Universiteit
Leiden
The Netherlands

Computational aspects of class group actions and applications to post-quantum cryptography

Houben, M.R.

Citation

Houben, M. R. (2024, February 28). *Computational aspects of class group actions and applications to post-quantum cryptography*. Retrieved from https://hdl.handle.net/1887/3721997

Version:	Publisher's Version
License:	Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden
Downloaded from:	https://hdl.handle.net/1887/3721997

Note: To cite this publication please use the final published version (if applicable).

Chapter 7

Horizontal racewalking using radical isogenies

This chapter consists of a paper written together with Wouter Castryck, Thomas Decru, and Frederik Vercauteren. It has been published as

Wouter Castryck, Thomas Decru, Marc Houben, and Frederik Vercauteren. Horizontal racewalking using radical isogenies. In *Advances in Cryptology – ASIACRYPT 2022*, pages 67–96, Lecture Notes in Computer Science, vol 13792. Springer, Cham. `https://doi.org/10.1007/978-3-031-22966-4_3`.

All authors of this paper contributed equally to the work.

Compared to the published version, we have corrected a few typos and mathematical errors, added a reference in the proof of Theorem 7.6.5 to code in the GitHub repository associated to Conjecture 7.6.4, and extended radical isogeny formulae up to degree $N = 41$ (previously up to $N = 37$). The numbering (of e.g. theorems and definitions) in the published version is different.

Abstract

We address three main open problems concerning the use of radical isogenies, as presented by Castryck, Decru and Vercauteren at Asiacrypt 2020, in the computation of long chains of isogenies of fixed, small degree between elliptic curves over finite fields. Firstly, we present an interpolation method for finding radical isogeny formulae in a given degree N, which by-passes the need for factoring division polynomials over large function fields. Using this method, we are able to push the range for which we have formulae at our disposal from $N \leq 13$ to $N \leq 41$ (where in the range $18 \leq N \leq 41$ we have restricted our attention to prime powers). Secondly, using a combination of known techniques and ad-hoc manipulations, we derive optimized versions of these formulae for $N \leq 19$, with some instances performing more than twice as fast as their counterparts from 2020. Thirdly, we solve the problem of understanding the correct choice of radical when walking along the surface between supersingular elliptic curves over $\mathbb{F}_p$ with $p \equiv 7 \bmod 8$; this is non-trivial for even N and was settled for $N = 2$ and $N = 4$ only, in the latter case by Onuki and Moriya at PKC 2022. We give a conjectural statement for all even N and prove it for $N \leq 14$. The speed-ups obtained from these techniques are substantial: using 16-isogenies, the computation of long chains of 2-isogenies over 512-bit prime fields can be accelerated by a factor 3, and the previous implementation of CSIDH using radical isogenies can be sped up by about 12%.

7.1 Introduction

One of the core operations in isogeny-based cryptography is the fast computation of the codomain curve of a cyclic chain of horizontal $\mathbf{F}_q$-isogenies of some fixed small-to-moderate degree $N \geq 2$ between elliptic curves over a finite field $\mathbf{F}_q$. Here, let us recall that an $\mathbf{F}_q$-isogeny between two elliptic curves over $\mathbf{F}_q$ is called horizontal if their $\mathbf{F}_q$-rational endomorphism rings are isomorphic imaginary quadratic orders. The primary use cases are CRS [10, 22] and CSIDH [7], which are proposals for post-quantum key exchange. However fast horizontal isogenies are also key to various other recent constructions, including digital signatures [2], oblivious transfer constructions [15], verifiable delay functions [12], and schemes for delay encryption [11].

This paper presents a speed-up of such computations. More concretely, we upgrade the radical isogeny approach from [6], where for any given N one produces an iterable formula for computing the elliptic curves in a cyclic chain of N-isogenies, with each step involving the extraction of an Nth root of some radicand $\rho_N \in \mathbf{F}_q$; whence the name "radical". Asymptotically, for fixed N and growing q, the cost of evaluating this formula is dominated by one exponentiation in $\mathbf{F}_q$. This should be compared to one scalar multiplication on an elliptic curve over $\mathbf{F}_q$, which is the dominant cost of the standard approach using Vélu's formulae [26]. In practice however, radical isogenies are useful for small N only, because they come with a large overhead; part of the goal of the current paper is to reduce this overhead.

A first problem is simply *finding* radical isogeny formulae. Indeed, while their existence was argued in [6, §3] by means of the Tate pairing, producing concrete instances is a non-trivial task. The method proposed in [6, §4] relies on finding a zero of the reduced N-division polynomial of a Vélu-type codomain curve over a certain modular function field over $\mathbf{Q}$. As N grows, not only the division polynomial but also this codomain curve and the function field become increasingly complicated, and one quickly reaches the point where this method becomes infeasible. Consequently, the GitHub repository accompanying [6] contains no radical isogeny formulae beyond $N = 13$.

A second problem is that radical isogeny formulae are highly non-unique, with freedom coming from the choice of curve-point model (e.g., the Tate normal form), from the choice of the radicand ρ_N, and from relations in the modular function field. Different radical isogeny formulae for the same value of N can have very different practical performances, and in view of the large overhead it is crucial to try and produce the most efficient version. Here we should mention recent work by Onuki and Moriya [17], who use Montgomery curves to find faster formulae in degrees $N = 3, 4$. Chi-Dominguez and Reijnders [9] have presented projective (= inversion-free) radical isogeny formulae in degrees $2 \leq N \leq 5$ and $N = 7, 9$, but these are constructed directly from the corresponding formulae from [6].

A third problem is that it is not always clear *which* Nth root of ρ_N needs to be chosen in order to walk horizontally. In the CSIDH setting of supersingular elliptic curves over a finite prime field $\mathbf{F}_p$, horizontality comes for free if N is odd; in this case ρ_N has exactly one Nth root in $\mathbf{F}_p$. But even-degree $\mathbf{F}_p$-isogenies, of which non-trivial cyclic chains exist when $p \equiv 7 \bmod 8$ only, are a concern. In this case ρ_N will

admit two Nth roots in $\mathbf{F}_p$, and selecting the wrong option will lead to a change of endomorphism ring and, as a result, in a breakdown of the iteration. This can be circumvented by an additional quadratic residuosity check at each step, but this is an annoying extra cost. In [4, Lem. 4] it was shown that this cost can be avoided when $N = 2$, because for the concrete radical isogeny formula presented there, the correct choice always turns out to be the principal square root, i.e. the unique square root which is again a square. This observation was extended to $N = 4$, now in terms of a principal fourth root, first as a conjecture [6, Conj. 2] and recently proved by Onuki and Moriya [17]. As mentioned in [6, §7], the correct generalization to arbitrary even N is not immediately apparent.

Contributions

We contribute significantly to each of the above open problems, which are listed explicitly in [6, §7]. Concretely, we address:

1. *Formula generation.* We develop an entirely different method for finding radical isogeny formulae in any given degree N, which avoids the need for factoring division polynomials over large function fields. The method uses interpolation over the modular curve $X_1(N)$ and is inspired by an alternative, Galois-theoretic proof of the existence of radical isogeny formulae along the lines of [5]. Using this method, we managed to generate radical isogeny formulae in degree as large as $N = 41$.

2. *Formula optimization.* The optimization and/or simplification of rational expressions modulo relations is an old and complicated problem, see for example [16]. In our case however, ad-hoc manipulations seem to yield the best results. We now believe to have found reasonably optimized formulae up to $N = 19$, with e.g. formulae for $N = 11, 13$ that can compete with our (optimized) version of $N = 7$. To highlight one example, for $N = 8$ we present the iteration

$$A \leftarrow \frac{-2A(A-2)\alpha^2 - A(A-2)}{(A-2)^2\alpha^4 - A(A-2)\alpha^2 - A(A-2)\alpha + A} \text{ with } \alpha = \sqrt[8]{\frac{-A^2(A-1)}{(A-2)^4}}$$

   whose counterpart from [6] spanned nearly a quarter of a page.

3. *Ensuring horizontality.* We believe to have found the correct generalization, at least conjecturally, of the observations from [4, Lem. 4], [6, Conj. 2] and [17, §5] for $N = 2, 4$ to arbitrary even N. The surprising new ingredient beyond $N = 4$ is that the principal Nth root needs to be tweaked by the Legendre symbol of a certain coefficient appearing in Tate's normal form; for $N = 4$ this Legendre symbol is always -1 so it goes unnoticed. With the aid of Magma we managed to prove this generalization up to $N = 14$.

One illustrative example where the three contributions resonate is the case $N = 16$. When computing long chains of 2-isogenies, e.g. as in the set-up phase of the delay

function from [11], we can use radical 16-isogenies to take 4 horizontal steps "at once", resulting in an asymptotic speed-up by a factor of 4. Experimentally, we observed a speed-up by a factor of about 3 over a 512-bit prime field.

As for CSIDH, we have generated a new prime CRAD-513 capable of handling radical 8- and 9-isogenies, and using our new and optimized formulae we obtained a speed-up of about 12% when compared to the implementation of CSURF-512 from [6]. Furthermore, comparing this to the pre-radical isogenies implementation of CSIDH-512, one sees that the overall speed-up caused by radical isogenies at the 512-bit prime level is about 35%. We expect that there remains room for pushing this quite a bit further, for example by optimizing formulae for $N > 19$.

7.2 Background

Throughout, we let K denote a field, unless otherwise specified. The base point (= neutral element) of an elliptic curve E/K is denoted by $\mathcal{O}_E$, or just $\mathcal{O}$ if E is clear from the context.

7.2.1 Division polynomials

For an elliptic curve $E : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$, $a_i \in K$ in long Weierstrass form we set $b_2 = a_1^2 + 4a_2$, $b_4 = 2a_4 + a_1a_3$, $b_6 = a_3^2 + 4a_6$, $b_8 = a_1^2a_6 + 4a_2a_6 - a_1a_3a_4 + a_2a_3^2 - a_4^2$. For each integer $N \geq 0$ we define the *N-division polynomial* as

$$\Psi_{E,0} = 0, \quad \Psi_{E,1} = 1, \quad \Psi_{E,2} = 2y + a_1x + a_3, \quad \Psi_{E,N} = t \cdot \prod_{Q \in (E[N] \setminus E[2])/\pm} (x - x(Q)),$$

where $t = N$ if N is odd and $t = \frac{N}{2} \cdot \Psi_{E,2}$ if N is even. Note that $\Psi_{E,2}^2 = 4x^3 + b_2x^2 + 2b_4x + b_6$ is a univariate polynomial in x. These division polynomials can be computed efficiently, thanks to the following recurrence relations:

$$\begin{aligned}
\Psi_{E,3} &= 3x^4 + b_2x^3 + 3b_4x^2 + 3b_6x + b_8, \\
\frac{\Psi_{E,4}}{\Psi_{E,2}} &= 2x^6 + b_2x^5 + 5b_4x^4 + 10b_6x^3 + 10b_8x^2 + (b_2b_8 - b_4b_6)x + b_4b_8 - b_6^2, \\
\Psi_{E,2N+1} &= \Psi_{E,N+2}\Psi_{E,N}^3 - \Psi_{E,N-1}\Psi_{E,N+1}^3 \text{ if } N \geq 2, \\
\Psi_{E,2N} &= \frac{\Psi_{E,N}}{\Psi_{E,2}}(\Psi_{E,N+2}\Psi_{E,N-1}^2 - \Psi_{E,N-2}\Psi_{E,N+1}^2) \text{ if } N \geq 3.
\end{aligned}$$

By definition, we have that $\Psi_{E,N}(P) = 0$ for any non-trivial $P \in E[N]$. If one is interested in the points of exact order N, then one can use the *reduced N-division polynomial* $\psi_{E,N}$ defined as $\Psi_{E,N}/\mathrm{lcm}_{d|N, d\neq N}\{\Psi_{E,d}\}$. For all primes ℓ, we simply have $\Psi_{E,\ell} = \psi_{E,\ell}$. Observe that for $N > 2$, the reduced N-division polynomial of E is a univariate polynomial in x.

Scalar multiplication by N on E can be expressed explicitly using division polynomials [20, Ex. 3.6]:

$$[N]P = \left(\frac{\phi_{E,N}(P)}{\Psi_{E,N}(P)^2}, \frac{\omega_{E,N}(P)}{\Psi_{E,N}(P)^3} \right), \tag{7.1}$$

with $\phi_{E,N} = x\Psi_{E,N}^2 - \Psi_{E,N+1}\Psi_{E,N-1}$ and $\omega_{E,N} = \frac{1}{2\Psi_{E,N}}(\Psi_{E,2N} - \Psi_{E,N}(a_1\phi_{E,N} + a_3\Psi_{E,N}^2))$.

7.2.2 Tate's normal form

We study elliptic curves E/K that are equipped with a distinguished K-rational point P of finite order N. For $N \geq 4$ such a curve-point pair (E, P) is isomorphic to a unique pair of the form

$$E_{b,c} : y^2 + (1-c)xy - by = x^3 - bx^2, \qquad P = (0,0), \tag{7.2}$$

for some $b, c \in K$. This distinguished model is called the *Tate normal form*. It is worth mentioning that the first few scalar multiples of $(0,0) \in E_{b,c}$ are easy expressions in terms of b and c, e.g.,

$$-(0,0) = (0,b),\ 2(0,0) = (b,bc),\ -2(0,0) = (b,0),$$
$$3(0,0) = (c,b-c),\ -3(0,0) = (c,c^2).$$

Expressions for higher multiples can be found using (7.1).

Furthermore, for every $N \geq 4$ one can write down a polynomial $F_N \in \mathbf{Z}[b,c]$ whose vanishing, along with the non-vanishing of the discriminant

$$\Delta(E_{b,c}) = b^3(16b^2 - 8bc^2 - 20bc + b + c(c-1)^3),$$

characterizes in any characteristic that the point $(0,0) \in E_{b,c}$ has exact order N. This polynomial can be found as a factor of the constant term of $\psi_{E_{b,c},N}(x) \in \mathbf{Z}[b,c][x]$, or by analyzing $N(0,0)$. It is uniquely determined up to sign. The first few instances are $F_4 = c$, $F_5 = c - b$, $F_6 = c^2 - b + c$, $F_7 = c^3 - b^2 + bc$, $F_8 = bc^2 - 2b^2 + 3bc - c^2$, see again [23, §2]. Thus, when viewing $E_{b,c}$ over the fraction field of $K[b,c]/(F_N)$, one can think of it as a "universal" curve-point pair from which all elliptic curves $E/\overline{K}$ equipped with a point $P \in E$ of order N are obtained through specialization at (unique) concrete values in $\overline{K}$ for b, c.

7.2.3 Radical isogenies

Vélu's formulae from [26] must be fed with the explicit coordinates of the points in $G = \ker \varphi$. In many applications, this kernel is a priori described in a more implicit form. For instance, in CSIDH it typically concerns the "unique subgroup of $E(\mathbf{F}_p)$ of order ℓ" for some odd prime number ℓ. An explicit generator of this subgroup can be found by repeatedly sampling $Q \leftarrow E(\mathbf{F}_p)$ and computing $\frac{p+1}{\ell}Q$ until its order is ℓ, but this scalar multiplication comes at a major cost which can dominate the application

of Vélu's formulae itself. Radical isogenies, as introduced in [6], are an attempt at mitigating this.

The key observation behind radical isogenies is that if $\ker \varphi$ is cyclic, say generated by a point $P \in E(K)$ of order $N \geq 2$ coprime to $\operatorname{char} K$, then Vélu's formulae for producing a defining equation of $E' = E/\langle P \rangle$ can be augmented with formulae yielding the coordinates of a point $P' \in E'$ such that

$$E \xrightarrow{\varphi} E' = E/\langle P \rangle \to E'/\langle P' \rangle$$

is cyclic of degree N^2. Consequently, when computing a non-backtracking chain of N-isogenies, from the second step onwards the formulae allow to bypass the scalar multiplication. The formulae depend on N and can be chosen to

- be *radical*, in that they are algebraic expressions in the coefficients of E, the coordinates of P and a radical $\sqrt[N]{\rho_N}$, where the radicand ρ_N is itself an algebraic expression in the coefficients of E and the coordinates of P,
- be *complete*, in that changing the choice of $\sqrt[N]{\rho_N}$, i.e., scaling it with Nth roots of unity, produces generators for the kernel of each N-isogeny that cyclically extends φ,
- have *good reduction*, in the sense that they have coefficients in $\mathbf{Z}[1/N]$ and they can be applied to any elliptic curve E, over any field K with $\operatorname{char} K \nmid N$, equipped with a point $P \in E(K)$ of order N.

In [6] the existence of such formulae is argued using properties of the Tate pairing. The good reduction property is in fact stated as a conjecture [6, Conj. 1].

Remark 7.2.1 When working over $K = \mathbf{F}_q$ for some prime power q satisfying $\gcd(q-1, N) = 1$, one usually wants to choose the unique instance of $\sqrt[N]{\rho_N}$ belonging to $\mathbf{F}_q$; see [6, §5.1]. This instance can be computed as ρ_N^{μ} with $\mu \in \mathbf{Z}$ a multiplicative inverse of N modulo $q-1$. So the cost of evaluating the formulae is asymptotically dominated by one field exponentiation. Unfortunately, the formulae come with a large overhead and, for fixed q, they outperform plain Vélu for small values of N only. The main goal of this paper is to push this crossover point to larger values of N. ◊

Example 7.2.2 (taken from [6, §4]) Consider an elliptic curve E with a point P of order $N = 5$. The Tate normal form of this curve-point pair is $E_{b,b} = y^2 + (1-b)xy - by = x^3 - bx^2$, $P = (0,0)$ for some $b \neq 0, (11 \pm 5\sqrt{5})/2$. Vélu's formulae produce the following equation for $E' = E/\langle P \rangle$:

$$y^2 + (1-b)xy - by = x^3 - bx^2 - 5b(b^2 + 2b - 1)x - b(b^4 + 10b^3 - 5b^2 + 15b - 1).$$

Analyzing the roots of $\psi_{E',5}(x)$ shows that for $\alpha = \sqrt[5]{\rho_5}$ with $\rho_5 = b$ the point

$$P' = \big(5\alpha^4 + (b-3)\alpha^3 + (b+2)\alpha^2 + (2b-1)\alpha - 2b, \\ 5\alpha^4 + (b-3)\alpha^3 + (b^2-10b+1)\alpha^2 + (13b-b^2)\alpha - b^2 - 11b \big)$$

on E' has order 5 and generates the kernel of a cyclic extension of φ (it is such that $\hat{\varphi}(P') = P$). There are five such cyclic extensions, corresponding to the five possible choices for α. Rewriting the curve-point pair (E', P') into Tate normal form produces the curve $E_{b',b'}$ where b' is given by the iterable formula

$$\rho_5' = b' = \alpha \frac{\alpha^4 + 3\alpha^3 + 4\alpha^2 + 2\alpha + 1}{\alpha^4 - 2\alpha^3 + 4\alpha^2 - 3\alpha + 1}. \tag{7.3}$$

☆

The above example illustrates the strategy from [6] for *finding* radical isogeny formulae. The cases $N = 2, 3$ are easy to handle [6, §4] so we assume that $N \geq 4$. One starts from the "universal" curve-point pair $E = E_{b,c}$, $P = (0, 0)$ over

$$\mathbf{Q}_N(b, c) := \mathrm{Frac}\, \frac{\mathbf{Q}[b, c]}{(F_N)}$$

and one computes a defining equation for $E' = E/\langle P \rangle$ using Vélu's formulae. One then computes the division polynomial $\psi_{E',N}(x)$ and, for a suitable radicand $\rho_N \in \mathbf{Q}_N(b, c)$, one finds the root $x_0' \in \mathbf{Q}_N(b, c)(\sqrt[N]{\rho_N})$ that is the x-coordinate of a point $P' \in E'$ such that $\hat{\varphi}(P') = P$, using a root-finding algorithm; this step is a severe bottleneck. If successful, then the corresponding y-coordinate $y_0' = y(P')$ can be found by solving a quadratic equation over $\mathbf{Q}_N(b, c)(\sqrt[N]{\rho_N})$. The coordinates x_0', y_0' are the radical isogeny formulae we are after; one hopes, and observes in practice, that the good reduction property comes for free. By writing the curve-point pair (E', P') back in Tate normal form $(E_{b',c'}, (0, 0))$ one obtains formulae for b', c' that can be applied iteratively, as in the case of (7.3).

Concerning the radicand ρ_N, it was argued in [6, §3] that $\rho_N = f_{N,P}(-P)$ works, where $f_{N,P}$ is the function on $E_{b,c}$ with divisor $N(P) - N(\mathcal{O})$ and having leading coefficient 1 when expanded in terms of the uniformizer x/y at $\mathcal{O}$, so that ρ_N is a representative of the Tate pairing $t_N(P, -P)$; see [14, Lem. 1].

7.3 Modular curves and Galois theory

This section recalls some of the theory of Galois coverings of modular curves. We mainly refer to [18] and [19]. Along the way we present an alternative proof of the existence of radical isogeny formulae [6, Thm. 5]. This closely resembles the discussion in [5, §3].

7.3.1 Congruence subgroups

Classically, as Riemann surfaces, modular curves are quotients $X = X_\Gamma = \mathbf{H}^*/\Gamma$ of the extended complex upper half plane $\mathbf{H}^* = \mathbf{H} \cup \mathbf{P}^1(\mathbf{Q})$ by a *congruence subgroup* $\Gamma \subset \mathrm{SL}_2(\mathbf{Z})$, i.e. a subgroup containing $\Gamma(N) \subset \mathrm{SL}_2(\mathbf{Z})$, the kernel of reduction modulo N, for some $N \in \mathbf{Z}_{>0}$. The minimal N for which this last property holds is called the *level* of X. The modular curve X admits a natural Zariski-open subset $Y = \mathbf{H}/\Gamma$, and the (finite collection of) points $X \setminus Y$ are called the *cusps* of X. Modular curves can be seen as irreducible smooth complex projective curves, and they always have a "moduli interpretation", in the sense that they (specifically, the non-cuspidal points) parametrize complex elliptic curves together with some additional structure on the N-torsion subgroup.

To make this latter viewpoint more precise, we will consider a different, slightly more general, method to construct "modular" curves. These modular curves will be more general in the sense that they may be reducible as complex projective curves; but they will be irreducible over $\mathbf{Q}$, and their geometrically irreducible components shall be modular curves in the classical sense. Let $N \geq 1$ be an integer and consider the "universal" elliptic curve

$$E_j : y^2 = 4x^3 - \frac{27j}{j-1728}x - \frac{27j}{j-1728}$$

over $\mathbf{Q}(j)$, whose j-invariant equals the indeterminate j. Let $\mathbf{Q}(j, E_j[N]) \subset \overline{\mathbf{Q}(j)}$ be the field obtained by adjoining the coordinates of all N-torsion points of E_j. Then this is a Galois extension, whose Galois automorphisms are completely determined by their action on $E[N]$. In particular, we have that the Galois group is isomorphic to the automorphism group $\mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$ of the N-torsion.

Let $H \subset \mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$ be a subgroup containing -1. The fixed field $\mathbf{Q}(j, E_j[N])^H$ is the function field of a smooth projective curve over $\mathbf{Q}$, which we will denote by X_H. This curve has a natural moduli interpretation, in the sense that away from a finite set its geometric points parametrize elliptic curves over $\overline{\mathbf{Q}}$ together with a certain structure on the N-torsion. More explicitly, it parametrizes pairs (E, α) up to H-isomorphism, where $\alpha : E[N] \to (\mathbf{Z}/N\mathbf{Z})^2$ is an isomorphism of abelian groups and two pairs (E_1, α_1) and (E_2, α_2) are called H-isomorphic if there exists an isomorphism $\varphi : E_1 \to E_2$ and an element $h \in H$ such that $\alpha_1 = h \circ \alpha_2 \circ \varphi$; see [19, §3] for more details. E.g. if we take for H the subgroup of $\mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$ of upper-diagonal matrices then X_H is the classical modular curve $X_0(N)$, which parametrizes elliptic curves together with a cyclic subgroup of order N.

The connection to modular curves in the classical sense is quite straightforward. If we denote by $\Gamma_H = \pi^{-1}(\mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})) \subset \mathrm{SL}_2(\mathbf{Z})$ the congruence subgroup that is the inverse image of H under the reduction modulo N map $\pi : \mathrm{SL}_2(\mathbf{Z}) \to \mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$, then we have that $X_H \cong X_{\Gamma_H}$ as complex projective curves if and only if $\det(H) = (\mathbf{Z}/N\mathbf{Z})^\times$; in general X_H will be geometrically isomorphic to the disjoint union of $[(\mathbf{Z}/N\mathbf{Z})^\times : \det(H)]$ copies of X_{Γ_H}.

7.3.2 The main suspects

Let $N \geq 3$. The subgroups $H \supset H'$ of $\mathrm{GL}_2(\mathbf{Z}/N^2\mathbf{Z})$ consisting of matrices having respective forms

$$\begin{pmatrix} \pm 1 \bmod N & * \\ 0 \bmod N & * \end{pmatrix}, \qquad \text{and} \qquad \begin{pmatrix} \pm 1 \bmod N & * \\ 0 & * \end{pmatrix}$$

correspond to the modular curves which we denote $X_1(N) = X_H$ and $X_1'(N) = X_{H'}$ respectively. The curve $X_1(N)$ is the classical modular curve parametrizing pairs (E, P) where E is an elliptic curve and $P \in E$ is an N-torsion point. The curve $X_1'(N)$ parametrizes triples (E, P, P') where P' is a *P-distinguished point*, i.e. a point $P' \in E/\langle P \rangle$ that maps to P under the dual isogeny $E/\langle P \rangle \to E$. Alternatively, it parametrizes pairs (E, C), where $C = \{Q, Q + P, \ldots, Q + (N-1)P\}$ is a coset on E modulo the order-N point P, where $NQ = P$.

Let us denote by $K \subset L$ the respective function fields over $\mathbf{Q}$ of these curves:

$$K := \mathbf{Q}(X_1(N)) = \mathbf{Q}(j, E_j[N])^H, \quad L := \mathbf{Q}(X_1'(N)) = \mathbf{Q}(j, E_j[N])^{H'}.$$

Then K, L are the fields $\mathbf{Q}_N(b, c)$ and $\mathbf{Q}_N(b, c, \sqrt[N]{\rho_N})$ from Section 7.2.3. The canonical inclusion $K \hookrightarrow L$ corresponds to the degree-N forgetful map $X_1'(N) \to X_1(N) : (E, P, P') \mapsto (E, P)$. As we will see in the next section, it is possible to deduce from a purely Galois-theoretic argument that the extension L/K is radical.

7.3.3 The Galois structure

Lemma 7.3.1 *Let $N \in \mathbf{Z}_{>0}$ and let $K \subset L$ be a degree N extension of fields whose characteristic does not divide N. Let $\zeta_N \in \overline{L}$ be a primitive Nth root of unity and assume that $L(\zeta_N)$ is Galois over K with Galois group*

$$\mathrm{Gal}(L(\zeta_N)/K) = \mathrm{Gal}(L(\zeta_N)/K(\zeta_N)) \rtimes \mathrm{Gal}(L(\zeta_N)/L),$$

where the first factor is cyclic of order N, say generated by σ, and where the semidirect product is according to the rule

$$\tau_j \circ \sigma^i \circ \tau_j^{-1} = \sigma^{ij} \tag{7.4}$$

for all $i = 0, 1, \ldots, N-1$ and all $\tau_j : \zeta_N \mapsto \zeta_N^j \in \mathrm{Gal}(L(\zeta_N)/L)$. Then there exists an $\alpha \in L$ such that $L = K(\alpha)$ and $\alpha^N \in K$.

Proof. The restricted maps $\sigma^i|_L : L \to L(\zeta_N)$ are pairwise distinct. Indeed, if $i, i' \in \{0, 1, \ldots, N-1\}$ are such that $\sigma^i|_L = \sigma^{i'}|_L$, then

$$\sigma^{i-i'} \in \mathrm{Gal}(L(\zeta_N)/K(\zeta_N)) \cap \mathrm{Gal}(L(\zeta_N)/L) = \{\mathrm{id}\},$$

which can only be true if $i = i'$. From [21, Lem. 0CKL] we get that these restricted maps are linearly independent over $L(\zeta_N)$. Thus there exists $\beta \in L$ such that $\alpha :=$

$\sum_{i=0}^{N-1} \zeta_N^i \sigma^i(\beta)$ is non-zero. From

$$\tau_j(\alpha) = \sum_i \zeta_N^{ij}(\tau_j \circ \sigma^i)(\beta) = \sum_i \zeta_N^{ij}(\sigma^{ij} \circ \tau_j)(\beta) = \sum_i \zeta_N^{ij}\sigma^{ij}(\beta) = \alpha$$

it follows that $\alpha \in L$ as well. Now observe that α was constructed in such a way that $\sigma^i(\alpha) = \zeta_N^{-i}\alpha$ for $i = 0, 1, \ldots, N-1$, which has two crucial consequences. On the one hand, it implies that $\mathrm{Gal}(L(\zeta_N)/L)$ is the exact group of automorphisms fixing $K(\alpha)$, or in other words $L = K(\alpha)$. On the other hand, it implies that $\sigma(\alpha^N) = \sigma(\alpha)^N = (\zeta_N\alpha)^N = \alpha^N$, so that α^N is fixed by the entire Galois group, i.e. $\alpha^N \in K$ as wanted. □

Now let K, L as in Section 7.3.2. Below we give an alternative proof of the fact that L/K is a radical extension. Our strategy is to apply Lemma 7.3.1, so we will first prove that $L(\zeta_N)/K$ is Galois, and then find explicitly elements $\sigma, \tau_j \in \mathrm{Gal}(L(\zeta_N)/K)$ satisfying (7.5).

Theorem 7.3.2 *The morphism $X_1'(N) \to X_1(N)$ is a simple radical extension, i.e. the degree N extension of function fields*

$$\mathbf{Q}(j, E_j[N^2])^H \subseteq \mathbf{Q}(j, E_j[N^2])^{H'}$$

can be realized by adjoining $\sqrt[N]{\rho}$ for some function ρ on $X_1(N)$.

Proof. Let $\mathcal{H} \subset H'$ be the subgroup consisting of matrices whose determinant is $\equiv 1 \pmod N$. Then the corresponding fixed field $\mathbf{Q}(j, E_j[N^2])^{\mathcal{H}}$ is $L(\zeta_N)$. One can verify that $\mathcal{H}$ is a normal subgroup of H, which implies that $L(\zeta_N)/K$ is Galois of degree $N\varphi(N)$ with Galois group $H/\mathcal{H}$.

In order to understand its structure, we first consider the intermediate extension $L \subseteq L(\zeta_N)$, which is just a cyclotomic extension with Galois group $\{\, \tau_j : \zeta_N \mapsto \zeta_N^j \mid 0 \leq j < N,\ \gcd(j, N) = 1 \,\} \cong (\mathbf{Z}/N)^*$. When viewed as elements of $H/\mathcal{H}$, these maps can be identified with

$$\tau_j = \begin{pmatrix} 1 & 0 \\ 0 & j \end{pmatrix} \bmod \mathcal{H}.$$

Next, we concentrate on the intermediate extension $K(\zeta_N) \subset L(\zeta_N)$ which is of degree N, and its Galois group can be identified with the cyclic group

$$\left\langle \sigma := \begin{pmatrix} 1 & 0 \\ N & 1 \end{pmatrix} \right\rangle = \left\{ \sigma^i = \begin{pmatrix} 1 & 0 \\ iN & 1 \end{pmatrix} \,\middle|\, i = 0, 1, \ldots, N-1 \right\},$$

which, as before, we consider modulo $\mathcal{H}$. It is easy to see that the elements $\tau_j \circ \sigma^i$ are pairwise distinct (e.g. because j is fully determined by the action of $\tau_j \circ \sigma^i$ on ζ_N, and then the uniqueness of i follows at once). Therefore these $N\varphi(N)$ elements must constitute the whole Galois group. The structure of the Galois group is then

determined by the rules $\sigma^N = 1$, $\tau_j^{\varphi(N)} = 1$, and

$$\sigma^i \circ \tau_j = \begin{pmatrix} 1 & 0 \\ iN & j \end{pmatrix} = \tau_j \circ \sigma^{ij^{-1}}; \tag{7.5}$$

matching (7.4). The result now indeed follows by applying Lemma 7.3.1. □

Remark 7.3.3 The subgroup $\mathcal{H} \subset \mathrm{GL}_2(\mathbf{Z}/N^2\mathbf{Z})$ introduced in the proof of the Theorem corresponds to a modular curve $\mathcal{X}_1'(N)$ over $\mathbf{Q}$ with function field $L(\zeta_N)$. Since $[(\mathbf{Z}/N^2\mathbf{Z})^\times : \det(\mathcal{H})] = \varphi(N)$ it consists geometrically of $\varphi(N)$ copies of $X_1'(N)$, labeled by the different primitive Nth roots of unity ζ_N.

The level structure induced by $\mathcal{H}$ yields the following moduli interpretation of $\mathcal{X}_1'(N)$: it parametrizes triples (E, C, R), where $(E, C) \in X_1'(N)$ is as in Section 7.3.2 and $R \in E[N]$ is an N-torsion point independent of P (i.e. such that $E[N] = \langle P, R\rangle$), where we identify two such points R_1 and R_2 if their Weil pairing with P yields the same (primitive) Nth root of unity, i.e. if $e_N(P, R_1) = e_N(P, R_2)$. Forgetting R leads to a covering $\mathcal{X}_1'(N) \to X_1'(N)$ of degree $\varphi(N)$.

One can make sense of the Galois action of $L(\zeta_N)/K$ in terms of this moduli interpretation. Given a triple $\mathcal{P} = (E, \{Q, Q+P, \ldots, Q+(N-1)P\}, R)$, the images under σ and τ_j are

$$\begin{aligned} \sigma(\mathcal{P}) &= (E, \{Q+R, Q+R+P, \ldots, Q+R+(N-1)P\}, R), \\ \tau_j(\mathcal{P}) &= (E, \{jQ, jQ+P, \ldots, jQ+(N-1)P, R). \end{aligned}$$

◊

7.4 Radical isogeny formulae through interpolation

We now describe the method we used to compute the radical isogeny formulae. Explicitly, starting from the universal Tate normal curve $E = E_{b,c}$ over $K = \mathbf{Q}_N(b, c)$ together with the point $P = (0, 0) \in E$ of order $N \geq 4$, we would like to find an expression for the coordinates of a P-distinguished point P' on the quotient curve $E' = E/\langle P\rangle$ (whose Weierstrass model, let us assume, is given by Vélu's formulae). According to Section 7.3, these coordinates live over some radical field extension L of K. For simplicity, we will mostly focus on computing the x-coordinate of P', as the computation of the y-coordinate is more or less analogous.

7.4.1 A linear system

Let us denote by $\overline{K}$ an algebraic closure of K, and let $Q \in E(\overline{K})$ be such that $NQ = P$. We would like to find an expression for

$$\beta_0 := \sum_{i=0}^{N-1} x(Q + iP),$$

since by Vélu's formulae this is equivalent to finding the x-coordinate of P'. If we define

$$\gamma_d := \sum_{S \in E[N]} e_N(P,S)^d x(Q+S),$$

then $\gamma_d^N \in K$ for all $d \in \mathbf{Z}$: indeed, let $R \in E(\overline{K})$ be an N-torsion point so that $E[N] = \langle P, R \rangle$ and denote by $e_N : E[N] \times E[N] \to \overline{K}$ the Weil pairing. Then $\zeta_N := e_N(P,R)$ is a primitive Nth root of unity. By Remark 7.3.3, it follows that

$$\gamma_d = \sum_{j=0}^{N-1} e_N(P,jR)^d \sum_{i=0}^{N-1} x(Q+jR+iP) = \sum_{j=0}^{N-1} \zeta_N^{jd} \sigma^j(\beta_0),$$

for some generator $\sigma \in \mathrm{Gal}(L(\zeta_N)/K)$ of $\mathrm{Gal}(L(\zeta_N)/K(\zeta_N))$. Following the last paragraph of the proof of Lemma 7.3.1 now shows that $\gamma_d^N \in K$.

Note that $\gamma_d \in L(\zeta_N)$ depends on the choice of Q. However all of them are related as follows:

Lemma 7.4.1 *Let $Q, Q' \in E[N^2]$ be such that $NQ = NQ' = P$. Then there exists an Nth root of unity $\zeta \in \overline{K}$ such that $\gamma_d(Q) = \zeta^d \gamma_d(Q')$ for all $d \in \mathbf{Z}$. Moreover, for all $d \in \mathbf{Z}$ we have that γ_d/γ_1^d is an element of K that is independent of the choice of Q.*

Proof. We have that Q' differs from Q by an N-torsion point. Note that adding multiples of P to Q clearly does not affect the value of γ_d while adding a multiple kR of R scales it by ζ_N^{-kd}. This shows the first statement with $\zeta = \zeta_N^{-k}$. For the second part, note that the independence on Q already follows from the first part. Now let σ be as above and let τ_j be a generator for the cyclotomic extension $L(\zeta_N)/K(\gamma_1)$. Then $\tau_j(\gamma_d) = \gamma_d$, whereas $\sigma(\gamma_d) = \zeta_N^{-d}\gamma_d$. Since σ, τ_j together generate $\mathrm{Gal}(L(\zeta_N)/K)$ we see that γ_d/γ_1^d is invariant under all Galois automorphisms of $L(\zeta_N)/K$ and we conclude that it is an element of K. □

Defining

$$\beta_j := \sigma^j(\beta_0) = \sum_{i=0}^{N-1} x(Q+jR+iP),$$

we now have the following linear system.

$$\begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \zeta_N & \zeta_N^2 & \cdots & \zeta_N^{N-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \zeta_N^{N-1} & \zeta_N^{2(N-1)} & \cdots & \zeta_N^{(N-1)^2} \end{pmatrix} \begin{pmatrix} \beta_0 \\ \beta_1 \\ \vdots \\ \beta_{N-1} \end{pmatrix} = \begin{pmatrix} \gamma_0 \\ \gamma_1 \\ \vdots \\ \gamma_{N-1} \end{pmatrix}.$$

In particular, if we set $\alpha := \gamma_1$ then we see that

$$\beta_0 = \frac{1}{N}\sum_{d=0}^{N-1}\gamma_d = \frac{1}{N}\sum_{d=0}^{N-1}\left(\frac{\gamma_d}{\gamma_1^d}\right)\alpha^d \in K(\alpha) = L. \tag{7.6}$$

We have now reduced the problem of finding radical isogeny formulae (at least the determination of the x-coordinate of P') to finding expressions for the elements $\gamma_d/\gamma_1^d \in K$ for all $d \in \{0, \ldots, N-1\}$. In the next subsection we will describe the method we used to do this. Before that we should point out one subtlety. To ensure that (7.6) is well defined we must have $\alpha \neq 0$; in fact, to be able to use the formula in practice, we should know exactly the value of $\alpha^N \in K$. Though, given N, this is not so difficult to establish (or even guess) in practice; a proof of a closed expression for α^N that works for all N can be found in the appendix (from which it also follows that α is never zero), see Theorem 7.7.1.

7.4.2 Finding the formulae

Expressions for $c_d := \gamma_d/\gamma_1^d$ will of course depend heavily on how one represents the field $K = \mathbf{Q}(X_1(N))$. It turns out that the representation $K = \mathbf{Q}_N(b, c)$ as presented in Section 7.2.3 is not always optimal. In order to minimize the complexity of the resulting formulae, as well as the running time complexity of the algorithm used to find them, we will instead employ Sutherland's optimized models of $X_1(N)$ [24]. These models are optimal in the sense that they write K as the fraction field, which we will denote $\mathbf{Q}_N(A, B)$, of $\mathbf{Q}[A, B]/G_N(A, B)$ for some modular polynomial $G_N(A, B)$ whose degree in B matches the gonality of $X_1(N)$ over $\mathbf{Q}$ (at least for $N \leq 40$). In particular, we can theoretically write every element of K, specifically the c_d we are after, as a polynomial in $\mathbf{Q}(A)[B]$, where the degree in B is as small as one could hope for. It is also possible, and relatively easy in fact, to find an explicit expression for $b, c \in K$ in terms of Sutherland's functions A, B, so one can also express the universal Tate normal curve $E_{b,c}$ as a curve $E_{A,B}$ over $\mathbf{Q}_N(A, B)$.

The idea is now to determine the reduction $\overline{c_d} \in \mathbf{F}_p(A)[B]$ of the coefficients c_d modulo several primes p, and then to lift the results to $\mathbf{Q}(A)[B]$ using the Chinese Remainder Theorem. To find the $\overline{c_d}$, we sample many curves $E_{A,B}$ over $\mathbf{F}_p$ for which Q, R, and ζ_N of the previous section are all defined over $\mathbf{F}_p$. For each of these curves, we explicitly compute the coefficients c_d as elements of $\mathbf{F}_p$. Then, as long as the number of samples is sufficiently large, we can determine an expression for $\overline{c_d} \in \mathbf{F}_p(A)[B]$ by means of rational interpolation (this last step can be achieved purely by linear algebra over $\mathbf{F}_p$).

The main problem that arises is how to efficiently generate suitable samples $(A, B) \in X_1(N)(\mathbf{F}_p)$. The requirement that ζ_N be defined over $\mathbf{F}_p$ is rather trivially met by demanding that $p \equiv 1 \pmod{N}$. The condition that $Q, R \in E_{A,B}(\mathbf{F}_p)$, however, is more intricate, and simply generating random curves turns out to be far too inefficient for large N. Instead, we rely on an approach based on the theory of complex multiplication.

The CM Method

The endomorphism ring of an elliptic curve $E/\mathbf{C}$ is isomorphic to either $\mathbf{Z}$ or an order $\mathcal{O}$ in an imaginary quadratic number field. In the latter case we say that E has complex multiplication (CM) by $\mathcal{O}$. The j-invariants of such elliptic curves are algebraic integers. The *Hilbert class polynomial* $H_D(X) \in \mathbf{Z}[X]$ is the minimal polynomial over $\mathbf{Q}$ of the j-invariant of an elliptic curve $E/\mathbf{C}$ with CM by the quadratic order of discriminant D.

Ordinary elliptic curves over a finite field always have CM. An ordinary elliptic curve $E/\mathbf{F}_q$ with CM by the imaginary quadratic order $\mathcal{O}$ of discriminant D exists if and only if there exist $t, u \in \mathbf{Z}$ such that $u^2 D = t^2 - 4q$ and $p \nmid t$ (where $p = \operatorname{char} \mathbf{F}_q$). In this case H_D splits completely over $\mathbf{F}_q$ and its roots are precisely the j-invariants of elliptic curves with CM by $\mathcal{O}$. The trace of Frobenius of such curves is $\pm t$, so they will have $q + 1 \pm t$ points. One can use this to find curves over $\mathbf{F}_q$ with a desired number of points; this is known as the *CM Method*.

Sampling curves with torsion

We now describe how to use the CM method to construct curves $E_{A,B}$ with full N^2-torsion over $\mathbf{F}_p$; this will certainly ensure that the desired points Q, R be defined over $\mathbf{F}_p$. We thus want to find curves with number of points divisible by N^4. One approach is to strengthen the requirement that $p \equiv 1 \pmod{N}$ to $p \equiv 1 \pmod{N^4}$ and construct curves of trace 2 using the CM method, i.e. with CM by an order whose discriminant D satisfies an equation of the form $u^2 D = 2^2 - 4p$ for some $u \in \mathbf{Z}_{>0}$. The structure of the $\mathbf{F}_p$-rational N^∞-torsion also be controlled by D; if we choose D to be a divisor of $(2^2 - 4p)/N^4$ then $E[N^2](\mathbf{F}_p) \cong (\mathbf{Z}/N^2\mathbf{Z})^2$, see e.g. [8, Thm. 7].

Algorithm

We summarize the above discussion in the following pseudo algorithm generating radical isogeny formulae for $N \geq 4$. The SageMath code we used can be found in the GitHub repository accompanying this paper.

(i) Find all prime numbers $p \equiv 1 \pmod{N^4}$ up to a certain bound.

(ii) For each prime number p, determine the roots j_i of the Hilbert class polynomials H_D modulo p for every imaginary quadratic discriminant D of the form $u^2 N^4 D = 4(p-1)$ for some $u \in \mathbf{Z}$.

(iii) For each root j_i, determine the $(A, B) \in X_1(N)(\mathbf{F}_p)$ for which $j(E_{A,B}) = j_i$.

(iv) For each pair (A, B), if $E_{A,B}$ has trace $+2$, determine $c_d \in \mathbf{F}_p$ for all $d \in \{0, \ldots, N-1\}$.

(v) For each d, find a formula for $c_d \in \mathbf{F}_p(A)[B]$ by rational interpolation.

(vi) Lift the formulae to $\mathbf{Q}(A)[B]$ by the Chinese Remainder Theorem.

7.4.3 Iterative formulae

The above describes how to find an expression for the x-coordinate of P' as an element of $L = K(\alpha)$. An analogous method can be used to find an expression for the y-coordinate. By transforming the pair (E', P') to Tate normal form one can then also determine explicit formulae for Sutherland's parameters $A', B' \in L$ corresponding to the point $(E', P') \in X_1(N)(L)$. In this way, we obtain radical isogeny formulae that can be applied iteratively. We list formulae for prime powers $16 < N \leq 41$ in our GitHub repository.[1]

7.5 Optimizing the formulae

When optimizing radical isogeny formulae, one needs to take into account all of the following choices.

- The radicand ρ_N is not unique: it can be scaled with Nth powers in $\mathbf{Q}_N(b, c)$, and it can be raised to exponents that are coprime with N. Switching from one radicand to another results in different radical isogeny formulae with different performances.
- It is not self-evident that the optimized representations of $X_1(N)$ by Sutherland from [24] will result in optimized radical isogeny formulae.
- Elements in $\mathbb{Q}_N(b, c, \alpha)$ can be expressed in several ways since we work modulo the two relations $F_N(b, c) = 0$ and $\alpha^N = \rho_N(b, c)$.
- It is a priori not clear what formulae we are trying to optimize; e.g. for $E' = E/\langle P \rangle$ we can try to find optimal expressions for a P-distinguished point P' on E', or we can try to write E' in Tate normal form immediately.

We will focus on finding efficient enough formulae in this setting, where it seems nigh impossible to prove that they are indeed the most optimal (especially for $N \geq 10$ as we will see further up ahead). Hence we do not claim they are optimal, but they should not be far off and at the very least in certain cases a big improvement compared to the work in [6].

For $N \in \{4, 5, \ldots, 10\} \cup \{12\}$, the Tate normal form can be parametrized by a single parameter, say A. This means that the codomain curve of a radical N-isogeny can be put into a (new) Tate normal form with a single parameter, say A', where we translated the P-distinguished point P' to $(0, 0)$. In practice, this new parameter seems a good candidate to try to optimize, as can be seen from the case of $N = 4, 5$ from [6]. The raw equation for A' can be easily obtained by any algebraic software package for these small N.

To find an efficient representation of A', consider the curve $X_1'(N)$ defined by $\alpha^N - \rho_N, F_N = 0$. Then A' can be seen as a function on this curve and we can compute its divisor. For $N < 10$, an algebraic software package has no issues checking

[1] `https://github.com/KULeuven-COSIC/Horizontal_Radical_Isogenies`

which linear combinations of places in its support constitute principal divisors, and we can use this to peel off (easy) factors from A'. For every $N \in \{4, \ldots, 9\}$, there are clear contenders for which factorization is most efficient. We list them all, skipping the case $N = 5$ which can be found in (7.3). Note that for $N \geq 6$, our "factorization" merely amounts to writing A' as the quotient of two easyish expressions in A and α.

N = 4. In this case we have $b = A$, $c = 0$ and for $\alpha^4 = A$ we have that

$$A' = \alpha \frac{4\alpha^2 + 1}{(2\alpha + 1)^4}. \tag{7.7}$$

N = 6. In this case we have $b = A(A-1)$, $c = A - 1$ and for $\alpha^6 = -A^2(A-1)$ we have that

$$A' = \frac{(-3A + 2)\alpha^4 + 3A^2\alpha^2 + 2A\alpha - 3A^3 + 4A^2}{\alpha^4 + 2A\alpha^2 + 3A\alpha + A^2}. \tag{7.8}$$

N = 7. In this case we have $b = A^2(A-1)$, $c = A(A-1)$ and for $\alpha^7 = A^4(A-1)$ we have that

$$A' = \frac{\alpha^6 + A\alpha^5 + 2A^3\alpha^2 - A^3\alpha + A^4}{-\alpha^6 + A\alpha^4 + A^3\alpha^2 - 2A^3\alpha + A^4}.$$

N = 8. In this case we have that $b = \frac{A(A-1)}{(A-2)^2}$, $c = \frac{-A(A-1)}{A-2}$ and for $\alpha^8 = \frac{-A^2(A-1)}{(A-2)^4}$ we have that

$$A' = \frac{-2A(A-2)\alpha^2 - A(A-2)}{(A-2)^2\alpha^4 - A(A-2)\alpha^2 - A(A-2)\alpha + A}.$$

N = 9. In this case we have that $b = A^2(A-1)(A^2 - A + 1)$, $c = A^2(A-1)$ and for $\alpha^9 = A^4(A-1)(A^2 - A + 1)^3$ we have that

$$A' = \frac{A(A^2 - A + 1)(\alpha^5 + A(A^2 - A + 1)\alpha^2 + A^2(A^2 - A + 1)^2)}{\alpha^7 - A(A^2 - A + 1)(A-1)\alpha^4 - A^3(A^2 - A + 1)^2\alpha + (A(A^2 - A + 1))^3}.$$

For $N \geq 10$, Magma struggles to efficiently verify whether a given divisor is principal, and those that do get found are less clean than the above factors, so we will optimize these two cases with the more general method for larger N.[2]

If we compute E' as $E/\langle P \rangle$ by means of Vélu's formulae, then E' is in (long) Weierstrass form and we still need to compute an isomorphism to put E' back in Tate normal form E'_t for certain $b', c' \in \mathbb{Q}_N(b, c, \alpha)$. By [20, Prop. 1.3(d)], the isomorphism $\iota : E'_t \to E'$ is determined by a 4-tuple (u, r, s, t), where $P' = (r, t)$ is the P-distinguished point and u is a unit. This u, when seen as a polynomial of degree $N-1$ in $\mathbb{Q}_N(b, c)[\alpha]$, seems to always be efficient to write down and evaluate. Furthermore, the expressions uc' and ub'/c' also enjoy this feature. In particular, a factor that arises in the coefficient of α^i has a high chance of also being there in the coefficient of α^j for $j > i$, which makes this efficient to evaluate in a Horner scheme with rising powers of α. We provide the concrete expressions for $N = 10$ and refer the reader

[2] We remark that for the smaller N it can be extremely fast to let a computer algebra software package verify that a given divisor is *not* principal, but to prove it is principal is harder in the majority of cases.

to our GitHub repository for larger N. Remark that for $N = 10$ we still work with a one-parameter family of curves and the expression uA' is just as efficient as uc' or ub'/c'. The operation counts for all formulae $N \in \{4, 5, \ldots, 17\} \cup \{19\}$ can be found in Table 7.1.

N = 10. In this case we have

$$b = \frac{A^3(A-1)(2A-1)}{(A^2-3A+1)^2}, \quad c = \frac{-A(A-1)(2A-1)}{(A^2-3A+1)}, \quad \alpha^{10} = \frac{A^9(A-1)(2A-1)^2}{(A^2-3A+1)^5},$$

and then $A' = v_{A'}/u$ with

$$\begin{aligned} u = 1 &+ 3\alpha + \frac{4A-1}{A}\alpha^2 + \frac{2c}{b}\alpha^3 - \frac{c(A-4)}{bA}\alpha^4 + \frac{(A-1)(4A-1)}{bA}\alpha^5 + \\ &\frac{(A+1)(A-1)}{bA^2}\alpha^6 + \frac{4c(A-1)}{b^2A}\alpha^7 + \frac{c(A-1)(4A-1)}{b^2A^2}\alpha^8 - \frac{c^2(A-1)}{b^3A}\alpha^9, \\ v_{A'} = A &+ 2\alpha + \frac{A+1}{A}\alpha^2 + \frac{3c}{b}\alpha^3 + \frac{c(A+1)}{bA}\alpha^4 + \frac{(A-1)(A+1)}{bA}\alpha^5 + \\ &\frac{(A+1)(4A-1)}{bA^2}\alpha^6 + \frac{c(A-1)}{b^2A}\alpha^7 + \frac{c(A+1)(A-1)}{b^2A^2}\alpha^8 + \frac{c^2(A-1)}{b^3A}\alpha^9. \end{aligned}$$

7.6 Ensuring horizontality

If both E and P are defined over a finite field $\mathbf{F}_q$ with $\gcd(q-1, N) = 1$ then, as discussed in [6, §5.1], the isogeny $\varphi : E \to E' = E/\langle P \rangle$ is necessarily horizontal. The radicand $\rho_N \in \mathbf{F}_q$ admits a unique Nth root $\alpha \in \mathbf{F}_q$, and for this choice of α the resulting point $P' \in E'$ is again defined over $\mathbf{F}_q$, so the argument repeats. Thus, if N and $q-1$ are coprime, then walking horizontally using radical isogenies is natural and easy. As explained in Remark 7.2.1, for any fixed N the cost of an iteration is dominated by this Nth root extraction, which amounts to one exponentiation in $\mathbf{F}_q$. But if $\gcd(q-1, N) > 1$ then maintaining horizontality is more subtle.

In the remainder of this section we focus on the CSIDH case of supersingular elliptic curves over a finite prime field $\mathbf{F}_p$, where this issue arises (only) if $p \equiv 7 \bmod 8$ and one navigates with cyclic isogenies of even degree N, see [13, Thm. 2.7]. In this case $\gcd(p-1, N) = 2$ because $N \mid \#E(\mathbf{F}_p) = p+1$. Let us recall that if $p \equiv 7 \bmod 8$ then supersingular elliptic curves over $\mathbf{F}_p$ come in two kinds: curves on the surface of their 2-isogeny volcano, and curves on the floor. The surface is characterized by the existence of three $\mathbf{F}_p$-rational points of order 2; more precisely, the group of $\mathbf{F}_p$-rational points is isomorphic to $\mathbf{Z}_2 \times \mathbf{Z}_{(p+1)/2}$. The points of order 2 can be classified as follows (see Figure 7.1):

- a point $P_{\rightarrow}$, whose halves are $\mathbf{F}_p$-rational,
- a point $P_{\leftarrow}$, whose halves are not $\mathbf{F}_p$-rational, but their x-coordinates are,
- a point $P_{\downarrow}$, the x-coordinates of whose halves are not $\mathbf{F}_p$-rational.

	Previous work [6]	This work	Cost per 2-isogeny
2-isogeny	-	$\mathbf{E} + \mathbf{M} + 3\mathbf{m} + 2\mathbf{A}$	1
3-isogeny	$\mathbf{E} + 6\mathbf{M} + 3\mathbf{A}$	$\mathbf{E} + 2\mathbf{M} + 3\mathbf{m} + 3\mathbf{A}$	1.023
4-isogeny	$\mathbf{E} + 4\mathbf{M} + 3\mathbf{A} + \mathbf{I}$	$\mathbf{E} + 3\mathbf{M} + \mathbf{m} + 3\mathbf{A} + \mathbf{I}$	1.008
5-isogeny	$\mathbf{E} + 7\mathbf{M} + 6\mathbf{A} + \mathbf{I}$	$\mathbf{E} + 6\mathbf{M} + \mathbf{m} + 6\mathbf{A} + \mathbf{I}$	1.034
6-isogeny	-	$\mathbf{E} + 9\mathbf{M} + 6\mathbf{m} + 9\mathbf{A} + \mathbf{I}$	1.090
7-isogeny	$\mathbf{E} + 24\mathbf{M} + 20\mathbf{A} + \mathbf{I}$	$\mathbf{E} + 12\mathbf{M} + 2\mathbf{m} + 9\mathbf{A} + \mathbf{I}$	1.043
8-isogeny	-	$\mathbf{E} + 11\mathbf{M} + \mathbf{m} + 9\mathbf{A} + 2\mathbf{I}$	1.151
9-isogeny	$\mathbf{E} + 69\mathbf{M} + 58\mathbf{A} + \mathbf{I}$	$\mathbf{E} + 17\mathbf{M} + 9\mathbf{A} + \mathbf{I}$	1.062
10-isogeny	-	$\mathbf{E} + 57\mathbf{M} + 5\mathbf{m} + 31\mathbf{A} + 3\mathbf{I}$	1.196
11-isogeny	$\mathbf{E} + 599\mathbf{M} + 610\mathbf{A} + \mathbf{I}$	$\mathbf{E} + 50\mathbf{M} + 21\mathbf{m} + 71\mathbf{A} + 2\mathbf{I}$	1.293
12-isogeny	-	$\mathbf{E} + 90\mathbf{M} + 8\mathbf{m} + 35\mathbf{A} + 3\mathbf{I}$	1.296
13-isogeny	$\mathbf{E} + 783\mathbf{M} + 776\mathbf{A} + \mathbf{I}$	$\mathbf{E} + 89\mathbf{M} + 33\mathbf{m} + 120\mathbf{A} + 2\mathbf{I}$	1.448
14-isogeny	-	$\mathbf{E} + 159\mathbf{M} + 16\mathbf{m} + 131\mathbf{A} + 4\mathbf{I}$	1.613
15-isogeny	-	$\mathbf{E} + 149\mathbf{M} + 32\mathbf{m} + 125\mathbf{A} + 2\mathbf{I}$	1.599
16-isogeny	-	$\mathbf{E} + 120\mathbf{M} + 4\mathbf{m} + 40\mathbf{A} + 3\mathbf{I}$	1.388
17-isogeny	-	$\mathbf{E} + 217\mathbf{M} + 55\mathbf{m} + 332\mathbf{A} + 3\mathbf{I}$	1.921
19-isogeny	-	$\mathbf{E} + 329\mathbf{M} + 125\mathbf{m} + 437\mathbf{A} + 3\mathbf{I}$	2.532

Table 7.1: The computational cost of radical N-isogenies for $N \in \{2, 3, \ldots, 17\} \cup \{19\}$ compared to previous work [6, Tbl. 3]. The letters $\mathbf{E}, \mathbf{M}, \mathbf{A}$ and $\mathbf{I}$ denote exponentiation, (full) multiplication (including squaring), addition and inversion respectively. The letter $\mathbf{m}$ denotes multiplication with a small constant. The last column expresses the cost of an N-isogeny relative to a 2-isogeny, based on the evaluation of a chain of 100 000 horizontal N-isogenies over $\mathbf{F}_p$, where p is the CRAD-513 prime from Section 7.7. Remark that the cost of $\mathbf{E}$ is approximately $(1.5 \log p)\mathbf{M}$ with the square-and-multiply algorithm. In particular, the last column would converge to 1 for larger values of p since the cost of a radical isogeny will be dominated by $\mathbf{E}$.

Each of these points spans the kernel of a 2-isogeny. The point $P_{\downarrow}$ takes us to the floor, while the other two isogenies are horizontal. It can be checked that the dual of an isogeny in the $P_{\rightarrow}$-direction is in the $P_{\leftarrow}$-direction, and vice versa. Therefore, non-backtracking chains of horizontal 2-isogenies necessarily happen on the surface and consistently walk in either of these two directions.

7.6.1 Horizontal vs. non-horizontal N-isogenies

Fix $N \geq 2$ even and assume that $p \equiv -1 \bmod \operatorname{lcm}(2N, 8)$, so that every curve E on the surface satisfies

$$E(\mathbf{F}_p)[N] \cong \mathbf{Z}_2 \times \mathbf{Z}_N. \tag{7.9}$$

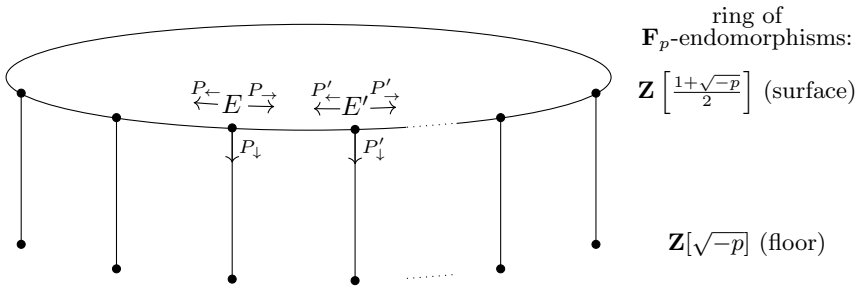

Figure 7.1: *Component of the 2-isogeny graph over $\mathbf{F}_p$ when $p \equiv 7 \bmod 8$. The top layer belongs to the surface; the bottom layer belongs to the floor; and $\sqrt{-p}$ is identified with the Frobenius endomorphism.*

Then $E(\mathbf{F}_p)$ has 2 or 3 cyclic subgroups of order N, depending on whether $t = \mathrm{ord}_2(N) > 1$ or $t = 1$ (see Lemma 7.6.1 below). Every corresponding isogeny $\varphi : E \to E/\langle P \rangle$ can be decomposed as $\varphi = \theta \circ \psi$, where ψ is the $N/2$-isogeny with kernel $\langle 2P \rangle$ and θ is the 2-isogeny with kernel $\langle \psi(P) \rangle$. The isogeny ψ is necessarily horizontal: indeed, if it would involve a vertical step, then composing with θ would necessarily involve backtracking, rendering φ non-cyclic. However, θ may take us to the floor.

Lemma 7.6.1 *Write $r = \mathrm{ord}_2(p+1) \geq \mathrm{ord}_2(2N) = t+1$.*

(i) If $t = 1$ then there are 3 options for $\langle P \rangle$, corresponding to θ being in the $P_{\rightarrow}$-direction, the $P_{\leftarrow}$-direction or the $P_{\downarrow}$-direction.

(ii) If $t \geq 2$ then there are 2 options for $\langle P \rangle$, corresponding to θ being in the $P_{\rightarrow}$-direction or the $P_{\downarrow}$-direction.

(iii) If $r \geq t+2$ (automatic if $t = 1$) then the group corresponding to θ being in the $P_{\rightarrow}$-direction can be characterized as follows: it is the unique group all of whose elements admit halves in $E(\mathbf{F}_p)$.

Proof. (i) Under the isomorphism (7.9), the cyclic subgroups of order N are generated by $(0,1)$, $(1,1)$ or $(1,2)$. Note that the group $\langle 2P \rangle$ does not depend on this choice, hence neither does ψ. Necessarily, the three groups must then correspond to the three stated options for θ.

(ii) If $t \geq 2$ then only the groups generated by $(0,1)$ or $(1,1)$ remain. Also note that we can further decompose $\psi = \theta' \circ \psi'$, where θ' is a 2-isogeny with kernel $\langle \psi'(2P) \rangle$. Since $\psi'(2P)$ is halvable over $\mathbf{F}_p$, this isogeny is necessarily in the $P_{\rightarrow}$-direction. But then θ cannot be in the $P_{\leftarrow}$-direction, otherwise φ would be non-cyclic.

(iii) If $r \geq t+2$ then $E(\mathbf{F}_p)[2N] \cong \mathbf{Z}_2 \times \mathbf{Z}_{2N}$ from which we see that the group generated by $(0,1)$ under the isomorphism (7.9) is uniquely characterized by its

elements being halvable over $\mathbf{F}_p$. But then $\psi(P)$ is also halvable over $\mathbf{F}_p$, from which the claim follows. □

The central question of Section 7.6 is: how do we avoid that $\ker\theta = \langle P_\downarrow \rangle$, within the framework of radical isogenies?

7.6.2 Square vs. non-square radicands

As explained in [6, §5.3], there is a simple algebraic criterion for determining whether quotienting out an order-N point $P \in E$ keeps us on the surface or takes us to the floor. Namely, we stay on the surface if and only if $\rho_N = f_{N,P}(-P)$ is a non-zero square in $\mathbf{F}_p$. In this case ρ_N admits two different Nth roots $\alpha \in \mathbf{F}_p$, which are each other's negatives. The challenge is to select the sign in such a way that the next radicand ρ'_N is again a square. Indeed, for this choice of Nth root the argument repeats and one keeps walking horizontally. Of course, one fallback is to make an arbitrary choice for α, at the cost of an exponentiation in $\mathbf{F}_q$ as before. One then computes the resulting ρ'_N and checks if it is a square. If it is not, then one switches to $-\alpha$.

It was observed in [4, Lem. 4] that for $N = 2$ the extra quadratic residuosity check can be avoided, because the correct choice of α admits an explicit description in terms of the "principal" square root of ρ_2, by which we mean the unique square root which is itself a square.

Remark 7.6.2 More generally, for any non-zero square $\rho \in \mathbf{F}_p$ we will refer to the unique Nth root of ρ that is a square as the principal Nth root. Note that when computing the Nth root through exponentiation, i.e., as $\rho^{(p+1)/2N}$, then it is automatically principal. ◊

Then, in more detail, the observation from [4, Lem. 4] was as follows: the radical isogeny iteration

$$E : y^2 = x^3 + Ax^2 + Bx \quad \to \quad E' : y^2 = x^3 + (A + 6\alpha)x^2 + 4\alpha(A + 2\alpha)x,$$

with $\alpha = \sqrt{B}$, repeatedly quotients out $(0, 0)$. If $(0, 0) \in E$ is the point $P_\rightarrow$, then $(0, 0) \in E'$ is the point $P'_\rightarrow$ if and only if α is the principal square root. This changes if $(0, 0) \in E$ is the point $P_\leftarrow$, in which case $(0, 0) \in E'$ is the point $P'_\leftarrow$ if and only if α is the non-principal square root.

This convenient fact was adapted to $N = 4$, first as a conjecture [6, Conj. 2] but recently this got proved by Onuki and Moriya [17, §5]. We will recall the precise statement of this adaptation in Section 7.6.4, where it will arise as an easy consequence to our generalization to arbitrary even N. But let us first highlight two takeaways that are already apparent from the case $N = 2$:

(i) When considering radical isogeny formulae for even N, then substituting $-\alpha$ for α produces formulae that are equally legitimate, e.g., because -1 is an Nth root of unity. Consequently, one cannot hope for a general rule saying that the $P_\rightarrow$-direction always corresponds to the principal Nth root.

(ii) Even worse, imagine that the rule does apply to some concrete choice of formulae, and now scale the radicand ρ_N with g^N for some arbitrary modular unit $g \in \mathbf{Q}_N(b,c)$, i.e. a function whose zeroes and poles are supported on the cuspidal part of $X_1(N)$; see [23]. The radical isogeny formulae transform into a version in which each occurrence of $\sqrt[N]{\rho_N}$ gets replaced by $\sqrt[N]{\rho_N}/g$. For these new formulae, the correct Nth root will depend on the Legendre symbol of the evaluation of g at the point $(E,P) \in X_1(N)$ under consideration.

7.6.3 Conjectural shape of ρ'_N modulo squares (proved for $N \leq 14$)

We ran into the following property of ρ'_N, which unfortunately we could not prove beyond $N = 14$, but which implies a generalization of the aforementioned observations for $N = 2, 4$ to arbitrary even N. Concretely, for every even $N \geq 4$ we can consider

$$\phi_{E,2}(x) = x^4 + b(1-c)x^2 - 2b^2x + b^3, \tag{7.10}$$

whose roots are the x-coordinates of the four halves of $P = (0,0)$ on $E = E_{b,c}$. Over $\mathbf{Q}_N(b,c)(\alpha^{N/2})$ this polynomial splits in two quadratic factors, with one quadratic factor corresponding to a pair of points

$$\frac{N}{2}Q,\ \frac{N}{2}Q + \frac{N}{2}P,$$

mapping to $\frac{N}{2}P'$ under φ. The discriminant of said quadratic factor is a modular unit of $X'_1(N)$ that we denote by Δ.

Example 7.6.3 Over $\mathbf{Q}_4(b,c)(\alpha^2)$ the polynomial (7.10) splits as $(x^2 - \alpha^2 x - \alpha^6)(x^2 + \alpha^2 x + \alpha^6)$. The roots of the first factor are the x-coordinates of two preimages of $2P'$. The discriminant of that factor is $\Delta = \alpha^4(1 + 4\alpha^2)$. ☆

Our conjecture is as follows:

Conjecture 7.6.4 *If the radicand $\rho_N = f_{N,P}(-P)$ was chosen, then one has*

$$\rho'_N \equiv \sigma\alpha b\Delta \tag{7.11}$$

modulo multiplication with a non-zero square in $\mathbf{Q}_N(b,c)(\alpha)$, for some $\sigma \in \{\pm 1\}$.

Here, we note:

- The sign σ should be viewed against our first takeaway message *(i)* above: substituting $-\alpha$ for α produces equally valid radical isogeny formulae but flips the sign.
- The congruence sign absorbs squares, so the conjecture is insensitive to replacing ρ'_N with any other representative of $t_N(P', -P')$, or even $t_N(P', \lambda P')$ for

whatever odd λ. However, as discussed in our second takeaway *(ii)* above, in the case of ρ_N the precise representative does matter. Interestingly, scaling with b^N would make the statement somewhat cleaner, as it would remove the mysterious factor b. This suggests that the radicand from Theorem 7.7.1 in the appendix is in fact a more natural choice than $f_{N,P}(-P)$.

It is exactly the presence of this factor b that made it difficult to guess how to go beyond the case $N = 4$; in the case $N = 4$ we have $b = -\alpha^4$ so that modulo squares this factor just appeared as a sign.

Theorem 7.6.5 *Conjecture 7.6.4 is true for $N \leq 14$.*

Proof sketch. From (7.7) and Example 7.6.3 we see that $\rho'_4 \equiv \alpha\Delta$ modulo squares, which matches with Conjecture 7.6.4 with $\sigma = -1$ because $-b = \alpha^4$ is a square. So the case $N = 4$ is immediate.

The case $N = 6$ is more illustrative. Take A', α, b as in (7.8) and let

$$\Delta = 4(1 - A)\alpha^3 + 3A^3 - 7A^2 + 4A$$

be the discriminant of the relevant quadratic factor of (7.10). One verifies, aided by the Magma command `IsPrincipal`, that for $\rho'_6 = f_{6,P'}(-P') = -A'^2(A' - 1)$ the function $-b\rho'_6/\alpha\Delta$ is a square in the function field of $X'_1(6) : \alpha^6 + A^2(A - 1) = 0$. So this again matches with Conjecture 7.6.4 (now with a minus sign).

In a similar way we have managed to deal with all even N up to 14, with further help coming from the observation that $\rho'_N = f_{N,-P'}(-P') \equiv f_{2,\frac{N}{2}P'}(P')$ modulo squares, see [3, Thm. IX.9(2)]. The right-hand side is a simpler function and therefore easier to handle by Magma. As an example, the Magma code for $N = 14$ can be found in the GitHub repository.[3] □

As mentioned, beyond $N = 14$ we were no longer able to verify Conjecture 7.6.4, although for $N = 16$ we gathered evidence by experimentally verifying Proposition 7.6.6 below for various concrete horizontal supersingular isogeny walks over finite prime fields.

7.6.4 Horizontal isogenies and principal Nth roots

Proposition 7.6.6 *Let $N \geq 4$ be even and consider radical isogeny formulae for computing chains of N-isogenies in terms of the radicand $\rho_N = f_{N,P}(-P)$. Assume that Conjecture 7.6.4 applies to these formulae and let $\sigma = \pm 1$ be the sign involved in its statement.*

Let $p \equiv -1 \bmod \operatorname{lcm}(2N, 8)$ and consider a supersingular elliptic curve $E/\mathbf{F}_p$ on the surface, along with a point $P \in E(\mathbf{F}_p)[N]$ such that the resulting isogeny $\varphi : E \to E' = E/\langle P\rangle$ is horizontal; let θ be the corresponding degree-2 component as in Section 7.6.1 and let $b, c \in \mathbf{F}_p$ be the corresponding Tate normal form coefficients. Let

[3] `https://github.com/KULeuven-COSIC/Horizontal_Radical_Isogenies`

$P' \in E'$ be the point produced by our radical isogeny formulae, where $\alpha = \sqrt[N]{\rho_N(b,c)}$ was computed as

$$\sigma \cdot s \cdot b^{(p-1)/2} \rho_N(b,c)^{(p+1)/2N}.$$

Here the sign s is determined as follows:

(i) if θ walks in the $P_\rightarrow$-direction and $r > t+1$ then $s = 1$,

(ii) if θ walks in the $P_\rightarrow$-direction and $r = t+1$ then $s = -1$,

(iii) if θ walks in the $P_\leftarrow$-direction (only possible if $t = 1$) then $s = -1$.

Then the isogeny $E' \to E'/\langle P' \rangle$ is horizontal.

Proof. Recall that the goal is to choose the instance of α that renders ρ'_N a square. Assuming Conjecture 7.6.4, this happens if and only if $\sigma\alpha b\Delta$ is a square.

In case *(i)* the point P is fully halvable over $\mathbf{F}_p$ thanks to Lemma 7.6.1(iii), so that Δ always evaluates to a square, regardless of the choice of α. So in order for ρ'_N to be a square, it is necessary and sufficient to choose α such that $\sigma\alpha b$ is a square: the claim follows.

If we are in cases *(ii)* or *(iii)* then none of the halves of P belong to $E(\mathbf{F}_p)$. Even stronger: none of these halves can have an $\mathbf{F}_p$-rational x-coordinate, because otherwise such a half H would satisfy $\pi_p(H) = -H$ and therefore $P = \pi_p(P) = -P$; a contradiction. This means that Δ is a non-square, regardless of the choice of α, and we can conclude as before. □

Example 7.6.7 For $N = 4$ we recover [6, Conj. 2], proved in [17]. Indeed, recall that $\sigma = -1$ and that b is always non-square in view of $\rho_4 = -b = \alpha^4$. Thus we have to compute $\alpha = s\rho_4^{(p+1)/8}$ with $s = -1$ if $p \equiv 7 \bmod 16$ and $s = 1$ if $p \equiv 15 \bmod 16$. ☆

We conclude by noting that $b^{(p-1)/2}\rho^{(p+1)/2N} = b^{-1}(b^N\rho_N)^{(p+1)/2N}$, effectively showing that the cost of root computation remains a single exponentiation.

7.7 Implementation

In this section we focus on N-isogenies between supersingular elliptic curves over prime fields $\mathbb{F}_p$ such that computing the required radical can be done deterministically by a single exponentiation. All tests were done in Magma v2.32-2 on an Intel(R) Xeon(R) CPU E5-2630 v2 @ 2.60GHz with 128 GB memory.

7.7.1 Isogeny chains

The main application of these radical isogeny formulae is that they can be used to efficiently compute a cyclic N^k-isogeny for small N and large k. This is similar to the work in [6], but we can now use larger N, have more efficient formulae for smaller N and are not restricted to odd N.

Remark that the radical 5-isogeny formulae from [6] were already optimized. Table 7.1 however shows a modest to strong speed up for radical N-isogenies for $N = 7, 9, 11, 13$. Over the field $\mathbb{F}_p$ with p the 513-bit CRAD-513 prime from Section 7.7.2, they provide a speed-up of respectively 4%, 13%, 55% and 57% compared to the work of [6].

The best known method to compute a chain of 17- or 19-isogenies so far was by sampling 17- or 19-torsion points and then applying Vélu-style formulae to compute the codomain. The cost of this is dominated by the computation of an appropriate torsion point. With the new radical formulae from Section 7.5, we only need to initialize the chain by computing such a torsion point once, and then can iteratively apply the radical isogeny formulae. Working over a prime field of roughly 512 bits, this results in an asymptotic speed-up of chaining 17-isogenies by a factor of 14, and a factor of 10 for chaining 19-isogenies. There is somewhat of a jump in complexity when going to optimized equations from $X_1(19)$ to $X_1(23)$ due to a jump in gonality. In particular, we do not expect radical 23-isogenies to be much of a speed-up over prime fields of characteristic roughly 512 bits,[4] so we did not try to optimize these. Nonetheless, for asymptotically large p the computational cost of a radical isogeny is expected to be dominated by a full exponentiation over $\mathbb{F}_p$.

For composite N, one can make a similar argument with regards to speed-up but the comparison is more subtle. For instance, the cost of computing a 15-isogeny is dominated by one exponentiation and 149 full multiplications according to Table 7.1. Alternatively, a 15-isogeny can also be computed by means of the concatenation of a 3- and 5-isogeny, the cost of which is dominated by two exponentiations and 8 full multiplications. Assuming we work over a prime field of cryptographic size - say at least 128 bits - the 15-isogeny will be the fastest method. However, assuming we have rational 9-torsion available, we have access to highly efficient radical 9-isogeny formulae, so asymptotically a 3-isogeny can be seen as half the cost of a 9-isogeny.

	512 bits	1024 bits	1536 bits
$2^{60,000}$-isogeny	23.38s	97.42s	264.59s
$4^{30,000}$-isogeny	11.93s	49.51s	133.12s
$8^{20,000}$-isogeny	8.77s	34.58s	91.33s
$16^{15,000}$-isogeny	7.92s	29.23s	75.01s
$3^{60,000}$-isogeny	23.39s	98.08s	266.31s
$9^{30,000}$-isogeny	12.77s	49.88s	134.61s

Table 7.2: Comparison in speed with regards to computing a chain of radical ℓ-isogenies over a prime field $\mathbb{F}_p$ for $\ell \in \{2, 3\}$ by means of different prime powers. The bit levels correspond to the size of p.

In general, composite N seem to yield more efficient formulae compared to prime

[4] Especially in the CSIDH setting from Section 7.7.2 where the initializing overhead is less negligible.

N as can be seen in Table 7.1. This stems from the fact that optimized equations for $X_1(N)$ typically have lower degree when N is composite, but also from the radical isogeny formulae themselves which appear to have parameterless integer coefficients (including zero) noticeably more often for composite N. These zero coefficients are even more frequently present in the radical isogeny formulae for prime-power N. In Table 7.2 one can see a comparison for computing low-degree prime-power chains of isogenies for three levels of prime bitsizes.

As can be seen, computing a chain of prime-power degree isogenies can be done more efficiently than a chain of prime degree isogenies for at least these values. The effect is more prominent for larger prime fields, since the exponentiation in those cases is more dominating in the overall cost of the radical isogeny formulae. We did not optimize the formulae for $N = 25$, since an optimal parametrization of $X_1(25)$ is already more complex than $X_1(19)$, and from Table 7.1 it is clear that computing chains of 5-isogenies would most likely be just as fast or faster (at least on the 512-bit level). Assuming the arithmetic for a radical ℓ^{k+1}-isogeny is always more complex than the arithmetic for a radical ℓ^k-isogeny, the asymptotic speed-up that can be gained from going to the next prime power is always bounded by $(k+1)/k$. For this reason, we expect that optimized radical 27- and 32-isogenies would be less efficient than radical 9- and 16-isogenies for all bitsizes in Table 7.2, though from a certain threshold onwards they would be the most efficient option again.

7.7.2 Impact on CSIDH

An application where chains of isogenies can be used is CSIDH [7]. We proceed just as in [6, §6], with the following differences:

- We make use of radical 17- and 19-isogenies.
- The optimzed formulae allow us to sample higher exponents of N-isogenies for $N = 7, 9, 11, 13$.
- We no longer use radical 4-isogenies, instead switching to radical 8-isogenies.

This last point may seem counterintuitive considering that chains of 16-isogenies are faster on the 512-bit prime level, as illustrated in Table 7.2. In CSIDH however, p is chosen such that $p+1$ is divisible by as many small primes as possible. If we want to make use of radical 16-isogenies, we would need to have that $32 \mid p+1$ (instead of $16 \mid p+1$ for radical 8-isogenies). This means that p would need to be roughly one bit larger, making all the other arithmetic more expensive. The trade-off in practice seems to be not worth it, considering the relative small gain from switching from chains of radical 8-isogenies to chains of radical 16-isogenies. The gap in efficiency between radical 4-isogenies and radical 8-isogenies does make a noticeable difference so we will use those. Nonetheless, we still need an extra factor of 2 that divides $p+1$ compared to the suggested prime in [6], so we choose CRAD-513 as the prime

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

The following sampling interval for the private key was determined heuristically, but can be considered (near) optimal:

$$\begin{aligned}&[-303;303] \times [-198;198] \times [-103;103] \times [-101;101] \times [-91;91] \\ &\times [-68;68] \times [-51;51] \times [-41;41] \times [-6;6]^{13} \times [-5;5]^{13} \\ &\times [-4;4]^{11} \times [-3;3]^{10} \times [-2;2]^{10} \times [-1;1]^{10}.\end{aligned}$$

Using these parameters, the class group action of the maximal private key can be computed 12% more efficiently than in the case of [6]. For an average private key, this speed-up will be roughly halved but from a constant-time implementation angle, the maximal private key is a more apt benchmark. This implementation in Magma is meant as a comparison to the work of [6], and can not be translated directly to other (constant-time) implementations such as CTIDH [1].

Appendix: an explicit radicand

The goal of the appendix is to prove the following result.

Theorem 7.7.1 *Let $N \in \mathbf{Z}_{>2}$. Let $K = \mathbf{Q}_N(b,c)$ as in Section 7.2.3. Let E/K be the elliptic curve given by $y^2 + (1-c)xy - by = x^3 - bx^2$. Let $P = (0,0) \in E$. Denote by Ψ_j the j-th division polynomial on E. Set $k = \lceil N/2 \rceil$. Then*

$$\left(\sum_{S \in E[N]} e_N(P,S) x(Q+S) \right)^N = N^{2N} \cdot \begin{cases} \frac{\Psi_k^2}{\Psi_{k-1}^2}(P) & \text{if } N \text{ is odd;} \\ \frac{\Psi_{k+1}}{\Psi_{k-1}}(P) & \text{if } N \text{ is even.} \end{cases}$$

Pairings and division polynomials

Let K be a field and let E/K be an elliptic curve. Suppose $P \in E(K)$ is of order N, such that $\operatorname{char} K \nmid N$. Let $Q \in E(\overline{K})$ satisfying $NQ = P$. Let $f \in K(E)$, $g \in \overline{K}(E)$ with respective divisors

$$\operatorname{div} f = N(P) - N(\mathcal{O}), \qquad \operatorname{div} g = \sum_{S \in E[N]} ((Q+S) - (S)).$$

Assume that g is such that $g^N = f \circ [N]$. Denote by $e_N : E[N] \times E[N] \to \mu_N$ the Weil pairing and by $t_N : E(K)[N] \times E(K)/NE(K) \to K^\times/(K^\times)^N$ the Tate pairing. For $\mathcal{P} \in E$, denote by $\tau_{\mathcal{P}} : E \to E$ the translation-by-$\mathcal{P}$ map. Let $\omega \in \Omega_E$ be an invariant differential and denote by $\operatorname{res}_{\mathcal{P}}(-) : \Omega_E \to \overline{K}$ the residue at $\mathcal{P}$ as defined in [25].

Lemma 7.7.2 *For every $Q \in E(K)$ we have*

$$t_N(P,Q) = \frac{\text{“Leading coefficient of } f \text{ at } Q\text{”}}{\text{“Leading coefficient of } f \text{ at } \mathcal{O}\text{”}} \in K^\times/(K^\times)^N.$$

Remark 7.7.3 Note that the leading coefficient of f (meaning the leading coefficient of the expansion of f with respect to a uniformizer) is everywhere well defined up to Nth powers, since the order of vanishing of f is at every point divisible by N (hence a different choice of uniformizer scales the leading coefficient by an Nth power). Also, the quotient in Lemma 7.7.2 is invariant under scaling f by an element of K, hence well-defines an element of $K^\times/(K^\times)^N$ given only the divisor of f. ◊

Proof. If $P = \mathcal{O}$ or $Q = \mathcal{O}$ then both sides are equal to 1, so assume $P \neq \mathcal{O} \neq Q$. We distinguish two cases.

Case $P = Q$. Let $h \in K(E)$ be any function such that $\operatorname{ord}_P(h) = -1$ and $\operatorname{ord}_{\mathcal{O}}(h) = 1$. Then $t_N(P,P) = f(\operatorname{div}(h) + (P) - (\mathcal{O}))$. By Weil reciprocity

$$\prod_R (-1)^{\operatorname{ord}_R(f)\operatorname{ord}_R(h)} \frac{f^{\operatorname{ord}_R(h)}}{g^{\operatorname{ord}_R(f)}}(R) = (-1)^{-2N} \frac{f^{-1}}{h^N} \frac{f^1}{h^{-N}}(P) \prod_{R \neq P,\mathcal{O}} f^{\operatorname{ord}_R(h)}(R).$$

equals 1. Hence

$$t_N(P,P) = \prod_{R \neq P,\mathcal{O}} f^{\mathrm{ord}_R(h)}(R) = \frac{h^N f(P)}{h^N f(\mathcal{O})} \in K^\times/(K^\times)^N.$$

Case $P \neq Q$. Let $h \in K(E)$ be any function such that $\mathrm{ord}_P(h) = 0$, $\mathrm{ord}_Q(h) = -1$, $\mathrm{ord}_\mathcal{O}(h) = 1$. Then $t_N(P,Q) = f(\mathrm{div}(h) + (Q) - (\mathcal{O}))$. By Weil reciprocity

$$1 = \prod_R (-1)^{\mathrm{ord}_R(f)\,\mathrm{ord}_R(h)} \frac{f^{\mathrm{ord}_R(h)}}{g^{\mathrm{ord}_R(f)}}(R) = (-1)^{-N} \frac{f}{h^{-N}}(\mathcal{O}) \frac{\prod_{R \neq \mathcal{O}} f^{\mathrm{ord}_R(h)}(R)}{h^N(P)}$$

Hence $t_N(P,Q)$ can be rewritten as

$$f(Q) \prod_{R \neq \mathcal{O}} f^{\mathrm{ord}_R(h)}(R) = (-1)^N \frac{h^N(P)}{(h^N f)(\mathcal{O})} f(Q) = \frac{f(Q)}{(h^N f)(\mathcal{O})} \in K^\times/(K^\times)^N.$$

□

Lemma 7.7.4 *Let $R \in E[N]$ such that P, R generate $E[N]$. We have*

$$t_N(P,P) = \left(\sum_{i,j=0}^{N-1} e_N(P,R)^i x(Q + iR + jP) \right)^N \quad \text{in } K^\times/(K^\times)^N.$$

Proof. We rely on the residue theorem [25, Thm. 3], whose use was suggested to us by Alexander Lemmens. This theorem implies that $\sum_{\mathcal{P} \in E} \mathrm{res}_\mathcal{P}(xg^{-1}\omega) = 0$, therefore

$$\begin{aligned} -\mathrm{res}_\mathcal{O}(xg^{-1}\omega) &= \sum_{S \in E[N]} \mathrm{res}_{Q+S}(xg^{-1}\omega) \\ &= \sum_{S \in E[N]} x(Q+S) \frac{g}{g \circ \tau_S}(Q)\, \mathrm{res}_Q(g^{-1}\omega) \\ &= \mathrm{res}_Q(g^{-1}\omega) \sum_{S \in E[N]} e_N(P,S) x(Q+S). \end{aligned}$$

It follows that (the last equivalence is due to Lemma 7.7.2)

$$\begin{aligned} \left(\sum_{S \in E[N]} e_N(P,S) x(Q+S) \right)^N &= (-1)^N \frac{x^N (g^N \circ \tau_Q)}{g^N}(\mathcal{O}) \\ &= (-1)^N \frac{x^N}{x^N \circ [N]} \frac{(x^N \circ [N])(f \circ [N] \circ \tau_Q)}{f \circ [N]}(\mathcal{O}) \\ &= (-1)^N N^{2N} \frac{x^N (f \circ \tau_P)}{f}(\mathcal{O}) \end{aligned}$$

which equals $t_N(P,P)$ in $K^\times/(K^\times)^N$. □

Now let $K = \mathbf{Q}(b,c)$, where b and c are both transcendental over $\mathbf{Q}$, though possibly algebraically dependent. Let E/K be the elliptic curve given by $y^2 + (1-c)xy - by = x^3 - bx^2$ and set $P := (0,0) \in E$.

For $Q \in E(K)$, we denote by $h_{P,Q} \in K(E)^\times$ any function with divisor $(P)+(Q)-(P+Q)-(\mathcal{O})$. For $j \in \mathbf{Z}$, we define

$$L_j := \left(\left(\frac{x}{y}\right)^{\mathrm{ord}_{\mathcal{O}}(h_{P,jP}) - \mathrm{ord}_P(h_{P,jP})} \cdot \frac{h_{P,jP} \circ \tau_P}{h_{P,jP}} \right)(\mathcal{O}).$$

In other words, L_j is the leading coefficient at $\mathcal{O}$ of the Laurent expansion of the function $(h_{P,jP} \circ \tau_P)/h_{P,jP}$ with respect to the uniformizer x/y. Note that, whereas $h_{P,Q}$ is only well-defined up to scalar multiplication, we have that L_j is a well-defined element of $K^\times$.

Lemma 7.7.5 *We have*

$$L_j = \begin{cases} b & \text{if } jP = -2P \text{ or } jP = -P; \\ 1 & \text{if } jP = \mathcal{O}; \\ -b & \text{if } jP = P; \\ b \cdot \dfrac{y_{jP}}{x_{jP} \cdot x_{(j+1)P}} & \text{else.} \end{cases}$$

Proof. Using (note that $h_{P,Q}$ as given by the formula below indeed has the desired divisor)

$$h_{P,Q} = \begin{cases} x & \text{if } Q = -P; \\ 1 & \text{if } Q = \mathcal{O}; \\ \dfrac{y}{x - x_{2P}} & \text{if } Q = P; \\ \dfrac{y - (y_Q/x_Q)x}{x - x_{P+Q}} & \text{else,} \end{cases}$$

this is a straightforward check for $Q \in \{-2P, -P, \mathcal{O}, P\}$. If $Q \notin \{-2P, -P, \mathcal{O}, P\}$ then in particular $x_{P+Q} \neq 0$. Let $u = x/y$. Then $x \circ \tau_P = bu + O(u^2)$ and $y \circ \tau_P = O(u^2)$, while $x = u^{-2} + O(u^{-1})$ and $y = u^{-3} + O(u^{-2})$. Thus the leading term at $\mathcal{O}$ of $(h_{P,Q} \circ \tau_P)/h_{P,Q}$ becomes

$$\frac{-y_Q/x_Q \cdot b}{-x_{P+Q}} = b \cdot \frac{y_Q}{x_Q \cdot x_{Q+P}}$$

as claimed. □

In what follows, $N > 2$ will always denote an integer and $k = \lceil N/2 \rceil$. We will assume that b, c are such that P has order at least $k+1$. Let $f \in K(E)$ be any

function with divisor $N(P) - N(\mathcal{O}) + ((k-N)P) - (kP)$.

Lemma 7.7.6 *We have*

$$\left(x^N \cdot \frac{f \circ \tau_P}{f}\right)(\mathcal{O}) = \prod_{j=-\lfloor N/2 \rfloor}^{\lfloor (N-1)/2 \rfloor} L_j.$$

Proof. This follows by noting that

$$\left(\left(\frac{x}{y}\right)^{2N} \cdot x^N\right)(\mathcal{O}) = 1.$$

and that $f = \prod_{j=-\lfloor N/2 \rfloor}^{\lfloor (N-1)/2 \rfloor} h_{P,jP}$ has the desired divisor. □

Define

$$\rho_N := \begin{cases} \dfrac{\Psi_k^2}{\Psi_{k-1}^2}(P) & \text{if } N \text{ is odd;} \\ \\ \dfrac{\Psi_{k+1}}{\Psi_{k-1}}(P) & \text{if } N \text{ is even,} \end{cases} \qquad \text{and} \qquad \pi(N) := \prod_{j=-\lfloor N/2 \rfloor}^{\lfloor (N-1)/2 \rfloor} L_j.$$

Lemma 7.7.7 *For all $N \in \mathbf{Z}_{>2}$, we have $\pi(N) = (-1)^N \rho_N$.*

Proof. We use induction on N. One easily verifies the claim for $N = 3, 4, 5$. Suppose $N = 2k \geq 6$ is even. Then

$$\pi(N)/\pi(N-1) = b \cdot \frac{y_{-kP}}{x_{-kP} \cdot x_{(-k+1)P}}, \quad \text{and} \quad \pi(N+1)/\pi(N) = b \cdot \frac{y_{kP}}{x_{kP} \cdot x_{(k+1)P}},$$

whereas $-\rho_N/\rho_{N-1} = -(\Psi_{k+1}\Psi_{k-1}/\Psi_k^2)(P) = -\rho_{N+1}/\rho_N$. But the middle term $-(\Psi_{k+1}\Psi_{k-1}/\Psi_k^2)(P)$ can be rewritten as $x_{kP} = x_{-kP}$ (from the multiplication-by-k formula using division polynomials; e.g. [20, Ex. 3.7]), so we can conclude using Lemma 7.7.8. □

Lemma 7.7.8 *For all $k \in \mathbf{Z} \setminus \{-1, -2\}$, we have $x_{kP}^2 x_{(k+1)P} = b \cdot y_{kP}$.*

Proof. Using the coordinate-wise addition formula for Weierstrass elliptic curves (e.g. [20, III.2.3]), we find $x_{kP}^2 x_{(k+1)P} = y_{kP}^2 + (1-c)x_{kP}y_{kP} + bx_{kP}^2 - x_{kP}^3 = by_{kP}$.

□

Proof of Theorem 7.7.1. In the proof of Lemma 7.7.4, we already saw that the left hand side equals $(-1)^N N^{2N} \left(x^N \cdot \frac{f \circ \tau_P}{f}\right)(\mathcal{O})$. The desired result now follows by combining Lemmas 7.7.6 and 7.7.7. □

7.8 Bibliography

[1] Gustavo Banegas, Daniel J. Bernstein, Fabio Campos, Tung Chou, Tanja Lange, Michael Meyer, Benjamin Smith, and Jana Sotáková. CTIDH: faster constant-time CSIDH. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(4):351–387, 2021.

[2] Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren. CSI-FiSh: Efficient isogeny based signatures through class group computations. In *Asiacrypt (1)*, volume 11921 of *Lecture Notes in Computer Science*, pages 227–247. Springer, 2019. `https://ia.cr/2018/485`.

[3] Ian F. Blake, Gadiel Seroussi, and Nigel P. Smart. *Elliptic Curves in Cryptography*. Cambridge University Press, United Kingdom, 1999.

[4] Wouter Castryck and Thomas Decru. CSIDH on the surface. In Jintai Ding and Jean-Pierre Tillich, editors, *PQCrypto 2020*, volume 12100 of *Lecture Notes in Computer Science*, pages 111–129. Springer, 2020.

[5] Wouter Castryck and Thomas Decru. Multiradical isogenies. In *AGC^{2}T-18*, Contemp. Math. (to appear). American Mathematical Society, 2022. `https://eprint.iacr.org/2021/1133`.

[6] Wouter Castryck, Thomas Decru, and Frederik Vercauteren. Radical isogenies. In *Proceedings of Asiacrypt 2020 Part II*, volume 12492 of *Lecture Notes in Computer Science*, pages 493–519. Springer, 2020.

[7] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In *Asiacrypt 2018 Pt. 3*, volume 11274 of *Lecture Notes in Computer Science*, pages 395–427. Springer, 2018.

[8] Wouter Castryck, Jana Sotáková, and Frederik Vercauteren. Breaking the decisional diffie-hellman problem for class group actions using genus theory. In *Advances in Cryptology – CRYPTO 2020: 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17–21, 2020, Proceedings, Part II*, page 92–120, Berlin, Heidelberg, 2020. Springer-Verlag.

[9] Jesus-Javier Chi-Dominguez and Krijn Reijnders. Fully projective radical isogenies in constant-time. In *Topics in Cryptology – CT-RSA 2022: Cryptographers' Track at the RSA Conference 2022, Virtual Event, March 1–2, 2022, Proceedings*, page 73–95, Berlin, Heidelberg, 2022. Springer-Verlag.

[10] Jean-Marc Couveignes. Hard homogeneous spaces, 2006. Unpublished article, available at `https://eprint.iacr.org/2006/291`.

[11] Luca De Feo and Jeffrey Burdges. Delay encryption. In *Proceedings of Eurocrypt 2021 Part I*, volume 12696 of *Lecture Notes in Computer Science*, pages 302–326. Springer, 2021.

[12] Luca De Feo, Simon Masson, Christophe Petit, and Antonio Sanso. Verifiable delay functions from supersingular isogenies and pairings. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019*, pages 248–277, Cham, 2019. Springer International Publishing.

[13] Christina Delfs and Steven D. Galbraith. Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. *Designs, Codes and Cryptography*, 78(2):425–440, 2016. `https://arxiv.org/abs/1310.7789`.

[14] Robert Granger, Florian Hess, Roger Oyono, Nicolas Thériault, and Frederik Vercauteren. Ate pairing on hyperelliptic curves. In Moni Naor, editor, *Advances in Cryptology - EUROCRYPT 2007*, pages 430–447, Berlin, Heidelberg, 2007. Springer.

[15] Yi-Fu Lai, Steven D. Galbraith, and Cyprien Delpech de Saint Guilhem. Compact, efficient and UC-secure isogeny-based oblivious transfer. In *Proceedings of Eurocrypt 2021 Part I*, volume 12696 of *Lecture Notes in Computer Science*, pages 213–241. Springer, 2021.

[16] Michael Monagan and Roman Pearce. Rational simplification modulo a polynomial ideal. In *ISSAC '06*, pages 239–245. ACM, 2006.

[17] Hiroshi Onuki and Tomoki Moriya. Radical isogenies on montgomery curves. In *Proceedings of PKC 2022 Part I*, volume 13177 of *Lecture Notes in Computer Science*, pages 473–497. Springer, 2022.

[18] David E. Rohrlich. Modular curves, Hecke correspondence, and L-functions. In *Modular forms and Fermat's last theorem*, pages 41–100. Springer, 1997.

[19] Samir Siksek. Explicit arithmetic of modular curves. Summer school notes, available at `https://homepages.warwick.ac.uk/staff/S.Siksek/teaching/modcurves/lecturenotes.pdf`, 2019.

[20] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, second edition, 2009.

[21] The Stacks project authors. The Stacks project. Available at `https://stacks.math.columbia.edu`, 2021.

[22] Anton Stolbunov. Public-key encryption based on cycles of isogenous elliptic curves. Master's thesis, Saint-Petersburg State Polytechnical University, 2004. In Russian.

[23] Marco Streng. Generators of the group of modular units for $\Gamma^1(N)$ over the rationals. *Ann. H. Lebesgue*, 6:95–116, 2023.

[24] Andrew V. Sutherland. Constructing elliptic curves over finite fields with prescribed torsion. *Mathematics of Computation*, 81:1131–1147, 2012.

[25] John Tate. Residues of differentials on curves. *Ann. Sci. École Norm. Sup. (4)*, 1:149–159, 1968.

[26] Jacques Vélu. Isogénies entre courbes elliptiques. *Comptes Rendus de l'Académie des Sciences, Série I*, 273:238–241, 1971.