Algorithmic Number Theory Symposium XVII, Groningen, July 6-10, 2026

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 separable 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.

Authors listed alphabetically; see https://ams.org/profession/leaders/CultureStatement04.pdf. Date of this document: 2026-06-05.

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., algebraic) evaluation algorithm for the isogeny.

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.)

1.1. **Previous work.** The idea of "automatic differentiation" (evaluating derivatives of functions that are given as algebraic algorithms) can be traced back to [Wen64]. The application of this technique to implicit Hensel lifting appears to be folklore: For instance, in [Bac09], Bach replaces in the Newton step the derivative by iterated differences. In [Mas20], Mascot uses implicit Hensel to p-adically lift n-torsion points on a Jacobian. As far as we know, using implicit Hensel lifting to lift isogenies is due to the second author in [Rob21, Remark 6.6.2]. This idea was developed further to compute canonical lifts of elliptic curves in [MRS25; Rob22b; KR24].

The main contribution of this paper is to first develop in greater generality the idea of implicit Henseling (which was used in a more ad-hoc manner in the previous articles), both from an algorithmic viewpoint (where we provide a very general approach to Hensel lifting for p-adically analytic constraints), but also from a geometric viewpoint (see Section 2.5). Secondly, in order to compute the trace of an arbitrary endomorphism rather than just the Frobenius, we explain how to lift arbitrary isogenies rather than just the Verschiebung as was done previously.

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, by $\mathbb{Q}_p = \operatorname{Frac}\mathbb{Z}_p$ the *field of p-adic numbers*. For $q = p^m$, we let $\mathbb{Q}_q$ denote the unique unramified extension of $\mathbb{Q}_p$ of degree m, and $\mathbb{Z}_q$ its valuation ring.

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 $a_n \in C$. 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 $p^k\mathbb{Z}_q$. Similar notation will (again slightly abusively) be used for tuples or vectors: $(a_1, ..., a_n) + O(p^k)$ is shorthand

for $(a_1 + O(p^k), ..., a_n + O(p^k))$. 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.

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 a 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.,* $F(\overline{x}) \in O(p^k)$, *assume that* F *is analytic (i.e., locally defined by a power series) around* $\overline{x}$ *and that* $\mathrm{D}F(\overline{x}) \bmod p$ *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 $(\mathbb{Z}_q/p^k\mathbb{Z}_q)$*-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, all partial derivatives $\partial_i F(\overline{x})$ for $i = 1, \ldots, |\mathsf{in}|$ can be computed in precision k by evaluating F at the two elements $\overline{x}$ and $\overline{x} + p^k e_i$ in precision $2k$, where e_i is a standard basis vector.

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 full-rank mod* p *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. 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.** Our Hensel lifting algorithm is based on an(y) algorithm to evaluate the analytic function $F\colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$ of which the lifted solution should be a zero. 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 at least one 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* $\mathsf{in} \neq \emptyset$ *and* $|\mathsf{in}| \geq |\mathsf{out}|$ *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.

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. In Algorithm 1, we compute the entries of the Jacobian matrix $\mathrm{D}F(\overline{a})_{l,i}$ with $l \in Y_j$, $i \in X_j$ using Proposition 2.3. To prove correctness, we point out that all other entries are zero and need not be computed: indeed, given our definition of

[1] A *linear-algebra exponent* is a value $\omega \in \mathbb{R}$ such that $n\times n$ matrix multiplication can be carried out using $O(n^\omega)$ base-ring operations. The best currently known algorithms achieve $\omega \approx 2.37$. A rectangular $m\times n$ linear system can be solved in $O(\max\{m,n\}^\omega)$ ring operations; see [DP12].

constraint systems, the outputs indexed by Y_j only depend on the inputs indexed by X_j.

Regarding the complexity: For each $j \in \{1, ..., r\}$, 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. For $k' = 2^{h'} \in \mathbb{Z}$, let $C_{k'}$ be the upper bound from Proposition 2.4 or Proposition 2.5 on the cost of lifting solutions from precision $k'/2$ to precision k'. Fix $k = 2^h$. 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 $2, 4, \ldots, 2^h$. Given that the bit-operation cost of arithmetic modulo $p^{k'}$ grows at least linearly in k', each lifting step costs at most half as much as the next one. In particular, we have $C_k + C_{k/2} + C_{k/4} + \ldots + C_2 \in O(C_k)$. □

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.

In the context of Hensel lifting, we're given a system of equations over $\mathbb{Z}_q$ of the form $F(x_1, \ldots, x_n) = 0$, where $F = (F_1, \ldots, F_s)$ is a polynomial map, and a solution $P = (a_1, \ldots, a_n)$ 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 $F(\widetilde{P}) = 0 \pmod{p^k}$, 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

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 $\operatorname{Spec} \mathbb{Z}_q[x_1, \ldots, x_n]/(F_1, \ldots, F_m)$, 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 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

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

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 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 2.8. 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 2.9. *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 2.8 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 2.9. The following Proposition is a reformulation of Proposition 2.3 in geometric language:

Proposition 2.10. *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 $F_j(Q)$ 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 2.11. 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.

    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.

(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 $y(P)$, $x(\widetilde{P})$, and the curve equation.

Remark 2.12. Using the notation of Proposition 2.10, 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 set of generators of our ideal of equations I). This allows to check if φ is smooth at P, and pick 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 (local) equations of $i'(X')$ and the equations defining $i(U)$ inside X', then we can apply the tools of Proposition 2.10 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.

Remark 2.13. 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 $4a^3+27b^2 \neq 0$. Throughout the paper, we assume $\mathrm{char}(k) \notin \{2,3\}$ and consider elliptic curves $E = E_{a,b}$ defined over k to be in short Weierstraß form.

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 $j(E_{a,b}) = 1728 \cdot 4a^3/(4a^3+27b^2)$.

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 $P \neq 0_E \in E$, 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 $E_{\mathbb{Q}_q}$. In particular, this generic fiber is enough to recover E.

Finally, by the rigidity lemma [MFK94, Proposition 6.1] two morphisms $\varphi_1, \varphi_2 : E_1 \to E_2$ of elliptic schemes over R that are equal over the special fiber $\mathbb{F}_q$ are equal everywhere. And furthermore by [MFK94, Corollary 6.4], $\varphi : E_1 \to E_2$ is a morphism if and only if it takes the section section of E_1 to the zero section of E_2. This rigidity lemma, combined with the fact that E/R is proper smooth allows to extend many of the algorithmic results on elliptic curves over a finite field $\mathbb{F}_q$ to the case of an elliptic scheme over R. Informally, as long as our formulas have 'good reduction' modulo p, they are automatically correct over R.

Isogenies and endomorphisms. An isogeny $\varphi\colon E_0 \to E_1$ is a non-constant 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 an imaginary-quadratic number field. If k is finite, 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 imaginary-quadratic 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 [Rob24] 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 $N \in \mathbb{Z}_{\geq 1}$ and $p \nmid N$, 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, where the products of elliptic curves are equipped with the product polarization. If $N = \prod_i \ell_i$, the isogeny Φ splits as a chain $\Phi = \Psi_n \circ \ldots \circ \Psi_1$ of ℓ_i-isogenies Ψ_i, 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. The *scaling factor* of φ on invariant differentials [Gal12, § 9.7] is defined as the unique quantity $c_\varphi \in k$ satisfying $\varphi^*(\mathrm{d}x/y) = c_\varphi \cdot \mathrm{d}x/y$ with $\mathrm{d}x/y$ the standard *invariant differentials* on E_0 resp. E_1. In case φ is separable, it can be written in the form $\varphi(x,y) = (f(x), c\,y f'(x))$ for some $f \in k(x)$ and $c \in k^\times$ [Gal12, Theorem 9.7.5]; the factor c then equals $1/c_\varphi$. An isogeny φ is said to be *normalized* if $c_\varphi = 1$. Given an elliptic curve E over k, the map $\mathrm{End}(E) \to k$ defined by $\varphi \mapsto c_\varphi$ is a ring homomorphism.

Retrieving the scaling factor from a given isogeny is immediate in some cases: if $\eta\colon E_0 \to E_1$, $(x,y) \mapsto (x/u^2, y/u^3)$ 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, by means of the formula $\varphi(x,y) = (f(x), \frac{1}{c_\varphi} y f'(x))$ for separable isogenies, resp. the fact that $c_\varphi = 0$ for inseparable isogenies, 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 around $x \in k$, 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 which takes as 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 x-only isogeny over k.

Input: An x-only evaluation algorithm eval_φ of a k-isogeny $\varphi\colon E_0 \to E_1$; a point $P = (x,y) \in E_0$ and its image $P_1 = \varphi(P) = (x_1, y_1) \in E_1$ under φ with $y_1 \neq 0$.

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 $(y \cdot b)/y_1$ // $(x_1, y_1) = (f(x), \frac{1}{c_\varphi} y f'(x)) = (a, \frac{1}{c_\varphi} y b)$

Remark 3.2. An x-only evaluation algorithm for φ inherently only determines φ up to a sign, i.e., up to post-composition with $[\pm 1]$, hence c_φ up to a sign: That is why the pair P, P_1 is needed in the algorithm to differentiate between $+\varphi$ and $-\varphi$.

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 fine moduli space of elliptic curves and, if Γ is a congruence subgroup, denote by $\mathfrak{A}_1(\Gamma)$ the fine moduli space of elliptic curves with Γ-level structure (cf. [DR72], technically these moduli spaces are algebraic stacks). 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 $\varphi\colon E \to E_1$ is an ℓ-isogeny over $\mathbb{F}_q$ with $p \nmid \ell$, then φ lifts uniquely to an isogeny $\widetilde{\varphi}\colon \widetilde{E} \to \widetilde{E_1}$. We can compute $\widetilde{\varphi}$ by lifting the points of the kernel of φ: since $\widetilde{E}[\ell]/\mathbb{Z}_q$ is étale, the points of ℓ-torsion of $E_{\mathbb{F}_q}$ lift uniquely.

Now let $\gamma\colon E \to E$ be a (non scalar) endomorphism over $\mathbb{F}_q$. If γ is of degree ℓ prime to p, then γ lifts to an isogeny $\widetilde{\gamma}\colon \widetilde{E} \to \widetilde{E}'$ for every $\widetilde{E}/\mathbb{Z}_q$ that lifts E, but in general $\widetilde{E}'$ is not isomorphic to $\widetilde{E}$. Only for an (essentially) unique choice of lift $\widetilde{E} = \hat{E}$ 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.

We introduce a few more definitions that we will use in the next proposition. Given $\gamma \in \mathrm{End}(E)$, its discriminant Δ_γ is the discriminant of the imaginary quadratic order $\mathbb{Z}[\gamma]$ it generates. A γ-oriented isogeny $\varphi : E = E_1 \to E_2$ is an isogeny whose kernel is stable by γ, then γ descends via φ to an endomorphism of E_2.

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 $\widehat{E}_\gamma$ 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 $\mathbb{Z}_q/p^k\mathbb{Z}_q$. 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 $\mathcal{A}/(\mathbb{Z}_q/p^k\mathbb{Z}_q)$ of A, the Hodge filtration on $H^1_{\mathrm{dR}}(\mathcal{A}/(\mathbb{Z}_q/p^k\mathbb{Z}_q)) \simeq H^1_{\mathrm{crys}}(A/\mathbb{Z}_q) \otimes_{\mathbb{Z}_q} \mathbb{Z}_q/p^k\mathbb{Z}_q$ 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. □

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 π_q 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 lift uniquely to endomorphisms of $\widehat{E}$, they may not lift uniquely to 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.

4. Efficient representations of isogenies

In this section, we review some widely used forms of representations isogenies, and we describe how to lift a given isogeny over $\mathbb{F}_q$, with $q = p^m$, to the ring $\mathbb{Z}_q$. Throughout the section, we will assume $p \geq 5$ and work with isogenies φ of degree coprime to p. We will use Notation 2.1 to describe the cost of arithmetic in $\mathbb{Z}_q/p^k\mathbb{Z}_q$.

An **isogeny representation** can be defined [Wes24, Definition 1.3] as a polynomial-time algorithm $\mathscr{A}$ taking as input a bit string $\mathcal{D}$ that describes an isogeny $\varphi_{\mathcal{D}}\colon E_{\mathcal{D}} \to E'_{\mathcal{D}}$, from which $\mathscr{A}$ is able to compute the degree $\deg(\varphi_{\mathcal{D}})$, domain $E_{\mathcal{D}}$, codomain $E'_{\mathcal{D}}$, and evaluations $\varphi_{\mathcal{D}}(P) \in E'_{\mathcal{D}}$ at given points $P \in E_{\mathcal{D}}$. In order to define a reasonable notion of *efficiency* of the representation, only isogenies over finite fields are usually considered.

For our purposes, we introduce a somewhat stricter, but related notion of *algebraic* isogeny representation that allows us to extend the finite-field definition to the p-adic context, under mild assumptions on the algorithm $\mathscr{A}$ that are satisfied by several widely-used examples of isogeny representations. The main differences to the definition from [Wes24, Definition 1.3] are: We work over (approximations of) $\mathbb{Z}_q$ instead of just $\mathbb{F}_q$; we require the "data" bit string $\mathcal{D}$ to contain only elements of

the base ring of the isogeny (by moving all remaining data into separate inputs), and we request that the specification of what constitutes a "valid" data tuple to be p-adically analytic in a neighbourhood of all valid inputs.

Definition 4.1. An **algebraic isogeny representation** consists of algorithms check, curves, and eval, as well as (for prime powers $q = p^m$ and positive integers d, k) pairs $(\mathsf{V}_{q,d,k}, \mathrm{isog}_{q,d,k})$ of subsets $\mathsf{V}_{q,d,k} \subseteq (\mathbb{Z}_q/p^k\mathbb{Z}_q)^{\mathsf{in}}$ and maps

$$\mathrm{isog}_{q,d,k} \colon \mathsf{V}_{q,d,k} \to \{d\text{-isogenies between elliptic curves over } \mathbb{Z}_q/p^k\mathbb{Z}_q\},$$

where the $\mathsf{V}_{q,d,k}$ are pairwise disjoint. The tuple $\mathcal{D} \in \mathsf{V}_{q,d,k}$ is referred to as the "data" of the isogeny $\mathrm{isog}_{q,d,k}(\mathcal{D})$. For brevity, we write $\varphi_{\mathcal{D}} \colon E_{\mathcal{D}} \to E'_{\mathcal{D}}$ for the isogeny $\mathrm{isog}_{q,d,k}(\mathcal{D})$ and its (co)domain curves.

- check takes (q, d, k) and a tuple $\mathcal{D} \in (\mathbb{Z}_q/p^k\mathbb{Z}_q)^{\mathsf{in}}$, and outputs a boolean indicating whether $\mathcal{D} \in \mathsf{V}_{q,d,k}$.
- curves takes (q, d, k) and $\mathcal{D} \in \mathsf{V}_{q,d,k}$ and outputs the curves $E_{\mathcal{D}}$ and $E'_{\mathcal{D}}$.
- eval takes (q, d, k) and $\mathcal{D}$ and some point $P \in E_{\mathcal{D}}(R)$, where R is a finite $\mathbb{Z}_q/p^k\mathbb{Z}_q$-algebra, and outputs the image point $\varphi(P) \in E'_{\mathcal{D}}(R)$.

It is required that all three algorithms run in time polynomial in $d, k, \log q$, and the degree of R over $\mathbb{Z}_q/p^k\mathbb{Z}_q$, and we say the representation is **efficient** if the runtime is polynomial in $\log d$. We require the following *compatibility* conditions: For all q, d and all $k' > k$, if $\mathcal{D}' \in \mathsf{V}_{q,d,k'}$ reduces to $\mathcal{D} \in \mathsf{V}_{q,d,k}$ modulo p^k, then $\mathsf{check}(q, d, k', \mathcal{D}') = \mathsf{check}(q, d, k, \mathcal{D})$, and the output of $\mathsf{algo}(q, d, k', \mathcal{D}')$ reduces to the output of $\mathsf{algo}(q, d, k, \mathcal{D})$ modulo p^k for $\mathsf{algo} \in \{\mathsf{curves}, \mathsf{eval}\}$. Aditionally, for all q, d there must be a function $F_{q,d} \colon \mathbb{Z}_q^{\mathsf{in}} \to \mathbb{Z}_q^{\mathsf{out}}$ which is analytic at all points $x \in \mathbb{Z}_q^{\mathsf{in}}$ with $x + O(p) \in \mathsf{V}_{q,d,1}$, such that $\mathsf{check}(q, d, k, \cdot) \colon (\mathbb{Z}_q/p^k\mathbb{Z}_q)^{\mathsf{in}} \to \{\mathtt{true}, \mathtt{false}\}$ is given by $\mathcal{D} \mapsto$ *"*$F_{q,d}(\mathcal{D}) \in O(p^k)$*?"*.[3]

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*

$$E_0 \xrightarrow{\varphi_1} E_1 \xrightarrow{\varphi_2} \cdots \xrightarrow{\varphi_{r-1}} E_{r-1} \xrightarrow{\varphi_r} E_r$$
$$\varphi$$

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 $\overline{\mathcal{D}} = (\overline{j}_0, \ldots, \overline{j}_n) + O(p^k) \in \mathbb{Z}_q^{n+1}$ that satisfies the polynomial constraints*

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

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

Assume that there exists only a single isogeny chain $E_0 \to \cdots \to E_n$ (up to isomorphism) such that each $j(E_i) = \overline{j}_i \bmod p$ and each step $E_{i-1} \to E_i$ has degree ℓ_i.

[3] These properties are satisfied by all isogeny representations whose algorithms are defined (for a fixed degree) by arithmetic circuits; however, in principle, non-algebraic operations are permitted as long as the outputs satisfy the compatibility conditions above.

Letting $\ell = \max_i \ell_i$ and assuming that $\ell < p$, we can compute a lift $\mathcal{D} = (j_0, \dots, j_n)$ of $\overline{\mathcal{D}}$ to precision $2k$ using $\widetilde{O}(\ell^2 n)\mathsf{M}_{2k}$ bit operations.

Proof. Let $X_0, ..., X_n$ denote variables. Write $F_i = \Phi_{\ell_i}(X_{i-1}, X_i)$; the constraints thus say $F(\overline{\mathcal{D}}) \in O(p^k)$ and we are looking for a lift $\mathcal{D}$ with $F(\mathcal{D}) \in O(p^{2k})$. The Jacobian matrix of F is clearly sparse; it is banded with only 2 nonzero diagonals. The uniqueness assumption on the isogeny path ensures that $\partial_{i-1} F_i(\overline{\mathcal{D}}) = \partial_X \Phi_{\ell_i}(\overline{j}_{i-1}, \overline{j}_i)$ and $\partial_i F_i(\overline{\mathcal{D}}) = \partial_Y \Phi_{\ell_i}(\overline{j}_{i-1}, \overline{j}_i)$ are nonzero modulo p for all i [MM24, Proposition 4.1], i.e., all elements on the nonzero diagonals are nonzero modulo p, hence why $\mathrm{D}F(\mathcal{D}) \bmod p$ has full rank.

To compute the lift, we apply Proposition 2.5. Note that each polynomial constraint depends at most 2 variables, and the complexity of evaluation of a modular polynomial of degree ℓ is $\widetilde{O}(\ell^2)$; finally, thanks to 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 isogeny 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.** For a cyclic isogeny $\varphi\colon E_0 \to E_1$ of small (typically prime) degree ℓ, there are several options for representing φ which take as input the domain E_0 and a single extra value x. For all common such representations, the check algorithm can be encoded as verifying an algebraic (hence analytic) constraint system, hence they're suited for Hensel lifting.

Vélu, $\sqrt{\text{élu}}$. Vélu's formulas [Vél71] and the $\sqrt{\text{élu}}$ algorithm [BDLS20] can be used to represent a separable normalized 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 $\ker \varphi = \langle P \rangle$ where $P \in E_0[\ell]$. Vélu's formulas (resp. the $\sqrt{\text{élu}}$ 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 $\mathbb{Z}_q/p^k\mathbb{Z}_q$, where E_i has short-Weierstraß coefficients $(\overline{a}_i, \overline{b}_i)$. For every lift $(a_0, b_0) + O(p^{2k})$, using Vélu's formulas (resp. the $\sqrt{\text{élu}}$ 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 such a 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 $(a_1, b_1) = \mathsf{BMSS}_\ell(a_0, b_0, j_1)$ 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:*

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

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 $E\colon y^2 + (1-c)xy - by = x^3 - bx^2$, with b, c such that the point $P = (0, 0)$ is of order ℓ. 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 [Gal12, Definition 26.3.4]. 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.

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 $\ell = \max_i \ell_i$ and assume $\ell < p$.*

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*

$$\begin{cases} Y_{\text{divpoly},i}\colon & \psi_{\overline{a}_i, \overline{b}_i, \ell_i}(\overline{x}_i) \in O(p^k) \qquad i = 1, \dots, n \\ Y_{\mathsf{Vélu},i}\colon & (\overline{a}_i, \overline{b}_i) - \mathsf{Vélu}(\overline{a}_{i-1}, \overline{b}_{i-1}, \overline{x}_i) \in O(p^k) \qquad i = 1, \dots, n \\ Y_{\text{endo}}\colon & (\overline{a}_n - \overline{a}_0 \cdot \overline{u}^4, \overline{b}_n - \overline{b}_0 \cdot \overline{u}^6) \in O(p^k). \end{cases}$$

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 $p \nmid N$. 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 (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 between Corollary 4.7 and Proposition 4.8. In the former, we see that for a general isogeny $\overline{\varphi}\colon \overline{E}_0 \to \overline{E}_1$, any arbitrary lift 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 $\overline{\varphi} \in \mathrm{End}(\overline{E}_0)$, for every E_0 lifting $\overline{E}_0$ there is a pair (E_0', φ) where E_0' is also a lift of $\overline{E}_0$ and $\varphi\colon E_0 \to E_0'$ lifts $\overline{\varphi}$ as an isogeny; but for *only one* choice of E_0 (up to isomorphism) we get $E_0' \cong E_0$.

Remark 4.10. We've focused so far on curves in short Weierstraß form, as this model is defined for all curves in characteristic $p \geq 5$. The algorithms we presented generalize to curves of different models though: one needs to replace (a_0, b_0) with the tuple of relevant coefficients and accordingly adjust the constraints defining isogeny steps. The full-rank condition on the Jacobian matrix is ensured by the fact that separable isogenies and endomorphisms lift uniquely.

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 higher-dimensional representations.** Suppose given a separable isogeny $\varphi\colon E \to E'$ over $\mathbb{F}_q$ of odd degree d, not necessarily polylog(qd)-smooth. Though φ cannot in general be split into a chain of small-degree factors, [Rob22a] describes a way to "embed" φ as a component of an 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 $T = (T_0, T_1) \in (E_0 \times E_1)(R)$ a point with $U_{11,11}(\Phi(T)) \neq 0 \pmod p$. Using $O(n \log n)\mathsf{M}_k$ bit operations, on input $((a_i, b_i, x_{P_i}, x_{Q_i}, x_{T_i})_{i=0,1})$, 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 $(\chi_{\text{split}}(A))^2 = U_{11,11}(0)^2$, but we need the non-squared value for our purposes. To compute it, we can use the duplication formula [Rob21, Eq. (2.9)] with $\widetilde{x} = T, \widetilde{y} = 0$ to get $\theta^A_{11,11}(0)$, using the fact that $\theta^A_{11,11}(T)$ 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_T as a constant, fixed at the beginning depending on the input curves and kernel modulo p. 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 2.13.

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. As in the previous sections, assume $p \geq 5$. 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. *Let $\overline{\varphi} \in \mathrm{End}(\overline{E})\backslash\mathbb{Z}$ be a nonscalar endomorphism defined over $\mathbb{F}_q$. Following Definition 4.1, suppose given an algebraic isogeny representation $(\mathsf{check}, \mathsf{curves}, \mathsf{eval}, (\mathsf{V}_{q,d,k}, \mathrm{isog}_{q,d,k})_{q,d,k}, (F_{q,d})_{q,d})$ and a data tuple $\mathcal{D} \in V_{q,d,1} \subseteq \mathbb{F}_q^{\mathsf{in}}$ such that $\overline{\varphi} = \mathrm{isog}_{q,d,1}(\mathcal{D})$. 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 $O(\log k)$ 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 of $(\overline{E}, \overline{\varphi})$ 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 degree d and 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

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

of $x^2 + (t \bmod p^k)x + (d \bmod p^k)$. 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 $\mathbb{Z} \cap [-p^k/2, p^k/2]$. □

Algorithm 3: Computing the trace of a separable endomorphism over $\mathbb{F}_q$.

Input: An algebraic isogeny representation $(\mathsf{check}, \mathsf{curves}, \mathsf{eval}, (\mathsf{V}_{q,d,k}, \mathrm{isog}_{q,d,k})_{d,k}, (F_{q,d})_{q,d})$, an integer $d \in \mathbb{Z}$ and data $\mathcal{D}_1$ such that $\mathrm{isog}_{q,d,1}(\mathcal{D}_1) = \varphi \in \mathrm{End}(E)$ is a separable endomorphism over $\mathbb{F}_q$.
Output: $\mathrm{tr}(\varphi) \in \mathbb{Z}$.
$k \leftarrow \lceil \log_p(4\sqrt{\deg \overline{\varphi}}) \rceil$.
For i **from** 1 **to** $\lceil \log_2(k) \rceil - 1$ **do**
  $\mathcal{D}_i \leftarrow$ Lift $\mathcal{D}_i$ from precision 2^{i-1} to precision 2^i. // Algorithm 1
$\mathcal{D} \leftarrow$ Lift $(\mathcal{D}_{\lceil \log_2(k) \rceil - 1} \bmod p^{k/2})$ to precision k. // Algorithm 1
(E, φ) over $\mathbb{Z}_q/p^k\mathbb{Z}_q \leftarrow \mathrm{isog}_{q,d,k}(\mathcal{D}_k)$
// we can compute E and evaluate φ from the representation
Sample $P \in E$ at random until $P_1 = \varphi(P)$ is not 2-torsion.
Compute the scaling factor $c_\varphi \in \mathbb{Z}_q/p^k$ from (E, φ). // Algorithm 2
Set $t \leftarrow c_\varphi + d/c_\varphi \in \mathbb{Z}/p^k$.
Return the representative of t in $\mathbb{Z} \cap [-p^k/2, p^k/2]$.

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 $\widetilde{O}(n^2 \log_p q + n \log q)$.

(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. Sampling a point P over $\mathbb{Z}_q/p^k\mathbb{Z}_q$ and checking $y(\varphi(P)) \neq 0$ amounts to an extra evaluation of φ; we might need to lift to double precision at most 2 more times so that there are enough points in $E\backslash\varphi^{-1}(E[2])$ to ensure we can find a suitable P with probability at least $1/2$; this does not change the asymptotic complexity. 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 $\mathbb{F}_q$), 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 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 an algebraic 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 $\mathrm{tr}\,\pi$.

However, it is possible to adapt Satoh's algorithm to use implicit lifting of the small Verschiebung $\widehat{\pi}_p$ rather than the modular equations Φ_p, see [MRS25; Rob22b]. (This is a bit different than our current setting, because $\widehat{\pi}_p$, although separable, has degree p and is not an endomorphism.)

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 (Corollary 5.2). 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

[Bac09] Eric Bach. 'Iterative root approximation in p-adic numerical analysis'. In: *Journal of Complexity* 25.6 (2009), pp. 511–529.

[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.

[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.

[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.

[DP12] Jean-Guillaume Dumas and Clément Pernet. *Computational linear algebra over finite fields*. 2012. arXiv: 1204.3735.

[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.

[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.

[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.

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

[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.

[Mas20] Nicolas Mascot. 'Hensel-lifting torsion points on Jacobians and Galois representations'. In: *Mathematics of Computation* 89.323 (2020), pp. 1417–1455.

[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.

[MFK94] David Mumford, John Fogarty and Frances Kirwan. *Geometric invariant theory*. Vol. 34. Springer Science & Business Media, 1994.

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

[MM24] William E. Mahaney and Travis Morrison. *Computing Isogenies at Singular Points of the Modular Polynomial*. 2024. arXiv: 2402.02038 [math.NT]. URL: https://arxiv.org/abs/2402.02038.

[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.

[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.

[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.

[Rob24] 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.

[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.

[Wen64] Robert E. Wengert. 'A simple automatic derivative evaluation program'. In: *Communications of the ACM* 7.8 (1964), pp. 463–464. DOI: 10.1145/355586.364791.

[Wes24] Benjamin Wesolowski. 'Random Walks in Number-theoretic Cryptology'. Habilitation à diriger des recherches. ENS Lyon, Aug. 2024. URL: https://hal.science/tel-04837478.

†Technische Universität München, Germany
Email address: lorenz@yx7.cc

*Centre Inria de l'Université de Bordeaux, France
Email address: damien.robert@inria.fr

Email address: alessandro.sferlazza@tum.de