}_q$ computed by Algorithm 3, and let $(\widetilde{E}, \widetilde{\varphi})$ be a lift over $\mathbb{Z}_q$. Since reduction modulo p is a ring homomorphism, both φ and $\widetilde{\varphi}$ satisfy the quadratic equation $x^2 - tx + d = 0$ given by the characteristic polynomial of $\overline{\varphi}$; in particular, they have the same trace $t \in \mathbb{Z}$. Similarly, the map $\varphi \mapsto c_\varphi$ is a ring homomorphism $\mathrm{End}(E) \to \mathbb{Z}_q/p^k\mathbb{Z}_q$, so the scaling factor $c_\varphi \neq 0 \in \mathbb{Z}_q/p^k\mathbb{Z}_q$ is a root of $x^2 + (t \bmod p^k) + (d \bmod p^k)$. In particular, we have $c_\varphi + d/c_\varphi \equiv t \in \mathbb{Z}/p^k\mathbb{Z}$. Finally, Hasse's theorem gives the bound $|t| \leq 2\sqrt{d} < p^k/2$. Thus, t equals the representative of $c_\varphi + d/c_\varphi \bmod p^k$ centered around 0, i.e., the unique representative in the integer interval $\mathbb{Z} \cap [-p^k/2, p^k/2]$. □

The following corollary gives the complexity of trace computation for the case of a smooth-degree endomorphism represented as a chain of small-degree isogenies, and for the case of HD representations.

Corollary 5.2. *Let $\varphi \in \mathrm{End}(E_0)\backslash\mathbb{Z}$ be a separable endomorphism defined over $\mathbb{F}_q$. Let $k = \lceil \log_p(4\sqrt{\deg\overline{\varphi}})\rceil$.*

(1) Suppose that $\deg\varphi = \prod_i^n \ell_i$, $\ell = \max_i \ell_i$, and that we're given an efficient representation of φ as a chain $\varphi = \eta \circ \varphi_n \circ \cdots \circ \varphi_1$, with φ_i an ℓ_i-isogeny computed via Vélu's formulas (resp. $\sqrt{\text{élu}}$) and η an isomorphism. Then we can compute $\mathrm{tr}(\varphi)$ using $O(n\ell)\mathsf{M}_k$) (resp. $n\widetilde{O}(\sqrt{\ell})\mathsf{M}_k$) bit operations. If $\ell \in O(1)$, this cost simplifies to $\widetilde{O}(n^2 \log_p q + n \log q)$.

[4] As noted in Corollary 2.6, the last lifting from precision $\lceil k/2\rceil$ to k dominates the cost.

Algorithm 3: Computing the trace of a separable finite-field endomorphism.

Input: An efficient representation $\mathcal{A}_\varphi$ of $\varphi \in \mathrm{End}(E)$ over $\mathbb{F}_q$ given as a zero of a constraint system over $\mathbb{Z}_q$ in precision 1; and $d = \deg \varphi \in \mathbb{Z}$.

Output: $\mathrm{tr}(\varphi) \in \mathbb{Z}$.

$k \leftarrow \lceil \log_p(4\sqrt{\deg \widehat{\varphi}}) \rceil$.

For i **from** 1 **to** $\lceil \log_2(k) \rceil - 1$ **do**

  Lift $\mathcal{A}_\varphi$ from precision 2^{i-1} to precision 2^i // Algorithm 1

Lift $\mathcal{A}_\varphi$ from precision $\lceil k/2 \rceil$ to precision k. // Algorithm 1

// now $\mathcal{A}_\varphi$ satisfies the constraints in precision k

(E, φ) over $\mathbb{Z}_q/p^k \leftarrow$ compute representation from $\mathcal{A}_\varphi$

Compute scaling factor $c_\varphi \in \mathbb{Z}_q/p^k$ from (E, φ) // Algorithm 2

$t \leftarrow c_\varphi + d/c_\varphi \in \mathbb{Z}/p^k$

Return the representative of t in $\mathbb{Z} \cap [-p^k/2, p^k/2]$.

(2) Suppose given an HD representation of φ embedded into a 2^n-isogeny in dimension $r \in \{2, 4, 8\}$. Then we can compute $\mathrm{tr}(\varphi)$ using $O(n \log n)\mathsf{M}_k$ bit operations. As above, this cost simplifies to $\widetilde{O}(n^2 \log_p q + n \log q)$.

Proof. In the case of a chain, Proposition 4.8 and Corollary 2.6 show that the complexity of lifting φ to precision k is $O(n\ell)\mathsf{M}_k$ (resp. $n\widetilde{O}(\sqrt{\ell})\mathsf{M}_k$ with $\sqrt{\text{élu}}$). To find the scaling factor, applying Algorithm 2 amounts to evaluating the isogeny over $S = R[\varepsilon]/(\varepsilon^2)$, with $R = \mathbb{Z}_q/p^k\mathbb{Z}_q$; the cost of arithmetic over S is the same as in $\mathbb{Z}_q/p^{2k}\mathbb{Z}_q$, that is, $\mathsf{M}_{2k} \lesssim 4 \cdot \mathsf{M}_k$, therefore Algorithm 2 has the same asymptotic cost as lifting. We remark that, if the isogeny steps computed via Vélu are normalized (or more generally, if their scaling factors over $\mathbb{Z}_q$ can be freely chosen as arbitrary lifts of those over $\mathbb{F}_q$; see Remark 2.7), then Algorithm 2 is not needed. Indeed, since composition multiplies scaling factors, the scaling factor of the chain is the product of $\prod_i c_{\varphi_i}$ with the scaling factor u of the last isomorphism; note that u is explicit in the representation as part of η.

If $\ell \in O(1)$, from $\deg \varphi \leq \ell^n$ we get $k \in O(n \log \ell / \log p)$. Substituting the cost of arithmetic $\mathsf{M}_k \in \widetilde{O}(k \log q)$ into the cost $O(n\ell)\mathsf{M}_k$ yields $\widetilde{O}(n^2 \log_p q + n \log q)$.

In the case of a HD representation, lifting costs $O(n \log n)\mathsf{M}_k$ by Proposition 4.13; repeating the above reasoning, Algorithm 2 costs $O(n \log n)\mathsf{M}_k$ bit operations too. As $\deg \varphi \leq 2^n$, we see as above that the cost simplifies to $\widetilde{O}(n^2 \log_p q + n \log q)$. □

Remark 5.3. It seems tempting, but equally difficult, to adapt the methods outlined in this section to point counting, i.e., to computing the trace of Frobenius. For example, given the (inseparable) Frobenius endomorphism $\pi \in \mathrm{End}(E)$, lifting a separable endomorphism γ encoding π, such as $\gamma = \pi - 1$, would allow us to compute $\mathrm{tr}(\gamma)$ and thus $\mathrm{tr}(\pi) = \mathrm{tr}(\gamma) + 2$. However, the first step of this strategy would be to construct a liftable representation for γ, which seems impossible without already solving the point-counting problem in the process: Indeed, knowledge of the degree of γ is equivalent to knowledge of $\mathrm{tr}\,\pi$.

An alternative approach would be to compute a lift of the separable Verschiebung isogeny $\widehat{\pi}_p$ (the Verschiebung is separable, but of degree p, so we need to be a bit careful when lifting it). We can represent $\widehat{\pi}_p$ from its kernel polynomial h, using Vélu's formula, but this cost $O(p)$. Also, computing h is not obvious: it can be recovered from the division polynomial ψ_p, but computing ψ_p costs $O(p^2)$. (In [MRS25] an alternative method is given which is quasi-linear: $\widetilde{O}(p)$ arithmetic operations.)

A better idea is to compute an HD representation of $\widehat{\pi}_p$, this can be done in polynomial time in $\log p$. However, this involves the computation of π_p on enough small torsion points, so this initial step is already essentially Schoof's algorithm. In particular, it already gives the trace of π_p modulo p.

If $q = p$, we obtain the full trace, and there is no need to apply the implicit lifting method to $\widehat{\pi}_p$. But if $q = p^n$ for a large n, implicit lifting of the Verschiebung works and yield substantial improvements to Satoh's point-counting algorithm. We refer to [MRS25] for the resulting point counting algorithm when $\widehat{\pi}_p$ is computed via Vélu's formula from its kernel polynomial, and to [Rob22b, § 5] for a sketch of the algorithm when using an HD representation of $\widehat{\pi}_p$.

Remark 5.4. The algorithm from [MPSW25] can be seen as adapting the SEA algorithm, which computes the trace of the Frobenius, to the computation of the trace of an arbitrary morphism. Our algorithm follows a similar approach, except that it adapts Satoh's algorithm. A similar method should work to adapt Kedlaya's algorithm [Ked01], and Harvey's variant [Har07].

5.1. **Implementation & experiments.** Our lifting algorithms were implemented in SageMath 10.8 for Vélu isogeny chains and HD-embedded endomorphisms. The software is available at:

https://gitlab.inria.fr/isogenies/lifting_traces

For comparability with prior work, we approximately replicated the benchmarking strategy employed in [MPSW25]: By means of the Deuring correspondence, we generated a number of Vélu 2-isogeny chains over $\mathbb{F}_{p^2}$ of length approximately $4\log_2(p)$ for various random choices p for a given bit length. We then applied both the trace-computation function of [MPSW25] and our own implementation to those endomorphisms. For this particular size imbalance (≈ 4) between degree and characteristic, our implementation consistently outperforms [MPSW25]; see Table 1, with apparently superior asymptotic scaling, as predicted by theory. However, note that the difference diminishes when smaller size imbalances are considered; for instance, for endomorphism degrees $\leq p^2/16$ the required precision k is just 1, so computing the trace modulo p alone already suffices (in both algorithms).

TABLE 1. Example timings (in wall-clock seconds) for computing the trace of a separable endomorphism of degree $2^L \approx p^4$ defined over $\mathbb{F}_{p^2}$. (Using SageMath 10.8 on a single Intel "Alder Lake" core running at 2.1 GHz.)

$\approx \log_2(p)$	16	24	32	40	48
[MPSW25]	1.9	3.2	9.5	16.4	33.1
This work	0.7	1.2	2.0	2.9	4.3

REFERENCES

[BBM82] Pierre Berthelot, Lawrence Breen and William Messing. *Théorie de Dieudonné cristalline II*. Vol. 930. Lecture Notes in Mathematics. Springer, 1982.

[BDLS20] Daniel J. Bernstein, Luca De Feo, Antonin Leroux and Benjamin Smith. 'Faster computation of isogenies of large prime degree'. In: *Proceedings of the Fourteenth Algorithmic Number Theory Symposium (ANTS XIV)*. Vol. 4. Open Book Series, 2020, pp. 39–55. URL: https://ia.cr/2020/341.

[BMSS08] Alin Bostan, François Morain, Bruno Salvy and Éric Schost. 'Fast algorithms for computing isogenies between elliptic curves'. In: *Mathematics of Computation* 77.263 (2008), pp. 1755–1778. arXiv: cs/0609020.

[BRS25] Razvan Barbulescu, Damien Robert and Nicolas Sarkis. *Models of Kummer lines and Galois representations*. Preprint. Mar. 2025. URL: https://ia.cr/2025/543.

[CD21] Wouter Castryck and Thomas Decru. 'Multiradical isogenies'. In: *Arithmetic, Geometry, Cryptography, and Coding Theory*. Vol. 779. Contemporary Mathematics. AMS, 2021, pp. 57–89. URL: https://ia.cr/2021/1133.

[CDHV22] Wouter Castryck, Thomas Decru, Marc Houben and Frederik Vercauteren. 'Horizontal racewalking using radical isogenies'. In: *ASIACRYPT (2)*. Vol. 13792. Lecture Notes in Computer Science. Springer, 2022, pp. 67–96. URL: https://ia.cr/2022/1259.

[CDV20] Wouter Castryck, Thomas Decru and Frederik Vercauteren. 'Radical isogenies'. In: *ASIACRYPT (2)*. Vol. 12492. Lecture Notes in Computer Science. Springer, 2020, pp. 493–519.

[CRV14] Xavier Caruso, David Roe and Tristan Vaccon. 'Tracking-adic precision'. In: *LMS Journal of Computation and Mathematics* 17.A (2014), pp. 274–294.

[DeJ95] Arthur J. De Jong. 'Crystalline Dieudonné module theory via formal and rigid geometry'. In: *Publications Mathématiques de l'IHÉS* 82 (1995), pp. 5–96.

[DJP14] Luca De Feo, David Jao and Jérôme Plût. 'Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies'. In: *Journal of Mathematical Cryptology* 8.3 (2014), pp. 209–247. URL: https://ia.cr/2011/506.

[DMPR24] 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'. In: *ASIACRYPT (3)*. Vol. 15486. Lecture Notes in Computer Science. Springer, 2024, pp. 304–338. URL: https://ia.cr/2023/1747.

[DR72] Pierre Deligne and Michael Rapoport. 'Les schémas de modules de courbes elliptiques'. In: *Modular Functions of One Variable II*. Vol. 349. Lecture Notes in Mathematics. Antwerp, 1972, pp. 143–316.

[Elk92] Noam D. Elkies. *Explicit isogenies*. Manuscript. Unpublished; superseded by [Elk98]. Boston, MA, 1992.

[Elk98] Noam D. Elkies. 'Elliptic and modular curves over finite fields and related computational issues'. In: *Computational Perspectives on Number Theory: Proceedings of a Conference in Honor of A.O.L. Atkin*. Studies in Advanced Mathematics. AMS/IP, 1998. URL: https://people.math.harvard.edu/~elkies/modular.pdf.

[FC90] Gerd Faltings and Ching-Li Chai. *Degeneration of abelian varieties*. Ergebnisse der Mathematik und ihrer Grenzgebiete, 3. Folge, Band 22. Springer, 1990.

[Gal12] Steven D. Galbraith. *Mathematics of Public-Key Cryptography*. Cambridge University Press, 2012. URL: https://math.auckland.ac.nz/~sgal018/crypto-book/crypto-book.html.

[GD67] Alexander Grothendieck and Jean Dieudonné. 'Eléments de géométrie algébrique'. In: *Publications Mathématiques de l'IHÉS* 4, 8, 11, 17, 20, 24, 28, 32 (1960–1967), p. 1965.

[GG13] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. 3rd ed. Cambridge University Press, 2013.

[Gro70] Alexander Grothendieck. 'Groupes de Barsotti–Tate et cristaux'. In: *Actes du Congrès International des Mathématiciens*. Vol. 1. Nice, 1970, pp. 431–436.

[Har07] D. Harvey. 'Kedlaya's algorithm in larger characteristic'. In: *Int. Math. Res. Notices* (2007).

[HHL17] David Harvey, Joris van der Hoeven and Grégoire Lecerf. 'Faster polynomial multiplication over finite fields'. In: *Journal of the ACM* 63.6 (2017), pp. 1–23. arXiv: 1407.3361.

[Igu62] Jun-Ichi Igusa. 'On Siegel Modular Forms of Genus Two'. In: *American Journal of Mathematics* 84.1 (1962), pp. 175–200. URL: http://www.jstor.org/stable/2372812 (visited on 21/01/2026).

[Kat06] Nicholas Katz. 'Serre–Tate local moduli'. In: *Surfaces Algébriques: Séminaire de Géométrie Algébrique d'Orsay 1976–78*. Springer, 2006, pp. 138–202.

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

[KPR25] Jean Kieffer, Aurel Page and Damien Robert. 'Computing isogenies from modular equations between Jacobians of genus 2 curves'. In: *Journal of Algebra* 666 (Mar. 2025), pp. 331–386. arXiv: 2001.04137.

[KR24] Sabrina Kunzweiler and Damien Robert. 'Computing modular polynomials by deformation'. In: *Proceedings of the Sixteenth Algorithmic Number Theory Symposium (ANTS XVI)*. Vol. 11. Springer, 2024. arXiv: 2408.06990.

[Kun+25] Sabrina Kunzweiler, Luciano Maino, Tomoki Moriya, Christophe Petit, Giacomo Pope, Damien Robert, Miha Stopar and Yan Bo Ti. 'Radical 2–isogenies and cryptographic hash functions in dimensions 1, 2 and 3'. In: *PKC 2025*. Vol. 15676. Lecture Notes in Computer Science. Springer, May 2025, pp. 265–299. URL: https://ia.cr/2024/1732.

[Mes72] William Messing. 'The crystals associated to Barsotti–Tate groups'. In: *The crystals associated to Barsotti–Tate groups: with applications to abelian schemes*. Springer, 1972, pp. 112–149.

[Mil86] James S. Milne. 'Abelian Varieties'. In: *Arithmetic Geometry*. Ed. by Gary Cornell and Joseph H. Silverman. Springer, 1986, pp. 103–150.

[MPSW25] Travis Morrison, Lorenz Panny, Jana Sotáková and Michael Wills. *The SEA algorithm for endomorphisms of supersingular elliptic curves*. Preprint. Jan. 2025. arXiv: 2501.16321.

[MRS25] Abdoulaye Maiga, Damien Robert and Djiby Sow. 'Towards computing canonical lifts of ordinary elliptic curves in medium characteristic'. In: *Designs, Codes and Cryptography* (Aug. 2025). DOI: https://doi.org/10.1007/s10623-025-01719-4.

[Oda69] Tadao Oda. 'The first de Rham cohomology group and Dieudonné modules'. In: *Annales scientifiques de l'École Normale Supérieure* 2.1 (1969), pp. 63–135.

[Onu21] Hiroshi Onuki. 'On oriented supersingular elliptic curves'. In: *Finite Fields and Their Applications* 69 (2021). DOI: 10.1016/j.ffa.2020.101777. URL: https://arxiv.org/abs/2002.09894.

[PAR21] The PARI Group. *PARI/GP*. Version 2.14. Université de Bordeaux, 2021. URL: https://pari.math.u-bordeaux.fr/.

[Ren18] Joost Renes. 'Computing isogenies between Montgomery curves using the action of $(0,0)$'. In: *Post-Quantum Cryptography (PQCrypto 2018)*. Springer, 2018, pp. 229–247. URL: https://ia.cr/2017/1198.

[Rob21] Damien Robert. 'Efficient algorithms for abelian varieties and their moduli spaces'. Habilitation à Diriger des Recherches. Université de Bordeaux, June 2021. URL: https://hal.science/tel-03498268.

[Rob22a] Damien Robert. *Evaluating isogenies in polylogarithmic time*. Preprint. Aug. 2022. URL: https://ia.cr/2022/1068.

[Rob22b] Damien Robert. 'Some applications of higher dimensional isogenies to elliptic curves (overview of results)'. Dec. 2022. eprint: 2022/1704.

[Rob23] Damien Robert. *The geometric interpretation of the Tate pairing and its applications*. Preprint. Feb. 2023. URL: https://ia.cr/2023/177.

[Rob24a] Damien Robert. *Fast pairings via biextensions and cubical arithmetic*. Preprint. Apr. 2024. URL: https://ia.cr/2024/517.

[Rob24b] Damien Robert. 'On the efficient representation of isogenies (a survey)'. In: *Number-Theoretic Methods in Cryptology (NuTMiC)*. Vol. 14966. Lecture Notes in Computer Science. Springer, 2024, pp. 3–84. URL: https://ia.cr/2024/1071.

[Sat00] Takakazu Satoh. 'The canonical lift of an ordinary elliptic curve over a finite field and its point counting'. In: *Journal of the Ramanujan Mathematical Society* 15 (2000), pp. 247–270.

[Sch95] René Schoof. 'Counting points on elliptic curves over finite fields'. In: *Journal de Théorie des Nombre de Bordeaux* 7.1 (1995), pp. 219–254. URL: https://www.mat.uniroma2.it/~schoof/ctg.pdf.

[Stacks] *The Stacks project*. https://stacks.math.columbia.edu. 2008–today.

[Vél71] Jacques Vélu. 'Isogénies entre courbes elliptiques'. In: *Comptes Rendus de l'Académie des Sciences de Paris* 273 (1971), pp. 238–241.

[VPV01] Frederik Vercauteren, Bart Preneel and Joos Vandewalle. 'A Memory Efficient Version of Satoh's Algorithm'. In: *EUROCRYPT*. Vol. 2045. Lecture Notes in Computer Science. Springer, 2001, pp. 1–13. DOI: 10.1007/3-540-44987-6_1.

Appendix A. Other isogeny representations

In Section 4 we saw different representations of chains of isogenies: via their kernels, generators of their kernels, via j-invariants and modular polynomials, via radical isogenies, or via the HD representation.

Then in Section 5 we saw how to recover the trace by computing the action on differentials, which can be done by any evaluation algorithm applied to a suitable infinitesimal point.

In this appendix, we explore an alternative approach, which relies on using modular forms. Although this is a classical topic, to the best of our knowledge Appendix A.4 explains for the first time how to combine radical isogenies with modular forms.

A.1. Modular invariants and radical isogenies.

In this section, we expand on Section 4.1 on the use of modular invariants (respectively, sequences of modular invariants) to represent an isogeny (resp. isogeny chains).

Using j-invariants. As mentioned in Section 4.1, an ℓ-isogeny $\varphi\colon E_1 \to E_2$ is uniquely determined by specifying just ℓ along with the j-invariants $j(E_1), j(E_2)$ of the domain and codomain elliptic curves, as long as the modular polynomial Φ_ℓ is smooth at $j(E_1), j(E_2)$ and the j-invariants are different from $0, 1728$. Indeed, under these conditions we can set up from $j(E_1), j(E_2)$ a differential equation that allows us to reconstruct the isogeny φ. This idea is due to Elkies [Elk92], and was exploited by [BMSS08] for fast isogeny computation algorithms.

Geometrically, the isogeny $\varphi\colon E_1 \to E_2$ is encoded by the associated moduli point $(E_1, \ker\varphi)$ in the algebraic stack $\mathfrak{A}_1(\Gamma^0(\ell))$ (see Example 2.12). There is an étale modular correspondence $\mathfrak{A}_1(\Gamma^0(\ell)) \to \mathfrak{A}_1 \times \mathfrak{A}_1, (E_1, \ker\varphi) \mapsto (E_1, E_1/\ker\varphi)$ over $\mathbb{Z}[1/\ell]$, which induces a modular correspondence $\Psi_\ell\colon \mathcal{A}_1(\Gamma^0(\ell)) \to \mathcal{A}_1 \times \mathcal{A}_1 \simeq \mathbb{A}^1 \times \mathbb{A}^1$ at the level of the coarse moduli space. (Furthermore, if $j(E_i) \neq 0, 1728$, the modular correspondence is strongly étale by the same argument as in [KPR25, § 4.2], so it suffices to look at the coarse modular correspondence.) The fact that φ can be recovered from the j-invariants is due to the fact that zeroes of the modular polynomial $\Phi_\ell(X, Y)$ describe the image of this modular correspondence. Indeed, whenever $(j(E_1), j(E_2))$ is a smooth point of the curve defined by Φ_ℓ, then Ψ_ℓ is a local isomorphism around this point.

Radical isogenies. The arguments above can be extended beyond Φ_ℓ to different modular polynomials stemming from different modular invariants, possibly with an extra level structure. Radical isogenies, first introduced in [CDV20], may be seen as such an extension.

Suppose given a curve E_1 and a rational point $P_1 \in E_1[\ell]$. The point P_1 generates the kernel of a cyclic isogeny $\varphi_1\colon E_1 \to E_2$ of degree ℓ. The idea of radical isogenies is to pick up a new ℓ-torsion point $P_2 \in E_2[\ell]$ on the codomain, such that the corresponding isogeny φ_2 is not backtracking with respect to φ_1. One can show that the set of such P_2 is parametrized by choices of ℓth root of the non reduced self Tate pairing $e_{T,\ell}(P_1, P_1)$; we refer to [CD21; CDHV22; Rob23] for more details.

Geometrically, radical isogenies encode extra level structure as follows. Consider the moduli space $\mathfrak{A}_1(\Gamma^1(\ell))$ that parametrizes pairs (E, P) with $P \in E[\ell]$ (unlike with $\Gamma^0(\ell)$, here the points of $\langle P\rangle$ are not identified together). Secondly, let $\Gamma'(\ell)$ be the congruence subgroup (containing $\Gamma^1(\ell^2)$) such that $\mathfrak{A}_1(\Gamma'(\ell))$ parametrizes pairs (E, P') with $P' \in E[\ell^2]$ modulo translation by $\ell P'$. For such a point P', the group $\langle \ell P'\rangle$ is well defined, and if φ is the corresponding isogeny, $\varphi(P')$ is also well defined. There is a modular correspondence $\mathfrak{A}_1(\Gamma'(\ell)) \to \mathfrak{A}_1(\Gamma^1(\ell)) \times \mathfrak{A}_1(\Gamma^1(\ell))$ that sends

$(E, P' + \langle \ell P' \rangle)$ to $\big((E, \ell P'), (\varphi(E), \varphi(P' + \langle \ell P' \rangle))\big)$. In this modular correspondence, the projection to the first variable defines a Kummer extension, hence why the preimages are parametrized by ℓth roots. This parametrization has a more precise geometric description (see [Rob23]): if (E, P) is the universal elliptic curve over $\mathfrak{A}_1(\Gamma^1(\ell))$, then $\mathfrak{A}_1(\Gamma'(\ell))$ is isomorphic to the μ_ℓ torsor associated to the Tate pairing $e_{T,\ell}(P, P)$.

Example A.1 (Radical isogenies in the Montgomery model). When searching for cycles in the supersingular isogeny graph, using 2-radical isogenies for Montgomery curves is more convenient than using the classical modular polynomial Φ_2. We haven't been able to find radical isogeny formulas for the Montgomery model in the literature, so we include them for completeness. They can be easily found by combining the 2-radical isogeny formulas in the theta model [Kun+25] and the conversion map between the theta model and the Montgomery model (see for instance [BRS25, § 4]).

Let $E\colon By^2 = x^3 + Ax^2 + x$ be a Montgomery curve, and $P = \big(x(P), 0\big) \in E[2]$ be a two-torsion point with $x(P) \neq 0$. Using the identity $A = -x(P) - 1/x(P)$, we see that the knowledge of $x(P) = \big(X(P) : Z(P)\big)$ is enough to recover the Kummer line $E/\pm 1$, as the latter does not depend on B. Let $(X_2 : Z_2)$ be given by

$$\begin{cases} X_2 = X(P) + \gamma \\ Z_2 = X(P) - \gamma \end{cases} \qquad \gamma = \sqrt{(X(P) + Z(P))(X(P) - Z(P))}.$$

Depending on the choice of square root, we get either $x_2 = X_2/Z_2$ or its inverse $1/x_2$. Either way, The point $P_2 = (X_2 : 0 : Z_2)$ defines a 2-torsion point on the Montgomery curve $E_2 = E/\langle P \rangle$. The assumption that $x(P) \neq 0$ ensures that E_2 is a well-defined Montgomery curve. The isogeny with kernel $T = (0, 0) \in E_2[2]$ gives the dual isogeny $E_2 \to E$.

A.2. **Modular forms.** Let $\varphi\colon E_1 \to E_2$ be an isogeny of elliptic curves. Given invariant differentials ω_1, ω_2 on E_1, E_2 respectively, the isogeny φ acts on differentials via its scaling factor:

$$\varphi^* \omega_2 = c_\varphi \cdot \omega_1.$$

We discussed in Section 3.2 how to compute this action on differentials for a general φ. Here, we discuss the special case where φ is encoded as in Appendix A.1 by the j-invariants $j(E_1), j(E_2)$. As mentioned in Section 4.1 and appendix A.1, we can use the modular polynomial $\Phi_\ell(X, Y)$ to reconstruct the rational equation of isogeny φ. While the equations of φ are clearly enough to compute the action on differentials, it is possible to skip the reconstruction step and compute the action directly from Φ_ℓ, using modular forms. We first review the theory of modular forms.

For elliptic curves over a general local ring (not necessarily $\mathbb{C}$), we can define a modular form of weight k as a function g on tuples (E, ω_E) with ω_E an invariant differential on E, such that if $\alpha\colon E' \to E$ is an isomorphism with $\alpha^* \omega_E = \lambda \omega_{E'}$, then $g(E', \omega_{E'}) = \lambda^k g(E, \omega_E)$.

More formally, let $\mathfrak{h}$ be the Hodge line bundle, defined as follows: if $\pi\colon E \to \mathfrak{A}_1$ is the universal elliptic curve, with section ε, then $\mathfrak{h} = \pi_* \Omega_{E/\mathfrak{A}_1} = \varepsilon^* \Omega_{E/\mathfrak{A}_1}$. A modular form g of weight k is a section of $\mathfrak{h}^{\otimes k}$ that extends to a suitable compactification of the moduli space $\mathfrak{A}_1$ ensuring g is well-defined on the Tate curve. This algebraic well-definedness condition corresponds to the analytic condition of holomorphy at the cusp.

Example A.2. Let a_4, a_6 be modular forms defined as follows over $\mathbb{Z}[1/6]$. To an elliptic curve E/k and global differential ω_E, pick up a short Weierstraß equation

$E\colon y^2 = x^3 + a_4x + a_6$ for E such that $\omega_E = \mathrm{d}x/2y$. Then we define $a_4(E, \omega_E) = a_4$ and $a_6(E, \omega_E) = a_6$.

We remark that the isomorphism $\eta\colon E_1 \to E_2$ defined by $(x, y) \mapsto (x/u^2, y/u^3)$ gives a new Weierstraß equation $E_2\colon y^2 = x^3 + a/u^4x + b/u^6$, $\frac{\mathrm{d}x}{2y}$ on E_2 pulls back to $u\frac{\mathrm{d}x}{2y}$ on E_1. We see that a_4, a_6 are of weight 4 and 6 respectively.

The modular forms a_4, a_6 are the algebraic version of the holomorphic Eisenstein series $G_4(\tau), G_6(\tau)$ from the analytic theory of modular curves. In what follows, we make the correspondence more precise.

Define the *normalized* Eisenstein series $E_4(\tau) = \frac{G_4(\tau)}{2\zeta(4)} = \frac{45}{\pi^4}G_4(\tau)$, $E_6(\tau) = \frac{G_6(\tau)}{2\zeta(6)} = \frac{945}{2\pi^6}G_6(\tau)$. The q-expansion of E_4, E_6 is given by primitive integral coefficients, hence they are defined over all $\mathbb{Z}$ by the q-expansion principle. Given τ, the elliptic curve $E_\tau = \mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$ satisfies the equation $\wp'^2(z) = 4\wp^3(z) - g_2(\tau)\wp(z) - g_3(\tau)$, with $g_2(\tau) = 60\,G_4(\tau), g_3(\tau) = 140\,G_6(\tau)$. Setting $x = \wp(z)$ and $y = \wp'(z)$ in this equation, the invariant differential $\mathrm{d}z$ on E_τ corresponds to $\mathrm{d}x/y$.

We now make a change of variable to obtain a curve defined over $\mathbb{Q}(E_4, E_6)$: via $(x, y) \mapsto ((2\pi\mathrm{i})^{-2}x, (2\pi\mathrm{i})^{-3}y)$, the curve E_τ is mapped to $E'_\tau\colon y^2 = x^3 - \frac{E_4(\tau)}{48}x + \frac{E_6(\tau)}{864}$, where this time $2\pi\mathrm{i}\,\mathrm{d}z$ corresponds to $\frac{\mathrm{d}x}{2y}$. If we see E_4, E_6 as algebraic modular forms given analytically by $E_4(E_\tau, 2\pi\mathrm{i}\,\mathrm{d}z) = E_4(\tau)$, $E_6(E_\tau, 2\pi\mathrm{i}\,\mathrm{d}z) = E_6(\tau)$, then we have

$$a_4(E, \omega_E) = -E_4(E, \omega_E)/48, \qquad a_6(E, \omega_E) = E_6(E, \omega_E)/864.$$

We see that the short Weierstraß coefficients a_4, a_6 have bad reduction at $2, 3$ as expected.

We refer to the discussion in [FC90, p. 141, 142] as to why it is the analytic differential $2\pi\mathrm{i}\,\mathrm{d}z$ which is algebraic, rather than the differential $\mathrm{d}z$. We will call the first one the canonical algebraic differential, and the second one the canonical analytic differential. We warn the reader that in the literature on modular forms, given an algebraic modular form g of weight k, then over the upper half plane, it is often the 'analytic' modular form $G_{\mathrm{an}}(\tau) = g(E_\tau, \mathrm{d}z)$ that is studied rather than the 'algebraic' one $G_{\mathrm{alg}}(\tau) = g(E_\tau, 2\pi\mathrm{i}\,\mathrm{d}z) = 1/(2\pi\mathrm{i})^k G_{\mathrm{an}}(\tau)$. For instance, the discriminant Δ is a weight 12 modular form defined algebraically on a curve $E\colon y^2 = x^3 + a_4x + a_6$ by the formula $\Delta(E, \mathrm{d}x/2y) = -16(4a_4^3 + 27a_6^2)$. Seen as a modular form on τ, the analytic version is usually given: $\Delta_{\mathrm{an}}(\tau) = \Delta(E_\tau, \mathrm{d}z) = \Delta\big(y^2 = x^3 - \frac{g_2(\tau)}{4}x - \frac{g_3(\tau)}{4}, \frac{\mathrm{d}x}{2y}\big) = g_2(\tau)^3 - 27g_3(\tau)^2$. Compare with the algebraic version:

$$\begin{aligned}\Delta_{\mathrm{alg}}(\tau) &= \Delta(E_\tau, 2\pi\mathrm{i}\,\mathrm{d}z)\\ &= \Delta\left(y^2 = x^3 - (E_4(\tau)/48)x + E_6(\tau)/864,\ \frac{\mathrm{d}x}{2y}\right)\\ &= (E_4^3 - E_6^2)/1728\end{aligned}$$

Indeed, since Δ is of weight 12, we have $\Delta_{\mathrm{an}}(\tau) = (2\pi\mathrm{i})^{12}\Delta_{\mathrm{alg}}(\tau)$.

Example A.3. Although $j(\tau)$ is of weight 0 (it gives an isomorphism invariant), we can recover a modular form of weight 2 from j as follows. Write the q-expansion $j(\tau) = J(q)$ where $q = e^{2\pi\mathrm{i}\tau}$, and consider its derivative $j'(\tau) = \frac{\mathrm{d}J}{\mathrm{d}q}(q)$. We give now an algebraic interpretation of j'. The Kodaira–Spencer isomorphism gives a canonical isomorphism between the tangent space at E in $\mathfrak{A}_1$ and $\mathrm{Sym}^2(\mathrm{Lie}(E))$. That is, given a differential ω_E on E/k, we can associate a point $T_E\mathfrak{A}_1$ on the tangent space, which depends only on $\omega_E^{\otimes 2}$. Pulling back the universal elliptic curve to the point on the tangent space, we obtain a deformation $\widetilde{E}$ of E to $\operatorname{Spec} k[\varepsilon]/(\varepsilon^2)$. Then $j(\widetilde{E}) = j(E) + \varepsilon j'(E, \omega_E)$.

By [Sch95, Proposition 7.1], we have $j'/j = -E_6/E_4$ and, by Example A.2, $j' = -18\frac{a_6}{a_4}j$. In particular, given a_4, a_6 we can recover j, j', and conversely if $j \neq 0, 1728$. We remark that j' is not an entire modular form, as it has a pole at the curve $j = 0$.

A.3. **The action on differentials from modular forms.** A modular correspondence gives a relation between modular *functions* (i.e., modular forms of weight 0) of isogenous elliptic curves. For instance, whenever E_1, E_2 are ℓ-isogenous, their j-invariants are tied by the relation $\Phi_\ell(j(E_1), j(E_2)) = 0$ given by the classical modular polynomial.

Now let instead g be a modular form of weight $k > 0$. We want to define an "affine" modular correspondence[5] that doesn't only take into account the isomorphism *class* of a curve E, but rather the curve E itself (in other words, we keep track of its invariant differential). More concretely, we look for a polynomial $\Phi_{g,\ell}$ such that whenever we have $(E_1, \omega_1), (E_2, \omega_2)$ and an isogeny $\varphi\colon E_1 \to E_2$ with $\varphi^*\omega_2 = \omega_1$ (we say that ω_2 is *normalized* with respect to ω_1), then

$$\Phi_{g,\ell}\big(g(E_1, \omega_1), g(E_2, \omega_2), j(E_1), j(E_2)\big) = 0.$$

To find $\Phi_{g,\ell}$, we first observe that, by linearity of φ^*, given λ a scalar, $\lambda\omega_2$ is normalized with respect to $\lambda\omega_1$, so $\Phi_{g,\ell}$ must be homogeneous with respect to the first two variables. Because of this, we are allowed to represent g in either its analytic or algebraic version: we only need to know g up to an unspecified scalar, as long as the same scalar is the same for E_1 and E_2. Analytically, the ℓ-isogeny $E_\tau \to E_{\tau/\ell}$ defined by $z \mapsto z$ is normalized, so we must have $\Phi_{g,\ell}(g(\tau), g(\tau/\ell), j(\tau), j(\tau/\ell)) = 0$.

Now suppose that we know $\Phi_{g,\ell}$, and that we have a way to evaluate g algebraically. Then we can use the affine modular polynomial $\Phi_{g,\ell}$ to recover the action on differentials of $\varphi\colon E_1 \to E_2$ given only $j(E_1), j(E_2)$ as follows. First, pick up some arbitrary short Weierstraß equation for E_1 with j-invariant $j(E_1)$, and compute $g_1 = g(E_1, \frac{\mathrm{d}x}{2y})$, and do the same for E_2 to get g_2. Then compute $\overline{g}_2$ by solving $\Phi_{g,\ell}(g_1, \overline{g}_2, j_1, j_2) = 0$. As g is of weight k, we must have $g_2 = \lambda^k\overline{g}_2$ for some λ; moreover, the differential on E_2 that is normalized with respect to φ is of the form $\zeta\lambda\frac{\mathrm{d}x}{2y}$ on E_2, for some unknown $\zeta \in \mu_k$.

More generally, given $G = (g_1, \dots, g_m)$ a tuple of modular forms of weights $(k_1, \dots, k_m)$, we could use a similar approach and consider an affine G-modular correspondence given by multivariate system of modular equations $\Phi_{G,\ell}(G(\tau), G(\tau/\ell)) = 0$. This would allow to recover the action on differential up to a $\zeta \in \mu_k$ with $k = \gcd\{k_i\}_i$. The case above is $G = (g, j)$.

Example A.4. Differentiating the modular equation $\Phi_\ell(j(\tau), j(\tau/\ell)) = 0$ gives an affine modular equation for (j, j'):

$$j'(\tau)\frac{\partial\Phi_\ell}{\partial X}\Big(j(\tau), j(\tau/\ell)\Big) + j'(\tau/\ell)\frac{1}{\ell}\frac{\partial\Phi_\ell}{\partial Y}\Big(j(\tau), j(\tau/\ell)\Big) = 0.$$

Using the method outlined above, if $\varphi_\ell(j_1, j_2) = 0$, we can compute normalized equations for the associated elliptic curve E_1, E_2 as follows.

First we pick up any short Weierstraß equation $y^2 = x^3 + a_4x + a_6$ for E_1, and compute from this $j_1' = j'(E_1, \frac{\mathrm{d}x}{2y}) = -18\frac{a_6}{a_4}j\big(E_1, \frac{\mathrm{d}x}{2y}\big)$ using Example A.3. Then we solve the linear equation $j_1'\frac{\partial\varphi_\ell}{\partial X}(j_1, j_2) + j_2'\frac{1}{\ell}\frac{\partial\varphi_\ell}{\partial Y}(j_1, j_2) = 0$ to recover $j_2' = j'(E_2, \omega_2)$ where $\omega_2 = \frac{\mathrm{d}x}{2y}$ is the normalized differential associated to a specific short Weierstraß equation of E_2. We use Example A.3 again to compute the a_4, a_6 coefficients of

[5] While the standard modular correspondence lives in the moduli spaces $\mathfrak{A}_1$, the 'affine' modular correspondence lives in the Hodge line bundle $\mathfrak{h}^{\otimes k}$ above it, where g is seen a coordinate on $\mathfrak{h}^{\otimes k}$.

this short Weierstraß equation for E_2 from j_2, j_2'. We note that we only recover the normalized differential on E_2 up to a sign, but this sign ambiguity is inherent due to the fact that $-\varphi$ is also an isogeny from E_1 to E_2.

As mentioned above, the idea that we can recover the action on differentials directly from $j(E_1), j(E_2)$ by differentiating the modular polynomial φ_ℓ is due to Elkies. The isogeny rational reconstruction of [BMSS08] first applies the method outlined above to find normalized equations for E_1, E_2, which then gives an explicit differential equation satisfied by φ, whose resolution allows to recover φ as a rational function in the Weierstraß coordinates of E_1. We refer to [KPR25] for a generalization of this method to dimension 2.

As mentioned in Section 4.2, the methods discussed above can be used to compute a normalized isogeny from a chain of j-invariants: we choose a short Weierstraß equation for the domain, then step by step we find an equation for the intermediate curves such that each intermediate isogeny of the resulting chain $E_i \to E_{i+1}$ is normalized. As mentioned in the proof of Corollary 5.2, whenever we have an endomorphism $\gamma \in \mathrm{End}(E)$ expressed as an isogeny chain steps followed by an isomorphism, having normalized steps is a computational advantage for trace computation. Indeed, the action on differentials (i.e., the scaling factor, hence the trace) of γ is then obtained for free: suppose E is given by the equation $y^2 = x^3 + ax + b$. The curve E_n obtained at the last step is isomorphic to E, so it will be given by an equation $y^2 = x^3 + au^4x + bu^6$. Then u (which we can recover up to a sign), gives the action on differentials of γ.

Example A.5. Satoh's p-adic lifting algorithm for point counting usually proceeds in two steps: first computing the canonical lift $\widehat{E}/\mathbb{Z}_q$ of an ordinary curve $E/\mathbb{F}_q$, using the modular polynomials Φ_p via Newton's method, and then lifting the kernel of the Verschiebung to compute the associated isogeny and its action on differentials. But, by Example A.4, the second step is unnecessary and we just need to compute the derivate of the modular polynomial Φ_p on $\widehat{E}$ to recover the action (in some p-adic precision). In practice this gives large speedups (see [Rob21, § 6.6]): a factor $\times 32$ less time and $\times 64$ less memory for point counting over $\mathbb{F}_{101^{418}}$ in Pari [PAR21].

A.4. **Modular forms and radical isogenies.** A similar approach to Example A.4 can be employed to recover the action on differentials from radical isogeny (without reconstructing the isogeny)! We illustrate this with the 2-radical formulas from Example A.1.

First, we need to introduce suitable modular forms. In isogeny based cryptography, it is customary to represent the Kummer line associated to a Montgomery curve $E\colon By^2 = x^3 + Ax^2 + x$ by the Montgomery coefficient represented as projective coordinates: $A = (\mathcal{A} : \mathcal{C})$. But in practice, in the computer, one represent these coefficients by actual values, i.e., we work with an affine point $(\mathcal{A}, \mathcal{C})$ rather than the associated projective point. The key idea is that we may interpret these values as modular forms. (This is highly reminiscent of the fact that interpreting the $(X(P) : Z(P))$ projective coordinates of a point P ont he Kummer line as affine coordinates of some affine point $\widetilde{P}$ can be interpreted in terms of cubical arithmetic, see [Rob24a]. In both cases we lift a point on a scheme to a point on a line bundle above it.) Note however that the Montgomery form requires a $\Gamma^0(4)$-level structure, hence $(\mathcal{A}, \mathcal{C})$ will be modular forms of level $\Gamma^0(4)$ (and weight 2).

One way to define them is as follows: let $(E, \pm T, \omega_E)$ be an elliptic curve, a point of 4-torsion up to sign (this is the same thing as specifying a cyclic kernel of degree 4, hence a $\Gamma^0(4)$ level-structure), and a differential. Then there is a unique

Montgomery equation $By^2 = x^3 + Ax^2 + x$ such that $x(\pm T) = 1$, $x(2 \pm T) = 0$, and $\omega_E = \mathrm{d}x/2y$. We define $\mathcal{A}(E, \pm T, \omega_E) = AB$ and $\mathcal{C}(E, \pm T, \omega_E) = B$.

Now, in the radical 2-isogeny formulas, we work with the projective coordinates $(X(P) : Z(P))$. We can also interpret them as modular forms (of weight 1 and level $\Gamma^0(4) \cap \Gamma(2)$), by asking that $\mathcal{A} = X^2 + Z^2$ and $\mathcal{C} = -XZ$. (This leaves a sign ambiguity for X, Z, but in practice we just use them as proxies for $\mathcal{A}, \mathcal{C}$).

We can now explain how to compute the following affine/modular radical 2-isogeny formulas.

Example A.6. Start with $X(P), Z(P)$ where X, Z are seen as modular forms as above. More precisely, these are the values they take on the tuple $\big(E, (\pm T', P), \omega_E\big)$ where $E\colon By^2 = x^3 + Ax^2 + x$ is our starting curve, $(\pm T', P)$ is our $\Gamma^0(4) \cap \Gamma(2)$ level structure given by $x(T') = 1$, and P our 2-torsion point, and $\omega_E = \mathrm{d}x/2y$.

Let γ a square root of $(X(P) + Z(P))(X(P) - Z(P))$, and define $X_2 = X(P) + \gamma$ and $Z_2 = X(P) - \gamma$ as in Example A.1. The values X_2, Z_2 define a two-torsion point P_2 on the codomain $E_2 = E/\langle P \rangle$, via $x(P_2) = X_2/Z_2$. We also get the coefficient $A' = -x(P_2) - 1/x(P_2)$ of the Montgomery curve equation $B'y^2 = x^3 + A'x^2 + x$ of E_2. Here the Montgomery equation comes from the level $\Gamma^0(4)$-structure given by $\varphi(T')$ where $\varphi\colon E \to E_2$ is the isogeny, and of course adding P_2 gives us a level $\Gamma^0(4) \cap \Gamma(2)$ on E_2. We now want to determine the coefficient B' of E_2 such that the associated differential $\mathrm{d}x/2y$ is (up to a sign) the normalized differential ω_2 for φ. For that, we need to compute $X_2' = X\big(E_2, (\varphi(T'), P_2), \omega_2\big)$ and $Z_2' = Z\big(E_2, (\varphi(T'), P_2), \omega_2\big)$ Using the normalized isogeny formulas from [Ren18], we can compute $X_2' = X_2 \cdot x(P)$ and $Z_2' = Z_2 \cdot x(P)$. From these, we recover A', B' as we wanted.

A remark, due to Sabrina Kunzweiler, is that we can rewrite X_2', Z_2' as $(X_2', Z_2') = (X(P)x(P_1), X(P)x(P_2))$ where $x(P_1) \neq x(P_2)$ are the two different x-coordinates of the 4-torsion points above P.

As in Example A.4, we can use this to compute the action on differentials of an endomorphism γ given by a cycle of 2-radical isogenies as follows. We start with P_1 and the modular form $X(P_1), Z(P_1)$ associated to the initial curve equation $E\colon By^2 = x^3 + Ax^2 + x$ and we use the modular radical isogeny formulas at each step to find at the end $X(P_n), Z(P_n)$, hence a curve equation $E_n\colon B'y^2 = x^3 + A'x^2 + x$. Now, although the coefficients B, B' are given by the modular form $\mathcal{C}$ of weight 2, unlike Example A.4 we do not necessarily have $B' = Bu^2$ where u is the action on differentials because γ needs not preserve the initial $\Gamma^0(4)$ structure on E. We first need to compute an automorphism of E that maps the initial level structure to the one[6] on E_n, i.e., find the isomorphism between the curve equations E and E_n. The action of this isomorphism on the differentials then gives us the action of γ on differentials.

Remark A.7. The way we defined the modular forms $\mathcal{A}, \mathcal{C}$ above is of course somewhat arbitrary. We could scale them by the same factor $\lambda \neq 0$, and still use them to compute normalized differentials (as long as we use the same λ for the domain and codomain of course).

The definition we chose is one that is quite simple to describe algebraically. Using the transformation formula from theta coordinates to the Montgomery model, we can also write the Montgomery coefficient A in terms of level 2 theta constants: $A = 2(\theta_0(0)^4 + \theta_1^4)/(\theta_0(0)^4 - \theta_1(0)^4)$ so it would be natural to take the numerator and denominator of this formula as modular forms. We have not tried to express

[6] Since $\Gamma^0(4)$ is of index 6 in $\Gamma = \mathrm{Sl}_2(\mathbb{Z})$, there are 6 possibilities, but due to the way radical isogenies are computed there are only 2 possibilities for the level structure on E_n. We refer to the implementation for details.

our $\mathcal{A},\mathcal{C}$ in terms of these theta constants, but it could be interesting. (Note that a level 2 theta structure corresponds in dimension 1 to a $\Gamma^0(4)\cap\Gamma_0(4)$ level structure, hence encodes more information than the Montgomery level structure. We thus have several different conversion formulas depending on which part of the level structure we map from theta to Montgomery.)

Finally, we note that the mention of cubical coordinates in Example A.6 is deeper than just an analogy. Indeed, if $\mathcal{L}$ is the universal principal line bundle on the universal elliptic curve $\pi\colon E\to\mathfrak{A}_1$, then $\pi_*\mathcal{L}^{\otimes -8}=\mathfrak{h}^{\otimes 4}$ by [FC90, Theorem I.5.1]. So as explained in [FC90, Theorem I.5.1], sections of $\mathcal{L}$ (or some power of $\mathcal{L}$) evaluated at point of n torsion give a way to defined modular forms. In the modular 2-radical formulas, we interpret $X(P),Z(P)$ as modular forms, but we could also interpret them as cubical coordinates of a suitable cubical point $\widetilde{P}$. It would be interesting to work out the link between the two aspects in more details.

Appendix B. Deformation of abelian varieties

The goal of this section is to give more details on the proof of theorem 3.4. We would like to describe the lifts of an abelian variety $A/\mathbb{F}_q$ to $\mathbb{Z}_q$. By standard inductive limit arguments and Grothendieck's algebraization, it suffices to understand the lifts (also called deformation) to $\mathbb{Z}_q/p^m\mathbb{Z}_q$ (see [Mes72, Theorem V.3.3] for more details).

The theory is well-understood even for the general case of an abelian scheme A/S. First, by Serre–Tate, if $S\to T$ is an infinitesimal thickening where p is locally nilpotent on S, then there is an equivalence of categories between abelian schemes $\widetilde{A}/T$ and tuples $(A/S,G,i)$ where $A=\widetilde{A}_S$, G/T is a p-divisible/Barsotti–Tate group and i is an isomorphism over S of G with the p-divisible group $A(p)$ of A: That is, $i\colon G/S\simeq A(p)$. (See [Kat06] for a proof.)

So deforming A/S to T amounts to deforming its p-divisible group $A(p)$ to T. We can now use Grothendieck–Messing theory [Mes72] to understand the deformations of G over a locally nilpotent divided power thickening $S\to T$: a p-divisible group G admits a Dieudonné crystal $\mathbb{D}(G)$ (a quasi-coherent sheaf in the crystalline site with an action of F and V), along with a Hodge filtration of $\mathbb{D}(G)_S$. And deforming G to S amounts to a choice of deforming this Hodge filtration to a direct summand of $\mathbb{D}(G)_{S\to T}$ ([Gro70, § 4]).

There are several ways to construct the crystal $\mathbb{D}(G)$, Messing in [Mes72] use the Lie algebra of the universal vectorial extension of G. We will instead describe the construction due to [BBM82]. Assume that G/S is a p-divisible group, where p is locally nilpotent on S. Let $\Sigma=\operatorname{Spec}\mathbb{Z}_p$ and $\operatorname{Crys}(S/\Sigma)$ be the big crystalline site on S [BBM82, § 1.1]. Then $\mathbb{D}(G)=\mathcal{EXT}^1_{S/\Sigma}(G,O_{S/\Sigma})$ is a crystal [BBM82, Théorème 3.3.3], where $\mathcal{EXT}$ is the internal ext sheaf. It is in fact a crystalline Dieudonné module, i.e., a crystal in quasi-coherent locally free module with an action of F,V. Furthermore, there is a Hodge filtration $0\to\omega_G\to\mathbb{D}(G)_S\to\operatorname{Lie}(G^*)\to 0$ [BBM82, Corollaire 3.3.5] where $\mathbb{D}(G)_S$ is the value of the crystal $\mathbb{D}(G)$ at $S\to S$, ω_G is the sheaf of invariant differentials and G^* is the dual p-divisible group of G. In particular, if $S\to T$ is a divided power thickening, and $\widetilde{G}/T$ a deformation of G to T, then since $\mathbb{D}(G)$ is a crystal, we have $\mathbb{D}(G)_{S\to T}\simeq\mathbb{D}(\widetilde{G})_T$, and the Hodge filtration on $\mathbb{D}(\widetilde{G})_T$ lift the one on $\mathbb{D}(G)_S$. By Grothendieck–Mazur, the converse is true, at least when the divided powers are locally nilpotent: any admissible filtration $\operatorname{Fil}^1\subset\mathbb{D}(\widetilde{G})_T$ comes from a deformation $\widetilde{G}/T$. Here an admissible filtration [Mes72, Definition V.1.4] means that Fil^1 is a locally-free vector sub-group with locally-free quotient which reduces to the Hodge filtration on $\mathbb{D}(G)_S$, and the result above extends to an equivalence of categories [Mes72, Theorem V.1.6].

If $S = \operatorname{Spec} R$ is affine where R is a perfect ring of characteristic p, then the crystal $\mathbb{D}(G)$ is associated to classical Dieudonné theory as follows. First, the global section functor $\Gamma(S/\Sigma, \ldots)$ induces an equivalence between crystals in quasi-coherent modules and $W(R)$-modules complete and separated, where $W(R)$ are the Witt vectors with value in R [BBM82, Théorème 1.2.9]. In our case of interest, $R = \mathbb{F}_q$ and $W(R) = \mathbb{Z}_q$. Secondly, classical Dieudonné theory associate to G/S the Dieudonné module $\mathcal{D}(G) = \operatorname{Hom}_S(G, CW)$ over the (non commutative) Dieudonné ring $W(R)\{F, V\}$, and where CW is Fontaine's sheaf of co-Witt vectors. Then by [BBM82, Théorème 4.2.14], we have a sesquilinear isomorphism $\mathcal{D}(G) \simeq \Gamma(S/\Sigma, \mathbb{D}(G))$ (i.e., $\Gamma(S/\Sigma, \mathbb{D}(G)) = \mathcal{D}(G)^\sigma$ where σ is the lift of the Frobenius to $W(R)$).

The crystal $\mathbb{D}(G)$ and the isomorphism with the Dieudonné module $\mathcal{D}(G)$ behaves as expected with duality [BBM82, § 5]. We remark that $G \mapsto \mathbb{D}(G)$ is an anti-equivalence of category between p-divisible groups and crystalline Dieudonné modules whenever S is regular over a perfect field k (or even just a field k of characteristic p having a finite p-basis) [DeJ95].

If $\pi\colon A \to S$ is an abelian scheme, the Hodge to de Rham spectral sequence degenerates at the E_1 page, so in particular we have a Hodge decomposition $0 \to \omega_A \to H^1_{\mathrm{dR}}(A/S) \to R^1\pi_* O_A \to 0$ [BBM82, Proposition 2.5.2 et Lemme 2.5.3]. Here $\omega_A = \pi_*(\Omega^1_{A/S})$ is the sheaf of invariant differential, it is isomorphic to $\varepsilon^*(\Omega^1_{A/S})$ where $\varepsilon\colon S \to A$ is the zero section. If $\pi^\vee\colon A^\vee \to S$ is the dual abelian scheme, we have a canonical isomorphism $H^1_{\mathrm{dR}}(A^\vee/S) \simeq H^1_{\mathrm{dR}}(A/S)^\vee$ which induces an isomorphism $R^1\pi_*(O_A) \simeq \pi^\vee_*(\Omega^1_{A^\vee/S})$. Now if $G = A(p)$ is the p-divisible group of A, we have canonical isomorphisms [BBM82, Théorème 2.5.6 et Proposition 3.3.7]

$$\mathbb{D}(G) \simeq \mathcal{EXT}^1_{S/\Sigma}(A, O_{S/\Sigma}) \simeq R^1\pi_{\mathrm{crys},*}(O_{A/\Sigma}) \tag{2}$$

compatible with duality. The isomorphism from Equation (2) induces an isomorphism $\mathbb{D}(G)_S \simeq H^1_{\mathrm{dR}}(A/S)$ compatible with the Hodge filtrations on G and A respectively [BBM82, Proposition 2.5.8]. In particular, H^1_{dR} is crystalline on an abelian scheme, and if $\widetilde{A}/T$ is a deformation of A to T, then the Hodge filtration on $H^1_{\mathrm{dR}}(A/T)$ induces a Hodge filtration on $\mathbb{D}(G)_{S\to T}$ lifting the one on $\mathbb{D}(G)_S$. The deformation $\widetilde{A}$ is precisely the abelian scheme predicted by combining Grothendieck-Mazur and Serre-Tate theory.

We now specialize this to $S = \operatorname{Spec} \mathbb{F}_q$, and we consider the crystalline site on S associated to $\operatorname{Spec} \mathbb{Z}_q$ along with its canonical divided power $p^n/n!$ (note that by [BBM82, pp. 1.1.10–1.1.15], this is the same thing as considering the crystalline site induced by the canonical divided powers on Σ). We also assume that $p > 2$ so that the divided powers are nilpotent. As explained above, since $\mathbb{F}_q$ is a perfect field, a crystal of quasi-coherent modules over S/Σ is determined by its global section functor $\Gamma(S/\Sigma)$, and applying this to the crystal $\mathbb{D}(G)$ of a p-divisible group, gives up to a Frobenius twist the Dieudonné module $\mathcal{D}(G)/\mathbb{Z}_q\{F, V\}$ of G. Here $\mathbb{Z}_q\{F, V\}$ is the Dieudonné ring (so that $FV - p = 0$), and $\mathcal{D}(G)$ is a free $\mathbb{Z}_q$-module of rank the height h of G. Furthermore, $G/\mathbb{F}_q \mapsto \mathcal{D}(G)$ is an anti-equivalence of category between p-divisible groups and $\mathbb{Z}_q\{F, V\}$-modules free of finite rank as $\mathbb{Z}_q$-modules. If $\pi\colon A \to \mathbb{F}_q$ is an abelian variety, we denote by $H^1_{\mathrm{crys}}(A/\mathbb{Z}_q)$ the global sections of the crystalline sheaf $R^1\pi_{\mathrm{crys},*}(O_{A/\mathbb{Z}_q})$. By Equation (2), $H^1_{\mathrm{crys}}(A/\mathbb{Z}_q) \simeq \mathcal{D}(A(p))$: we have a canonical sesquiliear isomorphism between its crystalline cohomology and the Dieudonné module of its p-divisible group. And furthermore, if $\mathcal{A}/\mathbb{Z}_q$ is a lift of A, we have an isomorphism $H^1_{\mathrm{dR}}(\mathcal{A}/\mathbb{Z}_q) \simeq H^1_{\mathrm{crys}}(A/\mathbb{Z}_q)$. And as we saw above, via these canonical isomorphisms, the Hodge filtration on $H^1_{\mathrm{crys}}(A/\mathbb{Z}_q)$ induced by $\mathcal{A}$ is

the same as the Hodge filtration induced on $\mathcal{D}(A(p))$ by the p-divisible group $\mathcal{A}(p)$ of the lift.

Combining all of this together, we obtain Theorem 3.4. As an aside, we also remark that there is a well developed theory of Dieudonné crystals/modules for finite commutative p-group schemes $G/\mathbb{F}_q$ (or for that matter finite locally free commutative p-group schemes over a general base scheme S) for which we refer to [BBM82] again. Since $\mathbb{F}_q$ is perfect, the Dieudonné module functor $G \mapsto \mathcal{D}(G)$ gives an anti-equivalence between finite locally free commutative p-group schemes $G/\mathbb{F}_q$ of length h and Dieudonné modules $M/\mathbb{Z}_q\{F, V\}$ of finite length h over $\mathbb{Z}_q$. For instance, if $A/\mathbb{F}_q$ is an abelian variety, $\mathcal{D}(A[p^m]) \simeq \mathcal{D}(A(p)) \otimes_{\mathbb{Z}_q} \mathbb{Z}_q/p^m\mathbb{Z}_q$. Since the Hodge filtration $H^1_{\mathrm{dR}}(A/\mathbb{F}_q)$ is induced by the action of the absolute Frobenius endomorphism F_A on H^1_{dR}, that is $\omega_{A/\mathbb{F}_q} = H^0(\mathbb{F}_q, \Omega^1_{A/\mathbb{F}_q}) = \mathrm{Ker}\, H^1_{\mathrm{dR}}(F_A)$, we see that the Frobenius filtration on $A[p]$ induces via the Dieudonné functor the Hodge filtration on $H^1_{\mathrm{dR}}(A) \simeq H^1_{\mathrm{crys}}(A/\mathbb{F}_q)$ (see [Oda69] for more details).