Radical $\sqrt[N]{\text{élu}}$ Isogeny Formulae

Thomas Decru

Université Libre de Bruxelles

Abstract. We provide explicit radical N-isogeny formulae for all odd integers N. The formulae are compact closed-form expressions which require one Nth root computation and $\mathcal{O}(N)$ basic field operations. The formulae are highly efficient to compute a long chain of N-isogenies, and have the potential to be extremely beneficial for speeding up certain cryptographic protocols such as CSIDH. Unfortunately, the formulae are conjectured, but we provide ample supporting evidence which strongly suggests their correctness.
For CSIDH-512, we notice an additional 35% speed-up when using radical isogenies up to $N = 199$, compared to the work by Castryck, Decru, Houben and Vercauteren, which uses radical isogenies up to $N = 19$ only. The addition of our radical isogenies also speeds up the computation of larger class group actions in a comparable fashion.

Keywords: Post-quantum cryptography · isogeny-based cryptography · radical isogenies · CSIDH

1 Introduction

Since the fall of SIDH [7, 19, 23], CSIDH [10] has become the most efficient isogeny-based public-key exchange protocol. The scheme is noninteractive and its mathematical framework has led to a variety of other cryptographic primitives such as a signature scheme [4] and a verifiable delay function [16]. This general flexibility comes at the cost of being slower by at least an order of magnitude compared to other post-quantum cryptographic schemes that do not rely on isogenies. In particular, CSIDH also does not scale well for higher security levels, which could be troublesome given that the initial security estimates may not suffice [20]. Fortunately, CSIDH has undergone several noteworthy speed-ups since its inception.

In 2020, Bernstein, De Feo, Leroux and Smith managed to improve the asymptotic complexity of computing an N-isogeny from a kernel generator from $\mathcal{O}(N)$ to $\tilde{\mathcal{O}}(\sqrt{N})$ field operations with their so-called $\sqrt{\text{élu}}$ isogeny formulae [3]. The original formulae requiring $\mathcal{O}(N)$ operations are classical results due to Vélu [28], but the ones used in CSIDH are more efficient versions on elliptic curves in Edwards or Montgomery form (see for example [21]). The hidden constants in the asymptotic complexity from the $\sqrt{\text{élu}}$ isogeny formulae are not too large; they start outperforming the fastest isogeny formulae for primes $N \approx 100$ already, depending

on the chosen programming language. This benefits CSIDH considerably, given that even on the most basic security level it already requires computing isogenies up to prime degree $N \approx 400$.

On the other end of the spectrum, improvements for computing the lowest-degree isogenies have also appeared. In [5], Castryck and Decru adjust the CSIDH-setting to allow using highly efficient 2-isogenies which only require one square-root computation together with a handful of additional arithmetic operations. This was then generalized in [9] to so-called radical N-isogenies for slightly larger primes $N > 2$. These radical isogenies allow an efficient computation of a cyclic N^k-isogeny by means of iteratively drawing Nth roots together with additional basic arithmetic operations. The main advantage of these radical isogenies is that only the initial N-isogeny requires the explicit computation of a point of order N, which is a costly operation in the CSIDH setting. Unfortunately, this additional overhead turns cumbersome rather quickly, and in a follow-up work, the authors of [8] manage to find all radical isogeny formulae up to $N = 37$, yet can only make them useful in the CSIDH setting for N at most 19. Unlike the $\sqrt{\text{élu}}$ isogeny formulae, these radical isogeny formulae have less of an impact for larger parameter sets, since more high prime-degree isogenies are required for those. Hitherto, all radical isogeny formulae were derived and optimized ad hoc from parametrizations of the modular curve $X_1(N)$; in particular the optimized equations from Sutherland [27] were used to obtain relatively compact expressions.

Our contributions

Let N be an odd positive integer. We provide formulae which, on input of an elliptic curve E and point P of order N, output an elliptic curve $E' = E/\langle P \rangle$ with a point $P' \in E'$ of order N such that P' is a kernel generator of an N-isogeny which extends the former isogeny cyclically to an N^2-isogeny. Furthermore, the formulae benefit the following properties:

- they impose no further restrictions (e.g. they work over any field and N need not be prime);
- they require $\mathcal{O}(N)$ basic field operations and one Nth root computation;
- by scaling these Nth roots with Nth roots of unity, they generate all N-isogenies which extend the isogeny with kernel $\langle P \rangle$ cyclically;
- they do not require a parametrization of the modular curve $X_1(N)$;
- they are an extension of classical Vélu formulae, whence can evaluate points.

Our formulae are unfortunately still conjectured, but we provide ample evidence supporting their correctness.

Our Magma implementation for CSIDH-512 benefits from radical isogenies for $N \leq 199$ and provides an additional 35% speed-up compared to the radical isogeny implementation of [8], which only uses radical isogenies for $N \leq 19$. Our implementation also shows that the fastest way to compute the class group action in CSIDH is a combination of radical isogenies and $\sqrt{\text{élu}}$ isogenies; i.e. "the gap is closed" and regular isogeny formulae need not be used anymore. Finally, we show that our radical

isogeny formulae scale well, providing a significant and fairly consistent speed-up across several sizes of class groups.

Acknowledgment

Thanks to Damien Robert for discussing the geometric approach attempt at proving the main conjecture, and thanks to the helpful feedback and suggestions from the anonymous CRYPTO reviewers. This work is partly funded by a postdoctoral grant from the Fund for Scientific Research, Belgium (F.R.S.–FNRS), partly supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement ISOCRYPT - No. 101020788) and partly by CyberSecurity Research Flanders with reference number VR20192203.

Notation

Throughout this entire discussion, we will assume N to be a positive odd integer strictly greater than three unless specified otherwise. The notations $x(P)$ and $y(P)$ refer to the x- respectively y-coordinate of an affine point P. We denote the logarithm with base two as log.

2 Preliminaries

In this section we provide some of the necessary background related to isogenies and cryptography. We refer the interested reader to the book by Silverman [25] for a staple reference regarding elliptic curves and isogenies in general, and to the lecture notes by De Feo [15] for isogenies with a focus on cryptographic applications.

2.1 Isomorphisms and isogenies

An elliptic curve E/K is a smooth projective algebraic curve of genus one over a field K. For the sake of clarity, elliptic curves are often given by an affine representation where the only point at infinity $\mathcal{O}_E$ is assumed to be the neutral element for its group law. If we want to specify that we consider L-rational points on E for a specific field $L \supseteq K$, we will denote this by $E(L)$. Elliptic curves come in several forms, such as the Montgomery form

$$E_{A,B}/K : By^2 = x^3 + Ax^2 + x,$$

where $A, B \in K$ and the discriminant $B(A^2 - 4) \neq 0$. Even though this form has many practical applications in cryptography, it does impose certain restrictions such as requiring K-rational 2-torsion and disallowing

the field characteristic being two. In the most general setting, any elliptic curve can be given by a Weierstraß equation

$$E/K : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6,$$

where all $a_i \in K$ and the discriminant $\Delta = -b_2^2b_8 - 8b_4^3 - 27b_6^2 + 9b_2b_4b_6$ is nonzero, with

$$b_2 = a_1^2 + 4a_2, \quad b_4 = 2a_4 + a_1a_3, \quad b_6 = a_3^2 + 4a_6,$$
$$b_8 = a_1^2a_6 + 4a_2a_6 - a_1a_3a_4 + a_2a_3^2 - a_4^2.$$

Two elliptic curves E/K and E'/K are $\overline{K}$-isomorphic iff they have the same j-invariant, which is an algebraic expression in the coordinates of the curve. For an elliptic curve E/K in Weierstraß form, we have

$$j(E) = \frac{(b_2^2 - 24b_4)^3}{\Delta}.$$

For an explicit isomorphism over K, we have the following proposition.

Proposition 1. *Let E/K and E'/K be elliptic curves in Weierstraß form. Then E and E' are K-isomorphic iff there exist $(u, r, s, t) \in K^\times \times K^3$ such that the change of coordinates*

$$(x, y) \mapsto (u^2x + r, u^3y + su^2x + t)$$

transforms E into E'. Such an isomorphism is called a Weierstraß isomorphism.

Proof. See [25, Proposition VII.1.3]. □

An isogeny $\varphi : E \to E'$ is a nonzero surjective morphism with finite kernel. An isogeny with the same domain as codomain is called an endomorphism. An example of this is the multiplication-by-N map $[N] : E \to E, P \mapsto [N]P$. The kernel of $[N]$ is denoted by $E[N]$ and is referred to as the N-torsion of E.
We will almost exclusively work with separable odd-degree cyclic isogenies, for which the following computational result exists due to Vélu.

Theorem 1. *Let $C = \langle P \rangle$ be a finite subgroup of an elliptic curve*

$$E/K : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6,$$

where P is a point of order N, with N odd. Fix a partition $C = \{\mathcal{O}_E\} \cup C^+ \cup C^-$ such that for any $Q \in C^+$ it holds that $-Q \in C^-$. For all $Q \in C^+$ define

$$g_Q^x = 3x(Q)^2 + 2a_2x(Q) + a_4 - a_1y(Q),$$
$$g_Q^y = -2y(Q) - a_1x(Q) - a_3,$$
$$u_Q = (g_Q^y)^2, \quad v_Q = 2g_Q^x - a_1g_Q^y$$
$$v = \sum_{Q \in C^+} v_Q, \quad w = \sum_{Q \in C^+} (u_Q + x(Q)v_Q),$$
$$A_4 = a_4 - 5v, \quad A_6 = a_6 - (a_1^2 + 4a_2) - 7w.$$

Then the separable isogeny φ with domain E and kernel C has codomain $E' = E/C$ with Weierstrass equation

$$E' : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + A_4x + A_6$$

over K. Furthermore, for $R \in E$ we can compute the image of R as

$$x(\varphi(R)) = x(R) + \sum_{Q \in C \setminus \{\mathcal{O}_E\}} (x(R+Q) - x(Q))$$
$$y(\varphi(R)) = y(R) + \sum_{Q \in C \setminus \{\mathcal{O}_E\}} (y(R+Q) - y(Q)).$$

Proof. This follows from the classical formulae by Vélu [28]. □

Remark 1. Since we assume N to be odd, we can always take C^+ to be $\left\{P, [2]P, \ldots, \left[\frac{N-1}{2}\right] P\right\}$ in Theorem 1, although this is not strictly needed. Furthermore, the only situation where y-coordinates matter on elliptic curves in Montgomery form, is if we want to compute the y-coordinate of the image of a specific point, i.e. $y(\varphi(R))$. In particular, A_4 and A_6 only depend on the x-coordinates of $[k]P$ for $1 \leq k \leq (N-1)/2$.

If $\varphi : E \to E'$ is a separable isogeny, then the degree of φ is $\#\ker\varphi$. For each such isogeny, one can construct the dual isogeny $\hat{\varphi} : E' \to E$ as the isogeny for which $\varphi \circ \hat{\varphi} = [\#\ker\varphi]$, as well as $\hat{\varphi} \circ \varphi = [\#\ker\varphi]$. Up to composition with an isomorphism, this dual isogeny is uniquely defined.

2.2 Tate normal form and Tate pairing

We are interested in elliptic curves E with a given point P of order N. We can translate P to $(0,0)$ and after rescaling this allows us to always assume E is given by a Tate normal form.

Lemma 1. *Let E/K be an elliptic curve and let $P \in E(K)$ be a point of order $N \geq 4$, then (E, P) is isomorphic to a unique pair of the form*

$$E_{b,c} : y^2 + (1-c)xy - by = x^3 - bx^2, \qquad P = (0,0)$$

with $b, c \in K$ and

$$\Delta(b,c) = b^3(c^4 - 8bc^2 - 3c^3 + 16b^2 - 20bc + 3c^2 + b - c) \neq 0\,.$$

Proof. See for example [26, Lemma 2.1]. □

By symbolically computing $x([N]P)$ and $y([N]P)$, we can require $[N]P$ to equal $\mathcal{O}_E$ for any given N. Together with the restriction that $[k]P \neq \mathcal{O}_E$ for all $1 \leq k < N$ and the nonvanishing of $\Delta(b,c)$, we obtain an irreducible polynomial $F_N(b,c) \in \mathbb{Z}[b,c]$ which expresses that P has exact order N.

Remark 2. The modular curve $X_1(N)$ parametrizes pairs (E, P), where E is an elliptic curve and P is a point of order N (for a more formal definition of $X_1(N)$, see for example [26]). Up to birational equivalence, F_N is a defining polynomial for $X_1(N)$. In [9, 8], they make extensive use of F_N, or - more specifically - the optimized versions of the defining polynomials for $X_1(N)$ by Sutherland [27]. Unfortunately, this comes with the drawback of working with massive equations rather quickly, making it hard to manipulate any type of symbolic expressions. We will not require explicitly computing F_N at all.

Let $N \geq 4$ be an integer and E/K an elliptic curve. The Tate pairing is a bilinear map

$$t_N : E(K)[N] \times E(K)/NE(K) \to K^\times/(K^\times)^N : (P_1, P_2) \mapsto t_N(P_1, P_2)$$

which can be computed by means of a Miller function (see for example [8]). The Tate pairing is compatible with isogenies; i.e. if $\varphi : E \to E'$ is an isogeny then $t_N(\varphi(P_1), P_2) = t_N(P_1, \hat{\varphi}(P_2))$ holds. For certain fields, in particular for $K = \mathbb{F}_q$, we have that the Tate pairing is nondegenerate. Remark that the Tate pairing is only defined up to Nth powers, so one must always choose a representative for concrete evaluations.

2.3 Class group actions and CSIDH

Let $E/\mathbb{F}_p$ be a supersingular elliptic curve with $p \equiv 3 \bmod 4$. Then the ring of $\mathbb{F}_p$-rational endomorphisms of E is isomorphic to either $\mathbb{Z}[\sqrt{-p}]$ or $\mathbb{Z}[(1 + \sqrt{-p})/2]$, with the isomorphism determined by mapping p-Frobenius π_p to $\sqrt{-p}$. In the former case, we say that E is on the *floor*, whereas in the latter case E is said to be on the *surface*. Isogenies which do not change the $\mathbb{F}_p$-rational endomorphism ring are said to be *horizontal*, whereas the ones that do are called *vertical* (for more info on this isogeny volcano terminology, see [18]).
For CSIDH [10], fix the set of all supersingular elliptic curves E which have an $\mathbb{F}_p$-rational endomorphism ring isomorphic to $\mathbb{Z}[\sqrt{-p}]$. For any such E and any $\mathfrak{a} \subset \mathbb{Z}[\sqrt{-p}]$, we can associate the subgroup

$$E[\mathfrak{a}] = \bigcap_{\alpha \in \mathfrak{a}} \{P \in E \mid \alpha(P) = \mathcal{O}_E\} \subset E,$$

where α must of course be seen as an endomorphism. The ideal-class group $\mathrm{Cl}(\mathbb{Z}[\sqrt{-p}])$ then acts freely and transitively on this fixed set of supersingular elliptic curves by mapping E to $E/E[\mathfrak{a}]$, where $[\mathfrak{a}]$ is the class of an invertible ideal of the class group.
The underlying assumption in CSIDH is that this turns our setup into a "hard homogenous space" when p is large enough, since $\#\mathrm{Cl}(\mathbb{Z}[\sqrt{-p}]) \approx \sqrt{p}$. For a formal definition of hard homogenous spaces, we refer to the paper by Couveignes [14]. Informally, it means that given E and $E/E[\mathfrak{a}]$, it is conjectured to be hard to find $[\mathfrak{a}]$ (or any equivalent ideal class). The CSIDH setup should thus not be seen as an analogue of SIDH, but rather as the supersingular variant of the CRS scheme [14, 24].

Ideally, we would be able to sample elements from $\mathrm{Cl}(\mathbb{Z}[\sqrt{-p}])$ uniformly at random to obtain a random curve from our fixed set, but in practice this is unfeasible. Since we are working with supersingular elliptic curves, however, we know that $\#E(\mathbb{F}_p) = p+1$ and can thus choose p such that $p+1$ has many small distinct prime factors ℓ_i. This in turn allows us to efficiently evaluate the action of $(\ell_i, \pi - 1)$ (as well as its inverse) with the help of ℓ_i-isogenies that have $\mathbb{F}_p$-rational kernel generators. By choosing sufficient small-degree isogenies, we can represent any element of the class group and efficiently compute its action on a given curve E. The CSIDH key exchange protocol then goes as follows. Alice chooses a secret exponent vector $(e_1, \ldots, e_n)$ corresponding to an ideal

$$\mathfrak{a} = (\ell_1, \pi - 1)^{e_1} \cdot \ldots \cdot (\ell_n, \pi - 1)^{e_n}.$$

From a given starting curve E, she computes $E_A = E/E[\mathfrak{a}]$ and sends it to Bob. Bob chooses his own secret exponent vector and sends $E_B = E/E[\mathfrak{b}]$ to Alice. They can now both compute $E_{AB} = E/E[\mathfrak{ab}]$ as their shared secret.
As a final remark, note that one can also perform CSIDH on the surface [5]; i.e. by using the supersingular elliptic curves with $\mathbb{F}_p$-rational endomorphism ring isomorphic to $\mathbb{Z}[(1+\sqrt{-p})/2]$. The main advantage is that in this setting, horizontal 2-isogenies are also available, allowing a small speed-up.

3 Radical isogeny formulae

Let $E_{b,c}/K$ be an elliptic curve in Tate normal form, with $P = (0,0)$ a point of order N and $N \geq 5$ odd.* Our goal is to efficiently compute a cyclic N^k-isogeny starting from $E_{b,c}$, which can be done iteratively as soon as we have an efficient method to extend the isogeny with kernel $\langle P \rangle$ to a cyclic N^2-isogeny. This extension may not be K-rational, but we have the following theorem which also explains the term "radical isogeny".

Theorem 2. *Let $E_{b,c}/K$ be an elliptic curve in Tate normal form with P a point of order N for some odd integer $N \geq 5$. Let P' be such that the composition*

$$E_{b,c} \xrightarrow{\varphi} E' = E/\langle P \rangle \xrightarrow{\varphi'} E'' = E'/\langle P' \rangle$$

is a cyclic N^2-isogeny. Then the field extension $K(P')$ equals $K(\sqrt[N]{\rho})$ for an appropriately chosen Nth root of the Tate pairing $\rho = t_N(P, -P)$.

Proof. See [9, Theorem 5]. □

Remark 3. For any point P' from Theorem 2 it holds that $\hat{\varphi}(P') = [\lambda]P$ for some $1 \leq \lambda \leq N-1$. Points for which $\lambda = 1$ are called *P-distinguished points* in [9] and we will also use this terminology.

*For the separate case $N = 3$, see [9].

In [9] they use explicit expressions of $X_1(N)$ to find the x-coordinates of P' symbolically. From that, they then compute a Tate normal form of E', where $P' = (0,0)$ is now a kernel generator for an isogeny which extends φ cyclically. This allows for an iterative function on the parameters b, c (or more precisely, on a pair of optimized birational equivalent parameters). In [8] they push this further by means of interpolation methods to find symbolic expressions for P', allowing explicit radical isogeny formulae up to $N = 37$. We will use a different approach and skip a direct formulation of P' altogether.
Consider

$$E_{b,c} \xrightarrow{\varphi} E' \xleftarrow[\sim]{\iota} E_{b',c'},$$

where

- φ is the isogeny with kernel $\langle(0,0)\rangle$, computed with the classical Vélu formulae from Theorem 1;
- ι^{-1} is the Weierstraß isomorphism putting E' into Tate normal form such that $\hat{\varphi}(\iota((0,0)) = P$. Or equivalently, $\iota((0,0)) = P'$ for some P-distinguished point $P' \in E'$.

From Theorem 1, it follows that the curve E' is given by

$$E' : y^2 + (1-c)xy - by = x^3 - bx^2 + a_4x + a_6,$$

where a_4, a_6 are obtained by formulae in function of the x-coordinates of $[k]P$ for $k \in \{1, 2, \ldots, (N-1)/2\}$. The isomorphism ι is defined by a tuple (u, r, s, t), which associates a point $(u^2x + r, u^3y + su^2x + t) \in E'$ to each $(x, y) \in E_{b',c'}$. The pair (r, t) determines the P-distinguished point P' of order N on E' since it is the image of $(0,0) \in E_{b',c'}$. The parameters (u, s) on the other hand rescale and rotate the curve E'. In general, it can be shown that any two out of (u, r, s, t) completely determine ι, and the following lemma provides us with the specific case which is of interest to us.

Lemma 2. *The Weierstraß isomorphism $\iota : E_{b',c'} \to E'$ from the prior discussion is completely determined by (u, s) and (b, c, N).*

Proof. Recall that E' is given by

$$E' : y^2 + (1-c)xy - by = x^3 - bx^2 + a_4x + a_6,$$

where a_4 and a_6 are explicit expressions in b, c, N obtained by the classical Vélu formulae from Theorem 1. We now want to use an isomorphism $\iota : E_{b',c'} \to E'$ to put this into the form

$$E_{b',c'} : y^2 + (1-c') - b'y = x^3 - b'x^2.$$

Applying the map $\iota(x, y) = (u^2x + r, u^3y + su^2x + t)$ to the equation of $E_{b',c'}$ results in the expression

$$\begin{aligned}
&u^6y^2 + (1 - c' + 2s)u^5xy - (b' + c'r - r - 2t)u^3y \\
&= u^6x^3 - (b' - c's - 3r + s^2 + s)u^4x^2 \\
&- (2b'r - b's - c'rs - c't - 3r^2 + rs + 2st + t)u^2x \\
&- (b'r^2 - b't - c'rt - r^3 + rt + t^2).
\end{aligned}$$

The statement then follows by equating the coefficients of this expression to the ones of E'. A condensed version of the resulting expressions are given by:

$$\begin{aligned}
a_1 &= 1-c, \quad f_1 = \frac{3u-a_1}{2}, \quad f_2 = \frac{b-u(b+s(a_1+s))}{2},\\
f_3 &= 2b+a_1 s, \quad f_4 = 2s+a_1, \quad f_5 = f_3+f_1 f_4,\\
f_6 &= bs+a_4-f_2 f_4, \quad f_7 = 3(f_1(a_1+f_1)+b)-f_5,\\
f_8 &= f_5 f_7 + 3f_6 - 9(a_4 - a_1 f_2 + f_1(b-2f_2)),\\
r &= \frac{f_6 f_7 + 9(a_6 + f_2(b-f_2))}{f_8},\\
c' &= 1-\frac{f_4}{u}, \quad b' = \frac{f_3 - b - 3r + s^2}{u^2}.
\end{aligned}$$

□

As mentioned earlier, instead of giving explicit formulae for $E_{b',c'}$ using $P' = (r,t)$, we will use (u,s) instead which is allowed due to Lemma 2. However, we will first settle on a choice of ρ since the Tate pairing is only defined modulo Nth powers. Define $\varpi_0 = 2$ and for all $i \geq 1$ define

$$\varpi_i = \prod_{k=1}^{i} x([k]P),$$

where we use the conventions $x(P) = 1 = x(-P)$ and $x([N]P) = b^2$.

Proposition 2. *Let $E_{b,c}$ be an elliptic curve in Tate normal form with $P = (0,0)$ a point of odd order N. Then the Tate pairing $t_N(P,-P)$ can be represented by*

$$\tau_N := -(b^2 \varpi_N)^{-1}.$$

Proof. This is a porism from [8, Theorem 14]. □

With the notation from above, we now formulate the main contribution of this paper, where we repeat some of the terminology to keep it more self-contained.

Conjecture 1 *Let $N \geq 5$ be an odd integer and $E_{b,c}/K$ an elliptic curve in Tate normal form with $P = (0,0)$ a point of order N. Let $\varphi : E_{b,c} \to E'$ be the isogeny with kernel $\langle P \rangle$, where E' is computed by means of the classical Vélu formulae (i.e. Theorem 1). Let α be an Nth root of the Tate pairing $t_N(P,-P)$, where we choose the representative τ_N from Proposition 2. Define*

$$\begin{aligned}
u &= 1 + 3b \sum_{i=1}^{N-2} \varpi_i \alpha^i - \sum_{i=1, i \neq N-3}^{N-1} \varpi_i \varpi_{i+1} \varpi_{i+2} \alpha^{3i},\\
s &= b \sum_{i=1}^{N-2} \varpi_i \alpha^i - b^3 \tau_N^2 \sum_{i=2}^{N-1} \varpi_{2i} \varpi_{2i+1} \varpi_{N-i-1} \varpi_{N-i} \alpha^{2i}.
\end{aligned}$$

Let $\iota : E_{b',c'} \to E'$ be the Weierstraßß isomorphism determined by (u, s) in the sense of Proposition 1 and Lemma 2, where $E_{b',c'}$ is a curve in Tate normal form. Then $\iota((0,0))$ is a P-distinguished point and - consequently - the isogeny $\varphi' : E_{b',c'} \to E''$ with kernel $\langle(0,0)\rangle$ is such that $\varphi' \circ \iota^{-1} \circ \varphi$ is a cyclic N^2-isogeny. Furthermore, varying the choice of α, i.e. scaling it with Nth roots of unity, provides the formulae for all other N-isogenies for which the kernel intersects $\ker \hat{\varphi}$ trivially.

From φ and ι, one can then easily deduce a map $(b, c) \mapsto (b', c')$, which - when applied k times iteratively - results in a cyclic N^k-isogeny. If one wants to also push points through this isogeny chain instead of merely computing the new codomain curves, it suffices to postcompose the $x(\phi(R))$ and $y(\phi(R))$ from Theorem 1 with the map $(x, y) \mapsto (u^2x + r, u^3y + su^2x + t)$ for every point $R \in E_{b,c}$ at each step of the isogeny chain.

Before providing supporting evidence for this conjecture, we first note its compactness, which likely comes from (at least) the following choices made along the way:

- The specific representation τ_N of the Tate pairing; this was already alluded on in [8] to conjure the most compact expressions.
- The use of classical Vélu formulae, which are normalized isogenies; i.e. $\varphi^*(\omega_{E'}) = \omega_{E_{b,c}}$ where ω denotes the invariant differential of an elliptic curve.
- The choice of a P-distinguished point P' to compute an isomorphism to put E' back into Tate normal form.

Remark 4. For N even, the expression for u in Conjecture 1 seems to also be correct, but the expression for s is not. Unfortunately, we were unable to find a correct expression for s for even N. For cryptographic purposes however, the most relevant composite even-degree radical isogenies are undoubtedly the powers of two, for which the most relevant ones are already covered in [9, 8].

Finding the formulae

Throughout this subsection, it is important to remark that u and s are elements in

$$\frac{\mathbb{Q}(b, c, \alpha)}{(F_N(b, c), \alpha^N - \tau_N)},$$

where the polynomial $F_N(b, c)$ expresses that P has exact order N. An algebraic software package will symbolically express such elements as polynomials of degree at most $N - 1$ in $\mathbb{Q}(b, c)[\alpha]$.[†] Furthermore, the coefficients of these polynomials are not uniquely defined, since they must be seen modulo $F_N(b, c)$.

[†]One can also express them as polynomials in $\mathbb{Q}(b, \alpha)[c]$ by constructing the extensions in a different order. This avenue did not result in fruitful insights.

In [9, 8], they obtain the parameters (r, t) by means of direct symbolical computation and interpolation methods. From this, they compute expressions for b' and c' and try to optimize those. Along the way however, they compute (u, s) since they make use of the isomorphism ι as well.
When printing out the expression for u for small N, the term

$$3b \sum_{i=1}^{N-2} \varpi_i \alpha^i$$

partly stands out due to the factor 3 at all coefficients. Initially it is most clear that by splitting u into these two terms, the coefficients of the separate terms of $u \in \mathbb{Q}(b, c)[\alpha]$ are supported on the modular units F_k as defined in [26], which are related to the irreducible polynomials $F_N(b, c)$.
While trying to spot a pattern for coefficients for increasing powers of α, it becomes clear that certain modular units almost always show up together. The ones that show up together most often have noncoprime indices; e.g. when F_{15} is a factor of a coefficient, F_3 and F_5 are as well. However, when computing the x-coordinate of $[k]P$, these combined factors are *exactly* the ones that appear since one must also take into account the nonvanishing of $[k]P$ for $k < N$. Indeed, for $N > 15$ we have that $[15]P$ must never equal $\mathcal{O}_{E_{b,c}}$, for any specification of the field, thus a factor $F_3F_5F_{15}$ in the denominator is to be expected.
With this in mind, the term with the factor 3 can easily be deduced since the coefficients are exactly the ϖ_i. The term

$$\sum_{i=1, i\neq N-3}^{N-1} \varpi_i \varpi_{i+1} \varpi_{i+2} \alpha^{3i}$$

is somewhat trickier however. Initially it seemed to depend on the value of the exponent modulo 3 when expressed in terms of modular units. But given that $\alpha^N = \tau_N$ is also supported on the modular units, the case distinction can be unified by allowing the exponents to "overflow" and express it as a polynomial of degree $3(N-1)$.
Finding the expression for s is easier once u is found, since the first term of s is the same apart from the factor 3. The other term is again trickier since the coefficients seem to require a case distinction between the odd and even powers of α. Similarly as for u however, it can be unified if we express the term as a polynomial of degree $2(N-1)$.
Remark that the splitting of u in those two terms is the "hard" part since little logic can be found in the coefficients of the combined sum. The coefficient at α^6 for example is $3b\varpi_6 - \varpi_2\varpi_3\varpi_4$, but this is not supported on the modular units and is a fairly complicated expression in $\mathbb{Q}(b, c)/(F_N(b, c))$ already.

Supporting evidence for Conjecture 1

To support the conjecture, we provide the following argumentation:

- These new radical formulae coincide with the radical formulae for all odd integers $N \leq 37$ from [8], in the sense that the (u, s) parameters give rise to the same (b', c') from [8], which were verified symbolically.

– In Section 4, these formulae are used to correctly compute class group actions in the CSIDH setting for all prime-degree isogenies up to $N = 1523$.
– The file `conjecture_test` in [17] furthermore provides verification for these formulae for distinct starting curves over various fields for all odd (not necessarily prime) integers $N \leq 3999$.

Attempts at proving Conjecture 1

We discuss two distinct avenues that may be used to try to prove Conjecture 1. Unfortunately, both of them come with separate issues preventing us from proving it.

Algebraic approach. From (u, s), one can determine the other two parameters (r, t) which fully define the isomorphism ι. In fact, the explicit expressions for b' and c' in Lemma 2 compute r along the way. Looking at the image of ι, it follows immediately that (r, t) is the P-distinguished point $P' \in E'$, which gets mapped to $(0, 0) \in E_{b',c'}$.
With this in mind, one can compute the N-division polynomial $\psi_N(x)$ of E'. This polynomial vanishes exactly on the x-coordinates of all points of $E'[N]$. Proving that $\psi_N(r) = 0$ would essentially prove Conjecture 1, with some minor caveats. One would also need to show that (r, t) and not $-(r, t)$ is the P-distinguished point, although from a cyclic N^k-isogeny point of view this doesn't matter since $\langle P' \rangle = \langle -P' \rangle$. Furthermore, one would need to show that (r, t) is not a kernel generator of the dual of $\varphi : E_{b,c} \to E'$. This follows from the fact that φ is rational, hence $\hat{\varphi}$ is as well, so a kernel generator for $\hat{\varphi}$ would be defined over $\mathbb{Q}(b, c, \zeta_N)/(F_N(b, c), \zeta_N^N - 1)$ instead of $\mathbb{Q}(b, c, \alpha)/(F_N(b, c), \alpha^N - \tau_N)$. Finally, one would also still need to argue that (r, t) is a point of (full) order N, and not a point of order $k \mid N$.
Unfortunately, proving that $\psi_N(r) = 0$ seems highly nontrivial. While the symbolic expressions for u and s are not too bad, the derived ones for r would include summations in the denominator stemming from u and s so verifying the polynomial vanishes is not easy. Furthermore, constructing $\psi_N(x)$ requires using induction, on top of the expression for E' already containing somewhat elaborate summations stemming from Vélu formulae. A first step in trying to use this approach to prove Conjecture 1 would be to find a clean expression for r itself, similar to the ones for u and s. Note that this may lead to faster arithmetic as well, since one could use an x-only arithmetic approach on the Kummer line $E_{b,c}/[-1]$.

Geometric approach. The parameters (u, r, s, t) of the isomorphism ι can also be interpreted geometrically. The P-distinguished point $P' = (r, t)$ is a translation in the plane, whereas u represents a scaling factor and s a rotation. Even though this may not lead to proving the formulae directly, our formulae may be compatible with another geometric interpretation, namely theta coordinates (see [22]).

For simplicity, assume that N is an odd prime and suppose we have a triplet $(E_{b,c}, P, Q)$, where Q is such that $\langle P, Q\rangle = E[N]$. On the one hand, our isogeny $\varphi : E_{b,c} \to E'$ with kernel $\langle P\rangle$ gives rise to a pair $(E', \varphi(Q))$, where $\varphi(Q) = \lambda P'$ for some nonzero scalar λ. On the other hand, we can also construct a map to "forget" P and end up with a pair $(E_{b,c}, Q)$. Starting from just $(E_{b,c}, P)$, the former can be seen as constructing the preimages $\hat{\varphi}^{-1}(P)$ in terms of theta coordinates of level $2N$; see for example [22, Proposition 5.2.2]. Remark that the constants C_{e_i,e_j} from the accompanying equation (5.8) in [22] can be shown to equal the Tate pairings $T_N(e_i, e_j)$ using [22, Corollary 3.3.3]. This argument essentially leads to a higher-dimensional analogue of Theorem 2. In terms of theta coordinates, the "forget" map can be constructed by descending from level $2N$ to level 2; i.e. express the points of $\hat{\varphi}^{-1}(P)$ in theta coordinates of level 2 as discussed in [22, Subsection 2.10.3].

Again, this approach has some subtleties that would need to be addressed, such as the case for composite N, as well as the translation from theta coordinates to Weierstraß equations. The biggest hurdle however, would stem from the fact that the descent part of the approach seems to fundamentally rely on whether N is a sum of two squares or not. If not, one would derive more convoluted expressions which at least differ by a constant factor in complexity. Assuming one could derive explicit formulae for radical isogenies for all odd N by using this geometric approach, they seem to be incompatible with the ones from Conjecture 1 due to this distinction.

4 Cryptographic applications

Our main focus for cryptographic applications will be CSIDH. Other protocols may benefit from these formulae as well of course; in particular CSIDH's ordinary variant CRS [14, 24] should yield similar results.

Operation count for cyclic N^k-isogenies

By means of iteration, our radical isogeny formulae allow efficient computations of long chains of cyclic N^k-isogenies for any odd integer $N \geq 5$. In general, however, it is not known whether τ_N admits an Nth root over the ground field altogether, possibly requiring us to go to a degree-N extension at every next iteration. For isogeny-based cryptographic purposes, we can often assume to work over a finite field $\mathbb{F}_q$ such that $\gcd(q-1, N) = 1$. In this case, the map $\mathbb{F}_q \to \mathbb{F}_q : a \mapsto a^N$ is a bijection, and thus every τ_N admits a *unique* Nth root over $\mathbb{F}_q$. Moreover, we can efficiently compute $\sqrt[N]{\tau_N}$ as τ_N^{μ}, where $\mu \equiv N^{-1} \bmod (q-1)$.

This condition is not rare since it holds for our main application CSIDH, or can be chosen to hold for various other applications as well (e.g. CRS [14, 24] and the VDF from [16]). Furthermore, in the context of the CSIDH framework, these are also the correct isogenies to compute, as shown in the following proposition.

Proposition 3. *Let $E/\mathbb{F}_q$ be an elliptic curve and $N \geq 5$ an odd integer such that $\gcd(q-1, N) = 1$ and $char(\mathbb{F}_q) \nmid N$, and assume that Conjecture 1 is true. Then the cyclic N^k-isogeny obtained by iteratively using the formulae from Conjecture 1 corresponds to the action of the ideal class $[(N, \pi_q - 1)^k]$. Furthermore, this N^k-isogeny can be computed in $2k\log(q) + \mathcal{O}(kN)$ basic $\mathbb{F}_q$-operations.*

Proof. The first part of the statement is simply [9, Lemma 8]. As for the operation count, the cost of each iteration is dominated by the summation and the computation of the Nth root. The former requires $\mathcal{O}(N)$ basic $\mathbb{F}_q$-operations, whereas the latter can be computed as a (full) exponentation given that $\gcd(q-1, N) = 1$. Using square-and-multiply, this can be done in $\log(q)$ squarings and (at most) $\log(q)$ multiplications. □

Remark 5. The realistic cost for the Nth root is closer to $1.5k\log(q)$ basic $\mathbb{F}_q$-operations since the Hamming weight of $N^{-1} \bmod (q-1)$ is typically $(q-1)/2$ for concrete parameters. In practice, this exponentation is often slightly faster, and one could even pick parameters (q, N) to minimize this, but we will not elaborate on this.

We will now provide a more precise count for the required number of arithmetic operations to compute a single radical N-isogeny starting from a given Tate normal form. We will simplify the expression by counting a squaring as a multiplication, and by ignoring additions as well as small scalar multiplications, the latter which we justify by only having single-digit constants in all our formulae. To limit the number of inversions, we will use projective coordinates as much as possible. To the best of our knowledge, no explicit formulae appear anywhere in the literature that describe the multiple scalars of $(0, 0)$ on a curve in Tate normal form, so we provide them here. We omit the formulae for the y-coordinates since these are not needed to compute an isogeny (see the remark following Theorem 1).

Lemma 3. *Let $E_{b,c}$ be an elliptic curve in Tate normal form with $P = (0 : 0 : 1)$ the point of order N. Writing $[k]P = (X_k : Y_k : Z_k)$, we have that $[2]P = (b : bc : 1)$ and for $3 \leq k \leq N-1$ it holds that*

$$X_k = bZ_{k-2}Z_{k-1}(bZ_{k-1} + (c-1)X_{k-1}) - X_{k-2}X_{k-1}^2, \quad Z_k = X_{k-1}^2 Z_{k-2}.$$

Proof. This easily follows from a direct computation using differential addition. □

We will write **E** for an exponentiation, **I** for an inversion and **M** for a (full) multiplication. The following steps are made and the corresponding arithmetic can be verified in [17]. For the sake of simplicity, we will assume $N > 9$ as to not have to deal with exceptional operation counts for $N \in \{5, 7, 9\}$.

- Computing the (X, Z)-coordinates of all multiples of $(0, 0)$ has a cost of $\frac{7N-65}{2}\mathbf{M}$.
- Computing all ϖ_i has a cost of $\mathbf{I} + \frac{7N-25}{2}\mathbf{M}$.
- Computing the Nth root of the Tate pairing τ_N costs **E**.

– Computing the u and s from Conjecture 1 has a cost of $\mathbf{I}+(7N-2)\mathbf{M}$.
– Computing the codomain curve E' or - more precisely - the a_4 and a_6 from Theorem 1 costs $(2N+2)\mathbf{M}$.
– Computing b' and c' from Lemma 2 has a cost of $\mathbf{I}+20\mathbf{M}$.

Combining all of this, we get a total cost of $\mathbf{E}+3\mathbf{I}+(16N-25)\mathbf{M}$ to compute a single radical N-isogeny. The hidden constant for the $\mathcal{O}(N)$ operation count is hence 16 for the number of multiplications. As can be seen in Figure 1, our new formulae start outperforming the work from [8] for $N \geq 19$. Unfortunately, there is currently too much overhead to outperform the optimizations for the smallest odd N. In particular, it seems unlikely that the radical 5-, 7- and 9-isogenies from [9] could ever be improved upon since $X_1(N)$ is one-parametrizable for those values of N.

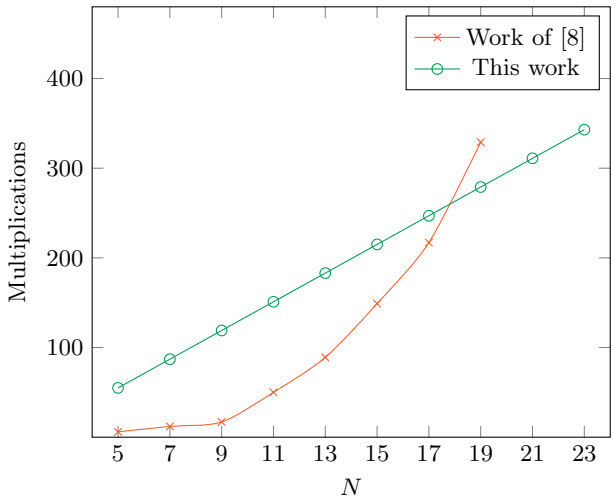

Fig. 1. Comparison of the number of multiplications that are needed to compute a radical N-isogeny between the previous state-of-the-art and this work. In [8], no general formulae were given but the required arithmetic clearly scales superlinearly.

Potential impact on CSIDH

We are now ready to discuss our improvements of CSIDH. We start by giving a brief overview of the algorithmic aspects of the class group action. Recall that Alice wants to compute the action of

$$\left(2, \frac{1+\sqrt{-p}}{2}\right)^{e_1} (3, \pi-1)^{e_2} (5, \pi-1)^{e_3} \dots (\ell_n, \pi-1)^{e_n}$$

on a supersingular elliptic curve $E/\mathbb{F}_p$ for some (secret) exponent vector $(e_1, e_2, \dots, e_n)$. Each of these ideals correspond to computing an $\mathbb{F}_p$-rational ℓ_i-isogeny. For the sake of simplicity, assume all $e_i \geq 0$ (if not,

one simply needs to compute an isogeny starting from the twist which is just a sign swap). Our implementation builds upon the work of [8] and computes the class group action as follows:

- We first compute the 2^{e_1}-isogeny as a chain of 8-isogenies, if necessary followed by a single 2- or 4-isogeny depending on e_1 mod 3. The formulae used here are the ones from [8].
- Similarly, we compute the 3^{e_2}-isogeny as a chain of 9-isogenies with the formulae from [8], followed by a single 3-isogeny if e_2 is odd.
- Next, we successively compute the 5^{e_3}-isogenies, 7^{e_4}-isogenies, 11^{e_5}-isogenies, 13^{e_6}-isogenies and 17^{e_7}-isogenies with the formulae from [8].
- Then, we successively compute the $\ell_i^{e_i}$-isogenies for all $19 \leq \ell_i \leq \ell_j$ by iterating the formulae from Conjecture 1, for some chosen j depending on the parameter set.
- Finally, we finish by computing the $\ell_{j+1}^{e_{j+1}} \dots \ell_n^{e_n}$-isogeny by means of the $\sqrt{\text{élu}}$ isogeny formulae from [3].

For each initial radical isogeny computation, we require a point of exact order ℓ_i to start our chain. Since sampling torsion points is by far most efficient on Montgomery curves, we swap to those forms in-between the Tate normal forms. At the start and end of the class group action we resort to curves of the form $E/\mathbb{F}_p : y^2 = x^3 + Ax^2 - x$ since these allow both easy verification in our CSIDH (on the surface) setting and the use of horizontal 2-isogenies (for more details on this, see [5]).
It is natural to wonder why we only use prime power degrees to compute the chains of 2- and 3-isogenies since the formulae from Conjecture 1 also work for - say - $N = 25$. While this is indeed the case, it also means that in the CSIDH setting it must hold that $25 \mid p + 1$ instead of just $5 \mid p + 1$. This implies that *all* arithmetic for the class group action now needs to be performed over a field with characteristic at least 2 bits more. In practice, the impact of this trade-off is negligible, even for the largest tested parameter set, so we opted to not use it. Computing the 3^{e_2}-isogeny as a chain of 27-isogenies instead of 9-isogenies is even worse, since this could only improve the speed for this part of the computation by a factor of 3/2 instead of 2.

We use the prime characteristic

$$p_{512} = 2^4 \cdot 3 \cdot \underbrace{(3 \cdot 5 \cdot \ldots \cdot 367)}_{72 \text{ consecutive primes}} \cdot 379 \cdot 409 - 1,$$

which is the same as in [8]. Note that $p_{512} \approx 2^{514}$ such that the corresponding class group has size approximately 2^{257}. To compare with the $\sqrt{\text{élu}}$ isogeny formulae, we will sample from the class group in that setting using integer exponents from $[-5; 5]$ since $\log(11^{74}) \approx 2^{256}$. We also make a second comparison with the $\sqrt{\text{élu}}$ isogeny formulae, using a skew sampling interval stemming from the optimal strategies of [12]. For the sake of simplicity, we used the same p_{512} for all benchmarks. Without making use of 2^k-isogenies, one could restrict to the original CSIDH-512 prime p for which $4 \mid p + 1$ but $8 \nmid p + 1$. This would come at a cost of computing 587-isogenies instead of 409-isogenies however, undoing any (very minor) gains.

To choose an optimal maximal exponent vector for our radical isogeny implementation we proceeded as follows. For each $\ell_i \mid p+1$ we computed a long chain of radical ℓ_i-isogenies to assign an averaged weight to a single radical ℓ_i-isogeny. Then, for every fixed ℓ_i and varying (integer) λ_j, we determined what the optimal maximal exponent vector would need to be based on these weights to reach 2^{λ_i} possible ending curves from a class group action exclusively using radical isogenies up to ℓ_i. For instance, if $\ell_i = 5$ and $\lambda_j = 10$, we would look at how to best reach 2^{10} ending curves using only radical 2-, 3- and 5-isogenies. The integer exponent vectors (e_1, e_2, e_3) could be drawn from

$$[-9; 9] \times [-5; 5] \times [-2; 2],$$

since $19 \cdot 11 \cdot 5 \approx 2^{10}$ and computing radical isogenies is faster for smaller ℓ_i given Proposition 3. For each pair (ℓ_i, λ_j), we then determined optimal maximal exponent vectors for the primes $\ell_{i+1}, \ldots, \ell_n$ such that these isogenies could reach $2^{\lambda-\lambda_j}$ possible ending curves, where 2^λ is the size of the class group. This was done by assigning weights to each prime according to the $\sqrt{\text{élu}}$ isogeny formulae. Finally, for each pair (ℓ_i, λ_j) we let our implementation compute the class group action for the maximal exponent vector and choose the fastest out of those options.

The choice for the maximal exponent vector instead of an average one is that it is a more apt benchmark. Indeed, due to side-channel attacks, a concrete implementation of CSIDH would need to run at least in constant time, for which the maximal exponent vector is an obvious lower bound since we can always pretend to be computing the maximal-degree isogeny with dummy variables. Unfortunately, the techniques used in CTIDH [1] do not translate to radical isogeny formulae since they batch isogeny degrees together, whereas we explicitly use long chains of fixed-degree isogenies starting from a curve in their respective Tate normal forms. For some more discussion concerning side-channel analysis regarding CSIDH, see for example [11, 2]. We currently see no way to turn this into a dummy-free constant-time implementation unfortunately, and leave this for future work.

Unsurprisingly, the optimal maximal exponent vector is heavily skewed towards computing many more small-degree isogenies. We found that over half of all isogenies are best computed as radical isogenies; more precisely, 46 out of the 75 prime degrees make use of radical isogeny formulae for the fastest class group action computation. The skew interval is displayed below.

```
[ 129, 90, 45, 44, 36, 33, 25, 21, 19, 16, 16, 14, 13, 12,
12, 11, 10, 10, 9, 9, 8, 8, 7, 7, 7, 6, 6, 6, 6, 6, 5, 5, 5,
5, 4, 4, 4, 4, 4, 4, 4, 4, 3, 3, 3, 3, 2, 2, 2, 2, 2, 2, 2,
2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1, 1,
1, 1 ]
```

Remark that the $\sqrt{\text{élu}}$ isogeny formulae are only used from $\ell_{47} = 211$ onward, and at most two isogenies for each of those primes are required

for the class group action computation. Given that the $\sqrt{\text{élu}}$ isogeny formulae already start outperforming regular isogeny formulae in the original CSIDH implementation from $N \approx 100$, it is clear that "the gap is closed" and the class group action computation is fastest with a combination of radical isogenies and $\sqrt{\text{élu}}$ isogeny formulae.

In Table 1 we show the results of our implementation. We chose for an implementation in Magma to make a fair comparison to the work already done in [9, 8] which was also done in Magma exclusively. As can be seen in Table 1, a sizeable part of the improvement for CSIDH-512 was due to their work, although our new radical isogeny formulae do improve it by another 35%, from 0.5705 to 0.3724. Since Magma uses a lot of internal machinery, we also display the number of multiplications used since those impact the running time most significantly. The multiplications include squarings, as well as the multiplications needed for the exponentiation (which were counted exactly as square-and-multiply and differ very little from the expected $512 + 256$). The multiplication count seems to follow the relative timing rather accurately, given that the number of inversions is limited (1231 for this work and 9352 for the work of [8]) and the inversion/multiplication ratio in this setting is only approximately 8 in Magma. The accompanying code can be found in [17], more precisely in the file `csidh_512`. It can be easily adapted for other prime characteristics, although in that case the exponent vectors would of course no longer be optimal (or even have the correct length for that matter).

We stress that our implementation is merely useful as a tool to compare to the radical isogenies already found and implemented in [9, 8], and is in no way meant as a practical version of CSIDH. Indeed, the constant-time implementation of for example [12] already outperforms our variable-time one with a factor of (approximately) two when it comes to multiplication count, without using any radical isogenies. This difference stems from the fact that certain subroutines in the implementation of [9, 8] are unoptimised; e.g. both their Vélu and $\sqrt{\text{élu}}$ isogeny implementations still require multiple inversions on top of several additional multiplications as well. To the best of our knowledge, there is no practical CSIDH implementation using radical isogenies, and it is unclear how well our speed-ups would carry over to a more optimised version (for instance, the cut-off for using radical isogenies for all $\ell_i \leq 199$ may differ). Ideally, such optimised version is a dummy-free and constant-time one, computing everything projectively and using at most one inversion at the end. Even though some initial attempts have been made in this regard in [11], they do not translate easily to our formulae, so we leave this part for future work.

Larger parameter sets

To show the scaling potential of our radical isogeny formulae on CSIDH, we elect to opt for 2 larger sets of parameters. The primes p_{1024} and p_{2048}

	Relative timing	Multiplications
$\sqrt{\text{élu}}$ + uniform interval [3]	1	4464062
$\sqrt{\text{élu}}$ + skew interval [12]	0.7877	3485937
$\sqrt{\text{élu}}$ + limited radicals [8]	0.5705	2481875
$\sqrt{\text{élu}}$ + $\sqrt[N]{\text{élu}}$ (this work)	0.3724	1615371

Table 1. Relative cost of computing a class group action with maximal coefficients in the CSIDH-512 setting by adding our $\sqrt[N]{\text{élu}}$ isogeny formulae compared to only using the $\sqrt{\text{élu}}$ isogeny formulae from [3]. For a more fair comparison, we also include a row to demonstrate how a skew sampling interval [12] can already benefit CSIDH with just $\sqrt{\text{élu}}$ isogeny formulae. The result using radical N-isogenies up to $N = 19$ from [8] is also displayed. All tests were done in Magma v2.28-5 using an Intel Xeon Gold 6248R CPU at 3.00GHz.

are as follows:

$$p_{1024} = 2^4 \cdot 3 \cdot \underbrace{(3 \cdot 5 \cdot \ldots \cdot 727)}_{\text{128 consecutive primes}} \cdot 743 \cdot 773 - 1,$$

$$p_{2048} = 2^4 \cdot 3 \cdot \underbrace{(3 \cdot 5 \cdot \ldots \cdot 1439)}_{\text{227 consecutive primes}} \cdot \underbrace{(1451 \cdot \ldots \cdot 1471)}_{\text{4 consecutive primes}} \cdot 1523 - 1.$$

Note that $p_{1024} \approx 2^{1023}$ and $p_{2048} \approx 2^{2051}$ resulting in corresponding class groups of size approximately 2^{512} and 2^{1025}. We do not try to make claims about the security for these larger parameter sets. For a discussion regarding that, see for instance [20]. We merely want to showcase how well our radical isogeny formulae scale when using them to compute the class group actions stemming from larger class groups, since the formulae from [9, 8] had been (rightfully) criticized to have less of an impact for larger parameters (see for instance [11]). In particular, the primes p_{1024} and p_{2048} do not target any specifc NIST security level.

To compare with the $\sqrt{\text{élu}}$ isogeny formulae, we sample from the class group in that setting by using integer exponents from respectively $[-7; 7]$ and $[-10; 10]$ since $\log(15^{132}) \approx 516$ and $\log(21^{232}) \approx 1019$. Given that [8] provides no higher parameter set implementations, we merely compare the uniform sampling interval applied to $\sqrt{\text{élu}}$ isogeny formulae with our new isogeny formulae. For CSIDH-1024, the optimal class group action computation uses our radical isogeny formulae for 74 out of 131 primes, whereas for CSIDH-2048 this is 133 out of 233 primes. Their respective relative speed-ups are $1 : 0.3667$ and $1 : 0.3623$. Overall, the radical isogeny formulae exhibit potential to scale up CSIDH relatively stably for all parameter sets. The accompanying code can be found in `csidh_1024` and `csidh_2048` in [17].

Again, it is unclear what the impact of radical isogenies would be on practical implementations of CSIDH for larger parameter sets. However, we are convinced that these formulae showcase that they are not merely somewhat useful for the lowest security level, but that they offer enough potential for further research for all security levels.

5 Conclusion

We provide radical N-isogeny formulae for all odd integers $N \geq 5$, which allow us to efficiently compute a cyclic N^k-isogeny as a chain of k N-isogenies. These contribute an additional 35% speed-up for the CSIDH-512 parameters compared to the partial radical isogeny results from [8]. In general, our new radical isogeny formulae scale well for higher parameter sets of CSIDH as well, providing a significant and consistent speed-up across all levels compared to not using radical isogenies.
Unfortunately, our formulae are only conjectured, so from a mathematical point of view, their proof is an open research question. Additionally, the case for radical N-isogeny formulae for even N is only partially answered, with one out of two parameters found. With the partial shift of focus of isogeny-based cryptographic protocols to higher dimensions, it is an interesting question whether similar results for the multiradical isogenies from [6] exist. Of course these would need to be within the framework of theta constants due to the lack of Vélu formulae in higher dimensions (see for instance [13]). Finally, a dummy-free constant-time implementation of these formulae in the CSIDH framework would be required to make them practical for concrete cryptographic implementations.

References

1. Banegas, G., Bernstein, D.J., Campos, F., Chou, T., Lange, T., Meyer, M., Smith, B., Sotáková, J.: CTIDH: faster constant-time CSIDH. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2021**(4), 351–387 (2021), https://doi.org/10.46586/tches.v2021.i4.351-387
2. Banegas, G., Krämer, J., Lange, T., Meyer, M., Panny, L., Reijnders, K., Sotáková, J., Trimoska, M.: Disorientation faults in CSIDH. In: Advances in Cryptology – EUROCRYPT 2023, part V. pp. 310–342. Springer Nature Switzerland (2023), https://doi.org/10.1007/978-3-031-30589-4_11
3. Bernstein, D.J., De Feo, L., Leroux, A., Smith, B.: Faster computation of isogenies of large prime degree. ANTS-XIV, The Open Book Series **4**(1), 39–55 (2020), https://doi.org/10.2140/obs.2020.4.39
4. Beullens, W., Kleinjung, T., Vercauteren, F.: CSI-FiSh: Efficient isogeny based signatures through class group computations. In: ASIACRYPT 2019, Part I. pp. 227–247. Springer (2019), https://doi.org/10.1007/978-3-030-34578-5_9
5. Castryck, W., Decru, T.: CSIDH on the surface. In: PQCrypto 2020. Lecture Notes in Computer Science, vol. 12100, pp. 111–129. Springer (2020), https://doi.org/10.1007/978-3-030-44223-1_7
6. Castryck, W., Decru, T.: Multiradical isogenies. In: 18th International Conference Arithmetic, Geometry, Cryptography, and Coding Theory, Contemporary mathematics, vol. 779, pp. 57–89. American Mathematical Society (2022), https://www.ams.org/books/conm/779/

7. Castryck, W., Decru, T.: An efficient key recovery attack on SIDH. In: EUROCRYPT 2023, Part II. pp. 423–447. Springer (2023), https://doi.org/10.1007/978-3-031-30589-4_15
8. Castryck, W., Decru, T., Houben, M., Vercauteren, F.: Horizontal racewalking using radical isogenies. In: Advances in Cryptology – ASIACRYPT 2022, Part II. pp. 67–96. Springer Nature Switzerland (2022), https://doi.org/10.1007/978-3-031-22966-4_3
9. Castryck, W., Decru, T., Vercauteren, F.: Radical isogenies. In: Advances in Cryptology – ASIACRYPT 2020, Part II. pp. 493–519. Springer International Publishing (2020), https://doi.org/10.1007/978-3-030-64834-3_17
10. Castryck, W., Lange, T., Martindale, C., Panny, L., Renes, J.: CSIDH: an efficient post-quantum commutative group action. In: ASIACRYPT 2018, Part III. pp. 395–427. Springer (2018), https://doi.org/10.1007/978-3-030-03332-3_15
11. Chi-Domínguez, J., Reijnders, K.: Fully projective radical isogenies in constant-time. In: Topics in Cryptology - CT-RSA 2022. Lecture Notes in Computer Science, vol. 13161, pp. 73–95. Springer (2022), https://doi.org/10.1007/978-3-030-95312-6_4
12. Chi-Domínguez, J., Rodríguez-Henríquez, F.: Optimal strategies for CSIDH. Adv. Math. Commun. **16**(2), 383–411 (2022), https://doi.org/10.3934/amc.2020116
13. Cosset, R., Robert, D.: Computing (ℓ, ℓ)–isogenies in polynomial time on Jacobians of genus 2 curves. Mathematics of Computation **84**(294), 1953–1975 (2015), http://www.jstor.org/stable/24489183
14. Couveignes, J.M.: Hard homogeneous spaces. Cryptology ePrint Archive, available at https://eprint.iacr.org/2006/291 (2006)
15. De Feo, L.: Mathematics of isogeny based cryptography, https://arxiv.org/abs/1711.04062
16. De Feo, L., Masson, S., Petit, C., Sanso, A.: Verifiable delay functions from supersingular isogenies and pairings. In: Advances in Cryptology – ASIACRYPT 2019, Part I. pp. 248–277. Springer International Publishing (2019), https://doi.org/10.1007/978-3-030-34578-5_10
17. Decru, T.: Online repository, https://gitlab.esat.kuleuven.be/Thomas.Decru/radical-velu-isogeny-formulae
18. Delfs, C., Galbraith, S.D.: Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. Des. Codes Cryptogr. **78**(2), 425–440 (2016), https://doi.org/10.1007/s10623-014-0010-1
19. Maino, L., Martindale, C., Panny, L., Pope, G., Wesolowski, B.: A direct key recovery attack on SIDH. In: EUROCRYPT 2023, Part II. pp. 448–471. Springer (2023), https://doi.org/10.1007/978-3-031-30589-4_16
20. Peikert, C.: He gives C-sieves on the CSIDH. In: Advances in Cryptology – EUROCRYPT 2020, Part II. pp. 463–492. Springer International Publishing (2020), https://doi.org/10.1007/978-3-030-45724-2_16
21. Renes, J.: Computing isogenies between Montgomery curves using the action of (0, 0). In: PQCrypto 2018. Lecture Notes in Computer

Science, vol. 10786, pp. 229–247. Springer (2018), https://doi.org/10.1007/978-3-319-79063-3_11
22. Robert, D.: Efficient algorithms for abelian varieties and their moduli spaces. Ph.D. thesis, Université de Bordeaux (UB) (2021)
23. Robert, D.: Breaking SIDH in polynomial time. In: EUROCRYPT 2023, Part II. pp. 472–503. Springer (2023), https://doi.org/10.1007/978-3-031-30589-4_17
24. Rostovtsev, A., Stolbunov, A.: Public-key cryptosystem based on isogenies. Cryptology ePrint Archive, Paper 2006/145 (2006), https://eprint.iacr.org/2006/145
25. Silverman, J.H.: The arithmetic of elliptic curves, vol. 106. Springer (2009), https://doi.org/10.1007/978-0-387-09494-6
26. Streng, M.: Generators of the group of modular units for $\Gamma_1(N)$ over the rationals. Cornell University **arXiv**, 1503.08127v2 (2019), https://arxiv.org/abs/1503.08127v2
27. Sutherland, A.: Constructing elliptic curves over finite fields with prescribed torsion. Mathematics of Computation **81**(278), 1131–1147 (2012), https://www.jstor.org/stable/23267989
28. Vélu, J.: Isogénies entre courbes elliptiques. Comptes-Rendus de l'Académie des Sciences, Série I **273**, 238–241 (1971), in French