Improved algorithms for ascending isogeny volcanoes, and applications

Steven D. Galbraith[1], Valerie Gilchrist[2], Damien Robert[3]

[1] University of Auckland, Auckland, New Zealand
[2] Université Libre de Bruxelles, Brussels, Belgium
[3] Inria Bordeaux, Institut de Mathématiques de Bordeaux, France

Abstract. Given two elliptic curves over $\mathbb{F}_q$, computing an isogeny mapping one to the other is conjectured to be classically and quantumly hard. This problem plays an important role in the security of elliptic curve cryptography. In 2024, Galbraith applied recently developed techniques for isogenies to improve the state-of-the-art for this problem.
In this work, we focus on computing ascending isogenies with respect to an orientation. Our results apply to both ordinary and supersingular curves. We give a simplified framework for computing self-pairings, and show how they can be used to improve upon the approach from Galbraith to recover these ascending isogenies and eliminate a heuristic assumption from his work. We show that this new approach gives an improvement to the overall isogeny recovery when the curves have a small crater (super-polynomial in size). We also study how these self-pairings affect the security of the (PEARL)SCALLOP group action, gaining an improvement over the state-of-the-art for some very particular parameter choices. The current SCALLOP parameters remain unaffected.

1 Introduction

Tate's isogeny theorem [30] states that two elliptic curves defined over $\mathbb{F}_q$ have the same number of points if and only if they are isogenous. To date, it is thought to be a hard problem to recover an isogeny between two fixed elliptic curves. The case where these two curves are supersingular greatly concerns modern-day isogeny-based cryptography, however, the case where they're ordinary still remains pertinent to elliptic curve cryptography and pairing cryptography. In 1999, Galbraith [13] gave algorithms solving the ordinary case, and then improved upon the worst case complexities in [14] using Kani's lemma.

In the later work, given elliptic curves on the floor of an isogeny volcano, Galbraith solves the problem by first computing paths up to the crater from

* Authors listed in alphabetical order: see https://www.ams.org/profession/leaders/CultureStatement04.pdf. Valerie Gilchrist is supported by a FRIA grant by the National Fund for Scientific Research (F.N.R.S.) of Belgium. Damien Robert is supported by the France 2030 program under grant agreement ANR-22-PETQ-0008 PQ-TLS.
Date of this document: 2025-08-15.

each curve, and then solving a meet-in-the-middle search on the crater to complete the path. In this work we focus on how to compute the paths up to the crater, these paths are called *vertical* or *ascending* isogenies. In Section 3.2 we give a new method to determine an ascending isogeny when the endomorphism ring is known. In Section 4, we make use of the work from Castryck et al [5], extended by Macula and Stange in [20], that shows how *self-pairings* can be used to recover unknown *horizontal* isogenies. These are pairings that can be evaluated using only one elliptic curve point and give a non-trivial output, i.e. $e(P, P) \neq 1$. We give a simplified framework for these pairings, and show how they can be used to recover vertical isogenies as well. In Section 4.2 we give a simplified proof of a special case of a result from [20] where we use only the standard Weil pairing rather than sesquilinear pairings. This proof is designed to be accessible to non-experts. We carefully characterize the cases when the order of the torsion subgroup does or does not divide the conductor, as well as the additional case when it divides the overall degree of the secret isogeny. This increases the available torsion subgroups on which we can gain *partial* information about how the isogeny acts. We then show how to encode the *missing* information into a conic equation, and treat the case of degenerate and non-degenerate conics individually.

With this new toolkit to compute unknown vertical isogenies, we consider the computational isogeny problem on certain curves with large volcanos (for example, pairing-friendly ordinary elliptic curves). In this case, the isogeny volcano is tall and the crater is very small (usually of size 1). The main result of [14] is to show that this case can be solved in $\tilde{O}(q^{1/4})$ field operations. In Section 5 we first show how to eliminate a heuristic assumption (about Elkies primes) from Galbraith's work using our pairing construction. We then detail an algorithm for solving the computational isogeny problem in volcanoes that are tall but with crater larger than polynomial size. Denote by Δ the discriminant of the maximal order and suppose $|\Delta| = q^a$ for some $0 < a < 1$. Write $t^2 - 4q = N^2\Delta$, where the curves in the isogeny class have $q + 1 - t$ points. The state-of-the-art algorithm from Galbraith [14] has complexity $\tilde{O}(h_0 N^{1/2})$ operations over $\mathbb{F}_q$ when the class number, h_0, is small enough, which works out as $\tilde{O}(q^{(1+a)/4})$ operations when a is small. Our new approach using self-pairings gives an improved complexity of $\tilde{O}(q^{(1-a)/4})$ operations, for most integers Δ.

In addition, in Section 2.4, we consider the case of volcanoes with small (but super-polynomial) conductor. The result is implicit in [13] but has not been explicitly stated anywhere. We show these two improvements compared to [14] in Figure 1.

Lastly, in Section 6, we demonstrate a second application of our approach to computing vertical isogenies. In this section our focus is on supersingular curves. Namely, we outline an attack on the SCALLOP group action for maximal orders with (almost) smooth discriminants. We show that if the fundamental discriminant Δ has a smooth factor of size at least $\Delta^{1/2+\epsilon}$, then our isogeny recovery approach from self-pairings potentially gives an improvement over the

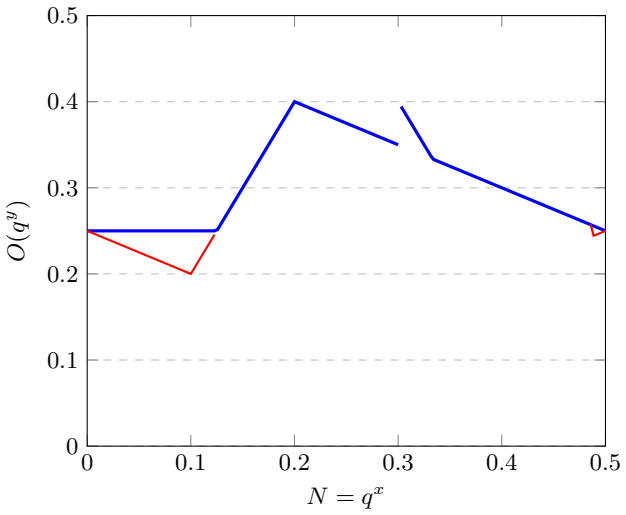

Fig. 1. The state-of-the-art from Galbraith is plotted in blue. The two improvements that will be presented in this paper are plotted in red.

state-of-the-art *classical* attacks. The exact complexity of our attack depends on the exact number and size of the prime factors of the smooth part of Δ.

Self pairings were introduced in [5] as a way to attack isogeny class group actions; they were then extended in [20]. However, they only apply when the degree d of the unknown isogeny $\phi : E_0 \to E_1$ is known. A key feature of our attacks in Section 5 and Section 6 is that we use the self-pairings only to ascend the volcano, and use this to compute an isogeny whose degree is not known. We can always recover the degree of the going up isogeny (and it is usually part of the SCALLOP parameter anyway).

2 Isogeny graphs and previous results on ascending isogenies

We study elliptic curves over a finite field $\mathbb{F}_q$. In some sections of the paper the focus is on ordinary curves and in some sections on supersingular curves.

In the ordinary case, the fundamental problem is to construct an isogeny $\phi : E_0 \to E_1$ between two given curves E_0, E_1 over $\mathbb{F}_q$. We are mainly interested in the case when ϕ is a vertical degree-N isogeny, which means one of the orders $\mathrm{End}(E_0), \mathrm{End}(E_1)$ has index N in the other.

We will sometimes use the notation that $\mathrm{End}(E_0)$ and $\mathrm{End}(E_1)$ have discriminants Δ_0, Δ_1 respectively, and write Δ for the discriminant of the imaginary quadratic field $K = \mathrm{End}(E_0) \otimes \mathbb{Q}$. Note that $\Delta_i = f_i^2 \Delta$ for $i = 0, 1$ and the integer f_i is called the conductor of the ring $\mathrm{End}(E_i)$. Sometimes we will abuse

notation and talk of the discriminant of a curve E when we mean the discriminant of $\mathrm{End}(E)$.

We write π_q for the q-power Frobenius map on E_0 and E_1. If R is a quadratic imaginary order, we will also often denote by $\Delta_R = \Delta(R)$ its discriminant, and likewise if $\gamma \in R$ we will denote by $\Delta_\gamma = \Delta(\gamma)$ the discriminant of its minimal polynomial (which is also the discriminant of $\mathbb{Z}[\gamma]$).

When we talk about supersingular curves we will always take them to be defined over $\mathbb{F}_{p^2}$ and we will assume that we are in the isogeny class of maximal curves, meaning that $\#E(\mathbb{F}_{p^2}) = (p+1)^2$.

2.1 Isogeny volcanos

Let E_0 be an elliptic curve over $\mathbb{F}_q$. Let S be a set of primes ℓ coprime to the characteristic of $\mathbb{F}_q$. The isogeny graph of E_0 with isogeny degrees in S is the graph whose vertex set is elliptic curves E over $\mathbb{F}_q$ that are isogenous to E_0, and there is an edge $\{E, E'\}$ if there exists an isogeny $\psi : E \to E'$ such that $\deg(\psi) \in S$. In the case $S = \{\ell\}$ we call the graph the ℓ-isogeny graph.

The connected components of the ℓ-isogeny graphs of ordinary elliptic curves have a specific structure. They constitute a cycle where each vertex of the cycle is the root of a tree. As can be seen in the example given in Figure 2, this structure resembles a volcano.

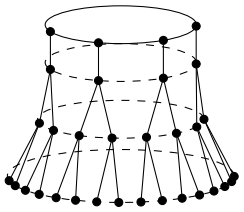

Fig. 2. Example 2-isogeny volcano of depth 3.

The vertices forming the top cycle of the volcano are often called the *crater*, or the *surface* of the volcano. The leaf vertices of the trees are called the *floor* of the volcano. A set of vertices at the same distance from the crater are called a level. The distance of a vertex on the floor from the cycle is called the *depth*. Vertices in the same level correspond to curves whose endomorphism ring has the same discriminant. Vertices on the crater correspond to curves whose discriminant is equal to the fundamental discriminant, Δ (or at least to a ring whose conductor is not divisible by ℓ). In an ℓ-isogeny volcano, a vertex in level k will correspond to a curve with discriminant $\Delta' = \ell^{2k}\Delta$.

In the ordinary case we will also consider volcanos with respect to a set S consisting of all primes dividing the conductor of the ring $\mathbb{Z}[\pi_q]$.

The Computational Isogeny Problem asks to recover an isogeny between two given elliptic curves, E_0, E_1. Taking a walk in the ℓ-isogeny volcano, where the

starting point is the vertex associated to E_0 and the ending point is the vertex associated to E_1, can be seen as a special case of this problem where the recovered isogeny is a power of ℓ.

2.2 The 1999 algorithm

In [13] Galbraith gave an algorithm that computes an isogeny between two ordinary elliptic curves. For two elliptic curves E_0, E_1, the approach used was to take a path starting at each of these curves that climbs up to the crater using modular polynomials. Once on the crater, a meet-in-the-middle computation determines how to join the two paths. If N is the largest prime divisor of the conductor and h is the class number of the maximal order then equation (5) of Section 6 of [13] states the complexity of the algorithm as

$$\begin{aligned} O(N^3(\log(N)+\log(q)) + \sqrt{h}(\log(h)^2 + \log(h)\log(q)^5) \\ + \log(h)\log(q)^6 + \log(q)^8/\log\log(q)) \end{aligned} \tag{1}$$

field operations. When N is constant or polynomial in $\log q$, the complexity is $\tilde{O}(q^{1/4})$. The worst case is when h is small and $N \approx q^{1/2}$, in which case the complexity is $\tilde{O}(q^{3/2})$ operations over $\mathbb{F}_q$.

2.3 The 2024 algorithm

Recently, Kani's lemma [18] was leveraged in some attacks [3, 21, 26] that were able to reconstruct a secret isogeny given some information about how the isogeny acts on a sufficiently large torsion group. Galbraith used this idea constructively in his recent paper [14] where he improves upon the worst case complexity of his previous work [13].

Instead of using modular polynomials to walk up the volcano from each of E_0, E_1 to the crater, Galbraith uses Kani's lemma to reconstruct the path. Suppose that N divides the conductor and that E_0 is directly above E_1 in the sense that there is a descending N-isogeny $\phi : E_0 \to E_1$ (see Figure 3). Then if we know ϕ on $E_0[M]$ for some suitable torsion group, the Kani construction can provide a representation of ϕ. Since we do not have any way to compute ϕ on $E_0[M]$, the method requires guessing how the isogeny acts on some torsion points. In other words, for several small primes $\ell \mid M$, given an ℓ-torsion basis (P_0, Q_0) on E_0, we need to guess $\phi(P_0), \phi(Q_0)$. It is explained in [14] that we need $M \approx N^{1/2}$.

To get the desired complexity we need the number of guesses for $\phi(P_0), \phi(Q_0)$ to be at most $O(M)$. Since there are M^2 choices for each M-torsion point, this seems to be a challenge. In [14] it is suggested to choose M to be a product of Elkies primes ℓ, namely primes such that $E_0[\ell]$ has a Frobenius eigenbasis with distinct eigenvalues, so $\pi_q(P_0) = uP_0$ and $\pi_q(Q_0) = qu^{-1}Q_0$ where π_q is the q-power Frobenius and $u \in (\mathbb{Z}/\ell\mathbb{Z})^*$. By choosing P_1 and Q_1 in the appropriate eigenspaces of $E_1[\ell]$, we know that $P_1 = [\lambda]\phi(P_0)$ and $Q_1 = [\nu]\phi(Q_0)$ for some integers λ, ν. Using the Weil pairing one can reduce to a single unknown integer,

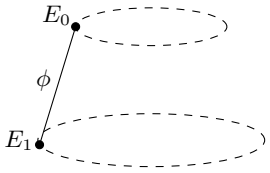

Fig. 3. Two curves at different levels of a volcano.

hence reducing the problem to trying $O(\ell)$ values. The drawback of this approach is that one needs to apply heuristic assumptions about the distribution of Elkies primes, see [14, Assumption 1] or [3, Section 10]).

For each guess, we run the Kani attack. If it returns a viable isogeny, then we can stop. Otherwise we repeat with a new guess. To run the Kani attack we have to choose our integer M so that $M - N$ is the sum of integer squares. One can write all positive integers as a sum of 4 squares, but the attack is faster if one can write $M - N$ as a sum of two squares.

2.4 Precise analysis in the case of small conductor

One of the observations in [14] is that it is sometimes more efficient to compute descending isogenies than ascending isogenies. Hence, [14] advocates descending to the floor, and solving the isogeny problem there using a meet-in-the-middle algorithm. This results in complexity $\tilde{O}(q^{1/4})$ operations for all conductors up to $q^{1/8}$. This is seen in Figure 1 with the flat blue line on the left hand side of the figure.

However, as already implied by equation (1), one can do better. We reproduce the analysis in this case.

First we recall the process for computing ascending and descending isogenies. In [13] ascending was done using modular polynomials and required cubic complexity. In [14] it is shown how to compute descending isogenies more efficiently (at least, asymptotically) when the group order is known by generating random kernel points. Given a curve over $\mathbb{F}_q$ with $q+1-t$ points and such that $\ell^4 \mid t^2-4q$, computing an ascending ℓ-isogeny may require $\tilde{O}(\ell^3)$ field operations, while computing a descending ℓ-isogeny takes $\tilde{O}(\ell^2)$ operations (see [14, Sect. 2.3] for details). It is possible, however, to compute the ascending isogeny in $\tilde{O}(\ell^2)$ operations as well. This will be detailed in Section 3.2. Now recall that we are trying to recover an isogeny, $\phi : E_0 \to E_1$ where $[\mathrm{End}(E_0) : \mathrm{End}(E_1)] = N$ and $\mathrm{End}(E_0) = \mathcal{O}_K$, i.e. E_0 lies on the crater. The approach from [13] is to take a walk up to the crater and then solve the problem there using meet-in-the-middle.

Suppose $N = q^b$ for a number $0 < b < 1/2$. As noted in [14, Section 2.3], walking up to the crater can always be done in $\tilde{O}(q^{2b})$ finite field operations (even with the $\tilde{O}(\ell^3)$ ascending algorithm from [14]). Call this ending curve E_2, so we have recovered the ascending isogeny $\varphi : E_1 \to E_2$ (see Figure 2.4 for reference).

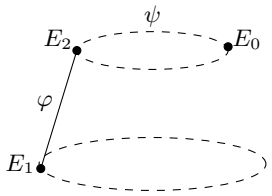

Fig. 4. Volcano where E_0 is not directly above E_1.

The class number of the crater is around $\sqrt{|\Delta_0|}$. Thus since $N^2\Delta_0 = t^2 - 4q$, we get that the class number will be $h_0 = \tilde{O}(\sqrt{q/N^2}) = \tilde{O}(q^{(1-2b)/2})$. Hence the meet-in-middle computation to recover the map $\psi : E_2 \to E_0$ on the crater will cost $\sqrt{h_0} = \tilde{O}(q^{(1-2b)/4})$.

The total complexity to recover $\hat{\varphi} \circ \hat{\psi} : E_0 \to E_1$ therefore comes from the cost to walk up to the crater and then conduct meet-in-the-middle which is $\tilde{O}(q^{2b}) + \tilde{O}(q^{(1-2b)/4})$ operations over $\mathbb{F}_q$. When $0 < b < 1/10$, we get that the meet-in-the-middle computation is dominating, giving a total complexity of $\tilde{O}(q^{(1-2b)/4})$. Otherwise, when $b \geq 1/10$, the total complexity will be $\tilde{O}(q^{2b})$ operations. The details are given in Algorithm 1.

Algorithm 1: Isogeny recovery in volcanoes with small conductor

Input : elliptic curves E_0, E_1 such that $\mathrm{End}(E_0) = \mathcal{O}_K$ and $[\mathrm{End}(E_0) : \mathrm{End}(E_1)] = N$
Output: $\phi : E_0 \to E_1$

1 Compute the unique ascending path $\varphi : E_1 \to E_2$ by sampling N-torsion points in E_1 and applying the formulas from Vélu ;
2 Compute an isogeny $\psi : E_2 \to E_1$ via meet-in-the-middle ;
3 Compute $\phi = \hat{\varphi} \circ \hat{\psi}$;
4 **return** ϕ ;

We gain an improvement over [13] when $0 < b < 1/8$. Suppose, for example, that $b = 1/14$. Then we get a complexity proportional to $q^{(1-2b)/4} \approx q^{0.21}$ operations over $\mathbb{F}_q$ which is better than $q^{0.25}$ operations. We outline this range of improvement in Figure 5.

3 Orientations

One can find a volcano structure in the supersingular isogeny graph over $\mathbb{F}_{p^2}$ when we endow the supersingular elliptic curves with an *orientation*. This terminology was introduced by Colo and Kohel [8].

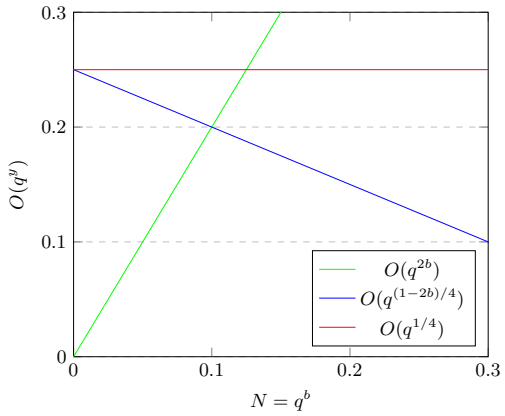

Fig. 5. The algorithmic complexity given in [14] is plotted in red. The complexity of the algorithm described in Section 2.4 is plotted in blue and green. Here we see that for $N < q^{1/8}$, the approach from Section 2.4 gives an improvement.

The terminology and notions in this section apply to both the ordinary and supersingular case. So let E be any elliptic curve over $\mathbb{F}_q$.

Definition 3.1 (K-orientation). *Let $K = \mathbb{Q}(\sqrt{d})$ be a quadratic imaginary field. We say an elliptic curve E is* K-oriented *if there exists a ring homomorphism*

$$\iota : K \hookrightarrow \mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}.$$

The pair (E, ι) is called a K-oriented elliptic curve.
If $R \subseteq K$ is a ring and $\iota(R) \subseteq \mathrm{End}(E)$ then we say E is R-oriented.

Suppose ι is the orientation on such a curve E. There exists a unique quadratic order $\mathcal{O} \subset K$ such that $\iota(\mathcal{O}) = \mathrm{End}(E) \cap \iota(K)$. Then we call ι a *primitive $\mathcal{O}$-orientation*, and we say that E is $\mathcal{O}$-orientable.

For any suborder $\mathcal{O}' \subset \mathcal{O}$, ι still induces an orientation on $\mathcal{O}'$, which is no longer primitive. Conversely, given an orientation ι on $\mathcal{O}'$, it extends uniquely to a K-orientation, and if $\mathcal{O}$ is the associated primitive orientation, we have $\mathcal{O}' \subset \mathcal{O}$. We call $\mathcal{O}$ the *saturation* of the orientation on E. We have $\Delta(\mathcal{O}') = \Delta(\mathcal{O})f^2$, and if m is an integer, we say that $\mathcal{O}'$ is *m-locally primitive* if m is coprime to f.

An oriented isogeny between two $\mathcal{O}$-oriented elliptic curves is an isogeny which commutes with the orientations. We will often drop the ι from our notation.

Example 3.1. The advantage of the above framework is that it applies just as well to ordinary curves.

If $E/\mathbb{F}_q$ is ordinary, there is a unique possible quadratic field K given by $K = \mathbb{Q}(\pi_q)$, and we will always use the natural orientation ι on K which sends π_q to the Frobenius endomorphism on E.

In particular, E is always oriented by the order $\mathbb{Z}[\pi_q]$, and the saturation of the orientation is simply given by $R = \mathrm{End}(E)$.

Finally, an oriented isogeny between ordinary elliptic curves with the natural Frobenius orientation is simply an $\mathbb{F}_q$-rational isogeny.

The class group $Cl(\mathcal{O})$ acts freely and transitively on the set of primitive $\mathcal{O}$-oriented curves up to isomorphisms (and Galois conjugacy in the special case where p is inert in $\mathcal{O}$, which can only happen in the supersingular case). A proof of this statement is given by Onuki in [23, Thm. 3.4]. This gives rise to a group action

$$\mathfrak{a} \star (E, \iota) = (E_\mathfrak{a}, \iota_\mathfrak{a})$$

where $\mathfrak{a}$ is an invertible $\mathcal{O}$-ideal coprime to the conductor of $\mathcal{O}$. To define it concretely, let $E[\mathfrak{a}] := \cap_{\alpha \in \mathfrak{a}} \ker \iota(\alpha)$, and denote by $\varphi_\mathfrak{a}^E$ the isogeny whose kernel is $E[\mathfrak{a}]$. Then

$$\varphi_\mathfrak{a}^E : E \to E_\mathfrak{a} = E/E[\mathfrak{a}] \text{ and } \iota_\mathfrak{a}(x) = \frac{1}{n(\mathfrak{a})} \varphi_\mathfrak{a}^E \circ \iota(x) \circ \hat{\varphi}_\mathfrak{a}^E.$$

If we also allow isogenies between $\mathcal{O}$-oriented curves, relaxing the primitive condition, then Colo and Kohel [8] show that there is a volcano structure just like the ordinary case: if E_1 is oriented by $\mathcal{O}$, and $\mathcal{O}_1$ is its saturation in E_1, and we have an order $\mathcal{O}_2$ containing $\mathcal{O}_1$ such that $\mathcal{O}_1$ is of conductor f in $\mathcal{O}_2$, then there is a unique "ascending" isogeny $\phi : E_1 \to E_2$ of degree f such that E_2 is primitively oriented by $\mathcal{O}_2$.

3.1 The module structure of the rational points of an oriented elliptic curve

In this section we let R be a quadratic imaginary ring, and we assume that we have an orientation $R \to \mathrm{End}(E)$ of an elliptic curve $E/\mathbb{F}_q$. Recall that the orientation is said to be primitive whenever R is saturated in $\mathrm{End}(E)$ (i.e., $R = (\mathrm{End}(E) \otimes_\mathbb{Z} \mathbb{Q}) \cap R$), or equivalently $\mathrm{End}(E)/R$ is torsion free. We call $E[n]$ a cyclic R-module if there is some $P \in E[n]$ such that $E[n] = RP = \{\phi(P) : \phi \in R\}$, in which case we call P and R-generator.

Theorem 3.1. *Let n be prime to the characteristic, and assume that the orientation R on E is primitive. Then $E[n]$ is a cyclic R-module.*

In particular, if $P \in E[n]$ is any R-generator, and $\gamma \in R$ is of discriminant $\Delta(\gamma) = f^2 \Delta_R$, then $\mathbb{Z}[\gamma]P$ is of cardinality $n^2/\gcd(n, f)$. So $E[n]$ is a cyclic $\mathbb{Z}[\gamma]$ module if and only if $\gcd(n, f) = 1$. It also follows that the Weil pairing $e_n(P, \gamma P)$ is a root of unity of exact order $n/\gcd(n, f)$.

Proof. The first statement is [20, Theorem 3]. For the second one, if $R = \mathbb{Z}[\omega]$, then we can write $\gamma = a + f\omega$. We have that $P, \omega P$ is a $\mathbb{Z}/n\mathbb{Z}$ basis of $E[n]$, and $\mathbb{Z}[\gamma]P$ is generated by P and $f\omega P$, and the latter is a point of order $n/\gcd(n, f)$. Finally $e_n(P, \gamma P) = e_n(P, \omega P)^f = \zeta^f$ is of order $n/\gcd(n, f)$ since $e_n(P, \omega P)$ is of exact order n. □

We can restate Theorem 3.1 as saying that $E[n]$ is R-cyclic if and only if R is n-locally primitive.

Corollary 3.1. *With the hypothesis of Theorem 3.1, $E[n]$ is isomorphic as an R-module to R/nR.*

Proof. By Theorem 3.1, $E[n]$ is a cyclic R-module, and if P is a generator, $E[n]$ is isomorphic to $R/\operatorname{ann}(P)$ where $\operatorname{ann}(P)$ is the annihilator of P. If $\gamma \in \operatorname{ann}(P)$, then since $\gamma(P) = 0$, R is commutative, and P generates $E[n]$ as an R-module, we find that $\gamma = 0$ on $E[n]$, so γ is divisible by $[n]$. The converse is obvious, so $\operatorname{ann}(P) = nR$, and $E[n] \simeq R/nR$. □

Corollary 3.2. *Let $n = \ell^e$ be prime to the characteristic, and assume that the orientation $R = \mathbb{Z}[\omega_R]$ on E is primitive. Let u be the degree of the field of definition of the geometric points of $E[\ell^e]$. Then $u \leq \ell^{e+1}$ in the ordinary case, and $u \leq \ell^e$ in the supersingular case.*

We can find an R-generator $P \in E[n]$, normalised such that $e_n(P, \omega_R P) = \zeta$, in time $\tilde{O}(\ell + u^2 \log^2 q + \log^{O(1)} q)$ (which can be improved to $\tilde{O}(\ell^{1/2} + u^2 \log^2 q + \log^{O(1)} q)$ in the supersingular case, or if ℓ is not inert in R).

For a general n, if R is n-locally primitive, we can find a basis (P_1, P_2) or an R-generator P of $E[n]$ in time $\tilde{O}(n^4 \log^2 q + \log^{O(1)} q)$ (resp. $\tilde{O}(n^2 \log^2 q + \log^{O(1)} q)$ in the supersingular case), by factorising n and applying the result above.

Proof. Assume for now that $E/\mathbb{F}_q$ is an ordinary curve (this is the harder case). We first explain how to find u. Recall that we can compute the discriminant of π_q acting on E in polynomial time in $\log q$ by point counting algorithms, and that we can use the endomorphism ring algorithm of [25] to also compute the ℓ-saturation R of $\mathbb{Z}[\pi_q]$ in E in polynomial time in $\log q$, so that R is ℓ-locally primitive in E. Now $E[\ell^e] \simeq R/\ell^e$ by Corollary 3.1, and u is the order of $\pi_q \in R/\ell^e$.

Finding u reduces to the computation of the order of elements in $\mathbb{F}_\ell$ if ℓ is ramified or split in R, and in $\mathbb{F}_{\ell^2}$ if ℓ is inert in R, hence costs $\tilde{O}(e \log \ell + \sqrt{\ell})$ or $\tilde{O}(e \log \ell + \ell)$ respectively. In the worst case $u \leq \ell^{e+1}$ (this can be improved to $u \leq \ell^e$ if ℓ is split in R).

Now $\pi_q^u - 1$ is divisible by ℓ^e in R. Write $\pi_q^u - 1 = \ell^e \gamma$, we have $N(\gamma) \leq q^u$. We can sample uniform points in $E[\ell^e]$ as follows. First we sample uniformly random points on $E(\mathbb{F}_{q^u})$ by computing square roots; this costs $\tilde{O}(u^2 \log^2 q)$ [24]. Then we apply γ to get uniformly random points on $E[\ell^e]$. This costs $\tilde{O}(u^2 \log^2 q + \log^{O(1)}(q))$. One then discards points that do not have the required order $n = \ell^e$. Finally, one takes pairs (P_1, P_2) and checks whether the Weil pairing $e_{\ell^e}(P_1, P_2)$ is of full order in μ_{ℓ^e}, in which case (P_1, P_2) is a basis. This requires generating $O(1)$ points and checking $O(1)$ conditions. Hence the overall cost is $\tilde{O}(\log(\ell^e) \log(q^u)) = \tilde{O}(eu \log \ell \log q)$ (because $\mu_{\ell^e} \subset \mathbb{F}_{q^u}$).

Once we have a basis, we know that either P_1 or P_2 is a generator of $E[\ell^e]$ as an R-module, and find it by evaluating ω_R followed by a Weil pairing. The total

cost, since $\ell^e < q^u$, is then $\tilde{O}(\ell + u^2 \log^2 q + \log^{O(1)} q)$ to find a basis (P_1, P_2) and an R-generator P of $E[\ell^e]$.

If $E/\mathbb{F}_{p^2}$ is a maximal supersingular curve, then $E(\mathbb{F}_{p^{2u}}) \simeq \mathbb{Z}/(p^u+1) \times \mathbb{Z}/(p^u+1)$, and so u is the order of p^2 modulo n. In particular, $u < n = \ell^e$.

We can also easily sample uniform points in $E[n]$ by sampling them in $E(\mathbb{F}_{q^u})$ and multiplying by the scalar cofactor (no need to use an endomorphism here). This costs $\tilde{O}(u^2 \log^2 q)$ operations. Then we use pairings to extract a basis (P_1, P_2), and then we apply ω_R and pairings again to find an R-generator R. The total cost is thus $\tilde{O}(\ell^{1/2} + u^2 \log^2 q)$ to find a basis (P_1, P_2) of $E[\ell^e]$, to which we need to add an endomorphism evaluation of $\tilde{O}(\log^{O(1)} q)$ to extract an R-generator P.

Finally we mention the case of general n. We simply handle each power of ℓ in turn and apply the Chinese remainder theorem to compute the point. The worst case is when $n = \ell$ is prime, in which case $u = O(n^2)$ and $u^2 = O(n^4)$, giving the stated complexity. □

Example 3.2. Let $E/\mathbb{F}_q$ be an ordinary curve. Specializing to the case $n = \ell$, we have three (well known) possibilities.

The first one is when ℓ splits in $\mathbb{Z}[\pi_q]$ (ℓ is Elkies), $\mathbb{Z}[\pi_q]/\ell\mathbb{Z}[\pi_q] \simeq \mathbb{F}_\ell^2$, π_q acts diagonally, and $u \mid \ell - 1$.

If ℓ is inert in $\mathbb{Z}[\pi_q]$ (ℓ is Atkin), $\mathbb{Z}[\pi_q]/\ell\mathbb{Z}[\pi_q] \simeq \mathbb{F}_{\ell^2}$, so $u \mid \ell^2 - 1$.

Finally, if ℓ is ramified in $\mathbb{Z}[\pi_q]$, then either π_q acts as a diagonal matrix $\begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix}$ (so $u \mid \ell - 1$) on $E[\ell]$, or as $\begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}$ (so $u \mid \ell(\ell-1)$). But the diagonal matrix case can happen if and only if $\pi_q - \lambda$ is divisible by ℓ in $\mathrm{End}(E)$, which is equivalent to $\mathbb{Z}[\pi_q]$ being of conductor divisible by ℓ in R. And the second case happens precisely when $\mathbb{Z}[\pi_q]$ is ℓ-locally primitive. Similar analysis holds for $n = \ell^e$, see for instance [22] for some results.

Remark 3.1. As a consequence of the proof above, we remark that in the ordinary case, if $\phi : E_0 \to E_1$ is an ascending oriented isogeny, so that $R_0 \subset R_1$, and u_i denotes the degree of the field of definition of the points in $E_i[n]$, then $u_1 < u_0$, because u_i is the order of π_q in R_i/nR_i.

Definition 3.2. *Let R be a quadratic order of discriminant Δ_R. We define the canonical (up to sign) imaginary element ω_R as follows. If $\Delta_R \equiv 1 \mod 4$, we let $\omega_R = \sqrt{\Delta_R}$; in that case $\mathbb{Z}[\omega_R]$ is of conductor 2 in R. Otherwise, $\Delta_R \equiv 0 \mod 4$, and we let $\omega_R = \sqrt{\Delta_R}/2$; which means $R = \mathbb{Z}[\omega_R]$.*

Corollary 3.3. *Let E be an elliptic curve oriented by R, $n \mid \Delta_R$, and assume that the orientation is n-locally primitive. We let ω_R be as in Definition 3.2.*

Then $E[n, \omega_R]$ is cyclic, and there exists a basis $(P, \omega_R(P))$ of $E[n]$.

Furthermore, if either n is odd, or $n \mid \Delta_R/4$, the matrix of ω_R on this basis is given by

$$\begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$$

In particular, for this n, $E[n, \omega_R] = \omega_R(E[n])$.

Proof. If R is a primitive orientation on E, then $E[\omega_R]$ is cyclic because ω_R is not divisible by any integer in R by construction, so $E[\omega_R]$ cannot contain any subgroups $E[\ell]$. It follows that $E[n, \omega_R]$ is still cyclic if R is just n-locally primitive.

If n is even, then Δ_R has to be divisible by 4, so $R = \mathbb{Z}[\omega_R]$. If n is odd, $\mathbb{Z}[\omega_R]$ might be of index 2 in R, but it is still n-locally primitive.

We deduce that there exists a basis $(P, \omega_R(P))$ of $E[n]$. Since ω_R^2 is either Δ_R or $\Delta_R/4$, under our assumptions $\omega_R^2 = 0 \mod n$, so we get the form of the matrix. □

3.2 The kernel of an ascending isogeny

This section gives new results that allow to compute an ascending isogeny in certain situations. These results extend previous work by Kohel [19] and Ionica and Joux [16].

Kohel [19] pointed out that for ordinary curves E over $\mathbb{F}_q$, if one is on the floor with respect to a prime ℓ, then there is only one $\mathbb{F}_q$-rational ℓ-isogeny from E and it is automatically ascending. This is linked to the fact that the ℓ-torsion is $\mathbb{Z}[\pi_q]$-cyclic for such a curve. Since the cyclic subgroup corresponding to the kernel of the isogeny is fixed by Frobenius, the kernel of the isogeny is $E[\ell, \pi_q - \lambda]$ for some integer λ. In other words, the ascending isogeny has kernel $E[\ell, \pi_q - \lambda]$, and this is a special case of our main result.

Ionica and Joux [16] extended this further, by showing how pairings can, in some situations, identify which subgroup of $E[\ell]$ leads to an ascending isogeny. Their results apply to ordinary curves, and require certain self-pairings to be non-trivial, which does not always hold.

Our result in this section is much more general. It shows that, if we are given a basis of $E[\ell]$, defined over some extension $\mathbb{F}_{q^d}$, we can find the kernel of the ascending isogeny in polynomial time in $\log \ell$ and $d \log q$. It applies in all ordinary settings. It also applies to the oriented supersingular case (provided we have an effective representation of the orientation). It is important to stress that the notion of "ascending" isogeny is not well-defined for supersingular curves in general. Instead, one must consider supersingular curves with an orientation, and it is the orientation that specifies whether or not an isogeny is ascending. This is exactly the information provided by the SCALLOP system.

Theorem 3.2. *Let $\phi : E_0 \to E_1$ be an ascending R-oriented cyclic isogeny of degree d. Let R_i be the saturation of R on E_i, let f be the conductor of R_0 in R_1. Then $f \mid d$, and ϕ factorizes uniquely as $\phi = \phi_2 \circ \phi_1$ where ϕ_1 is a purely ascending isogeny of degree f and ϕ_2 is an R_1-horizontal oriented isogeny.*

Furthermore, for any ω_1 such that $R_1 = \mathbb{Z}[\omega_1]$ (i.e. such that $\Delta_{\omega_1} = \Delta_{R_1}$), $R_0 = \mathbb{Z}[f\omega_1]$ and the kernel of ϕ_1 is given by $E_0[f\omega_1, f]$.

Proof. The unicity of the decomposition comes from the volcano structure.

The form of the kernel is an easy consequence of the general equivalence of categories described in [17, 27], which implies that the kernel of the ascending

isogeny is given by the action of the conductor ideal. We give a direct proof in a special case that is sufficient for our applications.

Let $\phi_1 : E_0 \to E_0'$ be our strictly ascending f-isogeny. Since ω_1 is an endomorphism on E_0', we have that $f\omega_1$ is 0 on $E_0'[f]$. Since $f\omega_1 \in R_0$ and ϕ_1 is R_0-oriented, if $P \in E_0$ is such that $\phi_1(P) \in E_0'[f]$, then $\phi(f\omega_1 P) = f\omega_1\phi(P) = 0$. In particular, $\ker\phi_1$ contains $(f\omega_1)E_0[f]$.

We also remark that $E_0[f, f\omega_1]$ does not depend on the choice of generator ω_1 for R_1: another generator will be of the form $\omega_1' = a \pm \omega_1$ with $a \in \mathbb{Z}$, so $E_0[f, f\omega_1] = E_0[f, f\omega_1']$. In fact, the argument shows that we can even take $\omega_1' = a + b\omega_1$, of discriminant $\Delta(\omega_1') = b^2\Delta_R$, as long as b is coprime to f.

So we can take ω_1 to be the canonical imaginary element of R_1 from Definition 3.2. Now for simplicity, we assume that f is odd or that Δ_0 is even. In that case, $f\omega_1$ is the canonical imaginary element of R_0, and by Corollary 3.3 there is a basis $(P, f\omega_1 P)$ of $E_0[f]$ and $(f\omega_1)^2$ is zero on $E_0[f]$. It follows that $(f\omega_1)E_0[f] = E_0[f, f\omega_1]$, and so $\ker\phi_1$ contains $E_0[f, f\omega_1]$. But ϕ_1 is of degree f, and the above basis shows that $E_0[f, f\omega_1]$ is too. □

Looking at the proof of the theorem above, we can give another description of a generator of the kernel of ϕ_1, more useful for algorithmic applications:

Corollary 3.4. *With the notations above, assume that R is f-locally primitive in E_0, and let ω_R be the canonical imaginary element from Definition 3.2.*

If either f is odd, or Δ_1 is even, then the kernel K of ϕ_1 is given by $E_0[\omega_R, f]$. If P is any generator of $E_0[f]$ as an R-module, then $\omega_R P$ is a generator of this kernel K.

We can compute the ascending isogeny $\phi_1 : E_0 \to E_0'$ in $\tilde{O}(u^2 \log f \log^2 q + uB_f^{1/2} \log f \log q + \log f \log^{O(1)} q) = \tilde{O}(f^2 \log^2 q + \log f \log^{O(1)} q)$ operations, where B_f is a bound on the largest prime divisor of f, and $u < f$ a bound on the degrees of the field extensions where the points of $(\ker\phi)[\ell_i]$ live, for $f = \prod \ell_i^{e_i}$.

Proof. If $\omega_0, \omega_1, \omega_R$ are the canonical imaginary elements of R_0, R_1, R respectively, then the proof of Theorem 3.2 shows that under our assumptions, we have $\omega_0 = f\omega_1$ and $\omega_R = u\omega_0$ for u coprime to f. From the R module structure of $E_0[f]$, we see that the kernel of ϕ_1, given by $E_0[f, \omega_R]$, is generated by $\omega_R P$ for any R-generator P of $E_0[f]$.

Now, to compute ϕ_1, we will work prime divisor by prime divisor of f (recall that we assumed that we know the factorisation of f_R, hence of f); there are at most $O(\log f)$ such primes. So in what follows we can assume that f is prime.

As in Corollary 3.2, in the ordinary case, we can determine the field extension where the points of the kernel of ϕ_1 are defined by determining the smallest u such that $\pi_q^u = 1$ in $R_0/(f\omega_1, f)$. Since the kernel is of order f, we have $u \mid f-1$. We also remark at this point that, as in Remark 3.1, since the saturation of R in E_0' will contain R_0, the degrees we will subsequently obtain on E_0' when we work prime by prime will be smaller than on E_0.

Now if $E_0[f](\mathbb{F}_{q^u})$ is cyclic, it is automatically equal to $\ker\phi_1$ since the kernel is rational. Sampling a point in $E_0(\mathbb{F}_{q^u})$ and multiplying by the (scalar) cofactor now gives a uniform point in the kernel.

The second case is that $E_0[f](\mathbb{F}_{q^u})$ is the full f-torsion. We can apply Corollary 3.2 to sample a generator P of $E_0[f]$ as an R-module. Then, in the situation of Corollary 3.4, we can apply ω_R to P to get a generator of $\ker \phi_1$.

Finally, computing ϕ_1 from a generator of the kernel can be done in $\tilde{O}(f^{1/2}u \log q)$ operations as shown in [2]. All in all, we can compute ϕ_1 in $\tilde{O}(u^2 \log f \log^2 q + uf^{1/2} \log q + \log^{O(1)} q)$ operations.

The supersingular case works as above, except that this time u is automatically also the degree of the field of definition of the full f-torsion. □

Remark 3.2. The only case not covered by Corollary 3.4 is when Δ_1 is odd and $f = 2f'$ is even. Then we can write ϕ_1 as an f_0-ascending isogeny $E_0 \to E_0''$ with R' the saturation of R in E_0'', of discriminant $4\Delta_1$, followed by a 2-isogeny $E_0'' \to E_0'$. We can apply Corollary 3.4 to the isogeny $E_0 \to E_0''$. And by Theorem 3.2, since $R' = \mathbb{Z}[\frac{1+\sqrt{\Delta_1}}{2}]$, the kernel of the 2-isogeny $E_0'' \to E_0'$ is given by $E_0''[1 + \sqrt{\Delta_1}, 2]$.

4 Recovering an oriented isogeny of known degree

In this section we study the following problem: let R be a quadratic imaginary order. Let $\phi : E_0 \to E_1$ be an R-oriented cyclic isogeny of known degree d between elliptic curves defined over $\mathbb{F}_q$. Our goal is, given R, E_0, E_1 and d, to recover ϕ. Following [14] we are going to do this by trying to determine ϕ on $E[n]$ for suitable n. In this paper the key tools are self-pairings.

We recall from Section 3 that this encompasses two cases. If $E_0/\mathbb{F}_q$ is ordinary, we can take $R = \mathbb{Z}[\pi_q]$ and we will always use the natural Frobenius orientation on E_0, E_1. This is the case we will consider for the (ordinary) computational isogeny problem. We remark that point counting algorithms give Δ_R in polynomial time in $\log q$. The other case is when E_0, E_1 are maximal supersingular elliptic curves defined over $\mathbb{F}_{p^2}$. In that case, ϕ is also automatically rational over $\mathbb{F}_{p^2}$. This is the case we will consider for the SCALLOP group action.

Let R_i be the saturation of R in $\mathrm{End}(E_i)$, and Δ_i the discriminant of R_i. We assume that ϕ is "ascending", meaning that $R_0 \subset R_1$, and we let $f = [R_1 : R_0]$ be the conductor of R_0 in R_1. We will say that ϕ is purely ascending if $f = d$. We remark that ϕ is automatically R_0-oriented: if $\gamma \in R_0$, there is some multiple $m\gamma$ in R, so since ϕ is R-oriented we have $[m]\gamma \circ \phi = \phi \circ [m]\gamma = [m]\phi \circ \gamma$. Dividing by $[m]$, we get $\gamma \circ \phi = \phi \circ \gamma$.

We will make the following assumptions for our complexity analysis:

- First, for simplicity, we will assume that $\log \Delta_R$ is polynomial in $\log q$. This is automatic in the ordinary case, since $|\Delta_{\pi_q}|$ is at most $4q$; and is also the case in all currently proposed instances of a supersingular oriented group action. All statements in this section are still valid in the general case if we replace the $O(\log q^{O(1)})$ in the complexity statements by $O((\log q + \log \Delta_R)^{O(1)})$.

– The orientation R is effective on both E_0 and E_1, meaning that we can evaluate any endomorphism $\gamma \in R$ on a point P using a polynomial in $\log N(\gamma)$ arithmetic operations over the field of definition of P.
We can refine this as follows: we fix $\omega_R = \sqrt{\Delta_R}/2$ if $\Delta_R \equiv 0 \mod 4$, and $\omega_R = (1 + \sqrt{\Delta_R})/2$ if $\Delta_R \equiv 1 \mod 4$. In both cases $N(\omega_R) \leq |\Delta_R|$, and if $\gamma = a + b\omega_R$, $a, b \leq \sqrt{2N(\gamma)}$. If we have an efficient representation of ω_R, then evaluating γ on P requires the evaluation of ω_R, scalar multiplication by $[a], [b]$ and a sum. This thus costs $O(\log(|a|) + \log(|b|) + \log^{O(1)} \Delta_R) = O(\log(N(\gamma)) + \log^{O(1)} \Delta_R)$ arithmetic operations over the field of definition of P.
As an aside, by adapting point counting algorithms, given an efficient representation of ω_R we can recover Δ_R in polynomial time. From now on, for our complexity analysis, whenever we pick up a generator ω_R of R, we will always assume that ω_R is a generator like the above, such that $N(\omega_R) = O(\Delta_R)$. This prevents silly situations like using $a + \omega_R$ for a very large a.
– As a consequence, the orientations by R_0 and R_1 are also effective. If γ_0 is an element of R_0, then some multiple n_0 of γ_0 is in R, hence can be efficiently evaluated (since $n_0 \mid \Delta_R$ cannot be too large). Now to evaluate γ_0 on a point P, we would like to find a point P' such that $P = [n_0]P'$; then $\gamma_0 P = (n_0\gamma_0)P'$. Such a point can be very expensive to compute for large, non-smooth n_0 since we would require a large extension field. Instead, Robert shows in [25, Sect. 2.3] how to efficiently divide an endomorphism by using higher dimensional isogenies.
If γ_0 is any generator of R_0, we can thus find an efficient representation of γ_0, and use this efficient representation to evaluate the other endomorphisms. Finding this efficient representation involves the evaluation of an endomorphism of R, but this is a precomputation that only needs to be done once.
Similar techniques also allow us to push forward the R-representation through an explicit R-isogeny $\phi' : E_0 \to E_0'$, by computing the image by ϕ' of $(P, \omega_R P)$ for several points P. This allows to find an efficient representation of the R-orientation on E_0' through a polynomial in $\log \Delta_R$ number of calls to the evaluation of ϕ'.
– Finally, we will assume that the factorisation of Δ_R is known. As a consequence, we can write $\Delta_R = \Delta f_R^2$ where Δ is a fundamental discriminant and the factorisation of f_R is known. We also have $\Delta_0 = \Delta f_0^2$ and $\Delta_1 = \Delta f_1^2$, with $f_0 = f f_1$. Then thanks to [25], we can recover f_0 and f_1, hence R_0 and R_1 in polynomial time in $\log \Delta_R + \log q$.

4.1 The general strategy

We will adapt the strategy of [14], which consists in guessing the image of ϕ on suitable subgroups of E_0, typically the image on the ℓ-torsion $E_0[\ell]$ for several small primes ℓ, and then reconstructing ϕ using a higher dimensional isogeny.

There are a number of different ways to reduce the number of guesses required to determine ϕ on $E_0[\ell]$ and we give a high level overview of them now.

- *The case of Elkies primes.* We can adapt the strategy of Section 2.3 to the general oriented case as follows. If $\ell \nmid \Delta_R$ splits as $\ell = \mathfrak{f}_1\mathfrak{f}_2$, then $E_0[\mathfrak{f}_i]$ is cyclic of order ℓ with generator P_i. Likewise, $E_1[\mathfrak{f}_i]$ is cyclic of order ℓ with generator Q_i. Since ϕ is R-oriented, $\phi(E_0[\mathfrak{f}_i]) \subset E_1[\mathfrak{f}_i]$, so we have $\phi(P_i) = a_iQ_i$, for some unknown scalars a_1, a_2 modulo ℓ. The Weil pairing gives us some information on a_1, a_2: we have $e_\ell(\phi(P_1), \phi(P_2)) = e_\ell(P_1, P_2)^d = e_\ell(Q_1, Q_2)^{a_1a_2}$. So, provided discrete logarithms are easy in the multiplicative group μ_ℓ of ℓ-th roots of unity in $\overline{\mathbb{F}}_q$, we can compute c such that $e_\ell(Q_1, Q_2) = e_\ell(P_1, P_2)^c$, and we know that $a_1a_2 = d/c$ modulo ℓ.
  In the special case that $\ell \mid d$, then either a_1 or a_2 is 0, and $E_0[\mathfrak{f}_1]$ or $E_0[\mathfrak{f}_2]$ is in the kernel of ϕ. So we recover part of the kernel of ϕ up to a choice. Otherwise, $\ell \nmid d$, and there are $\ell - 1$ possibilities for a_1, and then a_2 is completely determined by a_1.
- *Using self-pairings.* If $\ell \mid \Delta_R$, and R is a ℓ-locally primitive orientation on E_0, E_1, then we follow the insight of [5] that there exists a self-pairing that gives the image of ϕ up to a sign on a cyclic subgroup of $E_0[\ell]$. In [20, § 8], Macula and Stange give a more efficient construction of this self-pairing (see in particular [20, Example 3]), and we use a slight variant of their approach.
- *Using sesquilinear self-pairings.* If $\ell \nmid \Delta_R$, then $E[\ell]$ is a cyclic R-module, and Macula and Stange show in [20, Theorem 6] that there exists an R-sesquilinear self-pairing. While this gives less information than the previous case of self pairing, this still allows them to reduce the number of choices for the action of ϕ on $E_0[\ell]$ to $\ell - 1$ (if ℓ splits) or $\ell + 1$ (if ℓ is inert), see [20, Theorem 8 and 9]. In particular, this subsumes the previous special case of Elkies primes.
- The remaining case, at least if $R = R_0$ which we can assume from our hypotheses, is when $\ell \mid f$. In that case, R is not an ℓ-locally primitive orientation on E_1, and we cannot follow the pairing approach. Since we are speaking at the moment about small ℓ, this case is treated in [13, 14] as an easy case, because one can ascend in the volcano efficiently using the methods in those papers. We have shown in Section 3.2 an improved method to compute $(\ker \phi)[\ell]$ directly.

In summary, we will show that the self-pairings approaches of [5, 20], which were mainly used in the horizontal case, work just as well for the ascending case. The main difference compared to [20, Theorem 9] (which treats the case $\ell \nmid \Delta_R$) and [20, Theorem 11] (which treats the case $\ell \mid \Delta_R$) is that we also treat the case when ℓ divides d. This gives more flexibility in our choice of ℓ.

We also give in Section 4.2 a simplified and unified construction of "self-pairings", which only requires the standard Weil pairing (but is heavily inspired by the sesquilinear pairings used in [20]). By contrast, in [20] two different pairing constructions are given, depending on whether $\ell \mid \Delta_R$ or not. Our approach is similar to the use of the distortion map in pairing based cryptography. We refer to the end of Section 4.2 for a more detailed comparison of our approaches with the approach of [5, 20]. We also remark that the use of the distortion map was also explored in [12] as a way to determine whether an ℓ-isogeny was descending.

Finally, since we will apply self-pairings to large-ish ℓ, we will present some precise complexity statements that will be needed for the analysis in Section 5 and Section 6.

4.2 Pairings and distortion maps

We can reduce the computation of ϕ to the computation of the purely ascending isogeny ϕ_1 and an horizontal isogeny ϕ_2. We might not want to compute the purely ascending isogeny ϕ_1 fully, especially if the conductor f has large prime factors.

Section 3.2 can be seen as recovering partial information on the action of ϕ on $E[n]$ when $n \mid f$. But in this section, we focus on the case n coprime to f. More precisely, we explain how we can combine pairings with the orientation R to recover partial information about the image of ϕ on a basis of $E[n]$.

Recall that a pairing is a bilinear and non-degenerate map. Let E be an elliptic curve over $\mathbb{F}_q$ and m an integer co-prime to q. The most well-known and familiar pairing in elliptic curve cryptography is the Weil pairing

$$e_m : E[m] \times E[m] \to \mu_m \subseteq \bar{\mathbb{F}}_q^*$$

where μ_m is the multiplicative group of order m.

An important property of pairings, which has been widely used in isogeny cryptography and cryptanalysis, is that if $\phi : E_0 \to E_1$ is an isogeny of elliptic curves of degree d and $P, Q \in E_0[m]$ then

$$e_m(\phi(P), \phi(Q)) = e_m(P, Q)^d$$

where the first pairing is being computed on E_1 and the second pairing on E_0. Sometimes this property is used to learn information about the degree of an unknown isogeny. Other times, and this is how we use it, knowing d, one can use the above property to constrain the possible values for $\phi(P)$. These approaches have been used in a number of papers, including [5, 14, 20].

As mentioned in the introduction, an important set of techniques about self-pairings was given in [5], however in some contexts (including in our application) this requires a large field extension to be computed. Macula and Stange [20] showed a different approach that enables to obtain the same results without enlarging the fields over which one is working. Of particular importance for us will be [20, Theorem 8], which treats the case when n is coprime to Δ_R (this is a hidden hypothesis which comes from the invocation of Theorem 6 and Theorem 11 of [20] which treats the case when $n \mid \Delta_R$). Since [20] treats sesquilinear pairings and other advanced topics, the proof of this theorem may not be easily accessible to some researchers. Hence we take the opportunity to give an elementary proof of the result in a special case, and also note that the condition of supersingularity is not required.

Here is our version of "self-pairings".

Theorem 4.1. *Let $\phi : E_0 \to E_1$ be an R-oriented cyclic isogeny of degree d, let $n = \ell^e$ be an integer coprime to d, and assume that R is n-locally primitive on E_0 and E_1. Assume that we are also given $P_0 \in E_0[n], P_1 \in E_1[n]$ cyclic R-generators, and write $\phi(P_0) = \gamma(P_1)$ for some unknown endomorphism $\gamma \in R/nR$.*

Then we can compute $v = \deg(\gamma) \pmod n$ in time $\tilde{O}(ue \log \ell \log q + u'e\sqrt{\ell} \log q)$, where u is the degree of the field of definition of $E_0[\ell^e]$, and $u' \mid u$ the degree of the field of definition of μ_{ℓ^e}.

Proof. Take ω_R a generator of R, of norm bounded by $|\Delta_R|$ so that we can evaluate it efficiently. Since ϕ is oriented, it commutes with ω_R. Furthermore, since $RP_0 = E_0[n]$ it follows that $\{P_0, \omega_R(P_0)\}$ generates $E_0[n]$ and so $e_n(P_0, \omega_R(P_0))$ is an exact n-th root of unity. Similarly, $e_n(P_1, \omega_R(P_1))$ is an exact n-th root of unity. (The endomorphism ω_R here is performing the role of a "distortion map".) Suppose $\phi(P_0) = \gamma(P_1)$.

By compatibility of the Weil pairing with isogenies, we have

$$\begin{aligned} e_n(P_1, \omega_R(P_1))^{\deg(\gamma)} &= e_n(\gamma(P_1), \gamma\omega_R(P_1)) \\ &= e_n(\phi(P_0), \omega_R(\phi(P_0))) \\ &= e_n(P_0, \omega_R(P_0))^{\deg(\phi)}. \end{aligned}$$

Thus by computing these two pairings in $\tilde{O}(ue \log \ell \log q)$ and solving the discrete logarithm problem in $\tilde{O}(u'e\sqrt{\ell} \log q)$, one can compute $\deg(\gamma)$ modulo the order of $e_n(P_0, \omega_R(P_0))$, which is n. □

By combining Corollary 3.2 and Theorem 4.1, we see that sampling P_0, P_1 dominates the complexity:

Corollary 4.1. *With the notations above, we can sample cyclic R-generators $P_0 \in E_0[n], P_1 \in E_1[n]$, and compute $\deg \gamma \pmod n$ where $\gamma(P_1) = \phi(P_0)$ in $\tilde{O}(\ell + u^2 \log^2 q + \log^{O(1)} q)$.*

Let $\gamma = a + b\omega_R \in R \subseteq \mathrm{End}(E_1)$ be such that $\phi(P_0) = \gamma(P_1)$. Then $\deg(\gamma) = N(\gamma) = a^2 + ab \operatorname{Tr}(\omega_R) + b^2 N(\omega_R)$.

From now on, we will suppose that n is coprime to d, since the non-coprime case is easier and will be treated in Section 4.3, so that $v = \deg(\gamma) \neq 0 \pmod n$.

For our applications it will not be enough to know $\deg(\gamma)$. We will need to know the integers a and b, because we want to actually compute $\phi(P_0)$ by computing $\gamma(P_1)$. (From this we will also be able to compute $\phi(Q_0)$ where $\{P_0, Q_0\}$ is a basis for $E_0[n]$, by taking $Q_0 = \omega_R(P_0)$ and $\phi(Q_0) = \omega_R(\phi(P_0))$.) To find these integers we need to solve the conic equation

$$x^2 + xy \operatorname{Tr}(\omega_R) + y^2 N(\omega_R) = v \pmod n. \tag{2}$$

In practice it is usually simpler to separately solve the conic modulo each prime dividing n, and then use Hensel lifting and the Chinese remainder theorem to compute the set of solutions. We remark that this conic is of discriminant $\Delta(\omega_R) = \Delta_R$.

There are two cases of relevance in our paper. The first is when n is coprime to Δ_R and the conic is non-singular. In this case there are $O(n)$ solutions. To find solutions one calculates a rational parameterization of the conic and hence one can easily enumerate all solutions. The second case is when $n \mid \Delta_R$ and the conic degenerates, typically as the union of lines. In this case, as long as n does not have too many distinct prime factors, we can obtain a lot of useful information. We explain in more detail below.

The case of a non-degenerate conic We suppose here for simplicity that $n = \ell$ is a prime, and that $\ell \nmid \Delta_R$.

By assumption, we know there exists an endomorphism γ, hence a solution of Eq. (2) for $x, y \in \mathbb{Z}/\ell\mathbb{Z}$. Thus, the homogenised conic $X^2 + XY\,\mathrm{Tr}(\omega_R) + Y^2 N(\omega_R) - vZ^2$ is isomorphic to $\mathbb{P}^1$ over $\mathbb{Z}/\ell\mathbb{Z}$.

At infinity, $Z = 0$, we have 2 or 0 rational solutions $(X : Y : 0)$ depending on whether Δ_R is a square or not modulo ℓ. We deduce that if ℓ splits in R, the conic has $\ell - 1$ solutions (x, y), and if ℓ is inert in R, the conic has $\ell + 1$ solutions (x, y).

The case of a degenerate conic In this subsection, we assume that $n \mid \Delta_R$.

Theorem 4.2. *Let ω_R be the canonical imaginary element of R as defined in Definition 3.2. Let $n \mid \Delta_R$ be coprime to d, and such that $n \mid \Delta_R/4$ if n is even. Let $P_0 \in E_0[n]$ be a cyclic generator, and same for P_1. Let $Q_0 = \omega_R(P_0)$ and $Q_1 = \omega_R(P_1)$. Then $\phi(Q_0) = \alpha Q_1$, where $\alpha^2 = v \mod n$, where v is computed using pairings as in Theorem 4.1.*

Proof. As in Corollary 3.3, if n is even then $R = \mathbb{Z}[\omega_R]$, otherwise n is odd and $\mathbb{Z}[\omega_R]$ may be of index 2. In both cases, $\mathbb{Z}[\omega_R]/n\mathbb{Z}[\omega_R] = R/nR$, so we may assume that $\gamma = x + y\omega_R$ is such that $\phi(P_0) = \gamma(P_1)$. Now since ω_R is imaginary, the norm is $N(\gamma) = x^2 + \Delta_{\omega_R} y^2$. And by hypothesis, $n \mid \Delta_{\omega_R}$. So our conic degenerates to $x^2 = v \mod n$. The solutions are given by (α, y) for any α such that $\alpha^2 = v$.

We now recall that $\omega_R^2 = 0$ modulo n by Corollary 3.3. So if $\gamma = \alpha + y\omega_R$, then $\phi Q_0 = \phi\omega_R P_0 = \omega_R \phi P_0 = \omega_R(\gamma P_1) = \omega_R(\alpha P_1 + y\omega_R P_1) = \alpha\omega_R P_1 = \alpha Q_1$. In particular, the image of ϕ on Q_0 only depends on the possible solutions α. □

We remark that the equation $\alpha^2 = v$ has at most 2^{m+1} solutions, where m is the number of distinct prime factors of n. More precisely, since there exists at least one solution by hypothesis, there are exactly 2^{m-1}, 2^m or 2^{m+1} solutions according to whether $v_2(n) = 1$, $v_2(n) = 0, 2$, or $v_2(n) > 2$ respectively.

This means that while the image of Q_0 is fairly constrained (if n has not too many prime factors), the conic itself actually has $\approx 2^m n$ solutions, rather than just $O(n)$ (because it is not irreducible). We will call this version of self-pairing the cyclic version, because it gives information on a cyclic subgroup of $E_0[n]$.

Comparison with self-pairings As mentioned in [20], the sesquilinear pairings used in that article can be seen as a neat way to package together the pairing data $e_n(P,Q), e_n(\omega_R P, Q), e_n(P, \omega_R Q), e_n(\omega_R P, \omega_R Q)$ together into an R-sesquilinear Weil pairing. (In [20] Macula and Stange look at the R-sesquilinear Tate pairing rather than the sesquilinear Weil pairing, but the sesquilinear Weil pairing works similarly and was constructed in [29].)

Unfortunately, this sesqulinear pairing becomes degenerate when $n \mid \Delta_R$, so to handle this situation, in [20] the authors look at $(P,Q) \mapsto (t_n(\omega_R P, Q), t_n(P,Q))$ where t_n is the Tate pairing. Theorem 4.1 is almost the same, except that using the Weil pairing e_n instead of the Tate pairing allows us to only use $e_n(P, \omega_R Q)$, because the Weil pairing is alternating, and to use the same construction both for $n \mid \Delta_R$ and n coprime to Δ_R. We remark that the same phenomena was present in [6], where the authors remarked that it is often more convenient to use the Weil pairing than the Tate pairing to evaluate the "character" associated to an horizontal isogeny. This is not surprising: as explained in [5, § 6.2], this character evaluation can be done through self-pairings computation.

Finally, still for the case $n \mid \Delta_R$, the authors of [5] construct a self-pairing on the cyclic subgroup $E_0[n, \omega_R]$, using an "oriented" version of the Tate pairing [5, § 5.1]. By the discussion of [5, § 5.3], we see that a way to compute these self-pairings is via the distorted pairing $e_n(P, \omega_R Q)$ from Theorem 4.1. We refer to Section A for a generalisation of this result.

4.3 Putting everything together

In this section, we summarise the torsion information we can guess on ϕ, and how to use it to reconstruct ϕ. Recall that we assume $d = \deg(\phi)$ is known and is large, and we are choosing suitable $n = \ell^e$, where ℓ is a small prime, and attempting to deduce candidates for ϕ on $E_0[n]$.

We have developed tools for a number of cases.

1. The case $n = \ell$ is a prime dividing d. In this article, we will be interested in the case where ℓ furthermore divides the conductor f. Then we can use Section 3.2 to compute the associated ascending ℓ-isogeny $\phi_1 : E_0 \to E_0'$. We have $\phi = \phi_2 \circ \phi_1$, and we are reduced to recovering $\phi_2 : E_0' \to E_1$ of degree d/ℓ. We remark that the case where n divides d but not f is standard: we can write as above $\phi = \phi_2 \circ \phi_1$ where this time ϕ_1 is an horizontal ℓ-isogeny, and there are at most two such isogenies.
2. The case n is coprime to Δ_R and d. Then by Section 4.2, we have $O(n)$ possible choices for the action of ϕ on $E_0[n]$, which we can recover in $\tilde{O}(\ell + u^2 \log^2 q + eu \log \ell \log q + \log^{O(1)} q) = \tilde{O}(n^4 \log^2 q + \log^{O(1)} q)$. Here, u is the degree of the field extension of the points in $E_0[n]$. Indeed, this dominates both the cost of finding the conic equation Eq. (2), and the cost of finding all solutions of this conic.
3. The case $n \mid \Delta_R$, but with n coprime to d (hence also to f). For simplicity, we will assume that $n \mid \Delta_R/4$ if n is even. Then by Section 4.2, we have $O(2^m)$ possible choices for the action of ϕ on a specific point Q of $E_0[n]$

(with the notation of Theorem 4.2, the image by ω_R of an R-generator of $E_0[n]$), where m is the number of distinct prime divisors of n. These choices can be computed in $\tilde{O}(\ell + u^2 \log^2 q + eu \log \ell \log q + \log^{O(1)} q) = \tilde{O}(n^4 \log^2 q + \log^{O(1)} q)$.

We now explain how to use the torsion information from cases 2 and 3, which are the main cases of interest since d is coprime to n for these cases and so we are not directly getting information about the kernel of ϕ.

The method in [14] was based on the Kani attack [3, 21, 26], which needed the full torsion information (as provided by case 2). Namely, knowing the torsion information of ϕ on the n-torsion for $n^2 > 4d$ is enough to recover ϕ. This involves computing an n^2-isogeny $\Phi : E_0^r \times E_1^r \to E_0^r \times E_1^r$ in higher dimension $g = 2r$. The higher dimensional isogeny Φ is built out of an $(d, n^2 - d)$-isogeny diamond, involving the isogeny $\phi \,\mathrm{Id}_r$ with $r = 1, 2, 4$ and auxiliary isogenies ϕ', ρ, ρ' in dimension r of polarised degree $d, n^2 - d, n^2 - d$ respectively. Then Φ will be an isogeny of polarised degree n^2 in dimension $2r$.

The reason we need to move from dimension 1 to r is to be able to compute ψ, ψ' of the correct degree efficiently. Namely, if $n^2 - d = \sum_{i=1}^r a_i^2$ is the sum of $r = 1$, 2 or 4 squares, we can build the auxiliary isogenies in dimension r by using suitable $r \times r$ integer matrices. We know by Lagrange's theorem that we can always find a sum of 4 squares (and finding an explicit decomposition is efficient), so we can always work in dimension $g = 8$.

The Kani construction is well-documented in many papers so we do not give all the details. We remind the reader that $\ker \Phi$ may be computed as

$$\{(\hat{\phi}(P), -\rho'(P)) : P \in E_1^r[n^2]\}$$

where we view $\hat{\phi} : E_1^r \to E_0^r$ as a diagonal map, and where $\rho' : E_1^r \to E_1^r$ is coming from the matrix associated with the sum of squares. To compute $(\ker \Phi)[\ell]$ it suffices to know a basis (P, Q) for $E_0[\ell]$ and the values $(\phi(P), \phi(Q)) \in E_1[\ell]$ and everything else follows.

Finally, we mention that in practice, Φ is not computed directly, rather we write it as $\Phi = \Phi_2 \circ \Phi_1$ where Φ_i is of polarised degree n and we compute the two isogenies Φ_1 and $\tilde{\Phi}_2$ from $E_0^r \times E_1^r$ and $E_1^r \times E_0^r$ respectively and meet in the middle. It suffices to know ϕ on $E_0[n]$ to be able to compute the kernel of both these higher dimensional isogenies.

Remark 4.1. To handle case 3 we use an approach that was presented in an invited talk at ANTS in 2024 by Castryck [4]. At a high level, the result is that one can recover an unknown isogeny of degree d given interpolation data coming from a group of size $O(d)$.

We briefly recall the ideas from [4], which apply in a simplified form for our case 3. Reusing the notation and arguments from Theorem 4.2, we have points P_0, Q_0, P_1, Q_1 such that $\phi(P_0) = \gamma P_1$, for some $\gamma = \alpha + y\omega_R$, where α is constrained by an equation $\alpha^2 = v \pmod n$, but y can be arbitrary. From this we know that $\alpha Q_1 = \alpha \omega_R(P_1) = \phi(Q_0)$, but $\phi(P_0) = \alpha P_1 + y Q_1$. The idea is to

compose ϕ with an isogeny $\psi : E_1 \to E_2$ of kernel $\langle Q_1 \rangle$, so that

$$\psi(\phi(P_0)) = \psi(\alpha P_1 + yQ_1) = \alpha\psi(P_1).$$

We remark that $\ker \psi = E_1[n, \omega_R]$, so ψ is R-oriented, and does not depend on α. Now consider $\phi_2 = \psi \circ \phi$, of degree nd. Then we know that $\phi_2(P_0) = \psi(\alpha P_1 + y\omega_R P_1) = \alpha\psi(P_1)$. Since P_0 is an R-generator of $E_0[n]$, this allows to recover how ϕ_2 acts on the full n-torsion of $E_0[n]$.

In conclusion, we now have an isogeny ϕ_2 of degree dn and we know how ϕ_2 acts on the full group $E_0[n]$. We can then proceed with existing techniques.

All in all, let n_1 be the product of all prime powers we use for case 2, and n_2 the product of all prime powers we use for case 3. Let $n = n_1 n_2$. We build a n_2-isogeny $\psi : E_1 \to E_2$ using the method in the previous paragraph to obtain a $n_2 d$-isogeny $\phi_2 = \psi \circ \phi : E_0 \to E_2$ for which we know the action on the full $E_0[n]$-torsion. We need $n_1^2 n_2 > 4d$ to be able to recover ϕ_2 hence ϕ, via a $n_1^2 n_2^2$-isogeny $\Phi : E_0^r \times E_2^r \to E_0^r \times E_2^r$ in dimension $g = 2r$ (see Figure 6). Here the dimension r will depend on whether we can find a solution with $r = 1, 2, 4$ to $n_1^2 n_2^2 - n_2 d = \sum_{i=1}^{r} a_i^2$. Because of the common factor n_2, we require n_2 to be a sum of two squares to be able to work in dimension $g = 4$ (i.e., $r = 2$). In general we will need to work with $r = 4$ and so dimension 8 Abelian varieties.

Remark 4.2. Let E_0 be an ordinary elliptic curve over $\mathbb{F}_q$ with $q+1-t$ points and let $\mathrm{End}(E_0)$ have discriminant Δf^2. If t is odd then $\Delta f^2 = t^2 - 4q \equiv 1 \pmod 4$ and so $|\Delta| \equiv 3 \pmod 4$. So we know Δ is divisible by a prime congruent to 3 modulo 4. So a bad case is when t is odd and $|\Delta|$ is a prime.

Furthermore, while ψ does not depend on the torsion possibilities, we need to compute Φ for all the possibilities of the torsion we have. As explained before, in practice we compute two $n_1 n_2$-isogenies Φ_1, Φ_2 from $E_0^r \times E_2^r$. As a technical detail: to find the kernel of Φ_2 we need to know how $\tilde{\phi}_2$ acts on the n-torsion; this can be done from the knowledge of how ϕ_2 acts on the n-torsion through pairings and DLPs, because $\tilde{\phi}_2$ is the adjoint of ϕ_2 for the Weil pairing.

The complexity of the computation of the isogenies ψ, Φ itself will depend on the dimension where Φ lives, and if $n = \prod \ell_i^{e_i}$, of the largest prime ℓ_i and the degrees u_i of the field generated by the $\ell_i^{e_i}$-torsion points. More precisely, by [28, Lemma 5.7], if $n = \prod_{i=1}^{a} \ell_i^{e_i}$, computing an n-isogeny in dimension g can be done in time $\tilde{O}(ae\ell^g u \log q)$, where ℓ (resp. e) is a bound on the ℓ_i (resp. e_i), and u is a bound on the $lcm(u_i, u_j)$ where u_i is the degree of the field extension defined by the points of the $\ell_i^{e_i}$-torsion (so $u \leq n^2$).

The choice of prime power ℓ^e to use in practice will be a delicate balance between the size of ℓ, the degree of the field extension where the ℓ^e-torsion is defined, and the number of possibilities, mainly whether we are in case 2 or 3.

Remark 4.3. There is a non-trivial way to save time: rather than computing Φ from scratch from every possible choice of torsion information, we remark that we can reuse part of our construction of Φ.

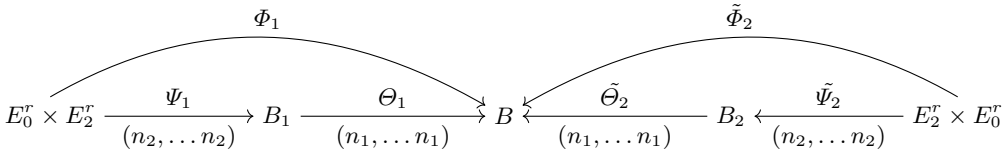

Fig. 6. This isogeny diagram corresponds to the explanation of how to address case 3. We use our knowledge of the $E_0[n]$ torsion group to compute the (probably non-smooth) $(n_2, \dots n_2)$-isogeny and guess the remaining (smooth) $(n_1, \dots n_1)$-isogeny ϕ_2 using a meet-in-the-middle computation (as detailed in Remark 4.3).

First, we can assume that we are given the torsion information of $\phi_2 : E_0 \to E_2$ in terms of matrices acting on basis elements (P_0, Q_0), (P_1', Q_1') of the $\ell_i^{e_i}$-torsion. We remark that the pairing information from Section 4.2 already gives this information on E_1, we have P_i a cyclic R-generator, and $Q_i = \omega_R P_i$, and the torsion information is expressed as a constraint on an endomorphism $\gamma \in R/\ell_i^{e_i}$, which we can recast as a matrix. We propagate this information through ψ, this involves sampling a new basis; we may pick P_1' a cyclic generator of $E_2[\ell_i^{e_i}]$ and let $Q_1' = \omega_R P_1'$ to simplify the computations. Then $\psi(P_1') = cQ_1'$, and we determine c by pairings and DLPs as in Section 4.2.

Now we have several choices of matrices for the action of ϕ_2 on $\ell_1^{e_1}$. We pick one choice, and use it to build the $\ell_1^{e_1}$-component of the kernel of Φ_1, Φ_2. We push the other $\ell_i^{e_i}$-basis points (using a product basis); this is where the compositum of the field of definition of the $\ell_1^{e_1}$ and $\ell_i^{e_i}$ torsion comes in. Then we make a choice of matrix for the action on $\ell_2^{e_2}$, determine the kernel of the next $\ell_2^{e_2}$ higher dimensional isogenies, and so on.

The key remark is that if we change our choice of matrix for $\ell_m^{e_m}$, we need to recompute the subsequent $\ell_i^{e_i}$ isogenies for $i \geq m$, but we can keep the ones already computed for $i < m$.

We illustrate this with a small example: assume that we want to use $n_1 = \ell_1 \ell_2$ with ℓ_1 split and ℓ_2 inert, and $n_2 = \ell_3 \ell_4$ a product of two different primes. We have $\ell_1 - 1$ possible choices for the action of ϕ on $E_0[\ell_1]$, $\ell_2 + 1$ possible choices for the action of ϕ on $E_0[\ell_1]$, and 2 possible choices for the action of ϕ on a subgroup of order ℓ_i of $E_0[\ell_i]$ for $i = 3, 4$. We can choose in which order to treat our primes, and it makes sense to pick them up by decreasing size. So suppose for instance that $\ell_3 > \ell_1 > \ell_2 > \ell_4$. Then via the above strategy, to compute all possibilities for Φ_1, we will compute in total $2(\ell_1 - 1)(\ell_2 + 1)2$ ℓ_4-isogenies, $2(\ell_1 - 1)(\ell_2 + 1)$ ℓ_2-isogenies, $2(\ell_1 - 1)$ ℓ_1-isogenies, and 2 ℓ_1-isogenies.

We have all the tools to bound the cost of recovering ϕ.

Theorem 4.3. *Let the notation and hypothesis be as at the beginning of Section 4. Specifically, let $\phi : E_0 \to E_1$ be an ascending R-oriented isogeny of degree d. Let $m \mid \Delta_R$ be a factor of Δ_R. Assume that R is m-locally primitive on E_0, and let $m = \prod \ell_i^{e_i}$. Denote the number of square roots of 1 modulo m by T, B*

a bound on the largest prime divisor ℓ_i of m, and u a bound on the $lcm(u_i, u_j)$, where u_i is the degree of the field of definitions of the points of $E_0[\ell_i^{e_i}]$.

Then we can recover $\phi : E_0 \to E_1$ in time $\tilde{O}((TuB^8 + T\sqrt{d/m})(\log q + \log d)^{O(1)})$.

Proof. First, dividing m by 2 or 4 if necessary, we can assume that $m \mid \Delta_R/4$. We can also assume that we know the factorisation of m, since we know the one of Δ_R.

If m is not coprime to d, we use standard tools such as case 1.

Following the notation used above, we take $n_2 = m$, and for n_1 a B'-powersmooth number coprime to Δ_R and of size $\Theta(\sqrt{d/m})$. We can always take a smoothness bound B' polynomial in $\log q + \log d$.

We compute the torsion information on the n_1 and n_2 torsion using Section 4.2. There are T possibilities for the n_2-torsion by Theorem 4.2, and $\tilde{O}(n_1)$ possibilities for the action on the n_1-torsion. By Corollaries 3.2 and 4.1 the cost of recovering this information is not the dominating step.

Computing the n_2-isogeny $\psi : E_1 \to E_2$ as in Remark 4.1 takes $\tilde{O}(Bu)$ field operations.

We apply the strategy of Remark 4.3 by computing the n_2-isogeny first, followed by the n_1-isogeny. We assume that we are in the worst case and that we need to work in dimension 8. The first n_2-isogeny costs $\tilde{O}(uB^8 \log^2(n_2) \log q)$. We also need to push the torsion information on the n_1-torsion, which costs $\tilde{O}(u'B^8 \log n_2 \log n_1 \log q)$. Here u' is a bound on the degree of the compositum field of the $\ell_i^{e_i}$-torsion and the $\ell_j^{e_j}$ torsion for $\ell_i \mid n_2$ and $\ell_j \mid n_1$, so since n_1 is B'-powersmooth, we have $u' \leq u(B')^2$. We need to perform these computations T times.

Then we need to compute the remaining n_1-isogeny and decide if we meet in the middle. This costs $\tilde{O}(B'^4 B'^8 \log^2(n_1) \log q)$. We need to do this computation $\tilde{O}(Tn_1)$ times. This gives us our final complexity result. □

Remark 4.4.

- A bound on u is given by $u = \min(O(m^2), O(B^4))$ in the ordinary case and $u = \min(O(m), O(B^2))$ in the supersingular case.
- In practice for our applications, the term $TuB^8 + T\sqrt{d/m}$ will be large enough to absorb the $\log q + \log d$ part into the $\tilde{O}$ factor, and we will use $\tilde{O}(TuB^8 + T\sqrt{d/m})$ as our complexity.
- If m is a sum of two squares, then heuristically we can work in dimension 4, and replace the term TuB^8 by TuB^4.
- If $\phi : E_0 \to E_1$ is an R-oriented cyclic descending isogeny of degree d, we can follow the same approach, *except* for case 2: ℓ is a prime dividing the conductor. Indeed, in that case, $K = (\ker \phi)[\ell]$ is a kernel of a strictly descending isogeny, but there are between $\ell-1$ and $\ell+1$ descending isogenies, so many more possibilities for K than in the case of a purely ascending isogeny when K is unique. Of course, the better strategy is to reconstruct the dual $\tilde{\phi} : E_1 \to E_0$ instead, which is ascending.

5 Improving the computational isogeny problem when the conductor is large

In this section we apply the techniques from Section 4 to the ordinary isogeny problem. In particular, we will show how to use the pairings from Section 4.2 to improve the algorithm from [14]. Our main application is the case of volcanoes with a small crater.

5.1 Improved computation of an ascending isogeny.

Let E_0 and E_1 be ordinary elliptic curves over $\mathbb{F}_q$ with $q+1-t$ points. Let $t^2-4q = f^2\Delta$, where Δ is a fundamental discriminant and $f > 1$. Suppose $\mathrm{End}(E_1)$ has discriminant $\Delta_1 = \Delta f_1^2$, $\mathrm{End}(E_0)$ has discriminant $\Delta_0 = \Delta f_0^2$ with $f_0 = Nf_1$. We want to recover the ascending $\mathbb{F}_q$-rational N-isogeny $\phi : E_0 \to E_1$ of degree $N|f$. Since ϕ is $\mathbb{F}_q$-rational we have $\phi \circ \pi_q = \pi_q \circ \phi$, where π_q denotes the q-power Frobenius maps on E_0 and E_1, so ϕ is $\mathbb{Z}[\pi_q]$-oriented. In this section we will neglect factors polynomial in $\log q$.

Galbraith in [14, Theorem 2] gives an algorithm that can compute ϕ in $\tilde{O}(N^{1/2})$ operations over $\mathbb{F}_q$. Our first improvement concerns a heuristic assumption used in the algorithm. As explained in Section 2.3, a key task of the 2024 algorithm is to guess $\phi(P_0), \phi(Q_0)$ for a suitable basis P_0, Q_0 of $E_0[\ell]$. This was done in [14] by choosing P_0 and Q_0 to be Elkies primes (i.e., eigenvectors for the q-power Frobenius). But by Theorem 4.1, we can use inert primes just as well as split primes. This removes the need for an assumption about the distribution of Elkies primes.

Our second main improvement will come from cyclic self-pairings (or from our point of view, the use of degenerate conics), coming from a factor $m \mid \Delta_1$. The details of the self-pairings algorithm is outlined in Algorithm 2.

In that algorithm we restrict for simplicity to the case where m is odd and $\gcd(m, f) = 1$. This is because in that case we can directly use the Frobenius π_q to get the self-pairing information. Indeed, define $\tau = 2\pi_q - t$, which will be interpreted depending on the context as lying in $\mathrm{End}(E_0)$ and $\mathrm{End}(E_1)$. Note that $\tau = \phi \circ \alpha \circ \hat{\phi}$, where $\alpha \in \mathrm{End}(E_0)$ is the endomorphism corresponding to $\sqrt{\Delta}$. We also have that $\mathrm{Tr}(\tau) = 2\,\mathrm{Tr}(\pi_q) - 2t = 0$ and $\Delta(\tau) = \tau^2 = t^2 - 4q$, $N(\tau) = 4q - t^2$. In particular τ is, up to a factor 2, our canonical imaginary element for $\mathbb{Z}[\pi_q]$ from Definition 3.2. We can thus apply Section 4.2, using $\omega_R = \tau$. Since $e_m(P, \tau Q) = e_m(P, \pi_q Q)^2$, the situation is simplified because we can simply evaluate the Frobenius rather than a more general endomorphism of the form τ/b. We refer to Theorem 4.2 for the general case, when m may have a common factor with f.

We then illustrate how to use the self-pairing information to recover ϕ in Algorithm 3. This algorithm calls two functions Guess() and Kani().

As already explained, we need to compute $(\ker \Phi)[\ell]$ for various primes ℓ, and this is fully determined by a basis (P, Q) for $E_0[\ell]$ and the corresponding image $(P', Q') \in E_2[\ell]$ under $\psi \circ \phi$.

In our applications we will have a point $P_0 \in E_0[\ell]$ such that $RP_0 = E_0[\ell]$ and a point $P_1 \in E_1[\ell]$ such that $RP_1 = E_1[\ell]$. We know that $\phi(P_0) = \gamma(P_1)$ for some $\gamma = a + b\omega_R$, but we don't know (a, b). Indeed, we will extend this to $P_2 \in E_2[\ell]$ such that $\psi \circ \phi(P_0) = \gamma(P_2)$, using the fact that ψ is also R-oriented. The function $\mathsf{Guess}()$ will produce a list of candidates for the unknown coefficients (a, b) in γ. To be precise, it will take as input torsion information as a sequence of data $(P_0, P_2) \in E_0[\ell] \times E_2[\ell]$, and it will output a list of items, each of which is a tuple of candidates for the correct $(a, b) \in (\mathbb{Z}/\ell\mathbb{Z})^2$.

The notation A_ℓ represents the data needed for the Kani construction. For example, $A_{n_2} = (n_2, (P_0, Q_0), (S, 0)))$ represents the information that (P_0, Q_0) is a basis for $E_0[n_2]$ and that $\psi \circ \phi(P_0) = S$ and $\psi \circ \phi(Q_0) = 0$.

The function $\mathsf{Kani}()$ takes as input the torsion data $\ell, (P_0, P_2) \in E_0[\ell] \times E_2[\ell]$ and a candidate (a, b) and constructs some prefix of the map Φ of degree $(\ell, \ell, \ldots, \ell)$, by computing the ℓ-part of $\ker \Phi$. An extra detail is that we will apply this to an abelian variety B_1 that is already part-way along the path of Φ. This is handled by calculating the image of $\psi \circ \phi(P_0) = (a + b\omega_R)(P_2)$ on B_1 by working with the images of the points P_0 and P_2 on B_1. We will also abuse notation and use the function $\mathsf{Kani}()$ for computing the ℓ-part of $\ker \tilde{\Phi}$ corresponding to the right hand side of Figure 6.

Algorithm 3 puts everything together. It describes how to recover an ascending isogeny between two fixed elliptic curves by making use of the partial information gained from self-pairings and then using $\mathsf{Guess}()$ and $\mathsf{Kani}()$ to recover the remaining part.

Applying Theorem 4.3 to our situation, we obtain the following complexity.

Proposition 5.1. *Let E_0 and E_1 be ordinary elliptic curves over $\mathbb{F}_q$. Suppose $\mathrm{End}(E_1)$ has discriminant Δ_1 and $\mathrm{End}(E_0)$ has discriminant $\Delta_0 = \Delta_1 N^2$. Suppose there is an N-isogeny $\phi : E_0 \to E_1$ (in other words E_1 is directly above E_0 in the volcano). Suppose Δ_1 has a large divisor $m = q^a$ which has $O(\log(\log(q)))$ distinct prime factors.*

Then we can recover ϕ using

$$\max\left(\tilde{O}(q^{10a}), \tilde{O}(q^{(1-3a)/4})\right)$$

operations over $\mathbb{F}_q$.

Proof. In the notation of Theorem 4.3, we have $d = N$, and since $N^2|\Delta| = O(q)$, we have $N = O(q^{(1-a)/2})$. We also have $m = O(q^a)$ since $m \mid \Delta$, so $\sqrt{d/m} = O(q^{(1-3a)/4})$. We also take the trivial upper bounds $B = m$ and $u = m^2$. Finally $T = O(\log q)$ because of our assumptions on the number of distinct prime factors of m. Then we can bound the term TuB^8 by $\tilde{O}(q^{10a})$, and the term $T\sqrt{d/m}$ by $\tilde{O}(q^{(1-3a)/4})$. □

Remark 5.1. It remains to address the question of whether $|\Delta_1|$ contains a large divisor m with $O(\log \log q)$ distinct prime factors. Unfortunately this may not

Algorithm 2: Torsion recovery via self-pairings

Input : Ordinary elliptic curves E_0, E_1 over $\mathbb{F}_q$ such that $[\mathrm{End}(E_0) : \mathrm{End}(E_1)] = N$, an odd integer m such that $m \mid \Delta_1$, $(m, f) = 1$;

Output: $S_0 \in E_0(\mathbb{F}_q), \{S_j\}_{j=1}^T$ such that $\phi(S_0) = S_j$ for one $j \in [1, T]$.

1 Sample $P_i \in E_i[m]$ that is a generator for the $\mathbb{Z}[\pi_q]$-module $E_i[m]$ for $i = 0, 1$;
2 Define $\tau = 2\pi_q - t$ and set $Q_i = \tau(P_i)$ for $i = 0, 1$;
3 Compute the following pairings: $T_i = e_m(P_i, Q_i), i = 0, 1$;
4 Compute all possible $\{a_j\}_{j=1}^T$ such that $(T_1)^{a_j^2} = (T_0)^N$;
5 **return** $\big((P_0, Q_0, P_1, Q_1), \{a_j\}_{j=1}^T\big)$;

always be the case. If $|\Delta_1|$ is a primorial then divisors of it, even of size only $|\Delta_1|^{1/2}$, will have too many distinct prime factors.

However, it is well-known (see Theorem 430 of Section 22.10 of Hardy and Wright [15]) that the average number of distinct prime factors of integers up to X is $\log\log(X)$. Hence the required condition on m applies to most integers Δ_1 and our claim holds on average.

We also consider the *very* special case that Δ_1 has a large powersmooth factor in the following corollary.

Proposition 5.2. *Let notation and hypotheses be as in Proposition 5.1. In addition, suppose the discriminant Δ_1 has a factor $m \mid \Delta_1$ of size q^a which is B-powersmooth. Denote the number of square roots of 1 modulo m by T.*

Then we can recover $\phi : E_0 \to E_1$ in time

$$\max\Big(\tilde{O}(TB^{12}), \tilde{O}(T \cdot q^{(1-3a)/4})\Big)$$

Proof. This is immediate from Theorem 4.3 using the bound $u = B^4$ and the analysis from Proposition 5.1. □

5.2 Volcanoes with small crater

Let E_0 and E_1 be ordinary elliptic curves over $\mathbb{F}_q$ with $q + 1 - t$ points where $t^2 - 4q = f^2\Delta$ where Δ is the fundamental discriminant, and $|\Delta| = q^a$ with $a > 0$ small. We want to find an isogeny between them, but this time we do not assume that E_0 is directly below E_1.

If E_0 is on the crater and E_1 on the floor then the approach in Galbraith [14] has complexity $\tilde{O}(h_0 f^{1/2})$ operations over $\mathbb{F}_q$, where h_0 is the class number of Δ. Since $f = O(q^{(1-a)/2})$ and $h_0 = \tilde{O}(\sqrt{|\Delta_0|}) = \tilde{O}(q^{a/2})$, we get that $\tilde{O}(h_0 f^{1/2}) = \tilde{O}(q^{a/2} q^{(1-a)/4}) = \tilde{O}(q^{(1+a)/4})$. The approach in [14] dealt with the cases when

Algorithm 3: Vertical isogeny recovery in volcanoes with small crater

Input : Ordinary elliptic curves E_0, E_1 such that $[\mathrm{End}(E_0) : \mathrm{End}(E_1)] = N$
Output: A representation of $\phi : E_0 \to E_1$

Choose $n_2 \in \mathbb{Z}$ dividing the discriminant Δ that satisfies the conditions in Algorithm 2;
Choose t small distinct primes $\ell_1, \dots \ell_t$ and a (small) integer s such that $N < 3^s \cdot n_2 \cdot \prod_{i\in[1,t]} \ell_i^2 < 2N(c + 2t\log(t))^2)$ for a small constant c;
Set $n_1 = 4 \cdot 3^{\lceil s/2 \rceil} \ell_1 \cdots \ell_t$;
Write $3^s n_2^2 (\ell_1 \cdots \ell_t)^2 - n_2 N > 0$ as a sum of squares;
Run Algorithm 2 to get $(P_0, Q_0, P_1, Q_1, \{a_j\}_{j=1}^T)$ such that $\phi(Q_0) = a_j Q_1$ for one of the $j \in [1, T]$;
Compute $\psi : E_1 \to E_2$ with kernel $\langle Q_1 \rangle$;
Set up torsion information $S_\ell = (P_{0,\ell}, P_{2,\ell}) \in E_0[\ell] \times E_2[\ell]$ for each ℓ_i, and also $S_{3^{\lceil s/2 \rceil}} \in E_0[3^{\lceil s/2 \rceil}] \times E_2[3^{\lceil s/2 \rceil}]$;
for $j \in [1, T]$ **do**
    Set $A_{n_2} = (n_2, (P_0, Q_0), (a_j \psi(P_1), 0)))$;
    Compute $\Psi_1, B_1 \leftarrow \mathsf{Kani}(E_0^4 \times E_2^4, A_{n_2})$; `// The partial Kani isogenies`
    Compute the kernel of the dual isogeny to Ψ_1 and use Kani to compute the image B_2 of the isogeny $\tilde{\Psi}_2 : E_2^4 \times E_0^4 \to B_2$;
    Compute $S'_\ell = \Psi_1((P_0, Q_0, S_j))$ and $S''_\ell = \tilde{\Psi}_2((P_0, Q_0, S_j))$ for each $i = 1, \dots t$; `// Push the torsion basis`
    Let $M = \tilde{O}(n_1)$ be the number of possible guesses for the images of ϕ on n_1-torsion;
    for $i \in [1, M]$ **do**
        Compute bases $A_{3^{s/2}}, A_{\ell_1}, \dots, A_{\ell_t}$ from $\mathsf{Guess}(S_{3^{\lceil s/2 \rceil}}, \{S_{\ell_1}, \dots S_{\ell_t}\})[i]$;
        $\Theta_1, B \leftarrow \mathsf{Kani}(B_1, n_1, A_{3^{s/2}}, A_{\ell_1}, \dots, A_{\ell_t})$;
        Use Kani to compute the corresponding $\tilde{\Theta}_2$ from B_2 and call the codomain B';
        `// Test if the two Kani isogenies have met in the middle:`
        if $B \cong B'$ **then**
            return $\Psi_2 \circ \Theta_2 \circ \Theta_1 \circ \Psi_1$ as a representation of ϕ;
return $\perp$;

E_0 and E_1 are not on the crater by descending to the floor and solving the isogeny problem in $\tilde{O}(q^{1/4})$ time.

We can improve these results using the results of Section 5.1, in the case when Δ has $O(\log\log(q))$ prime factors.

Theorem 5.1. *Assume that $|\Delta| = q^a$ has $O(\log\log(q))$ prime factors, and that $a < 1/41$. Then we can find a connecting isogeny between E_0 and E_1 in time $\tilde{O}(q^{(1-a)/4})$.*

Proof. We will show that we can climb up from E_0 to the crater in $\tilde{O}(q^{(1-a)/4})$ when a is small enough. Doing the same for E_1, we can then solve the isogeny problem in the crater in time $\tilde{O}(h_0^{1/2})$ where $h_0 = \tilde{O}(|\Delta|^{1/2}) = \tilde{O}(q^{a/2})$ is the number of curves on the crater, which is the class number of $\mathbb{Q}(\sqrt{\Delta})$. So $\tilde{O}(h_0^{1/2}) = \tilde{O}(q^{a/4})$, and the dominating step is the climbing up phase (since a is small).

We let $m = |\Delta|/4$ if $4 \mid \Delta$, and $m = |\Delta|$ otherwise. We still have $m = \Theta(q^a)$. If we knew the curve E_0' on the crater directly above E_0, we could apply Proposition 5.1 to climb up in time $\tilde{O}(q^{(1-3a)/4})$, when $a < 1/43$ (for larger a, the dominating complexity becomes $\tilde{O}(q^{10a})$). Since we do not know E_0', we need to test all the curves on the crater, for a total time of $\tilde{O}(h_0 q^{(1-3a)/4}) = \tilde{O}(q^{(1-a)/4})$. This is the dominating complexity over $\tilde{O}(q^{10a})$ when $a < 1/41$. □

Remark 5.2. If all primes dividing $f = O(q^{(1-a)/2})$ are smaller than $q^{a'}$ then one can compute ascending isogenies from E_0, E_1 to the crater in $\tilde{O}(q^{2a'})$ operations and solve the isogeny problem in the crater in $\tilde{O}(|\Delta|^{1/2}) = \tilde{O}(q^{a/2})$ operations, for a total cost of $\tilde{O}(q^{2a'} + q^{a/2})$. So Theorem 5.1 is mainly interesting in the case where f has at least one large enough prime divisor.

6 A special case attack attempt on SCALLOP

In what follows we outline how the general approach from Section 4 might potentially be leveraged to attack the SCALLOP group action when the fundamental discriminant has a large smooth factor. This indicates that any cryptographic protocol using this group action should exclude (almost) smooth discriminants from their parameter selection.

SCALLOP [11] is a group action on oriented supersingular elliptic curves, initially proposed with the purpose of improving upon the CSIDH group action. There have since been two variants: SCALLOP-HD [7] and PEARL-SCALLOP [1], that improve upon some subroutines and parameter choices but rely on similar security assumptions. Our attack will apply in the special case that the discriminant has a large smooth factor, so we will focus on PEARL-SCALLOP, whose parameter generation uses larger fundamental discriminants.

We consider the SCALLOP class group action by a non-maximal order $\mathcal{O} \subset K$ of discriminant $\Delta = \Delta_0 f^2$. Suppose we are given two oriented curves (E_0, ι_0) and $(E_\mathfrak{a}, \iota_\mathfrak{a})$ on the floor. We would like to recover an ideal $\mathfrak{a}$ in $Cl(\mathcal{O})$ such that $\mathfrak{a} \star (E_0, \iota_0) = (E_\mathfrak{a}, \iota_\mathfrak{a})$. In this section, we only look at classical attacks.

Standard attack. The standard classical attack for this problem is direct meet-in-the-middle [1, Sect. 3.5]. Since the class group is of order $\sqrt{|\Delta|}$ this gives a complexity of $\tilde{O}(|\Delta^{1/4}|) = \tilde{O}(|\Delta_0|^{1/4} f^{1/2})$ field operations. So for a security parameter, λ, we can expect $\log|\Delta| = 4\lambda$ in order to avoid this attack.

The "Climbing Up" approach. We consider the same attack from Galbraith [14] in the ordinary case, that uses the Kani machinery to climb up to the crater from each curve, and then use meet-in-the-middle on the crater to join the paths (see the figure below).

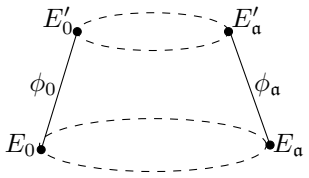

This is the approach we improved in Section 5.1 by using self pairings.

Recall, this approach involves guessing which curve on the crater is "directly above" the curves, i.e. the curve $E'_{\mathfrak{a}}$ such that $[\mathrm{End}(E'_{\mathfrak{a}}) : \mathrm{End}(E_{\mathfrak{a}})] = f$. Note that we usually already know E'_0. For brevity we outline how to recover $\phi_{\mathfrak{a}}$, though recovering ϕ_0 follows symmetrically. There are $\tilde{O}(\sqrt{\Delta_0})$ curves that lie on the crater. For each of these candidate curves, we must then guess the action of $\phi_{\mathfrak{a}}$ on some torsion points. Galbraith [14] shows that this guessing requires a complexity of $O(\sqrt{f})$; so in total the approach requires $(|\Delta_0|f)^{1/2}$ field operations to complete.

From here, we compute a meet-in-the-middle search on the crater to get an ideal mapping $E'_0 \to E'_{\mathfrak{a}}$. This requires $\tilde{O}(|\Delta_0|^{1/4})$ field operations. Let $\mathcal{O}_0$ be the maximal order of $\mathcal{O}$. So far we have recovered an ideal $\mathfrak{b}$ such that $\mathfrak{b} \star E'_0 = E'_{\mathfrak{a}}$, but where $[\mathfrak{b}]$ is in $Cl(\mathcal{O}_0)$ instead of $Cl(\mathcal{O})$. Notice that we have an exact sequence

$$0 \to G \to Cl(\mathcal{O}) \to Cl(\mathcal{O}_0) \to 0,$$

where $G := \ker(Cl(\mathcal{O}) \to Cl(\mathcal{O}_0))$ is of order $\Theta(f)$ that has a very explicit description (see Theorem 7.24 and Equation (7.25) of Cox [9]).

Taking an arbitrary preimage of $[\mathfrak{b}]$ in $Cl(\mathcal{O})$, say $[\mathfrak{b}']$, we may still be off from our target ideal by an element in G. With yet another meet-in-the-middle on the G action we recover this element with a complexity of $O(f^{1/2})$ operations.

Thus, to recover the ascending isogeny, compute the meet-in-the-middle on the crater, and recover the necessary element of G, the final complexity is

$$\tilde{O}(|\Delta_0|^{1/2} f^{1/2}) + \tilde{O}(|\Delta_0|^{1/4}) + \tilde{O}(f^{1/2}) = \tilde{O}(|\Delta_0|^{1/2} f^{1/2})$$

field operations. Note, this complexity is worse than the direct meet-in-the-middle attack from before.

Self-pairings attack. We can follow the same idea as the "Climbing Up" approach, but hope to gain some savings using self-pairings. In the following theorem, we compute the cost of recovering the ascending isogeny as outlined in Algorithm 3.

Theorem 6.1. *Fix an imaginary quadratic order K and a non-maximal order $\mathcal{O} \subset K$ with discriminant $\Delta = \Delta_0 f^2$, where Δ_0 is a fundamental discriminant. Further, suppose Δ_0 has a B-smooth factor m. Denote the number of square roots of 1 modulo m by T.*

Let (E_0, ι_0) and $(E_\mathfrak{a}, \iota_\mathfrak{a})$ be two oriented curves such that there exists an ideal $\mathfrak{a} \in Cl(\mathcal{O})$ such that $\mathfrak{a} \star (E_0, \iota_0) = (E_\mathfrak{a}, \iota_\mathfrak{a})$.

Then we can recover $\mathfrak{a}$ in

$$\tilde{O}\Big(T|\Delta_0|^{1/2}B^{10} + T(f|\Delta_0|/m)^{1/2}\Big)$$

field operations.

Proof. First, suppose we are given two supersingular curves $E_\mathfrak{a}, E'_\mathfrak{a}$ over $\bar{\mathbb{F}}_p$ such that $[\mathrm{End}(E'_\mathfrak{a}) : \mathrm{End}(E_\mathfrak{a})] = f$.

Then recovering the isogeny $\phi_\mathfrak{a} : E_\mathfrak{a} \to E'_\mathfrak{a}$ can be done in $\tilde{O}(T(B^{10} + f^{1/2}m^{-1/2})\log(q)^{O(1)})$ field operations by Theorem 4.3 (here we bound u by B^2).

Since we must also guess which curve $E'_\mathfrak{a}$ on the crater is above $E_\mathfrak{a}$, we must again multiply the complexity above by the number of guesses, which is $\tilde{O}\sqrt{|\Delta_0|}$. This gives a complexity of $\tilde{O}(T|\Delta_0|^{1/2}B^{10} + T(f|\Delta_0|/m)^{1/2})$ field operations.

From here, we have the same task as before of computing the precise ideal $\mathfrak{b}$, requiring yet another meet-in-the-middle computation in the action G. This costs $\tilde{O}(|\Delta_0|^{1/4} + f^{1/2})$ field operations, which is negligible compared to the climbing up phase. □

Let us assume that m has $O(\log\log q)$ distinct factors, so that $T = O(\log q)$. The complexity becomes $\tilde{O}(|\Delta_0|^{1/2}B^{10} + (f|\Delta_0|/m)^{1/2})$ field operations.

We see that for this method to be better than the direct attack on the floor which costs $\tilde{O}(|\Delta_0|^{1/4}f^{1/2})$, we need:

- $|\Delta_0|$ to be not too large compared to f so that $|\Delta_0|^{1/4}B^{10} \ll f^{1/2}$.
- m large enough compared to Δ_0, so that $m \gg |\Delta_0|^{1/2}$, but with few enough prime factors, which forces m not to be too large either: $m \leq B^{O(\log\log q)}$. In particular, we need $|\Delta_0|^{1/2} \ll B^{O(\log\log q)}$.

These make for tight restrictions on the choice of Δ_0 and f even when $|\Delta_0| \ll f$.

For instance, if $B^{10} = f^{1/2}|\Delta_0|^{-1/4-\epsilon}$, and $m = |\Delta_0|^{1/2+\epsilon}$ (with few enough prime factors), we obtain a complexity of $O(f^{1/2}|\Delta_0|^{1/4-\epsilon})$.

In the most favourable scenario for our attack, Δ_0 has a large smooth divisor $m = \prod \ell_i$, such that each ℓ_i is congruent to 1 modulo 4 so that we can work in dimension $g = 4$ rather than 8, and $\ell_i \mid p+1$ so the ℓ_i-torsion is already rational. The reason we need ℓ_i congruent to 1 modulo 4 to be able to work in dimension 4

is that this allows us to write a suitable multiple m' of m as a sum of two squares, which allows us to build a suitable m'-isogeny in dimension 2 given by a matrix of integers. But for SCALLOP we also have access to endomorphisms in $\mathcal{O}$, not just integers, and we can build an m'-isogeny in dimension 2 given by $\begin{pmatrix} \gamma_1 & \overline{\gamma_2} \\ -\gamma_2 & \overline{\gamma_1} \end{pmatrix}$ as long as we find $\gamma_1, \gamma_2 \in \mathcal{O}$ such that $m' = N(\gamma_1) + N(\gamma_2)$. This can allow us to relax the conditions on the ℓ_i. The total complexity of the attack would then be $\tilde{O}(|\Delta_0|^{1/2}B^4 + (f|\Delta_0|/m)^{1/2})$.

We finish by an example showing that our attack does not apply to the public SCALLOP parameter sets.

Example 6.1.

- The original version of SCALLOP uses $|\Delta_0| = O(1)$. This is too small to be of use, since it means that m is $O(1)$ and the self-pairings do not provide enough information.
- For $\lambda = 128$ bits of classical security, the authors of PEARL-SCALLOP use a discriminant $\Delta = \Delta_0 f^2$ of 512 bits, where Δ_0 has 6 factors, $\log(|\Delta_0|) = 256$, and $\log(f) = 128$. Here Δ_0 is too large compared to f for our attack to apply: we would have too many curves on the crater to try.
- For $\lambda = 256$ bits of security the authors of PEARL-SCALLOP [1, Sect. 4.1] suggest a discriminant of 1024 bits where Δ_0 has 4 factors, $\log(|\Delta_0|) = 258$, and $\log(f) = 390$.
  The self-pairings attack described above would give an improvement over state-of-the-art for a smoothness bound B such that
  $$|\Delta_0|^{1/2}B^{10} < |\Delta_0|^{1/4}f^{1/2}.$$
  This gives $\log B < 13$ (of course this is a back of the envelope computation, the exact value would depend on the exact hidden factors). Their Δ_0 contains a 2^{16}-smooth factor m of size $\log(m) \approx 33 \ll 258/2$. So Δ_0 is not smooth enough for our attack to apply either.

References

1. Bill Allombert, Jean-François Biasse, Jonathan Komada Eriksen, Péter Kutas, Chris Leonardi, Aurel Page, Renate Scheidler, and Márton Tot Bagi. Faster SCALLOP from non-prime conductor suborders in medium sized quadratic fields. In Tibor Jager and Jiaxin Pan, editors, *PKC 2025*, volume 15676 of *Lecture Notes in Computer Science*, pages 333–363. Springer, 2025.
2. Daniel J. Bernstein, Luca De Feo, Antonin Leroux, and Benjamin Smith. Faster computation of isogenies of large prime degree. MSP Open Book Series, 2020.
3. Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023*, volume 14008 of *Lecture Notes in Computer Science*, pages 423–447. Springer, 2023.
4. Wouter Castryck, Thomas Decru, Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. Isogeny interpolation for elliptic curves, and applications. Presented by Wouter Castryck at

the Sixteenth Algorithmic Number Theory Symposium, Massachusetts Intitute of Technology, July 15, 2024. https://antsmath.org/ANTSXVI/schedule.html.
5. Wouter Castryck, Marc Houben, Simon-Philipp Merz, Marzio Mula, Sam van Buuren, and Frederik Vercauteren. Weak instances of class group action based cryptography via self-pairings. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023*, volume 14083 of *Lecture Notes in Computer Science*, pages 762–792. Springer, 2023.
6. Wouter Castryck, Marc Houben, Frederik Vercauteren, and Benjamin Wesolowski. On the decisional Diffie–Hellman problem for class group actions on oriented elliptic curves. *Research in Number Theory*, 8(4):99, 2022.
7. Mingjie Chen, Antonin Leroux, and Lorenz Panny. SCALLOP-HD: group action from 2-dimensional isogenies. In *PKC 2024*, volume 14603 of *Lecture Notes in Computer Science*, pages 190–216. Springer, 2024.
8. Leonardo Colò and David Kohel. Orienting supersingular isogeny graphs. *Journal of Mathematical Cryptology*, 14(1):414–437, 2020.
9. David A. Cox. *Primes of the form $x^2 + ny^2$: Fermat, Class Field Theory, and Complex Multiplication*. AMS Chelsea Publishing/American Mathematical Society, third edition, 2022.
10. B. Edixhoven, G. van der Geer, and B. Moonen. Abelian varieties. Book project, 2012. http://van-der-geer.nl/~gerard/AV.pdf.
11. Luca De Feo, Tako Boris Fouotsa, Péter Kutas, Antonin Leroux, Simon-Philipp Merz, Lorenz Panny, and Benjamin Wesolowski. SCALLOP: scaling the CSI-FiSh. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023*, volume 13940 of *Lecture Notes in Computer Science*, pages 345–375. Springer, 2023.
12. Mireille Fouquet, Josep M. Miret, and Javier Valera. Distorting the volcano. *Finite Fields and Their Applications*, 49:108–125, 2018.
13. Steven D. Galbraith. Constructing isogenies between elliptic curves over finite fields. *LMS J. Comput. Math.*, 2:118–138, 1999.
14. Steven D. Galbraith. Climbing and descending tall isogeny volcanos. *Research in Number Theory*, 11(7), 2025.
15. G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers (5th edition)*. Oxford, 1979.
16. Sorina Ionica and Antoine Joux. Pairing the volcano. *Math. Comput.*, 82(281):581–603, 2013.
17. Bruce W Jordan, Allan G Keeton, Bjorn Poonen, Eric M Rains, Nicholas Shepherd-Barron, and John T Tate. Abelian varieties isogenous to a power of an elliptic curve. *Compositio Mathematica*, 154(5):934–959, 2018.
18. Ernst Kani. The number of curves of genus two with elliptic differentials. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1997:122 – 93, 1997.
19. D. R. Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California, Berkeley, 1996.
20. Joseph Macula and Katherine E. Stange. Extending class group action attacks via sesquilinear pairings. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024*, volume 15486 of *Lecture Notes in Computer Science*, pages 371–395. Springer, 2024.
21. Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023*, volume 14008 of *Lecture Notes in Computer Science*, pages 448–471. Springer, 2023.

22. Josep M. Miret, Ramiro Moreno, Daniel Sadornil, Juan Tena, and Magda Valls. An algorithm to compute volcanoes of 2-isogenies of elliptic curves over finite fields. *Applied Mathematics and Computation*, 176(2):739–750, 2006.
23. Hiroshi Onuki. On oriented supersingular elliptic curves. *Finite Fields Their Appl.*, 69:101777, 2021.
24. Thomas Pornin. Optimized discrete logarithm computation for faster square roots in finite fields. Cryptology ePrint Archive, Paper 2023/828, 2023.
25. Damien Robert. Some applications of higher dimensional isogenies to elliptic curves (preliminary version). *IACR Cryptol. ePrint Arch.*, page 1704, 2022.
26. Damien Robert. Breaking SIDH in polynomial time. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023*, volume 14008 of *Lecture Notes in Computer Science*, pages 472–503. Springer, 2023.
27. Damien Robert. The module action for isogeny based cryptography. IACR Eprint 2024/1556, October 2024.
28. Damien Robert. On the efficient representation of isogenies (a survey). IACR Eprint 2024/1071, June 2024.
29. Katherine E Stange. Sesquilinear pairings on elliptic curves. *arXiv preprint arXiv:2405.14167*, 2024.
30. John Tate. Endomorphisms of abelian varieties over finite fields. *Inventiones mathematicae*, 2:134–144, 1966.

A Tate-Weil-Cartier pairings

The goal of this section is to revisit the construction of self pairings from [5].

Consider a commutative diagram of isogenies of abelian varieties:

$$\begin{array}{ccc} A_1 & \xrightarrow{\sigma_1} & A_2 \\ \downarrow{\scriptstyle \psi_1} & & \downarrow{\scriptstyle \psi_2} \\ B_1 & \xrightarrow{\sigma_2} & B_2 \end{array}$$

We know that the Weil-Cartier pairing [10, Ch. 11] is non degenerate: $e_{\sigma_1} : A_1[\sigma_1] \times \hat{A}_2[\hat{\sigma}_1] \to \mathbb{G}_m$, where we denote by $A_1[\sigma_1]$ the kernel of σ_1. By analogy with the standard Tate pairing, we can define a Tate-Weil-Cartier pairing by restricting to the subgroup $A_1[\sigma_1, \psi_1] = A_1[\sigma_1] \cap A_1[\psi_1]$ of $A_1[\sigma_1]$.

Theorem A.1. *There is a well defined Tate-Weil-Cartier pairing:*

$$T_{\psi_1}^{\sigma_1} : A_1[\sigma_1, \psi_1] \times \hat{A}_2[\hat{\sigma}_1]/\hat{\psi}_2(\hat{B}_2[\hat{\sigma}_2]) \to \mathbb{G}_m$$

which can be computed for $P_1 \in A_1[\sigma_1, \psi_1]$ and $Q_2 \in \hat{A}_2[\hat{\sigma}_1]$ by

$$T_{\psi_1}^{\sigma_1}(P_1, Q_2) = e_{\sigma_1}(P_1, Q_2) = e_{\psi_1}(P_1, \hat{\sigma}_2 Q') \tag{3}$$

where $\hat{\psi}_2 Q' = Q$.

Proof. Let

$$A_1[\sigma_1, \psi_1]^\perp = \{P \in \hat{A}_2[\hat{\sigma}_1] : e_{\sigma_1}(Q, P) = 1 \ \forall \ Q \in A_1[\sigma_1, \psi_1]\}$$

be the orthogonal of $A_1[\sigma_1, \psi_1]$ in $\hat{A}_2[\hat{\sigma}_1]$ for the Weil-Cartier pairing. We then have a well defined pairing:

$$A_1[\sigma_1, \psi_1] \times \hat{A}_2[\hat{\sigma}_1]/A_1[\sigma_1, \psi_1]^{\perp} \to \mathbb{G}_m$$

It is not hard to see that $\hat{\psi}_2(\hat{B}_2[\hat{\sigma}_2])$ is in this orthogonal, because $e_{\sigma_1}(P_1, \hat{\psi}_2 Q_2) = e_{\psi_2 \circ \sigma_1}(P_1, Q_2) = e_{\sigma_2 \circ \psi_1}(P_1, Q_2) = e_{\sigma_2}(\psi_1 P_1, Q_2) = 1$, since $\psi_1(P_1) = 0$ when $P_1 \in A_1[\sigma_1, \psi_1]$. Hence $T^{\sigma_1}_{\psi_1}$ is well defined.

We can use the compatibility of the Weil-Cartier pairings with isogenies to prove Eq. (3). Indeed, we have $\hat{\psi}_1 \hat{\sigma}_2 Q' = \hat{\sigma}_1 \hat{\psi}_2 Q' = 0$, so $e_{\psi_1}(P_1, \hat{\sigma}_2 Q')$ is well defined. And we compute: $e_{\psi_1}(P_1, \hat{\sigma}_2 Q') = e_{\sigma_2 \psi_1}(P_1, Q') = e_{\psi_2 \sigma_1}(P_1, Q') = e_{\sigma_1}(P_1, \hat{\psi}_2 Q') = e_{\sigma_1}(P_1, Q_2)$. □

Remark A.1.

- The non degenerate pairing on the left $A_1[\sigma_1, \psi_1] \times \hat{A}_2[\hat{\sigma}_1]/\hat{\psi}_2(\hat{B}_2[\hat{\sigma}_2]) \to \mathbb{G}_m$ is also non degenerate on the right whenever $\hat{\psi}_2(\hat{B}_2[\hat{\sigma}_2])$ is the full orthogonal, e.g., when $\#\hat{B}_2[\hat{\sigma}_2] = \#\hat{A}_2[\hat{\sigma}_1]$ and $\#\hat{B}_2[\hat{\sigma}_2, \hat{\psi}_2] = \#A_1[\sigma_1, \psi_1]$.
- The proof of Eq. (3) shows that $Q' \in B_2[\sigma_2 \psi_1]$ and that $T^{\sigma_1}_{\psi_1}(P_1, Q_2) = e_{\sigma_2 \psi_1}(Q') = e_{\sigma_2 \psi_1}(Q')$ which shows that $T^{\sigma_1}_{\psi_1}$ is also induced by the Weil-Cartier pairing $e_{\sigma_2 \psi_1}$, which highlight the symmetric role of the ψ_i, σ_i.
- Still by the compatibility of the Weil-Cartier pairing and isogenies, whenever σ_1 is a d-isogeny, that is we have a contragradient isogeny $\tilde{\sigma}_1$ satisfying $\sigma_1 \circ \tilde{\sigma}_1 = \tilde{\sigma}_1 \circ \sigma_1 = [d]$, we have: $e_{\sigma_1}(P, Q) = e_d(P, Q') = e_d(P', Q)$ where $\sigma_1(Q') = Q$ and $\tilde{\sigma}_1(P') = P$.

Example A.1.

- Take $E_1/\mathbb{F}_q$ an elliptic curve defined over a finite field $\mathbb{F}_q$, $\psi_1 = \psi_2 = \psi : E_1 \to E_2$ a rational isogeny, and $\sigma_i = \hat{\pi}_q - 1$ on E_i. Since $E[\hat{\pi}_q - 1] = E[\pi_q - q]$ and $E[\pi_q - 1] = E(\mathbb{F}_q)$, we obtain a non degenerate pairing on the left $E_1[\pi_q - q, \psi] \times E_1(\mathbb{F}_q)/\hat{\psi}(E_2(\mathbb{F}_q)) \to \mathbb{G}_m$. We remark that if $E_1[\psi] \subset E_1[m]$ with $q \equiv 1 \mod m$, then $E_1[\pi_q - q, \psi] = E_1(\mathbb{F}_q)[\psi]$.
  As a further special case, taking $\psi = [m]$ (where m satisfy the above condition), we obtain a non degenerate pairing on the left $e : E_1(\mathbb{F}_q)[m] \times E_1(\mathbb{F}_q)/m(E_1(\mathbb{F}_q)) \to \mathbb{G}_m$, which can be computed as $e(P, Q) = e_{\hat{\pi_q} - 1}(P, Q) = e_m(P, (\pi_q - 1)Q')$ where $mQ' = Q$. We recover the standard Tate pairing.
- Take $E/\mathbb{F}_q$ an elliptic curve, and R a primitive orientation on E by a quadratic imaginary order of discriminant Δ_R. Let ω_R be the canonical imaginary element of Definition 3.2. Since $\hat{\omega} = -\omega$, we have a non degenerate Cartier-Weil pairing $e_{\omega_R} : E[\omega_R] \times E[\omega_R] \to \mathbb{G}_m$. Since $E[\omega_R]$ is cyclic, we see that e_{ω_R} gives a cyclic self pairing.
  Now if $N(\omega_R) = m_1 m_2$, then we can construct a commutative diagram as above, by taking $\sigma_1 = \omega_R, \psi_1 = [m_1], \psi_2 = \hat{\omega_R}, \sigma_2 = [m_2]$, to obtain a non degenerate pairing on the left (hence also on the right) $E[\omega_R, m_1] \times E[\omega_R]/\omega_R E[m_2] \to \mathbb{G}_m$, which can be computed as $e_\omega(P, Q) = e_{m_1}(P, m_2 Q')$ where $\omega_R Q' = Q$.

Let P_0 be a $\mathbb{Z}[\omega_R]$-generator of $E[m_1 m_2]$. Then $Q_0 = \omega_R P_0$ is a generator of $E[\omega_R]$, $P = m_2 P_0$ is a $\mathbb{Z}[\omega_R]$-generator of $E[m]$, and $Q = \omega_R P = m_2 \omega_R P_0$ a generator of $E[\omega_R, m_1]$. By the result above, we find that $e_\omega(Q, Q) = e_{m_1}(\omega_R P, P)$. We recover the distorted pairing from Theorem 4.1 (up to a sign), and this is the self pairing built in [5].