CLIMBING AND DESCENDING TALL VOLCANOS

STEVEN D. GALBRAITH

ABSTRACT. We revisit the question of relating the elliptic curve discrete logarithm problem (ECDLP) between ordinary elliptic curves over finite fields with the same number of points. This problem was considered in 1999 by Galbraith and in 2005 by Jao, Miller, and Venkatesan. We apply recent results from isogeny cryptography and cryptanalysis, especially the Kani construction, to this problem. We improve the worst case bound in Galbraith's 1999 paper from $\tilde{O}(q^{1.5})$ to (heuristically) $\tilde{O}(q^{0.4})$ operations.

The two cases of main interest for discrete logarithm cryptography are random curves (flat volcanoes) and pairing-based crypto (tall volcanoes with crater of constant or polynomial size). In both cases we show a rigorous $\tilde{O}(q^{1/4})$ algorithm to compute an isogeny between any two curves in the isogeny class. We stress that this paper is motivated by pre-quantum elliptic curve cryptography using ordinary elliptic curves, which is not yet obsolete.

1. INTRODUCTION

Let E_0 and E_1 be ordinary elliptic curves over a finite field $\mathbb{F}_q$ with the same number of $\mathbb{F}_q$-rational points. It is a long-standing and important problem to determine if the discrete logarithm problem (ECDLP) is equally hard on the two curves. A natural approach to compare instances of ECDLP is to construct an efficiently computable group homomorphism between the two groups. This approach was studied by Galbraith [Gal99], building on work of Kohel [Koh96] for computing endomorphism rings of elliptic curves. These works made clear that a major obstruction to the equivalence of the ECDLP is the conductor gap, which we will now recall.

Let $\#E_0(\mathbb{F}_q) = q + 1 - t$ and suppose $D = t^2 - 4q = f^2 D_0$ where D_0 is the discriminant of the imaginary quadratic field $K = \mathbb{Q}(\sqrt{t^2 - 4q})$. The ring of integers of K is $\mathcal{O}_K = \mathbb{Z}[(D_0 + \sqrt{D_0})/2]$. The endomorphism rings of E_0 and E_1 are orders in K that contain $\mathbb{Z}[(D + \sqrt{D})/2]$. The integer f is called the *conductor* of the ring $\mathbb{Z}[(D + \sqrt{D})/2]$, and it is the index $[\mathcal{O}_K : \mathbb{Z}[(D + \sqrt{D})/2]]$. Let $f_i = [\mathcal{O}_K : \mathrm{End}(E_i)]$ for $i = 1, 2$. We have $f_i \mid f$. If f is divisible by a large prime N then, as noted by Kohel, the hard case of the isogeny problem for ordinary curves is when f_1 and f_2 are divisible by different powers of N, as in this case any isogeny from E_0 to E_1 (or vice versa) has degree divisible by N. We call this the case of *large prime conductor gap*. Note that $N^2 \mid |t^2 - 4q| \leq 4q$ implies $N \leq 2\sqrt{q}$.

A practical case of this question arises when constructing curves with a given number of points using the CM method, such as for pairing-based cryptography (for examples see Appendix A). In these situations, one generates a "special" elliptic curve whose endomorphism ring is $\mathbb{Z}[(D_0 + \sqrt{D_0})/2]$, but the "general"

Submitted to Algor. Num. Th. Symp.

elliptic curve with the same number of points would more likely have endomorphism ring $\mathbb{Z}[(D+\sqrt{D})/2]$. We might believe that the general case of the problem is hard, but we might have some concerns about the special case. For example, the German ECC guidelines[1] recommend that the class number of the maximal order in the endomorphism ring to be at least 200 [BSI]. However, if the conductor gap is a large prime then it is not feasible to compute any of the general elliptic curves, and we are stuck with the special curve.

Of course, in most real-world implementations of elliptic curve cryptosystems this situation does not arise, as for a randomly chosen elliptic curve we do not expect $t^2 - 4q$ to have a large square factor (we call this a *flat volcano*). However, the situation does arise in pairing-based cryptography (see Appendix A for examples of *tall volcanos*). It is also of theoretical interest to find tools to deal with the conductor gap.

The main result of Galbraith [Gal99] was an algorithm[2] to compute an isogeny between any two curves in $\tilde{O}(q^{3/2})$ field operations in the worst case (the worst case being when the conductor is divisible by a prime $N \approx \sqrt{q}$). This result doesn't allow us to conclude anything about the relative difficulty of the ECDLP, since the Pollard rho algorithm for the ECDLP runs in time $\tilde{O}(q^{1/2})$ field operations, and so clearly all instances of ECDLP are equivalent within that complexity bound. However, the merit of [Gal99] was an algorithm running in time $\tilde{O}(q^{1/4})$ field operations to compute an isogeny, as long as the conductor is *not* divisible by any large primes (which is the average case). In [Gal99] the complexity is rigorous, since a deterministic time-memory tradeoff is used, but in practice we prefer to use heuristic methods based on pseudorandom walks [GHS02, GS13]. We don't expect to do better than $\tilde{O}(q^{1/4})$ for flat volcanos without a major breakthrough.[3]

The dream result of this paper would be to compute isogenies in $\tilde{O}(q^{1/4})$ time for all cases, but anything strictly better than $\tilde{O}(q^{1/2})$ tells us something non-trivial about the (classical) hardness of ECDLP in an isogeny class.

Jao, Miller, and Venkatesan [JMV05] used these ideas to obtain a polynomial-time equivalence of the ECDLP among curves, again in the case where there is no large prime dividing the conductor gap. More precisely, they show that if one can efficiently solve ECDLP on some fixed positive proportion of curves in a given level of the isogeny volcano, one can probabilistically solve ECDLP efficiently on any given curve in that same level.

The question of the conductor gap was revisited in 2011 by Koblitz, Koblitz and Menezes [KKM11]. In Section 11.2 they summarise the conductor gap and, in opposition to the comments above about the CM method, conjecture that these ideas may "make a generic curve less secure than a special curve".

The aim of this paper is to revisit such problems. We employ recently developed tools such as the square root Vélu formula and the Kani construction, to see if we

[1] These guidelines concern traditional ECC and not pairing-based cryptography.

[2] All algorithms in this paper are probabilistic, so running times are expected values.

[3] This paper is about classical algorithms. With quantum computers, Kuperberg's algorithm solves this problem in subexponential time.

Submitted to Algor. Num. Th. Symp.

can improve on the previous results. In particular, we are influenced by Robert's observation that the Kani construction can be used to obtain an efficient representation of an isogeny of large prime degree [Rob22a].

As a taster of our result, consider the following problem (which is Problem 2 in our list in Section 5): Let E_0 and E_1 be elliptic curves over $\mathbb{F}_q$ with conductor gap of large prime degree N such that $q^{1/4} < N < \sqrt{q}$, and that are connected by an N-isogeny over $\mathbb{F}_q$. We want to compute a representation of an isogeny $\phi : E_0 \to E_1$. Traditional approaches would have complexity at least linear in N, but we show an algorithm with complexity $\tilde{O}(N^{1/2}) = \tilde{O}(q^{1/4})$ operations in $\mathbb{F}_q$. This means that if there is an efficient algorithm (e.g., complexity $\tilde{O}(q^{1/4})$) for the ECDLP on one of the curves, then there is also an algorithm for the ECDLP on the other that beats Pollard rho.

Building on that, we re-consider the general isogeny problem for ordinary curves. Our main result (see Section 8) is to reduce from $\tilde{O}(q^{3/2})$ to $\tilde{O}(q^{2/5})$ the heuristic worst-case expected cost of finding an isogeny between two given ordinary curves E_0 and E_1. We also show that for the case of main interest, namely pairing-friendly elliptic curves generated by the CM method, the isogeny problem can be solved in $\tilde{O}(q^{1/4})$ operations, which unifies this case with the typical case in elliptic curve cryptography. This latter claim is fully rigorous.

Finally, in Sections 10 and 11 we give an updated commentary on the papers by Jao, Miller, and Venkatesan [JMV05] and Koblitz, Koblitz and Menezes [KKM11].

1.1. **Acknowledgements.** This research was funded by the Ministry for Business, Innvovation and Employment in New Zealand. I thank Luca de Feo for suggestions during the early stages of the research. I also thank Damien Robert, Valerie Gilchrist, and the anonymous reviewers for constructive comments and suggestions.

2. Isogenies and volcanos

We assume the reader has basic knowledge of elliptic curves and isogenies. We use standard terminology and notation, such as in [Koh96, Gal99, Sut13]. An isogeny $\phi : E_0 \to E_1$ such that $\ker(\hat{\phi} \circ \phi) = E_0[N]$ is called an N-isogeny, where $\hat{\phi} : E_1 \to E_0$ is the dual isogeny.

We focus on ordinary elliptic curves over $\mathbb{F}_q$, namely those whose endomorphism ring is an order in an imaginary quadratic field. Many of the results also apply in the context of supersingular elliptic curves E over $\mathbb{F}_q$ with an *orientation*, which is an embedding of an order in an imaginary quadratic field into the endomorphism ring of E.

If E_0 is an ordinary elliptic curve then we consistently use the notation $\#E_0(\mathbb{F}_q) = q + 1 - t$, D_0 is the discriminant of $\mathbb{Q}(\sqrt{t^2 - 4q})$, h_0 is the class number of $K = \mathbb{Q}(\sqrt{t^2 - 4q})$, and $t^2 - 4q = f^2 D_0$ for some integer f (called the *conductor*). It is well-known (see Exercise 5.27 of Cohen [Coh93]) that $h_0 = O(\sqrt{|D_0|}\log(|D_0|))$. The class number of the order $\mathcal{O}$ in $\mathbb{Q}(\sqrt{t^2 - 4q})$ of discriminant $f^2 D_0$ is given

by

$$\frac{h_0 f}{[\mathcal{O}_K^* : \mathcal{O}^*]} \prod_{\ell | f} \left(1 - \left(\tfrac{D_0}{\ell}\right)\tfrac{1}{\ell}\right) \tag{1}$$

(see Theorem 7.24 of Cox [Cox89]).

A *level* of the isogeny volcano[4] is the set of all isomorphism classes of curves E such that $\mathrm{End}(E)$ has the same discriminant. If $\mathrm{End}(E)$ has discriminant D_0 (respectively. $t^2 - 4q$) then E is said to be on the *crater* (respectively. *floor*).

Given an ordinary elliptic curve E over $\mathbb{F}_q$, the problem of determining its level (equivalently, computing $\mathrm{End}(E)$) was first considered by Kohel [Koh96], who gave an $O(q^{1/3})$ algorithm under GRH. Bisson-Sutherland [BS11] gave a heuristic subexponential-time algorithm to compute $\mathrm{End}(E)$, and recently Robert [Rob22b] sketched a polynomial-time method once the conductor is factored. Hence, in this paper we assume that it is easy to compute the level of any given curve.

If E_0 is an ordinary elliptic curve with endomorphism ring $\mathcal{O}$ then an isogeny $\phi : E_0 \to E_1$ is *ascending* (respectively, *descending*) if the endomorphism ring of E_1 is a strict superset (respectively, subset) of $\mathcal{O}$. An isogeny is *horizontal* if the endomorphism ring of both curves is the same.

One usually considers volcano structures in an isogeny graph where the edges correspond to isogenies of a fixed degree ℓ, but we will be a bit more general and allow edges for all divisors of the conductor f.

2.1. **Elkies primes.** An *Elkies prime* for an elliptic curve E_0 over $\mathbb{F}_q$ with $q+1-t$ points is a (small) odd prime ℓ that splits (and is non-ramified) in $\mathbb{Q}(\sqrt{t^2-4q})$ (i.e., $(\frac{D_0}{\ell}) = 1$). Recall that the characteristic polynomial of the q-power Frobenius map π on E_0 is $x^2 - tx + q$. If ℓ is an Elkies prime then

$$x^2 - tx + q \equiv (x - \alpha)(x - \beta) \pmod{\ell}$$

for some distinct $\alpha, \beta \in \mathbb{Z}_\ell^*$. Then there is a pair of points $(P_0, Q_0) \in E_0^2$ that generate $E_0[\ell]$ and are such that $\pi(P_0) = [\alpha]P_0$ and $\pi(Q_0) = [\beta]Q_0$.

We know that for fixed (q, t), asymptotically half the primes are Elkies primes. Hence it is natural to conjecture that there are about $\frac{1}{2}X/\log(X)$ Elkies primes up to X for any given (q, t). Some results on the distribution of Elkies primes are given by Shparlinski and Sutherland [SS15]. They show that for "most (q, t)", at least $1/3$ of the primes in a dyadic interval $[X, 2X]$ are Elkies primes.

In one special case, which arises in our applications, one can prove rigorous results about the number of Elkies primes of size polynomial in $\log(q)$. Specifically, suppose the discriminant D_0 is bounded, but we wish to have many Elkies primes of size $O(\log(q))$ for large q (e.g., $D_0 = -3$ for pairing-based cryptography). Then it suffices to find small primes ℓ in an arithmetic progression such that $(\frac{D_0}{\ell}) = 1$. This condition is defined by congruences modulo a power of 2 and modulo the primes dividing $|D_0|$. For example when $D_0 = -3$ then $\ell \equiv 1 \pmod{12}$ always satisfies $(\frac{D_0}{\ell}) = 1$. Rigorous results on primes in arithmetic progressions,

[4] In this paper we are concerned with computing isogenies between any two given curves in an isogeny class, so we take as many isogeny degrees as needed so our volcanos are connected.

such as Bennett et al [BMO+18] and references therein, show that there exist sufficiently many Elkies primes $\ell = O(\log(q))$ in this case. Precisely, Corollary 1.7 of [BMO+18] implies that, when $|D_0| < 10^5$, taking $X = O(\log(N)\log\log(N)^2)$ suffices to ensure there are at least $2\log(N)$ Elkies primes up to X.

To simplify the calculations in our main result we make the strong heuristic assumption that the i-th Elkies prime is upper bounded by $c+2i\log(i)$ for some constant c, but this could be relaxed to an upper bound of the form $c_1 + c_2 i^{c_3}\log(i)^{c_4}$ for some constants $c_1, c_2, c_3, c_4 > 1$ and we would still get meaningful results.

2.2. **Computing large degree isogenies.** Given an elliptic curve E_0 over a finite field $\mathbb{F}_q$ and a large integer N coprime to q, a natural problem is to compute one or more $\mathbb{F}_q$-rational cyclic N-isogenies from E_0 (if they exist). It is known from the elliptic curve point counting literature (see for example Proposition VII.2 of [BSS99]) that for an ordinary elliptic curve E_0 with $j(E_0) \notin \{0, 1728\}$ and a prime N co-prime to q then there are either $0, 1, 2$ or $N+1$ $\mathbb{F}_q$-rational cyclic N-isogenies from E_0. The first and third cases are when $(\frac{t^2-4q}{N}) = -1$ or $+1$ respectively, and the last case occurs only when N divides the index $[\mathrm{End}(E_0) : \mathbb{Z}[\pi]]$

Lemma 1. *Let $q > 3$ be prime. The kernel of an $\mathbb{F}_q$-rational isogeny of degree N coprime to q is defined over a field of degree at most N.*

A method to compute an N-isogeny is to generate a random N-torsion point on E_0 over some extension. First, work out the smallest k such that $N \mid \#E_0(\mathbb{F}_{q^k})$, then generate a random point $P \in E_0(\mathbb{F}_{q^k})$, and finally multiply by a suitable cofactor to get a point of order N. One can then compute the isogeny using Vélu's algorithm. Since we know the isogeny is $\mathbb{F}_q$-rational we have $k < N$ by Lemma 1.

When there are $N+1$ possible isogenies then the above method generates any one of them, and with probability at least $(N-1)/(N+1)$ it is descending. When there is only one possible isogeny then there is a unique subgroup of order N for the smallest k. When there are two possible isogenies then the point must be an eigenvector for Frobenius, and this is arranged by applying $\pi - \lambda$ or $\pi - \mu$ where λ, μ are the distinct eigenvalues of Frobenius on the N-torsion.

We now determine the complexity of this method. A simple way to compute k in $O(N)$ operations modulo N is to use the recurrence formula to compute $\#E_0(\mathbb{F}_{q^k}) \pmod N$ sequentially and to stop when $N \mid \#E_0(\mathbb{F}_{q^k})$.

Computing random $P \in E_0(\mathbb{F}_{q^k})$ is dominated by the cost of a square-root in $\mathbb{F}_{q^k}$, which can be done in $\tilde{O}(k^2\log(q))$ operations in $\mathbb{F}_q$, using quasi-linear field multiplication algorithms for $\mathbb{F}_{q^k}$. When $N, k < \sqrt{q}$ this complexity is $\tilde{O}(q)$. Here $\tilde{O}(x)$ means $O(x\log(x)^c)$ for some constant c.

Then we compute $[\#E_0(\mathbb{F}_{q^k})/N]P$, which again requires $\tilde{O}(k^2\log(q))$ operations in $\mathbb{F}_q$. Finally, we run Vélu's algorithm on E_0 using the point of order N in $E_0(\mathbb{F}_{q^k})$. This takes $O(N)$ operations in $\mathbb{F}_{q^k}$, which is again $\tilde{O}(k^2\log(q))$ operations in $\mathbb{F}_q$. Square-root Vélu [BDFLS20] doesn't change the overall complexity, as it only reduces the second stage to $\tilde{O}(N^{1.5})$ operations, and so the cost is still dominated by finding the kernel generator.

3. THE KANI CONSTRUCTION

Let E_0 be an elliptic curve with finite subgroups H_1 and H_2 of co-prime order. Let $\phi : E_0 \to E_1 = E_0/H_1$ and $\gamma : E_0 \to E_2 = E_0/H_2$. We have $\deg(\phi) = \#H_1$ and $\deg(\gamma) = \#H_2$. Let ϕ' have kernel $\gamma(H_1)$ and γ' have kernel $\phi(H_2)$. Then ϕ' and γ' both map to a curve which we call E_3 and we can choose ϕ' and γ' such that $\gamma' \circ \phi = \phi' \circ \gamma$. This gives the standard SIDH square, which Kani calls an "isogeny diamond configuration". Let $\widehat{\phi'}$ and $\widehat{\gamma'}$ be the dual isogenies. We have the following diagram.

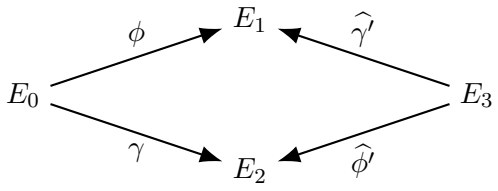

Kani defines a map

$$F : E_0 \times E_3 \to E_1 \times E_2$$

by

$$F(X, Y) = (\phi(X) - \widehat{\gamma'}(Y), \gamma(X) + \widehat{\phi'}(Y)).$$

One may represent F as the matrix

$$\begin{pmatrix} \phi & -\widehat{\gamma'} \\ \gamma & \widehat{\phi'} \end{pmatrix}.$$

Let $M = \#H_1 + \#H_2$ and let P_0, Q_0 be a basis for $E_0[M]$. Note that $\gcd(M, \#H_1) = \gcd(M, \#H_2) = \gcd(\#H_1, \#H_2) = 1$. Kani [Kan97] proves that F is an isogeny of polarized abelian varieties. Indeed, it is an (M, M)-isogeny, meaning the kernel is a Weil-isotropic subgroup of order M^2 and exponent M. We refer to Kani [Kan97] and Robert [Rob23] for details about polarizations. For algorithms to evaluate F we refer to Lubicz and Robert [LR22] and the references therein.

One can slo consider the adjoint map $\bar{F} : E_1 \times E_2 \to E_0 \times E_3$ given by

$$\bar{F}(X, Y) = (\hat{\phi}(X) + \hat{\gamma}(Y), -\gamma'(X) + \phi'(Y)).$$

This also is an isogeny of polarized abelian varieties. The map $\bar{F}$ is represented by the matrix

$$\begin{pmatrix} \hat{\phi} & \hat{\gamma} \\ -\gamma' & \phi' \end{pmatrix}.$$

It follows that $\bar{F} \circ F$ maps (X, Y) to (MX, MY). Indeed, as isogenies, $\bar{F} \circ F = [M]$, and we call F an M-isogeny.

It follows that

$$\begin{aligned} \ker(F) &= \bar{F}((E_1 \times E_2)[M]) \\ &= \{(\hat{\phi}(X) + \hat{\gamma}(Y), -\gamma'(X) + \phi'(Y)) : X \in E_1[M], Y \in E_2[M]\}. \end{aligned}$$

There are lots of alternative ways to compute the kernel.

Submitted to Algor. Num. Th. Symp.

Lemma 2. *Let $f = \gamma' \circ \phi = \phi' \circ \gamma$. Then*

$$\ker(F) = \{(\hat{\phi}(P), -\gamma'(P)) : P \in E_1[M]\},$$

$$\ker(F) = \{([\#H_1]P, -f(P)) : P \in E_0[M]\},$$

and, writing $\psi = [-\#H_2^{-1}]\hat{f}$

$$\ker(F) = \{(\psi(P), P) : P \in E_3[M]\}.$$

Proof. The first of these is stated by [Rob23]. It is just a specialisation of the earlier formula $\ker(F) = \bar{F}((E_1 \times E_2)[M])$ together with noting that the restriction of $\bar{F}$ to $E_1[M] \times \{0\}$ is injective.

The second is proved by noting that $\phi : E_0[M] \to E_1[M]$ is injective (since ϕ has degree co-prime to M) and substituting $P = \phi(P')$ where $P' \in E_0[M]$ into the previous formula. This is also given as Lemma 4 of Maino et al [MMP+23].

A version of the third is given in Kani [Kan97] in the proof of Corollary 2.4. It follows by substituting $P = [-\#H_2^{-1}]\widehat{\gamma'}(P')$ into the first formula where $P' \in E_3[M]$. □

We make a remark about the field of definition. Suppose E_0, E_1 and $\phi : E_0 \to E_1$ are all defined over $\mathbb{F}_q$. We will always choose γ to be over $\mathbb{F}_q$, so E_2 is over $\mathbb{F}_q$. It follows that E_3 is also over $\mathbb{F}_q$, and since F is built from ϕ and γ it is also defined over $\mathbb{F}_q$. This also follows by considering the Galois group acting on the kernel of F.

One can extend the ideas of this section to higher dimensions. It is common that γ needs to have some fixed degree m. In order to efficiently find such a map one writes m as a sum of $g \in \{2, 3, 4\}$ squares, and chooses $\gamma : E_0^g \to E_0^g$ to be an endomorphism on a product of E_0. For example, if $m = m_1^2 + m_2^2$ then $\gamma = \left(\begin{smallmatrix} m_1 & -m_2 \\ m_2 & m_1 \end{smallmatrix}\right) : E_0^2 \to E_0^2$ is an isogeny. Writing $\bar{\gamma} = \left(\begin{smallmatrix} m_1 & m_2 \\ -m_2 & m_1 \end{smallmatrix}\right)$ gives $\bar{\gamma} \circ \gamma = [m]$ on E_0^2.

For the Kani construction to be applied then it is necessary to diagonally extend ϕ to be a map from E_0^g to E_1^g. The formulas for the kernel extend naturally. We give the details for the case $g = 2$. Let $\gamma' : E_1^2 \to E_1^2$ be defined on E_1^2 by the same matrix as γ. The map $F : E_0^2 \times E_1^2 \to E_1^2 \times E_0^2$ is an (M, M, M, M)-isogeny with kernel

$$\ker(F) = \{(\hat{\phi}(P), \hat{\phi}(Q), -\gamma'(P, Q)) : P, Q \in E_1[M]\} \subseteq (E_0^2 \times E_1^2)[M]. \quad (2)$$

Similarly, the kernel of the dual map to F is given by

$$\ker(\bar{F}) = \{(\phi(P), \phi(Q), \gamma(P, Q)) : P, Q \in E_0[M]\}.$$

For future reference we note that setting $P = \widehat{\phi}(P')$ and $Q = \widehat{\phi}(Q')$ for $P', Q' \in E_1[M]$ in the above leads to the formula

$$\ker(\bar{F}) = \{([\#H_1]P', [\#H_1]Q', \gamma(\widehat{\phi}(P'), \widehat{\phi}(Q'))) : P', Q' \in E_1[M]\}. \quad (3)$$

4. MAIN TOOL

We now present the main tool that is used to obtain our results. This is inspired by the work of Robert [Rob22a] on computing isogenies of any degree, combined with various techniques that have been recently used in cryptanalysis of isogeny-based cryptosystems. It also builds on ideas from [QKL+21, DLRW23] and is a variant of a result by Castryck et al [CHM+23].

Let E_0 and E_1 be elliptic curves over $\mathbb{F}_q$ that are connected by an isogeny $\phi : E_0 \to E_1$ over $\mathbb{F}_q$ of degree N. The integer N need not be prime, but we require that N is not divisible by any small primes (say, primes smaller than $4\log(N)\log\log(N)$). The problem that we wish to solve is to obtain a representation of ϕ given only the two curves E_0 and E_1 and the integer N. The phrase "representation of ϕ" is the same as explained by Leroux [Ler22] and Robert [Rob22a]: it means that we can evaluate ϕ on any given point $P \in E_0(\mathbb{F}_{q^t})$ in time polynomial in $t\log(q)$ and $\log(N)$. This is in contrast to Vélu's algorithm, which has complexity at least linear in N.

We require exponential time to compute this representation of ϕ. However the complexity grows proportionally to $N^{1/2}$, which is better than anything one would expect from Vélu-like techniques.

The basic idea is to choose an integer m that is a sum of two or four squares and is such that $M = N + m$ is smooth. Then we follow the Kani/Robert technique to construct an M-isogeny $F : E_0^g \times E_1^g \to E_1^g \times E_0^g$ where $g \in \{2, 3, 4\}$. To determine the kernel of F we need to know $\phi(P)$ and $\phi(Q)$ where P, Q generate $E_0[M]$. The main challenge is how to determine this information. First, as done by Castryck et al [CHM+23], we choose P and Q to be Frobenius eigenspaces. Since ϕ is defined over $\mathbb{F}_q$, the images $\phi(P)$ and $\phi(Q)$ will also be Frobenius eigenspaces. By using the Weil pairing we can determine $\phi(P)$ and $\phi(Q)$ up to a single unknown scalar. However this scalar is still too large to guess, so the second trick is to split F into two parts following the "meet-in-middle" trick used in several works including [QKL+21, DLRW23]. It then suffices to guess one integer of size around $\sqrt{N}$, and the isogeny diagram will be consistent if and only if we have the correct guess. Everything else is polynomial in $\log(N)$ and $\log(p)$.

As is standard, to evaluate $\phi(P)$ on any point $P \in E_0(\mathbb{F}_{q^l})$ one computes $F(P, 0, \ldots, 0)$ and projects the image point to the first component. This works because F acts as a $2g \times 2g$ matrix, with ϕ in the top left entry.

For simplicity (and to minimise the number of heuristics) we present the four squares case, but in practice one would prefer to do the two squares version.

The result relies on a heuristic assumption about the size of Elkies primes (which, as we have seen, can be avoided when $|D_0|$ is bounded). We give an example in Appendix A as evidence that such heuristics are not a problem in practice.

Assumption 1: There exists a constant c such that for all prime powers q and integers t such that $|t| < 2\sqrt{q}$. Let $N < q$ be an integer and $s = \lfloor 2\log(N) \rfloor$. Let $3 < \ell_1 < \ell_2 < \cdots < \ell_s$ be the smallest s elements in the sequence of distinct Elkies primes (split primes in $\mathbb{Q}(\sqrt{t^2 - 4q})$) strictly greater than 3. Then $\ell_i < c + 2i\log(i)$ for $1 \leq i \leq s$.

Submitted to Algor. Num. Th. Symp.

Let $\{\ell_1, \ell_2, \ldots, \ell_s\}$ be as in Assumption 1. We will consider subsets $S \subseteq [s]$ such that

$$A_S = \prod_{i \in S} \ell_i$$

satisfies $A_S^2 \leq N/2$. Since each ℓ_i has average size roughly $s\log(s)$ we expect $|S|$ to have size around $\frac{1}{2}\log(N)/\log(s\log(s))$. By assumption 1, we can choose S such that

$$N/(2(c + 2s\log(s))^2) < A_S^2 < N/2.$$

Note that $\gcd(N, A_S) = 1$ for all choices of S.

For each such S choose the smallest n such that $3^n A_S^2 > N$. Note that $3^n A_S^2 < 3N$ and so $m = 3^n A_S^2 - N < 2N$. We have $3^n < 6(c + 4\log(N)\log\log(N))^2$ and that $m = 3^n A_S^2 - N$ is an even integer that is not divisible by 3. One can write m as a sum of four squares. For the two squares version, one can write m as a sum of two squares with probability $O(1/\sqrt{\log(N)})$ (assuming m is distributed close to uniformly), and hence choosing $O(\sqrt{\log(N)})$ subsets S allows to find such an m. We omit the analysis (see discussion and references in Section 10 of Castryck and Decru [CD23]).

We now give the details. The main application is ordinary curves, since better methods are available in the supersingular case. But we state the result in the general case.[5]

Theorem 1. *Let E_0 and E_1 be elliptic curves over $\mathbb{F}_q$ that are connected by an isogeny $\phi : E_0 \to E_1$ over $\mathbb{F}_q$ of degree N. Let $N > 1000$ be such that N is not divisible by any prime smaller than $4\log(N)\log\log(N)$. Then there is a (heuristic) algorithm to compute a representation of ϕ that can be evaluated on points in time polynomial in $\log(N)$ and the size of the representation of the points. The expected complexity of the algorithm to compute the representation is $\tilde{O}(N^{1/2})$ operations in $\mathbb{F}_q$.*

Proof. We give the algorithm as Algorithm 1.

The first step is to determine m and M such that $M = N + m$ is smooth and m can be written as a sum of four integer squares. This is the heuristic part and it has been discussed above. Note that N is coprime to 6. Let $\ell_1, \ell_2, \ldots, \ell_s$ be the smallest $s = \lfloor 2\log(N) \rfloor$ elements in the sequence of distinct Elkies primes. Let $S \subseteq [s]$ be such that

$$A_S = \prod_{i \in S} \ell_i$$

satisfies $N/(2(c+4\log(N)\log\log(N))^2) < A_S^2 < N/2$. Note that $\gcd(N, A_S) = 1$ for all choices of S. Then there is an integer n such that $3^n A_S^2 > N$ and $3^n = O((\log(N)\log\log(N))^2)$. Then $m = 3^n A_S^2 - N > 0$ can be written as a sum of four integer squares $m = m_1^2 + m_2^2 + m_3^2 + m_4^2$. Let $\gamma : E_0^4 \to E_0^4$ be the isogeny defined by the matrix representing mutiplication by the Hamilton quaternion $m_1 + m_2 i + m_3 j + m_4 k$; see Robert [Rob23]. The same matrix also defines an isogeny

[5] In the supersingular case we require the characteristic polynomial of Frobenius to be irreducible, which is the case when q is a prime, otherwise there are no Elkies primes.

$\gamma : E_1^4 \to E_1^4$. Write $\bar{\gamma} : E_0^4 \to E_0^4$ for the isogeny given by the transpose of the matrix, so that $\bar{\gamma} \circ \gamma = [m]$.

We now explain the meet-in-the-middle step.

We have $M = N + m = 3^n A_S^2$. Let $M_1 = 3^{\lceil n/2 \rceil} A_S$ and $M_2 = 3^{\lfloor n/2 \rfloor} A_S$. Then $M = M_1 M_2$ and $M_2 \mid M_1$. We have $M_1 \leq 3M_2$ and $M_1 M_2 \leq 3N$, so $M_1 \leq 3\sqrt{N}$.

The unknown N-isogeny $\phi : E_0 \to E_1$ extends to a diagonal (N, N, N, N)-isogeny $\Phi : E_0^4 \to E_1^4$, and we have $\Phi \circ \gamma = \gamma \circ \Phi$. The Kani machinery therefore shows the existence of an M-isogeny

$$F : E_0^4 \times E_1^4 \to E_1^4 \times E_0^4$$

with kernel

$$\{(\widehat{\phi}(P), \widehat{\phi}(Q), \widehat{\phi}(R), \widehat{\phi}(S), -\gamma(P, Q, R, S)) : P, Q, R, S \in E_1[M]\}.$$

As mentioned in equation (3), we also have

$$\begin{aligned} \ker(\bar{F}) \quad = \quad & \{([N]P, [N]Q, [N]R, [N]S, \gamma(\widehat{\phi}(P), \widehat{\phi}(Q), \widehat{\phi}(R), \widehat{\phi}(S))) \\ & : P, Q, R, S \in E_1[M]\}. \end{aligned}$$

We have $F : E_0^4 \times E_1^4 \to E_1^4 \times E_0^4$ being a M-isogeny. We can write F as

$$E_0^4 \times E_1^4 \xrightarrow{F_1} B \xleftarrow{\bar{F}_2} E_1^4 \times E_0^4$$

for some abelian variety B, where $F = F_2 \circ F_1$ and each F_i is an M_i-isogeny. Then

$$\ker(F_1) = \ker(F) \cap (E_0^4 \times E_1^4)[M_1]$$

and

$$\ker(\bar{F}_2) = \ker(\bar{F}) \cap (E_1^4 \times E_0^4)[M_2].$$

Since ϕ and γ are defined over $\mathbb{F}_q$, it follows that F is defined over $\mathbb{F}_q$ and so F_1, F_2 and the abelian variety B are defined over $\mathbb{F}_q$.

By equations (2) and (3), to compute $\ker(F_1)$ and $\ker(\bar{F}_2)$ it suffices to know $\widehat{\phi}(P_1), \widehat{\phi}(Q_1)$ for some basis P_1, Q_1 of $E_1[M_1]$. For each prime $\ell \mid A_S$ let $P_{1,\ell}, Q_{1,\ell}$ be a Frobenius eigenbasis for $E_1[\ell]$. (The primes 2 and 3 are not necessarily Elkies primes, and are handled differently.) So there are some integers α_ℓ, β_ℓ such that $\pi(P_{1,\ell}) = [\alpha_\ell] P_{1,\ell}$ and $\pi(Q_{1,\ell}) = [\beta_\ell] Q_{1,\ell}$, where π is the q-power Frobenius map on E_1. Note that $P_{1,\ell}$ and $Q_{1,\ell}$ are defined over an extension of $\mathbb{F}_q$ of degree at most $O(\ell)$, which is polynomially-bounded since we are assuming $\ell < c + 2s\log(s)$ with $s = O(\log(N))$. Since $\widehat{\phi}$ commutes with π we have $\pi(\widehat{\phi}(P_{1,\ell})) = [\alpha_\ell]\widehat{\phi}(P_{1,\ell})$ and $\pi(\widehat{\phi}(Q_{1,\ell})) = [\beta_\ell]\widehat{\phi}(Q_{1,\ell})$.

Let $P_{0,\ell}, Q_{0,\ell}$ be a Frobenius eigenbasis for $E_0[\ell]$, with $\pi(P_{0,\ell}) = [\alpha_\ell] P_{0,\ell}$ and $\pi(Q_{0,\ell}) = [\beta_\ell] Q_{0,\ell}$. It follows that $P_{0,\ell} = [\lambda]\widehat{\phi}(P_{1,\ell})$ and $Q_{0,\ell} = [\mu]\widehat{\phi}(Q_{1,\ell})$, for some $\lambda, \mu \in \mathbb{Z}_\ell^*$. Using the Weil pairing we have

$$e_\ell(P_{1,\ell}, Q_{1,\ell}) = e_\ell(P_{0,\ell}, Q_{0,\ell})^{\lambda\mu \deg(\phi)}.$$

Hence we can compute $\lambda\mu \pmod{\ell}$. Multiplying $P_{0,\ell}$ by $(\lambda\mu)^{-1} \pmod{\ell}$ gives $P_{0,\ell} = \mu^{-1}\widehat{\phi}(P_{1,\ell})$ and $Q_{0,\ell} = \mu\widehat{\phi}(Q_{1,\ell})$. All these computations are performed in

time polynomial in $\log(N)$. It suffices to guess the value μ, and there are at most $\varphi(\ell) = \ell - 1$ choices. This idea has been used in several papers, such as Castryck et al [CHM+23].[6]

We also need to address the 2-power and 3-power parts of M_1. Note that M_1 is odd, but following Appendix F.3 of Dartois et al [DLRW23] it is convenient to know the image of $\widehat{\phi}$ on $E_1[4]$ in order to efficiently determine whether or not the isogenies F_1 and $\bar{F}_2$ do actually meet-in-the-middle. We cannot assume that 2 and 3 are Elkies primes, but there are only polynomially-many guesses required to determine the image under ϕ of the points of 2^2-power and $3^{\lceil n/2 \rceil}$-power order dividing M.

If the guess is correct then the images of F_1 and $\bar{F}_2$ will be the same abelian variety. If the images of the isogenies do not match, then we repeat with another guess.[7]

Once we have a diagram

$$E_0^4 \times E_1^4 \xrightarrow{F_1} B \xleftarrow{\bar{F}_2} E_1^4 \times E_0^4$$

then we already have an efficient representation of F and can hence evaluate ϕ. The complexity of evaluating F is $O(\log(N)^c)$ operations for some constant c, since we need to compute $O(\log(N))$ consecutive ℓ-isogenies, where $\ell = O(\log(N))$.

Finally we address the complexity of Algorithm 1. As we have explained, we need to determine $\widehat{\phi}(P_{1,\ell})$ and $\widehat{\phi}(Q_{1,\ell})$ correctly for each $\ell | A_S$ (and some polynomially-bounded powers of 2 and 3). This requires $O(\varphi(M_1)) = O(\sqrt{N})$ iterations/guesses. Each iteration involves operations on points of order $\tilde{O}(\log(N))$ and defined over a field extension of degree $\tilde{O}(\log(N))$. Then we compute the sequence of ℓ-isogenies to get a sequence of Abelian varieties over $\mathbb{F}_q$. The isogeny computations are polynomial-time in $\log(N)$. If the isogenies meet in the middle then we have an isogeny $F : E_0^4 \times E_1^4 \to E_1^4 \times E_0^4$ as desired. Hence the total number of operations in $\mathbb{F}_q$ is $\sqrt{N}$ times a polynomial in $\log(N)$ and $\log(q)$. $\square$

Note that most of Algorithm 1 is deterministic. The only places where probabilistic algorithms are used are to generate random points on elliptic curves and to factorise polynomials (e.g., in lines 3, 4, 10 and 11). With $\tilde{O}(q^{1/4})$ field operations it may be possible to make the algorithm deterministic, but we do not consider this question further.

The pair $F_1, \bar{F}_2$ is a representation of F and hence of ϕ. However, one might want to compute F in one step, without having to handle general Abelian varieties B. Here is how to determine $\ker(F)$ directly: For each $\ell \mid M_2$ compute a basis for $(E_1^4 \times E_0^4)[\ell]$ and map it through $\bar{F}_2$ to get a set of generators for $\ker(F_2) : B \to E_1^4 \times E_0^4$. Then pull these generators back through F_1 (since ℓ is small this can be

[6] Note that Castryck et al [CHM+23] use a self-pairing $e(Q_1, Q_1) = e(Q_0, Q_0)^{\mu^2}$ to determine μ exactly. Self-pairings can be used in our setting to get a small improvement in some cases, but they do not change the worst-case result.

[7] In the applications we have in mind, there will be a unique solution and hence only one match. But in theory one could have more than one solution.

Algorithm 1 Construct representation of an unknown N-isogeny

INPUT: E_0, E_1, N

OUTPUT: Representation of an N-isogeny $F : E_0^4 \times E_1^4 \to E_1^4 \times E_0^4$.

Compute m, m_1, m_2, m_3, m_4, S, A_S, n, M, M_1 and M_2 as specified in the proof, so $M = 3^n A_S^2 = M_1 M_2$ and $m = M - N = m_1^2 + m_2^2 + m_3^2 + m_4^2$.
for each prime $\ell > 3$, $\ell \mid A_S$ **do**
    Compute eigenbases $P_{1,\ell}, Q_{1,\ell}$ for $E_1[\ell]$ and $P_{0,\ell}, Q_{0,\ell}$ for $E_0[\ell]$
    Determine eigenvalues α_ℓ, β_ℓ
    Compute pairings $e_\ell(P_{1,\ell}, Q_{1,\ell})$ and $e_\ell(P_{0,\ell}, Q_{0,\ell})$
    Solve for $\lambda\mu$ and set $P_{0,\ell} = [(\lambda\mu)^{-1} \pmod{\ell}]P_{0,\ell}$
end for
Set $P_1 = \sum_{\ell \mid A_S} P_{1,\ell}$, $Q_1 = \sum_{\ell \mid A_S} Q_{1,\ell} \in E_1[A_S]$
Set $P_0 = \sum_{\ell \mid A_S} P_{0,\ell}$, $Q_0 = \sum_{\ell \mid A_S} Q_{0,\ell} \in E_0[A_S]$
Choose a basis $\{R_1, R_2\}$ of $E_1[2^2 3^n]$ and set $P_1 = P_1 + R_1, Q_1 = Q_1 + R_2$
for each $\mu \in \mathbb{Z}_{A_S}^*$ and each basis $\{R_1, R_2\}$ of $E_0[2^2 3^n]$ **do**
    Set $P_0' = [\mu^{-1}]P_0 + R_1$ and $Q_0' = [\mu]Q_1 + R_2$
    Compute $\ker(F_1)$ and $\ker(\bar{F}_2)$ using P_0', Q_0', P_1, Q_1
    Compute $F_1 : E_0^4 \times E_1^4 \to B_1$ and $\bar{F}_2 : E_1^4 \times E_0^4 \to B_2$.
    if $B_1 \cong B_2$ **then**
        return F_1, F_2
    end if
end for

done by choosing random points in $(E_0^4 \times E_1^4)[\ell^2]$ and mapping them through F_1 and then solving multidimensional discrete log problems in B).

5. COMPUTATIONAL PROBLEMS

The main goal of this paper is to give an algorithm to compute an isogeny between two given ordinary elliptic curves with the same number of points, however there are a number of special cases and related computational problems that we study as warm-up and applications of our main result:

Problem 1: Given $E_0/\mathbb{F}_q$ on the crater (e.g., $j(E_0) = 0$) and a large prime N dividing the conductor, compute the image curve E_1 of a descending N-isogeny $\phi : E_0 \to E_1$ and be able to evaluate ϕ on chosen points.

Problem 2: Given $E_0, E_1/\mathbb{F}_q$ connected by an isogeny of large prime degree $N > 2q^{1/4}$ dividing the conductor, compute a representation of $\phi : E_0 \to E_1$. (Note that $N > 2q^{1/4}$ means N^2 does not divide the conductor.)

Problem 3: Same as previous problem, but when N^2 divides the conductor. (One can also consider variant of problem 1 in this setting, where the goal is to compute E_1 in the middle.)

Problem 4: Given $E_0/\mathbb{F}_q$ on the crater and $E_1/\mathbb{F}_q$ with large prime N dividing the conductor gap, compute a representation of an isogeny $\phi : E_0 \to E_1$. Here $N \mid \deg(\phi)$ but the degree is not necessarily equal to N as E_0 might not be "directly above" E_1 in the volcano.

We picture the four problems below.

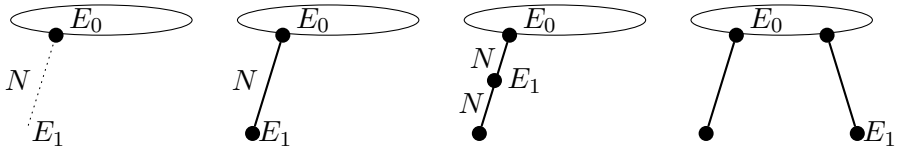

We defer Problem 1 to last. Problems 2, 3 and 4 are warm-ups to our main result.

Since we are focussed in this paper on exponential-time algorithms, we may assume the conductor has been factored (factoring t^2-4q can be done in $O(q^{1/4+o(1)})$ time using exponential-time methods [Pol74, Str77]). We also assume that $\mathrm{End}(E)$ is determined for all curves under consideration, as this is also easy relative to the problems we study. Bisson-Sutherland [BS11] give a heuristic subexponential-time algorithm to compute $\mathrm{End}(E)$, and recently Robert [Rob22b] sketched a polynomial-time method once the conductor is factored.

6. Warm-up: Solving Problems 2 and 3

For Problem 2 we have ordinary elliptic curves E_0, E_1 over $\mathbb{F}_q$ with conductor gap a large prime N, and such that there is an (unknown) isogeny $\phi : E_0 \to E_1$ defined over $\mathbb{F}_q$ of degree N. Note that $N < 2\sqrt{q}$. We assume $N > 2q^{1/4}$ so that N^2 does not divide the conductor. Our goal is to compute a representation of ϕ.

We consider a direct approach based on the Kani construction and our main tool Theorem 1. As explained, the basic idea is to choose a sum of $g \in \{2, 3, 4\}$ squares m such that $M = N + m$ is smooth, and then to follow the Kani technique to construct an isogeny $F : E_0^g \times E_1^g \to E_1^g \times E_0^g$.

If N is small (say $N < 1000$) or divisible by small primes, then we can use standard methods. For the hard case, N satisfies the requirements of Theorem 1. Hence, we can compute a representation of the isogeny $\phi : E_0 \to E_1$ in $\tilde{O}(N^{1/2})$ operations in $\mathbb{F}_q$. Since $N = O(q^{1/2})$ this method has complexity $\tilde{O}(q^{1/4})$. The isogeny F, and hence ϕ, can be evaluated in time polynomial in $\log(N)$ and $\log(q)$.

We now consider Problem 3. This is when N^2 divides the conductor and so there are three levels in the volcano. The three cases of interest are pictured below.

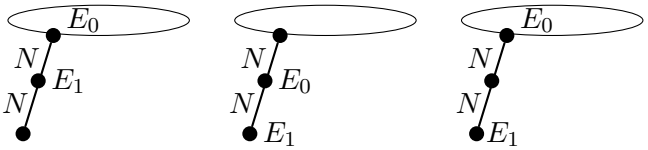

Since $N = O(q^{1/4})$ the first and second cases are easily handled using Theorem 1 as in our solution to Problem 2. The complexity is $\tilde{O}(q^{1/8})$ operations in $\mathbb{F}_q$.

For the third case we do not know the intermediate curve. Instead of computing two N-isogenies we just go all the way with an N^2-isogeny, using Theorem 1. This gives cost $\tilde{O}((N^2)^{1/2}) = \tilde{O}(q^{1/4})$ operations in $\mathbb{F}_q$. For our general result we handle these cases differently.

7. Solving Problem 4: Walking the crater

Now we are given $E_0/\mathbb{F}_q$ on the crater and $E_1/\mathbb{F}_q$ with large prime dividing the conductor gap, but E_0 is not directly "above" E_1. Let $N = [\mathrm{End}(E_0) : \mathrm{End}(E_1)]$, which we do not assume to be prime. Then an isogeny $\phi : E_0 \to E_1$ has $N \mid \deg(\phi)$ but the degree is not necessarily equal to N. This problem is considered in [Gal99] and the solution is to go "up" from E_1 to the crater and then apply a meet-in-the-middle algorithm on the crater. The complexity stated in [Gal99] is $\tilde{O}(N^3 + h_0^{1/2})$ operations in $\mathbb{F}_q$, but as explained in Section 2.2 one can actually do it in $\tilde{O}(N^2 + h_0^{1/2})$ operations in $\mathbb{F}_q$ (as long as there is no large prime $\ell \mid N$ such that $\ell^2 \mid f$).

To get improved results we introduce a new approach to dealing with this problem. Instead of going up from E_1, we try all curves E_0' on the crater until we find the one that is directly above E_1. Let h_0 be the class number of the maximal order, which is the size of the crater. In the worst case we will need to try h_0 curves.

The precise algorithm is as follows: Starting with E_0 we enumerate all h_0 curves E_0' on the crater. To do this efficiently one needs to pre-compute the class group structure and the relations among the generators.[8] An easy case is when the class group is cyclic and generated by a small prime ideal. More generally one can use tools from the CSIDH cryptosystem [CLM+18] and the recent Clapoti algorithm by Page and Robert [PR23]. Then, for each candidate curve E_0', apply our main tool Theorem 1 to the pair (E_0', E_1) with degree N. When E_0' is directly above E_1 the tool succeeds and we have a representation of an isogeny from E_0' to E_1. When E_0' is not connected to E_1 by an N-isogeny then the method fails. Finally, compute an isogeny from E_0 to E_0' using standard methods [Gal99, GS13].

The (heuristic) complexity is $\tilde{O}(h_0 N^{1/2})$ operations in $\mathbb{F}_q$ to find E_0' and the isogeny from E_0' to E_1. Here we are using the fact that we are provided with a curve E_0 on the crater, and so in $O(h_0 \log(q)^{O(1)})$ operations in $\mathbb{F}_q$ we can list all curves on the crater by using the action of the class group. The complexity is $\tilde{O}(h_0^{1/2})$ for computing the isogeny from E_0 to E_0', which is always less than the complexity of the first part.

In summary, we have proved the following.

Theorem 2. *Let E_0 be a curve on the crater (i.e., End(E) has discriminant D_0). Let h_0 be the class number of the maximal order. Let E_1 be such that End(E) has discriminant $N^2 D_0$ where $N > \log(q)$. There is an algorithm to compute an isogeny from E_0 to E_1 with (heuristic) complexity $\tilde{O}(h_0 N^{1/2} \log(q)^{O(1)})$ operations in $\mathbb{F}_q$.*

When $h_0 = O(1)$ then $\tilde{O}(h_0 N^{1/2}) = \tilde{O}(q^{1/4})$ operations in $\mathbb{F}_q$. The case where the class number is bounded or grows polynomially in $\log(p)$ is the most important case for applications in elliptic curve cryptography. But for the sake of completeness we now discuss the complexity of the general case.

[8] For our applications we have exponential-time for this precomputation if needed.

Walking the crater is clearly not the best method when N is small and h_0 is large. In practice we would run the $\tilde{O}(h_0 N^{1/2})$ algorithm when h_0 is smaller than some bound, and the $\tilde{O}(N^2 + h_0^{1/2})$ algorithm from [Gal99] when h_0 is bigger than some bound. We now determine the cross-over point.

Recall that the maximal order has discriminant D_0 and class number

$$h_0 = O(\sqrt{|D_0|}\log(|D_0|)) = \tilde{O}(|D_0|^{1/2}).$$

Also $t^2 - 4q = f^2 D_0$ where $N \mid f$. The worst case is when $f = N$, so from now on we assume this is the case. Hence we can approximate h_0 as $q^{1/2}/N$. The algorithm from [Gal99] has complexity growing as $N^2 + q^{1/4}/\sqrt{N}$ while the $\tilde{O}(h_0 N^{1/2})$ algorithm has complexity growing as $q^{1/2}/\sqrt{N}$. Note that $q^{1/4}/\sqrt{N}$ is always smaller than $q^{1/2}/\sqrt{N}$, so the crossover is roughly when $N^2 = q^{1/2}/\sqrt{N}$, namely $N = q^{1/5}$. Hence, when $q^{1/5} \leq N \leq q^{1/2}$ we apply the algorithm with complexity $\tilde{O}(h_0 N^{1/2}) = \tilde{O}(q^{1/2}/\sqrt{N}) = \tilde{O}(q^{2/5})$. When $N \leq q^{1/5}$ we apply the algorithm from [Gal99], with complexity $\tilde{O}(N^2 + h_0^{1/2}) = \tilde{O}(q^{2/5})$. Hence in all cases we obtain a method that is better than $\tilde{O}(q^{1/2})$.

In the case of pairing-based cryptography, the crater typically has size 1. Hence the results of this section show that one can reduce the ECDLP between curves in the isogeny class in $\tilde{O}(q^{1/4})$ operations in $\mathbb{F}_q$. This result unifies the complexity of the problem with the result from [Gal99] in the average case for elliptic curve cryptography.

8. General algorithm

We now consider the general isogeny problem as studied in [Gal99]: Given E_0 and E_1 over $\mathbb{F}_q$ to compute an isogeny between them. We can't expect to do better than the $\tilde{O}(q^{2/5})$ result of Section 7. But we show that all cases can be handled within this bound.

The strategy from [Gal99] was to always ascend to the crater, but this is not optimal for all cases. In fact, in many cases descending is a better idea. Indeed, if we are not given a curve on the crater, and if the class number h_0 is large enough, then it is not feasible to compute a curve on the crater.

Theorem 3. *There is an algorithm that takes as input two ordinary elliptic curves E_0 and E_1 over $\mathbb{F}_q$, runs in (heuristic) expected time $\tilde{O}(q^{2/5})$ operations in $\mathbb{F}_q$, and outputs a representation of an isogeny from E_0 to E_1 that can be evaluated in (heuristic) polynomial-time.*

Proof. We can factor $t^2 - 4q$ in time at most $O(q^{1/4})$ using exponential methods, and then we can determine the endomorphism rings of E_0 and E_1 in subexponential or polynomial time using [BS11, Rob22b]. Let $t^2 - 4q = f^2 D_0$ and let N be the largest prime dividing f.

Let f_1 and f_2 be the conductors of the curves E_0 and E_1 that we are trying to relate. If $f_1 = f_2$ we can compute an isogeny between them in $\tilde{O}(|t^2 - 4q|^{1/4}) = \tilde{O}(q^{1/4})$ operations in $\mathbb{F}_q$ using meet-in-middle approaches, as done by [Gal99, GS13].

There are four cases of the problem, each requiring a different algorithmic solution. The first case is when the conductor f does not have large prime factors. More precisely, if $N < q^{1/5}$ then we can descend to the floor in $\tilde{O}(q^{2/5})$ operations in $\mathbb{F}_q$ and then solve the isogeny problem in $\tilde{O}(q^{1/4})$ operations.

The second case is when the crater is small enough to use the class polynomial to construct a curve E on the crater, and then apply the $\tilde{O}(h_0 N^{1/2})$ algorithm Theorem 2 from Section 7 twice to find isogenies from E to both curves E_0 and E_1. If $|D_0| < q^{2/5}$ then we can construct a curve on the crater in $|D_0|^{1+o(1)} = \tilde{O}(q^{2/5})$ operations in $\mathbb{F}_q$. We then have $h_0 = \tilde{O}(q^{1/5})$ and $f > q^{3/10}$. Since $f^2|D_0| = O(q)$ we have $fh_0 = O(\sqrt{q}\log(q))$ and so $\sqrt{f}h_0 = O(\sqrt{q/f}\log(q))$. Since $f > q^{3/10}$ it follows that $\sqrt{N}h_0 \leq \sqrt{f}h_0 = O(q^{7/20}\log(q)) < \tilde{O}(q^{2/5})$. Hence the total cost to handle this case is bounded above $\tilde{O}(q^{2/5})$.

The remaining challenge is to deal with $|D_0| \geq q^{2/5}$ and $N \geq q^{1/5}$. Write $f = Nu$ for some integer u. Since $|D_0|f^2 < 4q$ we have $u^2 < 4q/(|D_0|N^2) < 4q^{1/5}$ and so $u = O(q^{1/10})$.

The third case for the proof is when either f_1 or f_2 is not divisible by N. Then we can efficiently go up to the crater in $\tilde{O}(u^3) = \tilde{O}(q^{3/10})$ operations. Now that we have a curve on the crater, we can list all curves on the crater and solve the isogeny problem using Theorem 2 from Section 7. Since $N \geq q^{1/5}$ we have $|D_0| \leq q^{3/5}$ and so $h_0 \leq \tilde{O}(q^{3/10})$. The complexity is $O(h_0 N^{1/2}) = \tilde{O}(h_0^{1/2} q^{1/4}) = \tilde{O}(q^{2/5})$.

The fourth and final case is when f_1 and f_2 are both divisible by N. This means the conductor is divisible by a large prime but the curves E_0 and E_1 are both "close to the floor". The idea is to go down to the floor of the volcano (not up to the crater as in [Gal99]) and then solve the isogeny problem. Precisely, one computes isogenies of degree dividing u from E_0 and E_1 to curves E_0' and E_1' both with endomorphism rings of discriminant $f^2 D_0$. This is done in $\tilde{O}(u^2) = \tilde{O}(q^{1/5})$ operations. We now have curves at the same level and can solve the isogeny problem using a meet-in-the-middle algorithm in $\tilde{O}(q^{1/4})$ operations.

This completes the proof. □

We stress that the worst case of the algorithm is when the conductor $f \approx q^{1/5}$ is prime and the class number of the maximal order is $h_0 \approx q^{3/10}$. In both the case of randomly chosen ordinary elliptic curves (where we have $f = O(\log(q))$) and pairing-friendly curves (where typically $D_0 = -3$ and $h_0 = 1$) the complexity is $\tilde{O}(q^{1/4})$, and no heuristics are needed.

Corollary 1. *Let E_0 and E_1 be two ordinary elliptic curves over $\mathbb{F}_q$ with $q + 1 - t$ points where the discriminant of $\mathbb{Q}(\sqrt{t^2 - 4q})$ is -3. Then there is a (rigorous) probabilistic algorithm that runs in expected $\tilde{O}(q^{1/4})$ operations in $\mathbb{F}_q$ that outputs a representation of an isogeny from E_0 to E_1 that can be evaluated in polynomial-time.*

Proof. As we explained in Section 2.1, when $|D_0|$ is bounded then it is proven that Elkies primes are distributed as needed for heuristic assumption 1 (with some small change to the constants). Applying the methods in the proof of Theorem 3 in

this case (since the class number is 1 there are no random walk methods needed) yields a rigorous algorithm. Since $h_0 = 1$, the cost is $\tilde{O}(N^{1/2}) = \tilde{O}(q^{1/4})$. $\square$

9. Solving Problem 1

For Problem 1 we are given $E_0/\mathbb{F}_q$ on the crater and there is a large prime N dividing the conductor. We want to compute the image curve E_1 of a descending N-isogeny $\phi : E_0 \to E_1$ and be able to evaluate ϕ on chosen points.

9.1. **Using standard techniques.** Let $E_0/\mathbb{F}_q$ be on the crater and let N be a large prime such that $N^2 \mid t^2 - 4q$. There are either $N-1$, N or $N+1$ $\mathbb{F}_q$-rational isogenies from E_0 to curves with endomorphism ring of index N. A rational N-isogeny has a kernel defined by a polynomial of degree $(N-1)/2$, so the x-coordinates of kernel points themselves are defined over an extension of degree at most $(N-1)/2$.

Following the work of Kohel, [Gal99] states that one can solve Problem 1 in $\tilde{O}(N^3) = \tilde{O}(q^{1.5})$ operations in $\mathbb{F}_q$ using modular polynomials. We now explain that this is not optimal.

The cubic complexity in [Gal99] arises from the cost of computing modular polynomials. However, even in 1999 one could have achieved $\tilde{O}(N^2)$ complexity by using the method explained in Section 2.2.

Taking $N = O(\sqrt{q})$ this gives complexity $\tilde{O}(q)$ operations in $\mathbb{F}_q$. If $N^4 \mid (t^2 - 4q)$, meaning that the conductor is divisible by N^2, then $N = O(q^{1/4})$ and then the complexity is $\tilde{O}(q^{1/2})$ operations in $\mathbb{F}_q$.

Alternatively one can consider computing E_1 directly using class polynomials. Sutherland [Sut11, Sut12b] states the cost is $|D|^{1+o(1)}$, which is $O(q^{1+o(1)})$ and so also no better than $\tilde{O}(N^2)$.

Hence, current techniques solve this problem in $\tilde{O}(N^2) = \tilde{O}(q)$ operations in $\mathbb{F}_q$. This is boring from the point of view of ECC.

9.2. **Guessing E_1.** We now show how to beat the $\tilde{O}(q)$ bound in the case when the conductor is equal to N and when the class number of the maximal order is 1.

If we ignore for the moment the problem of computing an isogeny $\phi : E_0 \to E_1$, one can find a curve E_1 by guessing. Indeed, the Hasse interval has length at most $4\sqrt{q} + 1$ and (apart from some exceptional "extreme" cases that do not arise for curves whose endomorphism ring has large conductor) a randomly chosen elliptic curve $E_1/\mathbb{F}_q$ typically has number of points equal to $\#E_1(\mathbb{F}_q)$ with probability $O(1/\sqrt{q})$. Hence one can choose random E_1 and compute $\#E_1(\mathbb{F}_q)$ using Schoof-Atkin-Elkies and stop when $\#E_1(\mathbb{F}_q) = \#E_0(\mathbb{F}_q)$. Indeed, since our focus is the curves used in classical cryptography, we have $\#E_0(\mathbb{F}_q)$ divisible by a large prime, and so one can test a curve E_1 by choosing a random point $R \in E_1(\mathbb{F}_q)$ and checking that $[\#E_0(\mathbb{F}_q)]R = 0$.

This gives an algorithm to compute E_1 in $\tilde{O}(\sqrt{q})$ operations in $\mathbb{F}_q$. With overwhelming probability E_1 will be on (or close to) the floor. We then apply the results from Section 8 to compute the isogeny $\phi : E_0 \to E_1$ in $\tilde{O}(q^{2/5})$ time. Hence we have done better than the $\tilde{O}(q)$ algorithm previously known for this problem.

9.3. **General case.** We now sketch a general solution to Problem 1.

Theorem 4. *One can solve Problem 1 in $\tilde{O}(q^{1/2})$ operations in $\mathbb{F}_q$.*

Proof. If $N < 2q^{1/4}$ then the standard methods (note that one has to try at most 3 random N-isogenies in order to have a descending one) solve the problem in time $\tilde{O}(q^{1/2})$, which is acceptable within the context of this section. So we consider the case when $2q^{1/4} < N < q^{1/2}$.

We wish to extend the approach given in Section 9.2. We have to handle two issues. The first is that the class number of the maximal order is in general not equal to one. The second is that N may be a large prime dividing the conductor f, but $N \neq f$. Note that $t^2 - 4q = f^2 D_0$ and so $f < 2q^{1/2}$.

The first problem arises since we generate a random elliptic curve on the floor of the volcano, and E_0 might not be directly above it. As in Section 7 we solve this problem by trying all curves E_0' on the crater, since one of them will be directly above the chosen curve. This will multiply the total cost by the class number h_0 of the maximal order. Note that $h_0 = \tilde{O}(\sqrt{|D_0|}) = \tilde{O}(\sqrt{q/f^2})$ and so it is small when f is large. Since we are assuming $f > q^{1/4}$ we have $h_0 = \tilde{O}(q^{1/4})$.

For the second problem, write $f = NN'$ for some N'. Since $N > q^{1/4}$ and $f < q^{1/2}$ we have $N' \leq f/N \leq q^{1/4} < N$. We proceed as in Section 9.2 by trying random curves E_1' until we get an elliptic curve that is on the floor. This requires $O(q^{1/2})$ operations in $\mathbb{F}_q$. We then compute a representation of an f-isogeny $\phi : E_0' \to E_1'$ for the corresponding elliptic curve E_0' above E_1' on the crater. This works since Theorem 1 applies for isogenies of any degree, not only prime degree. The cost is $\tilde{O}(h_0 f^{1/2})$ operations in $\mathbb{F}_q$.

At this point we have a representation of a descending f-isogeny ϕ from E_0', but we want a descending N-isogeny from E_0. In $O(\sqrt{h_0})$ time one computes an isogeny $\psi : E_0' \to E_0$ using standard algorithms for solving the group action problem [GS13]. We may assume the degree of ψ is coprime to f, by constructing ψ using prime ideals coprime to the conductor. By evaluating ϕ on the kernel of ψ we can push ψ to an isogeny $E_1' \to E_1''$ for some curve E_1 and, since the isogenies commute, it follows that there is an f-isogeny from E_0 to E_1''. Hence we now have a representation of a descending f-isogeny from the desired curve E_0.

The last step is to compute a descending N-isogeny, where $N \mid f$. Since the descending isogeny goes through the intermediate curves it is natural to think one can compute the equations of the intermediate curves. But we do not "directly" have the descending isogeny, as we have a representation of it from the Kani/Robert machinery. So it is not clear how to do this.

Instead, we work back up from the bottom, by computing an N'-isogeny. One can work with each prime $\ell \mid N'$ separately. When coming up from the floor there is a unique cyclic subgroup of order ℓ. Since $N' < N$, there are no primes in common between N and N'. We use the method in Section 9.1 above, by generating a kernel point P of order ℓ and applying Vélu. One can therefore compute the image of an ascending isogeny in $O((N')^2 \log(q))$ operations in $\mathbb{F}_q$, which is $\tilde{O}(q^{1/2})$ operations.

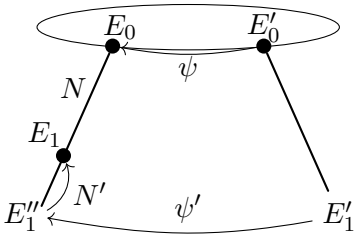

FIGURE 1. Illustration of the proof of Theorem 4.

The process is pictured in Figure 1.

Putting everything together, we obtain a representation of an N-isogeny from E_0 to E_1, as the composition of an f-isogeny from E_0 to E_1'' followed by an N'-isogeny from E_1'' to E_1. The total cost is

$$\tilde{O}(q^{1/2}) + \tilde{O}(h_0 f^{1/2}) + \tilde{O}(h_0^{1/2}) + \tilde{O}((N')^2)$$

operations in $\mathbb{F}_q$. It follows from our earlier analysis that the dominant cost is $\tilde{O}(q^{1/2})$. □

We remark that a special case of this problem arises in the algorithm by Bröker, Lauter and Sutherland for computing modular polynomials [BLS12], but our methods do not improve the complexity of their algorithm.

Finally, one of the reviewers pointed out that one can also compute a prime degree ascending isogeny from the floor in $\tilde{O}(q^{1/2})$ operations: When $N \leq q^{1/4}$ then one computes the unique N-isogeny using the method from Section 2.2, and when $N > q^{1/4}$ then the discriminant of of the maximal order is $O(q^{1/2})$ and one can use CM to compute a curve on the crater and then proceed similar to the above.

10. UPDATING JAO, MILLER, AND VENKATESAN

Jao, Miller, and Venkatesan [JMV05] showed a polynomial-time equivalence of the ECDLP among curves in the case where there is no large prime dividing the conductor gap. The main result of their paper is that, given an algorithm A that solves ECDLP on some fixed positive proportion of elliptic curves over $\mathbb{F}_q$ in "fixed level" (meaning: with some endomorphism ring $\mathcal{O}$), one can probabilistically solve ECDLP on any given curve with the same endomorphism ring with polynomially in $\log(q)$ expected queries to A with random inputs. The main tool is showing how one can efficiently compute random horizontal walks in the isogeny graph that give curves close to uniformly distributed among curves with endomorphism ring $\mathcal{O}$.

To get a more general result requires being able to jump between levels, after which the uniform mixing in the levels is enough to give the result. Large primes dividing the conductor gap are an obstacle to getting this to work.

Our techniques don't directly help with this problem, as they are about finding an isogeny connecting two given elliptic curves, rather than random self-reducing

the isogeny problem. However, we do have a new angle on this issue that does not seem to have been previously noted.

Jao, Miller, and Venkatesan consider the case where the algorithm A solves ECDLP on a fixed positive proportion of elliptic curves, and wish to solve ECDLP on *any* curve in the isogeny class. We relax this to solving ECDLP on *all but negligibly many* curves in the isogeny class. Our main observation is this: Suppose there is a large prime $N > q^{1/5}$ dividing the conductor gap, and partition the isogeny class into two sets: S_1 is the set of isomorphism classes of elliptic curves E over $\mathbb{F}_q$ with $\#E(\mathbb{F}_q) = q + t - t$ and such that $\mathrm{End}(E)$ has conductor coprime to N, while S_2 is the set for which $\mathrm{End}(E)$ has conductor divisible by N. Equation (1) shows that $\#S_2 \approx N\#S_1$, so S_2 is exponentially larger than S_1, and S_1 is exponentially small. Hence, the algorithm A solves ECDLP on a positive proportion of elliptic curves in S_2, and a randomly chosen curve in the isogeny class is in S_2 with overwhelming probability. So we just ignore the curves in S_1.

This allows to prove this theorem.

Theorem 5. *Let A be an algorithm that solves ECDLP on a fixed positive proportion of elliptic curves over $\mathbb{F}_q$ with n points. Then one can solve ECDLP for a random elliptic curve over $\mathbb{F}_q$ with n points, with overwhelming probability and in time bounded by $\tilde{O}(q^{2/5})$ operations in $\mathbb{F}_q$ plus the time taken for polynomially in $\log(q)$ queries to A with random inputs.*

Proof. Consider the conductor of the isogeny class. If the conductor does not have any large prime factors, then the result is essentially already in [JMV05] (one should descend close to the floor, and random self-reduce inside the large levels of the isogeny graph). As long as all primes dividing the conductor are bounded by $q^{1/5}$ then all the costs are within $\tilde{O}(q^{2/5})$ operations in $\mathbb{F}_q$.

So it suffices to consider the case when the conductor is divisible by a prime $N \geq q^{1/5}$. Let $\ell_1, \ldots, \ell_k$ be the prime divisors of the conductor that satisfy $\ell_i > q^{1/5}$ (there can be at most two of them, so define $\ell_2 = 1$ if $k = 1$).

Write the conductor as $f = \ell_1\ell_2 u$ where u is $q^{1/5}$-smooth. For each $w \mid u$ let $\mathcal{O}_w$ be the order of discriminant $w^2 D_0$. Equation (1) shows that the number of elliptic curves in the isogeny class with conductor dividing u is given by

$$N_0 = \sum_{w|u} \frac{h_0 w}{[\mathcal{O}_K^* : \mathcal{O}_w^*]} \prod_{\ell|w} \left(\ell - \left(\tfrac{D_0}{\ell}\right)\right).$$

Let S_1 be the set of isomorphism classes of elliptic curves E over $\mathbb{F}_q$ with $\#E(\mathbb{F}_q) = n$ and such that $\mathrm{End}(E)$ has conductor not divisible by $\ell_1\ell_2$, and S_2 the set for which $\mathrm{End}(E)$ has conductor divisible by $\ell_1\ell_2$. When $k = 1$ then $\#S_1 = N_0$ and $\#S_2 = (\ell_1 \pm 1)N_0$. When $k = 2$ then $\#S_1 = (1 + (\ell_1 \pm 1) + (\ell_2 \pm 1))N_0$ and $\#S_2 = (\ell_1 \pm 1)(\ell_2 \pm 1)N_0$. In both cases $\#S_2/\#S_1 \geq q^{1/5}/2$.

Let E be a randomly chosen curve in the isogeny class. Then with overwhelming probability E lies in S_2.

Similarly, consider the set S_A of elliptic curves in the isogeny class for which A solves ECDLP efficiently. Since S_2 is exponentially larger than S_1, the intersection $S_A \cap S_2$ must be a positive proportion of S_2.

Finally, since all curves in S_2 have (relative to each other) conductor gaps divisible by primes of size at most $q^{1/5}$, we can compute pseudo-random walks in S_2 in time $\tilde{O}(q^{2/5})$, and hence sample pseudo-randomly from S_2. Hence we will hit a curve in S_A after a polynomial number of steps in the walk and with polynomially many calls to A. □

11. The Koblitz, Koblitz, Menezes Speculation

Koblitz, Koblitz and Menezes [KKM11] introduced a bizarre consequence of the difficulty to compute isogenies across the conductor gap. They argued that this issue might imply that curves on the floor of the volcano are *less* secure (i.e., have easier discrete logarithm problem) than curves on the crater.

There is no direct evidence for this conjecture, in the sense that there is no known algorithm that would solve the discrete logarithm problem on curves on the floor in fewer than $O(q^{1/2})$ operations.

We briefly explain that the results in our paper do not refute this argument.

Let $E_0/\mathbb{F}_q$ be a curve on the crater, such that there is a large prime $N > q^{1/4}$ dividing the conductor of the isogeny class. Suppose there is an algorithm that solves ECDLP for curves on the floor in fewer than $O(q^{1/2})$ operations. Can we use this algorithm to solve the ECDLP on E_0 in fewer than $O(q^{1/2})$ operations? The problem is that we need to construct an isogeny $\phi : E_0 \to E_1$ where E_1 is on the floor. As we have seen, we only know two ways to do this: Either construct the isogeny ϕ in at least $N^2 > q^{1/2}$ operations, or "guess" E_1, which also takes at least $q^{1/2}$ attempts. Either way, the cost to transfer the ECDLP from E_0 to E_1 already takes at least $q^{1/2}$ operations and so the KKM speculation is not contradicted.

On the other hand, if the ECDLP is easier on the crater than the floor, and if E_1 is an elliptic curve on the floor, then we have shown that an isogeny from E_1 to a curve on the crater can be constructed in time at most $O(q^{1/4})$. Hence, when the class number of the maximal order is small, then the ECDLP on E_1 can be reduced to an instance of the ECDLP on the crater and it follows that if the ECDLP can be solved in fewer than $q^{1/2}$ operations on the crater then it can be solved in fewer than $q^{1/2}$ operations for all curves in the isogeny class.

In summary, our results are consistent with the argument by Koblitz, Koblitz and Menezes.

12. Conclusions and open problems

The paper is about understanding the hardness of the ECDLP among curves with the same number of points. The fundamental question is whether one can efficiently transfer instances of the discrete logarithm problem from one curve to another, for any two curves with the same number of points.

We have improved the algorithm by Galbraith [Gal99] in the worst-case from $\tilde{O}(q^{3/2})$ to $\tilde{O}(q^{2/5})$. The result can be made fully rigorous in some cases (e.g., when $|D_0|$ is bounded).

It would be very interesting to have a better solution to Problem 1, namely to construct a curve on the floor of the isogeny volcano in the case where the conductor is divisible by a large prime. One possible approach would be to use modular curves to replace guessing random E_1 by guessing curves that are already biased towards having the same number of points as E_0. Recall that we try to find a random curve E_1 such that $\#E_1(\mathbb{F}_q) = \#E_0(\mathbb{F}_q)$. Suppose $\ell \mid \#E_0(\mathbb{F}_q)$. Then if we can restrict our search to sampling random curves E_1 such that $\ell \mid \#E_1(\mathbb{F}_q)$, then we should only need to make $\sqrt{q}/\ell$ guesses. This approach works in general, and has been developed in [Sut11, Sut12a], but we are mainly interested in curves with a prime number of points, so this approach is not immediately applicable to the main cases of interest. If the quadratic twist of the curve E_0 has order with a suitable factor ℓ then one can attempt this idea. We leave this for future work, as we have not been able to get a full solution to the problem by this approach. Ideally, there will be a new idea that can solve the problem without such naive methods as guessing E_1.

REFERENCES

[BMO+18] Michael A. Bennett, Greg Martin, Kevin O'Bryant, and Andrew Rechnitzer, Explicit bounds for primes in arithmetic progressions, *Illinois Journal of Mathematics*, Vol. 62, No. 1-4 (2018) 427–532.

[BDFLS20] Daniel J. Bernstein, Luca De Feo, Antonin Leroux, and Benjamin Smith. Faster computation of isogenies of large prime degree. *The Open Book Series*, 4(1):39–55, 2020.

[BS11] Gaetan Bisson and Andrew V. Sutherland, Computing the endomorphism ring of an ordinary elliptic curve over a finite field, *Journal of Number Theory*, vol. 131, no. 5 (2011) 815–831.

[BSS99] I. F. Blake, G. Seroussi, and N. P. Smart, *Elliptic curves in cryptography*, Cambridge, 1999.

[BLS12] Reinier Bröker, Kristin Lauter, Andrew V. Sutherland, Modular polynomials via isogeny volcanoes, Mathematics of Computation, vol. 81, no. 278 (2012) 1201–1231.

[BSI] Federal Office for Information Security (BSI), Technical Guideline BSI TR-03111 Elliptic Curve Cryptography, Version 2.10 (2018)

[CLM+18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny and Joost Renes, CSIDH: An Efficient Post-Quantum Commutative Group Action, in T. Peyrin and S.D. Galbraith (eds.), ASIACRYPT 2018, Springer LNCS 11274 (2018) 395–427.

[CHM+23] 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 H. Handschuh and A. Lysyanskaya, CRYPTO 2023, Springer LNCS 14083 (2023) 762–792.

[CD23] Wouter Castryck and Thomas Decru, An Efficient Key Recovery Attack on SIDH, in C. Hazay and M. Stam (eds.), EUROCRYPT 2023, Springer LNCS 14008 (2023) 423–447.

[Coh93] H. Cohen, *A course in computational algebraic number theory*, GTM 138, Springer, 1993.

[Cox89] D. A. Cox, *Primes of the form* $x^2 + ny^2$, Wiley, 1989.

[DLRW23] Pierrick Dartois, Antonin Leroux, Damien Robert and Benjamin Wesolowski, SQISignHD: New Dimensions in Cryptography, eprint 2023/436 (2023).

[QKL+21] Victoria de Quehen, Péter Kutas, Chris Leonardi, Chloe Martindale, Lorenz Panny, Christophe Petit and Katherine E. Stange, Improved Torsion-Point Attacks on SIDH Variants, in T. Malkin and C. Peikert (eds), CRYPTO 2021, Springer LNCS 12827 (2021) 432–470.

[Gal99] Steven D. Galbraith. Constructing isogenies between elliptic curves over finite fields. *LMS Journal of Computation and Mathematics*, 2:118–138, 1999.

[GHS02] Steven D. Galbraith, Florian Hess, and Nigel Smart, Extending the GHS Weil descent attack, In L. Knudsen (ed.), EUROCRYPT 2002, Springer LNCS 2332 (2002) 29–44.

[GS13] Steven D. Galbraith and Anton Stolbunov, Improved Algorithm for the Isogeny Problem for Ordinary Elliptic Curves, Applicable Algebra in Engineering, Communication and Computing, 24(2) (2013) 107–131.

[JMV05] David Jao, Steven D. Miller, and Ramarathnam Venkatesan. Do all elliptic curves of the same order have the same difficulty of discrete log? In ASIACRYPT 2005, pages 21–40. Springer, 2005.

[Kan97] Ernst Kani, The number of curves of genus two with elliptic differentials, *J. reine angew.*, vol. 485 (1997) pp. 93-122.
[KKM11] A. H. Koblitz, N. Koblitz, A. Menezes, Elliptic Curve Cryptography: The Serpentine Course of a Paradigm Shift *Journal of Number Theory*, Vol. 131, No. 5 (2011) 781–814.
[Koh96] David Kohel, *Endomorphism rings of elliptic curves over finite fields*, PhD thesis, University of California at Berkeley, 1996.
[Ler22] Antonin Leroux, A New Isogeny Representation and Applications to Cryptography, in S. Agrawal and D. Lin (eds.), ASIACRYPT 2022, Springer LNCS 13792 (2022) 3–35.
[LR22] David Lubicz and Damien Robert, Fast change of level and applications to isogenies, *Research in Number Theory*, Volume 9, No. 7 (2023)
[MMP+23] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope and Benjamin Wesolowski, A Direct Key Recovery Attack on SIDH, in C. Hazy and M. Stam (eds.), EUROCRYPT 2023, Springer LNCS 14008 (2023) 448–471.
[PR23] Aurel Page and Damien Robert, Introducing Clapoti(s): Evaluating the isogeny class group action in polynomial time, *IACR Cryptol. ePrint Arch. 2023/1766*.
[Pol74] John M. Pollard, Theorems on factorization and primality testing, Proc. Cambridge Philos. Soc. vol. 76 (1974), 521–528.
[Rob23] 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, Cham, 2023. Springer.
[Rob22a] Damien Robert, Evaluating isogenies in polylogarithmic time, eprint 2022/1068
[Rob22b] Damien Robert. Some applications of higher dimensional isogenies to elliptic curves: Overview of results. eprint 2022.
[Sil09] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, second edition, 2009.
[SS15] Igor E. Shparlinski and Andrew V. Sutherland, On the distribution of Atkin and Elkies primes for reductions of elliptic curves on average, LMS J. Comput. Math. 18 (1) (2015) 308–322.
[Str77] Volker Strassen, Einige Resultate über Berechnungskomplexität, Jber. Deutsch. Math. Verein. vol. 78 (1976/77), no. 1, 1–8.
[Sut11] Andrew V. Sutherland, Computing Hilbert class polynomials with the Chinese remainder theorem, Math. Comp. 80 (2011) 501–538.
[Sut12a] Andrew V. Sutherland, Constructing elliptic curves over finite fields with prescribed torsion, Math. Comp. 81 (2012) 1131–1147.
[Sut12b] Andrew V. Sutherland, Accelerating the CM method, LMS J. Comput. Math. 15 (2012) 172–204.
[Sut13] Andrew V. Sutherland, Isogeny volcanoes, in Proceedings of the Tenth Algorithmic Number Theory Symposium, The Open Book Series, Vol. 1, No. 1 (2013) 507–530.

APPENDIX A. EXAMPLES OF PAIRING-FRIENDLY CURVES

We give some real-world examples of elliptic curves with large primes dividing the conductor. These curves have been proposed for certain pairing-based cryptosystems. The curve E_0 in applications has class number one, since the non-trivial automorphism is used to speed up computing the pairing. Hence, for the applications it is not necessary to construct curves on the floor. The interest in these curves from the point of view of this article is the question of whether the discrete logarithm problem might be easier or harder for the curve E_0 on the crater than the general case of curves in the isogeny class. The ability to compute isogenies across the conductor gap merely shows that the discrete logarithm problem cannot have dramatically different complexity at different levels in the volcano.

Curve	Conductor
Pasta	$3 \cdot 210890879 \cdot 31052728481172930447028584034 1$
Pluto-Eris	$3 \cdot 17387 \cdot 178307 \cdot 1531668177584969 \cdot 1166712141901129507$ $\cdot 60828567289519414610683 3$
BN382-plain	$3 \cdot 554633$ $\cdot 14150452518435797091709219053060069021096350876494 27$
Geppetto	$99609175647210028388479 3$ $\cdot 33728034835887799224372269381656381850708127921979643$
BLS12-381	$11 \cdot 31 \cdot 503 \cdot 10177 \cdot 64223 \cdot 859267 \cdot 52437899$ $\cdot 24305087161 \cdot 18815978399361992833$

The natural question is to relate the hardness of ECDLP for curves on the crater and curves on the floor. The obstacle is problem 1: We don't know how to compute an equation for a curve on the floor for any of these examples in fewer than $q^{1/2}$ operations in $\mathbb{F}_q$.

Even with BLS12-381 and Pluto-Eris (the smoothest conductors in the above list), doing things the naive way has cost at least

$$18815978399361992833^2 \approx 2^{128} \quad \text{and} \quad 60828567289519414610683 3^2 \approx 2^{160}$$

operations. For BLS-12-381 this beats $\sqrt{q}$, but is not as good as the $q^{1/4}$ that we can now achieve. For Pluto-Eris, $q \approx 2^{256}$, so $\sqrt{q} \approx 2^{128}$ and we are no better than $\sqrt{q}$.

As an example of parameters for Theorem 1, consider the case

$$N = 14150452518435797091709219053060069021096350876494 27$$

from the BN382-plain example. Here $D_0 = -3$ and the list of Elkies primes begins

$$7, 13, 19, 31, 37, 43, 61, 67, 73, 79, 97, 103, 109, 127, 139, 151, \ldots$$

These primes all satisfy $\ell_i < 7.5 + 4i\log(i)$. We ran a toy program to choose random subsets of Elkies primes and to compute the number A_S as defined in our theorem. Taking $S = \{8, 14, 21, 32, 43, 56, 78, 104, 121\}$ gives

$$A_S = 3465673716706679666094 37.$$

Taking $n = 9$ gives $3^n A_S^2 > N$. We have

$$m = 3^n A_S^2 - N = 2^3 \cdot 5^2 \cdot 4745295376629449267991147374952934684744279789357,$$

which satisfies $m = m_1^2 + m_2^2$ for $m_1 = 27609211712694305185354750$ and $m_2 = 13667132249653306290470330$.

MATHEMATICS DEPARTMENT, UNIVERSITY OF AUCKLAND, NZ.
Email address: s.galbraith@auckland.ac.nz