Extending the GHS Weil Descent Attack

S.D. Galbraith, F. Hess and N.P. Smart

Department of Computer Science,
University of Bristol,
Merchant Venturers Building,
Woodland Road,
Bristol, BS8 1UB,
United Kingdom.
{steven, florian, nigel}@cs.bris.ac.uk

Abstract. In this paper we extend the Weil descent attack due to Gaudry, Hess and Smart (GHS) to a much larger class of elliptic curves. This extended attack still only works for fields of composite degree over $\mathbb{F}_2$. The principle behind the extended attack is to use isogenies to find a new elliptic curve for which the GHS attack is effective. The discrete logarithm problem on the target curve can be transformed into a discrete logarithm problem on the new isogenous curve. One contribution of the paper is to give an improvement to an algorithm of Galbraith for constructing isogenies between elliptic curves, and this is of independent interest in elliptic curve cryptography. We conclude that fields of the form $\mathbb{F}_{q^7}$ should be considered weaker from a cryptographic standpoint than other fields. In addition we show that a larger proportion than previously thought of elliptic curves over $\mathbb{F}_{2^{155}}$ should be considered weak.

1 Introduction

The technique of Weil descent to solve the ECDLP was first proposed by Frey [5]. This strategy was elaborated on further by Galbraith and Smart [8]. The work of Gaudry, Hess and Smart [9] gave a very efficient algorithm to reduce the ECDLP to the discrete logarithm in a Jacobian of a hyperelliptic curve over $\mathbb{F}_q$. Since subexponential algorithms exist for the discrete logarithm problem in higher genus curves, this gives a possible method of attack against the ECDLP. We refer to the method of [9] as the GHS attack.

Menezes and Qu [14] analysed the GHS attack in some detail and demonstrated that it did not apply to the case when $q = 2$ and n is prime. Since this is the common case in real world applications, the work of Menezes and Qu means that the GHS attack does not apply to most deployed systems. However, there are a few fielded elliptic curve systems which use the fields $\mathbb{F}_{2^{155}}$ and $\mathbb{F}_{2^{185}}$. Hence there is considerable interest as to whether the GHS attack makes all curves over these fields vulnerable. In [17] Smart examined the GHS attack for elliptic curves with respect to the field extension $\mathbb{F}_{2^{155}}/\mathbb{F}_{2^{31}}$ and concluded that such a technique was unlikely to work for any curve defined over $\mathbb{F}_{2^{155}}$.

Jacobson, Menezes and Stein [10] also examined the field $\mathbb{F}_{2^{155}}$, this time using the GHS attack down to the subfield $\mathbb{F}_{2^5}$. They concluded that such a strategy could be used in practice to attack around 2^{33} isomorphism classes of elliptic curves defined over $\mathbb{F}_{2^{155}}$. Since there are about 2^{156} isomorphism classes of elliptic curves defined over $\mathbb{F}_{2^{155}}$, the probability of finding one where the GHS attack is applicable is negligible.

In this paper we extend the GHS attack to a much larger number of elliptic curves over certain composite fields of even characteristic.

The main principle behind the paper is the following. Let E_1 be an elliptic curve over a finite field $\mathbb{F}_{q^n}$ and suppose that the GHS attack transforms the discrete logarithm problem into one on a curve of genus g over $\mathbb{F}_q$. Now let E_2 be an elliptic curve over $\mathbb{F}_{q^n}$ which is isogenous to E_1 (i.e., $\#E_1(\mathbb{F}_{q^n}) = \#E_2(\mathbb{F}_{q^n})$). The GHS method is not usually invariant under isogeny, so the value for g which arises for E_2 can be different to the one for E_1. There are two ways this property might be exploited:

- To solve a discrete logarithm problem on an elliptic curve E_1 over $\mathbb{F}_{q^n}$ for which the GHS attack is not effective, one could try to find an isogenous curve E_2 for which the GHS attack is effective.
- It is often possible to construct a 'weak' elliptic curve E_2 over $\mathbb{F}_{q^n}$ for which the GHS attack is particularly successful (this is essentially what was done by [10]). One might 'hide' such a curve by taking an isogeny to a curve E_1 for which the GHS attack is not effective. Knowledge of the 'trapdoor' (i.e., the isogeny) would enable one to solve the discrete logarithm. This approach might have both malicious and beneficial applications.

We achieve the first point as follows. Given an elliptic curve E_1 over $\mathbb{F}_{q^n}$ with $N = \#E_1(\mathbb{F}_{q^n})$ the strategy is to search over all elliptic curves with small m (using the method of Section 4) until one is found which has N points (this is checked by exponentiating a random point). Once such an 'easy' curve is found one can construct an isogeny explicitly using the method of Section 3, which is an improved version of the algorithm of Galbraith [7].

This process extends the power of the GHS attack considerably. For instance, with $K = \mathbb{F}_{2^{155}}$ and $k = \mathbb{F}_{2^5}$, Jacobson, Menezes and Stein [10] found that there are only 2^{33} curves with a small value of g (i.e., the value m is small). Using our techniques the number of isomorphism classes of curves which are vulnerable to attack is increased to around 2^{104}. This is a significant breakthrough in the power of the GHS attack. Furthermore, we show that for certain fields (in particular, $\mathbb{F}_{q^7}$ where q is a power of 2) almost all elliptic curves are vulnerable to the extended GHS attack.

Regarding the second point, suppose an elliptic curve E_1 over $\mathbb{F}_{q^n}$ with large m is given. The same process as above gives a method to check whether an elliptic curve E_2 exists which is isogenous to E_1 over $\mathbb{F}_{q^n}$ and which has small m. Hence, we show that it is not possible to use this trapdoor for either malicious or constructive applications.

As an aside, we note that the methods of this paper give a way to unify the treatment of subfield curves (sometimes called Koblitz curves) with the general

case. Given an elliptic curve E_1 defined over $k = \mathbb{F}_q$ then special techniques are required to perform Weil descent with respect to K/k. However, by taking an isogeny $\phi : E_1 \to E_2$ such that E_2 is now defined over K one can proceed in the usual manner. We comment that subfield curves are only used when the extension degree n is a large prime, and, as shown in [14], the GHS method is not usually successful in this case.

We only describe the extended GHS strategy in the case of fields of characteristic two, though the principles can of course be easily adapted to the general case. We stress that, in the case of characteristic two, our results only apply to extension fields of composite degree.

The remainder of the paper is organised as follows. In section 2 we explain the GHS attack and the analysis of Menezes and Qu. In section 3 we discuss the method of Galbraith for finding isogenies between elliptic curves, in addition we sketch a new version of Galbraith's algorithm which requires much less memory. In section 4 we describe how to obtain an explicit list of isomorphism classes of elliptic curves for which the GHS attack can be successfully applied. In section 5 we examine the implications of using isogenies to extend the GHS attack.

2 The GHS Attack

Let us first set up some notation. Throughout this paper we let E denote an elliptic curve over the field $K = \mathbb{F}_{q^n}$ where $q = 2^r$. Let k denote the subfield $\mathbb{F}_q$. To simplify the discussion, and since those cases are the most important, we always assume that r and n are odd. We also assume that n is a prime. We stress that it is easy to obtain analogous results in the more general case.

Define $\sigma : K \to K$ to be the q-power Frobenius automorphism, and let $\pi : K \to K$ denote the absolute Frobenius automorphism $\pi : \alpha \to \alpha^2$. Therefore, $\sigma = \pi^r$.

The elliptic curve discrete logarithm problem (ECDLP) is the following: given $P \in E(K)$ and $Q \in \langle P \rangle$ find an integer λ such that $Q = [\lambda]P$. The apparent intractability of the ECDLP forms the basis for the security of cryptographic schemes based on elliptic curves.

We let l denote the order of the point P. To avoid various well known attacks, namely those described in [15], [16], [13] and [6], one chooses the curve such that l is a prime of size $l \approx q^n$. We also need to ensure that l does not divide $q^{ni} - 1$, for all "small" values of i.

The GHS attack is as follows. One takes an elliptic curve defined, as above, over $\mathbb{F}_{q^n}$, with a large subgroup of prime order l. We assume the curve is given by an equation of the form

$$E : Y^2 + XY = X^3 + aX^2 + b \text{ where } a \in \{0, 1\}, b, \in K.$$

We may assume that $a \in \{0, 1\}$ since r and n are odd. Then one constructs the Weil restriction of scalars $W_{E/k}$ of E over k, this is an n-dimensional abelian variety over k. The variety $W_{E/k}$ is then intersected with $n - 1$ carefully chosen

hyperplanes so as to obtain a hyperelliptic curve C. Let g denote the genus of C.

In addition, the GHS attack gives an explicit and efficient group homomorphism from $E(K)$ to $J_C(k)$. Assuming some mild conditions, the Jacobian of C will contain a subgroup of order l and the image of the subgroup of order l in $E(K)$ will be a non-trivial subgroup of order l in $J_C(k)$.

The genus of C is equal to either 2^{m-1} or $2^{m-1}-1$, where m is determined as follows.

Theorem 1 ([9]). *Let $b_i = \sigma^i(b)$, then m is given by*

$$m = m(b) = \dim_{\mathbb{F}_2}\left(\mathrm{Span}_{\mathbb{F}_2}\left\{(1, b_0^{1/2}), \ldots, (1, b_{n-1}^{1/2})\right\}\right).$$

In particular we have $1 \leq m \leq n$. If m is too small then the size of $J_C(k)$, which is $\approx q^g$, will be too small to contain a subgroup of size l. If m is too large, although we can translate discrete logarithm problems to the hyperelliptic setting, this does not help us to solve the original ECDLP in practice.

Menezes and Qu proved the following theorem which characterises the smallest value of $m > 1$ and the elliptic curves which give rise to such m.

Theorem 2 ([14]). *Keeping the notation as above, and considering the GHS technique for Weil restriction of E from K down to k. Let n denote an odd prime, let t denote the multiplicative order of two modulo n and let $s = (n-1)/t$. Then*

1. *The polynomial $x^n - 1$ factors over $\mathbb{F}_2$ as $(x-1)f_1 f_2 \cdots f_s$ where the f_i's are distinct irreducible polynomials of degree t. For $1 \leq i \leq s$ define*

$$B_i = \{b \in \mathbb{F}_{q^n} : (\sigma - 1) f_i(\sigma) b = 0\}.$$

2. *For all $1 \leq i \leq s$ and all $b \in B_i$ the elliptic curves*

$$Y^2 + XY = X^3 + b,$$
$$Y^2 + XY = X^3 + \alpha X^2 + b$$

   have $m(b) \leq t+1$, where α is a fixed element of K of trace one.
3. *If $m(b) = t+1$ then E must be one of the previous curves for some i and some $b \in B_i$.*
4. *The cardinality of the set $\cup_{i=1}^{s} B_i$ is $qs(q^t - 1) + q$.*

In particular $m(b) = t+1$ is the smallest attainable value of $m(b)$, which is greater than one, for the field $\mathbb{F}_{q^n}$ using the GHS technique for Weil restriction down to $\mathbb{F}_q$.

Menezes and Qu use the above theorem to show that if n is a prime in the range $160 \leq n \leq 600$ and $q = 2$ then the GHS attack will be infeasible.

If we look at smaller prime values of n we see that $n = 31$ is particularly interesting, since we obtain a particularly low value of $t = 5$ and $s = 6$. Thus for the field $\mathbb{F}_{2^{155}}$ there are around 2^{33} elliptic curves whose Weil restriction down to

$\mathbb{F}_{2^5}$ contains a hyperelliptic curve of genus 31 or 32. However, the next admissible size of t is 10, which would correspond to hyperelliptic curves of genus 2^{11} or $2^{11}-1$. The algorithms for solving the discrete logarithm problem on curves of genus 2^{11} over $\mathbb{F}_{2^5}$ do have subexponential complexity, but the problem has grown to such a size (10000 bits) that this is not an efficient way to solve a 155 bit elliptic curve discrete logarithm problem.

We also need to take into account the Weil restriction from $\mathbb{F}_{2^{155}}$ down to $\mathbb{F}_{2^{31}}$. This will always lead to values of m equal to 1 or 5, thus the most useful hyperelliptic curves have genus 15 or 16. It was shown in [17] that solving a discrete logarithm problem on a curve of genus 16 over the field $\mathbb{F}_{2^{31}}$ is infeasible using current technology.

It would therefore appear that, for the field $\mathbb{F}_{2^{155}}$, by avoiding the $\approx 2^{33}$ curves which gives rise to $t = 5$ means one need not worry about the GHS attack. However, as we have explained in the introduction, our new results show that this is argument is not true.

3 Constructing Isogenies

Let $K = \mathbb{F}_{q^n}$ be a finite field. Let $\Sigma = \sigma^n$ be the q^n-th power Frobenius. Let E_1 and E_2 be two non-supersingular elliptic curves over K which are isogenous over K (i.e., $\#E_1(K) = \#E_2(K)$). We wish to find an explicit representation for an isogeny

$$\phi : E_1 \to E_2.$$

In [7] the following result was proved.

Theorem 3. *There is an algorithm to compute ϕ which in the worst case takes $O(q^{3n/2+\epsilon})$ operations in $\mathbb{F}_{q^n}$ and requires at worst $O(q^{n+\epsilon})$ space. The average case complexity is $O(q^{n/4+\epsilon})$ operations.*

Galbraith in [7] gives the algorithm only in the case of prime fields of large characteristic. However, he also discusses how to extend the algorithm to arbitrary finite fields using the techniques of Couveignes [4] and Lercier [12]. It is clear that the complexity estimates are the same for the different types of finite fields.

In this section we sketch a version of Galbraith's algorithm which has polynomial storage requirement. The technique to obtain an algorithm with reduced storage requirement is inspired by ideas of Pollard [16]. This new version has expected average case running time $O(q^{n/4+\epsilon})$. The sketched algorithm applies over any base field.

Let t denote the common trace of Frobenius of E_1 and E_2. Set

$$\Delta = t^2 - 4q^n.$$

The endomorphism rings $\mathrm{End}(E_i)$ are orders in the imaginary quadratic field $\mathbb{Q}(\sqrt{\Delta})$. The maximal order of $\mathbb{Q}(\sqrt{\Delta})$ we shall denote by $\mathcal{O}$, and its class number by h_Δ. Let $\mathbb{Z}[\Sigma]$ denote the equation order given by the characteristic polynomial of Frobenius,

$$\Sigma^2 - t\Sigma + q^n = 0.$$

Then $\mathbb{Z}[\Sigma] \subseteq \mathrm{End}(E_i) \subseteq \mathcal{O}$.

The algorithm consists of a number of stages.

Stage 0: Reduce to considering finding an isogeny between two curves whose endomorphism ring is the maximal order.

Stage 1: Using a random walk determine an ideal of $\mathcal{O}$ corresponding to an isogeny between the elliptic curves.

Stage 2: Smooth the ideal (using ideas from index calculus algorithms for ideal class groups in quadratic fields).

Stage 3: Extract an isogeny corresponding to the smooth ideal output by the previous stage.

In the next subsections we outline the various stages.

3.1 Isogenies and Ideals

In the course of our algorithm we will need a map to pass from isogenies to prime ideals and vice-versa. We shall now explain how to perform this subtask. The ideas in this section are based on subprocedures of the Schoof-Elkies-Atkin (SEA) algorithm. For an overview of this we refer to Chapter VII of [3].

Let l denote an element in our factor base. We shall concentrate on the more complicated case where l splits, leaving the ramified case to the reader.

Let j denote the j-invariant of an elliptic curve E over K such that $\mathrm{End}(E) \cong \mathcal{O}$. Let l denote a prime which splits in $\mathcal{O}$. The characteristic polynomial of Frobenius factorises as

$$X^2 - tX + q^n \equiv (X - \mu)(X - \lambda) \pmod{l},$$

where $\mu, \lambda \in K$. By Dedekind's Theorem the prime l splits in $\mathcal{O}$ into the ideals

$$\mathfrak{l}_1 = (l, \Sigma - \mu), \mathfrak{l}_2 = (l, \Sigma - \lambda).$$

In addition, the modular polynomial $\Phi_l(j, X)$ factors in $K[X]$ as a product of two linear factors and s factors of degree r, where r divides $l - 1$. The roots of $\Phi_l(j, X)$ in K correspond to two j invariants j_1 and j_2, these are the j-invariants of elliptic curves E_1 and E_2 that are l-isogenous to E.

We wish to determine the correct association between $\{j_1, j_2\}$ and $\{\mathfrak{l}_1, \mathfrak{l}_2\}$. To do this we use techniques from the Elkies variant of Schoof's algorithm. Fix a j-invariant, say j_1, and determine the subgroup C_1 of $E[l]$ which lies in the kernel of the isogeny from E to E_1. We then determine whether μ or λ is an eigenvalue for this isogeny by checking whether

$$\Sigma(P) = [\mu]P \text{ or } [\lambda]P \text{ for } P \in C_1.$$

If μ is an eigenvalue then j_1 corresponds to $\mathfrak{l}_1$ and j_2 corresponds to $\mathfrak{l}_2$, otherwise the correspondence is the opposite.

Using the above techniques one can also solve the following inverse problem. Given j and a prime ideal $\mathfrak{l}_1$, determine the j-invariant of the isogenous curve corresponding to the isogeny determined by $\mathfrak{l}_1$.

In either direction the method requires operations on polynomials of degree $O(l)$. Hence, the total the complexity will be $O((\log q^n)l^2)$ field operations, since the main bottleneck is computing x^{q^n} modulo the a polynomial in x of degree $(l-1)/2$.

3.2 Stage 0

Using Kohel's algorithm [11] we find, for each i, a chain of isogenies from E_i to an elliptic curve E_i' whose endomorphism ring is the maximal order $\mathcal{O}$.

This is the part of the procedure which gives us the worst case running time. If we let c denote the square part of Δ and if c contains a large prime factor then this stage can be particularly difficult. This will lead to a large worst case complexity for our algorithm. However, on average c turns out to be both small and smooth, and so this stage is particularly simple. In fact if $c = 1$ (i.e., the order $\mathbb{Z}[\Sigma]$ is the maximal order) then **Stage 0** can be eliminated completely.

By abuse of notation for the rest of the description we shall set $E_i = E_i'$ and $j_i = j(E_i)$.

3.3 Stage 1

We define a random walk on the j-invariants of elliptic curves. More specifically, we will consider pairs of the form $(j, \mathfrak{a})$, where j is the j-invariant of some elliptic curve and $\mathfrak{a}$ is an element of the ideal class group of $\mathcal{O}$. The random walk only depends on the value of j. We require a function

$$f : K \to \mathcal{F} \times \{0, 1\}$$

which should be deterministic but have a distribution close to uniform. The function f will be used to define the random walk.

Recall we have two j invariants j_1 and j_2 of two isogenous elliptic curves and we wish to determine the isogeny between them, using as small amount of memory as possible. For this we use the ideas of Pollard.

We define a step of our random walk given a j-invariant $j_k^{(i)}$ as follows: First compute $(l, b) = f(j_k^{(i)})$. Then factor $\Phi_l(j_k^{(i)}, X)$ to obtain one or two new j-invariants. Using the bit b select one of the j-invariants in a deterministic manner and call it $j_k^{(i+1)}$. Use the technique from the previous subsection to determine the prime ideal $\mathfrak{l}$ corresponding to the isogeny from $j_k^{(i)}$ to $j_k^{(i+1)}$. Finally, update the original pair $(j_k^{(i)}, \mathfrak{a}_k^{(i)})$ to $(j_k^{(i+1)}, \mathfrak{a}_k^{(i+1)})$ where

$$\mathfrak{a}_k^{(i+1)} = \text{Reduce}\left(\mathfrak{a}_k^{(i)} \cdot \mathfrak{l}\right).$$

The whole algorithm then proceeds in the following way. We take a random walk of $T = O(\sqrt{h_\Delta}) = O(q^{n/4})$ steps, starting with the initial value $(j_1^{(0)} = j_1, \mathfrak{a}_1^{(0)} = (1))$. Only the final position $(j_1^{(T)}, \mathfrak{a}_1^{(T)})$ is stored. Then start a second

random walk from the initial value ($j_2^{(0)} = j_2, \mathfrak{a}_2^{(0)} = (1)$). Eventually, after an expected T steps, we will find a value of S such that

$$j_1^{(T)} = j_2^{(S)}.$$

If such a collision is not found then one can easily 'randomise' the starting points of the random walk using isogenies and repeat. Once a collision is found we know that the isogeny from j_1 to j_2 is represented by the ideal

$$\mathfrak{a} = \mathfrak{a}_1^{(T)} / \mathfrak{a}_2^{(S)}.$$

It is also possible to construct a chain of isogenies from E_1 to E_2 by following the paths in the random walk, but this is much longer than necessary. Instead, as we will show in the next section, one can obtain an isogeny which can be easily represented in a short and compact format.

To analyse the complexity of **Stage 1** we notice that since the random walk is on a set of size h_Δ then we expect a collision to occur after $\sqrt{\pi h_\Delta / 2}$ steps, by the birthday paradox. In the unlikely event that a collision does not occur after this many steps, we start again with related initial j-invariants, or one could repeat the process using a different function f. Since each step of the walk requires at most $O((\log q^n) L^2)$ field operations, and $\Delta \approx q^n$ we obtain a final complexity for **Stage 1** of

$$O((\log q^n) L^2 \sqrt{h_\Delta}) = O\left((\log q^n)^6 q^{n/4}\right) = O(q^{n/4+\epsilon}).$$

3.4 Stage 2

Now we have two j-invariants j_1 and j_2 and an ideal $\mathfrak{a}$ representing an isogeny between j_1 and j_2. We can assume that $\mathfrak{a}$ is a reduced ideal. In this stage we will replace $\mathfrak{a}$ by a smooth ideal. Of course, the ideal $\mathfrak{a}$ was originally constructed as a smooth product of ideals, but this representation has enormous (exponential) length. Hence we desire a representation which is more suitable for computation. This is accomplished using techniques from index calculus algorithms for imaginary quadratic fields. We let $\mathcal{F}'$ denote the set of prime ideals of $\mathcal{O}$ lying above those primes in $\mathcal{F}$.

The factor base $\mathcal{F}$ we will choose, for complexity theoretic purposes, to consist of those primes l which either split or ramify in $\mathcal{O}$ whose size is less than $L = 6(\log \Delta)^2 = O((\log q^n)^2)$. In real life, since we will need to use the modular polynomials of degree less than L, we would select $L = 1000$.

We repeatedly compute the following reduced ideal

$$\mathfrak{b} = \text{Reduce}\left(\mathfrak{a} \prod_{\mathfrak{l} \in \mathcal{F}'} \mathfrak{l}_i^{a_i}\right),$$

where a_i is chosen randomly, until the ideal $\mathfrak{b}$ factorises over the factor base $\mathcal{F}$ as

$$\mathfrak{b} = \prod_{\mathfrak{l} \in \mathcal{F}'} \mathfrak{l}_i^{b_i}.$$

We then have

$$\mathfrak{a} \equiv \prod_{\mathfrak{l} \in \mathcal{F}'} \mathfrak{l}_i^{b_i - a_i}. \tag{1}$$

The size of the b_i's are bounded since the ideal $\mathfrak{b}$ is reduced. Heuristically (see below) we expect to require at most $q^{n/4}$ choices of the a_i before we obtain a value of $\mathfrak{b}$ which is sufficiently smooth. Hence, if we assume $|a_i| \leq t$ then we will require

$$t^{\#\mathcal{F}} \geq q^{n/4}.$$

In other words the value of t can be taken to be of polynomial size in $n\log(q)$. Hence, we require polynomial storage to hold the isogeny as a smooth ideal.

To estimate the running time we need to examine the probability of obtaining a smooth number. We are essentially testing whether an integer of size $\sqrt{\Delta}$, i.e. the norm of a reduced ideal, factors over a factor base of integers less than $L = 6(\log(q^n))^2$. Standard estimates give this probability asymptotically as u^{-u} where $u = \log(\Delta)/\log(L)$. We, hence, obtain the probability as

$$u^{-u} \approx q^{n(-1+c/(\log\log q))/4}.$$

Therefore the complexity of this step will be $q^{n/4+\epsilon}$.

For real life applications we actually try to factor an integer of size 2^{80} over a factor base of primes less than 1000. The probability of this working is 2^{-24}, hence the complexity in practical examples is still less than $q^{n/4} = 2^{160/4} = 2^{40}$.

The probability is of finding a factorisation of this form is $\Psi(x, 1000)/x$. For completeness we give the following table on approximate values of $\Psi(x, 1000)$ for values of x of interest in our situation. These values have been computed by Dan Bernstein [2]. By examining the table we see that, for all values of n in the range $100 \leq n \leq 240$, **Stage 2** of our algorithm will run in time $O(q^{n/4})$.

$x = q^{n/2}$	$\approx$ $\log_2 \Psi(x, 1000)$	$\approx$ $\log_2(\Psi(x, 1000)/x)$
2^{50}	39	-11
2^{60}	45	-15
2^{70}	51	-19
2^{80}	56	-24
2^{90}	61	-29
2^{100}	66	-34
2^{110}	71	-39
2^{120}	75	-45

3.5 Stage 3

Finally we can write down the isogeny between E_1 and E_2. This is done by taking each prime ideal in equation (1) and applying the method previously given to determine the associated j-invariant. We comment that negative powers of a prime $\mathfrak{l}$ correspond to positive powers of the complex conjugate ideal $\overline{\mathfrak{l}}$. The actual isogeny is determined by Vélu's formulae [18].

Chaining these isogenies together we obtain the final desired map from E_1 to E_2. In practice we do not actually write down the isogeny but simply evaluate the isogeny on the points of interest.

The ideal $\mathfrak{b}$ in **Stage 2** will have norm at most $O(\sqrt{\Delta})$. The smooth representation of the ideal equivalent to $\mathfrak{a}$ will have at most $O(\log \Delta)$ not necessarily distinct factors in it, each factor corresponding to an isogeny of degree at most L. Hence, the mapping of points from E_1 to E_2, given the smooth representation of the ideal equivalent to $\mathfrak{a}$, can be performed in time polynomial in $\log q$.

4 Searching for Isogenous Curves

Theorem 2 shows that there is a set of possible values for b such that the elliptic curves $Y^2 + XY = X^3 + b$ and $Y^2 + XY = X^3 + X^2 + b$ have a specific small value for m. For our application it is necessary to be able to generate at least one representative for each isogeny class of elliptic curves with this small m. In this section we discuss methods to achieve this.

By Theorem 2, the set of possible values b is equal to the union of the sets

$$B_i = \{b \in \mathbb{F}_{q^n} : (\sigma - 1) f_i(\sigma) b = 0\}.$$

for $1 \leq i \leq s$.

An element b lies in two of these sets if there are indices i and j such that $(\sigma - 1) f_i(\sigma) b = (\sigma - 1) f_j(\sigma) b = 0$. Since $\gcd(f_i(x), f_j(x)) = 1$ it follows that $(\sigma - 1) b = 0$. Therefore $B_i \cap B_j = \mathbb{F}_q$ for all $i \neq j$.

Each of these sets will be handled in turn, so from now on we fix an index i and define $f(x) = (x - 1) f_i(x)$ and $B = \{b \in \mathbb{F}_{q^n} : f(\sigma) b = 0\}$.

Let α be a normal basis generator for $K = \mathbb{F}_{2^{rn}}$ over $\mathbb{F}_2$ In other words $\{\alpha, \alpha^2, \ldots, \alpha^{2^{nr-1}}\}$ is a vector space basis for K over $\mathbb{F}_2$. It is a fundamental fact that such an element α exists and that $K = \{g(\pi)\alpha : g(x) \in \mathbb{F}_2[x]\} = \{g(\sigma)\alpha : g(x) \in \mathbb{F}_q[x]\}$.

The following result is an easy exercise.

Lemma 1. *Let the notation be as above. Write $h(x) = (x^n - 1)/f(x)$ and define $\alpha' = h(\sigma)\alpha$. Then*

$$B = \{g(\sigma)\alpha' : g(x) \in \mathbb{F}_q[x]\}. \tag{2}$$

Indeed, in the above it is enough to let the $g(x)$ run over a set of representatives for the quotient ring $\mathbb{F}_q[x]/(f(x))$ (i.e., over all elements of $\mathbb{F}_q[x]$ of degree less than $\deg(f(x))$).

Lemma 1 gives an efficient algorithm to compute representatives for each set B which has running time $O(\#B)$. Clearly $\#B = q^{\deg(f(x))}$.

By taking the union of these sets we obtain all the values for b which we require. In the notation of Theorem 2 we have s values for i and $\deg(f_i(x)) = t$, therefore $\#B_i = q^{t+1}$ for each index i and, since the intersection of any two B_i has size q it follows that $\#(\cup_i B_i) = qs(q^t - 1) + q$ as claimed in Theorem 2.

4.1 A more refined approach

Recall that the goal is to produce a representative of each isogeny class of elliptic curves with small values for m. Since r and n are odd we may assume that all elliptic curves have the form $E : y^2 + xy = x^3 + ax^2 + b$ with $a \in \{0, 1\}$. If E is an elliptic curve which has a small value for m then $E^\pi : y^2 + xy = x^3 + ax^2 + \pi(b)$ is an isogenous elliptic curve. Since $f(\sigma)\pi b = \pi f(\sigma) b = 0$ it follows that E^π necessarily has the same small value for m. Rather than listing all $b \in B$ it would be better to find a set $B' \subset B$ such that $B = \{\pi^j(b) : b \in B', 1 \leq j \leq rn\}$.

From Lemma 1 we have $B = \{g(\sigma)\alpha' : g(x) \in \mathbb{F}_q[x]\}$. We want to work in terms of π rather than $\sigma = \pi^r$. An analogous argument to that used to prove Lemma 1 gives

$$B = \{g(\pi)\alpha' : g(x) \in \mathbb{F}_2[x]\}. \tag{3}$$

In this case the polynomials $g(x) \in \mathbb{F}_2[x]$ may be taken to have degree less than $\deg(f(x^r))$. Given any two elements $b_j = g_j(\pi)\alpha'$ for $j \in \{1, 2\}$ and any two elements $c_j \in \mathbb{F}_2$ we clearly have $c_1 b_1 + c_2 b_2 = (c_1 g_1 + c_2 g_2)(\pi)\alpha' \in B$. It is easy to show that B is a module over $\mathbb{F}_2[\pi]$ and that the following result holds.

Lemma 2. *Let notation be as above. There is an isomorphism of rings from $\mathbb{F}_2[\pi]$ to $\mathbb{F}_2[x]$ and there is a corresponding isomorphism of modules from the $\mathbb{F}_2[\pi]$-module B to the $\mathbb{F}_2[x]$-module $\mathbb{F}_2[x]/(f(x^r))$. The isomorphism is given by*

$$g(\pi)\alpha' \longmapsto g(x).$$

Consider the factorisation $f(x^r) = \prod_{j=1}^{l} f_j(x)$ into irreducibles. The polynomials $f_j(x)$ are all distinct since $x^{rn} - 1$ has no repeated roots (when r and n are both odd then $\gcd(x^{rn} - 1, rnx^{rn-1}) = 1$). The Chinese remainder theorem for polynomials implies that we have the following isomorphism of $\mathbb{F}_2[x]$-modules

$$\mathbb{F}_2[x]/(f(x^r)) \cong \bigoplus_{j=1}^{l} \mathbb{F}_2[x]/(f_j(x)). \tag{4}$$

The terms on the right hand side are finite fields $K_j = \mathbb{F}_2[x]/(f_j(x))$.

Combining Lemma 2 and equation (4) we see that the $\mathbb{F}_2[\pi]$-module B is isomorphic to the $\mathbb{F}_2[x]$-module $(\oplus_j K_j)$. We need to understand the action of π on elements of B, and this corresponds to multiplication of elements of $(\oplus_j K_j)$ by x.

Hence, write N_j for the normal subgroup of K_j^* generated by x, so that $x(gN_j) = gN_j$ for any $g \in K_j^*$. It follows that the cosets of K_j^*/N_j give representatives for the x-orbits of elements of K_j^* (the zero element of K_j must be handled separately). We write $c_j = [K_j^* : N_j]$ for the index (i.e., the number of distinct cosets). Let ζ_j be a generator for the cyclic group K_j^*, then $N_j = \langle \zeta_j^{c_j} \rangle$ (i.e., $x = \zeta_j^{dc_j}$ for some d which depends on ζ_j).

However, it is not possible to study the fields K_j individually. Instead, we have to consider the product $\oplus_j K_j$ and have to consider that the action of x on this product is multiplication by x on every coordinate.

The following result gives an explicit set of representatives for $(\oplus_j K_j)/\langle x\rangle$. This is essentially achieved by forming a diagonalisation (under the action of x) of the set $\oplus_j K_j$.

Lemma 3. *The set*

$$\{g_1 \oplus \cdots \oplus g_{a-1} \oplus \zeta_a^b : 1 \le a \le l, 0 \le b < c_a, g_j \in K_j\} \cup \{0\}$$

is a set of representatives for $(\oplus_j K_j)/\langle x\rangle$.

Proof. The result follows from two facts. First, every element of $\oplus_j K_j$ is of the given form if we disregard the condition $0 \le b < c_a$. Second, x acts as

$$x^i(g_1 \oplus \cdots \oplus g_{a-1} \oplus \zeta_a^b) = x^i(g_1) \oplus \cdots \oplus x^i(g_{a-1}) \oplus x^i(\zeta_a^b),$$

and $x^i(\zeta_a^b) = \zeta_a^{b+idc_a}$. □

This set of representatives can be easily mapped to B as follows. For any $g_1 \oplus \cdots \oplus g_l$ we can represent each $g_j \in K_j = \mathbb{F}_2[x]/(f_j(x))$ as a polynomial $g_j(x) \in \mathbb{F}_2[x]$. The Chinese remainder algorithm gives $g(x) \in \mathbb{F}_2[x]/(f(x))$ which reduces to each $g_j(x)$ modulo $(f_j(x))$. We then obtain the corresponding value $b = g(\pi)\alpha' \in B$.

Thus we obtain an algorithm to compute a set B' of elements whose π-orbits generate B.

4.2 Example

We give an example of the refined approach in a simplified way. Let us consider $K = \mathbb{F}_{155}$, $n = 31$, $r = 5$, $q = 2^5$ and $m = 6$. The factorisation of $x^{31} - 1$ over $\mathbb{F}_2$ is

$$(x-1)\prod_{i=1}^{6} f_i(x)$$

with $f_i(x)$ of degree 5. For brevity we skip the actual values of the $f_i(x)$ and subsequent polynomials. We obtain $s = 6$, $t = 5$ and $B_i \cap B_j = \mathbb{F}_q$ for $i \ne j$.

We now write $f(x)$ and B instead of $f_i(x)$ and B_i for all i in turn (note the difference in defining $f(x)$ compared to before). Over $\mathbb{F}_2$ we have the factorisation

$$f(x^r) = g_1(x)g_2(x)$$

with $\deg(g_1(x)) = 5$ and $\deg(g_2(x)) = 20$. Using the Chinese remainder theorem again yields

$$\begin{aligned} B &\cong \mathbb{F}_2[x]/((x^r-1)f(x^r)) \\ &\cong \mathbb{F}_2[x]/((x^r-1)g_1(x)) \ \oplus \ \mathbb{F}_2[x]/(g_2(x)). \end{aligned}$$

The first quotient ring above contains $2^{\deg((x^r-1)g_1(x))} = 1024$ elements. The second quotient ring $\mathbb{F}_2[x]/(g_2(x))$ is isomorphic to $\mathbb{F}_{2^{20}}$ and the index of the group

generated by the element $x + (g_2(x))$ in the full multiplicative group equals $(2^{20} - 1)/nr = 6765$. Let $\zeta \in \mathbb{F}_2[x]/(g_2(x))$ be a generator of the full multiplicative group. Mapping all elements of $\mathbb{F}_2[x]/((x^r - 1)g_1(x))$ and the elements $\{\, \zeta^j \,|\, 0 \leq j < 6765 \,\} \cup \{0\}$ under the above isomorphisms to K gives sets of elements $B_1, B_2 \subseteq K$ such that

$$B = \bigcup_{j=0}^{rn-1} \pi^j \big(B_1 + B_2\big).$$

In $B_1 + B_2$ clearly only the 1024 elements of B_1 can possibly be conjugated under π which makes up only a very small fraction of all the $1024 \cdot 6676$ elements of $B_1 + B_2$.

We conclude that altogether for $1 \leq i \leq 6$ we obtain a set of $6 \cdot 1024 \cdot 6766$ representatives of classes under the action of powers of π in $\cup_i B_i$ with only small redundancy due to double occurrences or the action of powers of π.

5 Implications

We first state that for curves over large prime fields the techniques of Weil restriction do not apply. Secondly for curves over fields of the form $\mathbb{F}_{2^p}$, where p is prime, Menezes and Qu [14] showed that for curves of cryptographic interest the GHS attack does not apply. Hence, for the rest of this section we will concentrate on the case where $K = \mathbb{F}_{q^n}$ where q is a non-trivial power of two and n is a prime such that $5 \leq n < 40$.

Assume all isomorphism classes are distributed uniformly over the isogeny classes (this is in accordance with the Sato-Tate conjecture). Hence, when choosing a curve uniformly at random we are essentially choosing an isogeny class at random. The number of isogeny classes is around $2q^{n/2}$ due to the Hasse-Weil bounds and the fact that we are in characteristic two.

Let s and t be defined as in Theorem 2, then this Theorem states that an upper bound on the number of isomorphism classes with the smallest non-trivial value of m is given roughly by $2sq^{t+1}$. However, since the Frobenius operator maps curves with small m to curves with the same value of m, we really need to consider the number of non-isomorphic curves with the smallest non-trivial value of m up to the trivial isogeny of the Frobenius operator. The number of these is roughly

$$\frac{2sq^{t+1}}{nr}.$$

Finally, we can deduce that probability that a random elliptic curve over K lies in an isogeny class which contains a curve with the smallest non-trivial value of m is approximated by

$$p = \min(1, \mathfrak{p})$$

where

$$\mathfrak{p} = \frac{sq^{t+1-n/2}}{nr}.$$

In the case $s = 1$ then this is almost every isomorphism class hence our extension does not produce significantly more weak curves. The only values of n under consideration with $s \neq 1$ are $n = 7, 17, 23$ and 31.

For values of $s = 2$ we obtain probably the most interesting case, since if n is fixed and $q \to \infty$ then eventually we should be able to consider all curves using the new method. This is particularly interesting when $n = 7$, since in this case we obtain $m = 4$ and $s = 2$. Hence, using isogenies it is perfectly possible that every curve over $\mathbb{F}_{q^7}$ can be attacked by reducing to a hyperelliptic curve of genus 7 or 8 over $\mathbb{F}_q$. The work in [17] shows that hyperelliptic curves of genus 8 should be considered weak for cryptographic purposes. However, before we can eliminate $n = 7$ completely we need to consider the complexity of the extended GHS attack in more detail.

Recall our extended GHS attack first enumerates the set B of Theorem 2 using the techniques of Section 4, for the case of $n = 7$ this will take time $O(\#B)$. Then having, hopefully, found an isogenous curve we need to determine the explicit isogeny. Using the techniques of Section 3 this takes time $O(q^{7/4+\epsilon})$. Gaudry's algorithm on the resulting curve of genus 7 or 8 is practical and runs in time $O(q^2)$. Hence, for $n = 7$ we find the extended GHS attack should be faster than using Pollard rho on $E(\mathbb{F}_{q^7})$, for almost all curves as long as q is large enough.

To show that our heuristic probability argument is correct we enumerated all the curves for small values of q and $n = 7$ to obtain the following comparison between the value of our formula for p and the actual value of p.

q	2	2^2	2^3	2^4	2^5
Formula p	0.404	0.286	0.269	0.286	0.323
Actual p	0.470	0.290	0.280	≈ 0.285	≈ 0.331

The actual p is either exact (computed via enumeration of all possibilities) or is an approximation computed via a Monte Carlo experiment. When $r = 10$ (i.e., $q = 1024$) one finds that the formula for p gives the value 0.914 and that for $r \geq 11$ our formula gives $p = 1$. This suggests that all cryptographic sized fields with $n = 7$ are weak with respect to the extended GHS attack.

The only prime value of n with $s \geq 3$ in the range $5 \leq n < 40$ is $n = 31$. This is particularly interesting given that the field $\mathbb{F}_{2^{155}}$ is used for an elliptic curve group in the IPSEC set of protocols for key agreement. See [1] for a description of the curve used and [17] for a previous analysis of this curve using the GHS attack with $n = 5$.

When $n = 31$ and $q = 2^r$ the proportion of all curves which succumb to the basic GHS attack is approximately $6q^{-25}$. For the extended attack, assuming the values of m are distributed evenly over the isogeny classes, the proportion of all affected curves is

$$\approx \frac{6q^{-9.5}}{31r}.$$

By the work of Jacobson, Menezes and Stein [10] we know that the GHS method is feasible for curves of cryptographic interest when $m = 6$ and $q = 2^5$.

For the IPSEC curve over $\mathbb{F}_{2^{155}}$ this means there is roughly a one in 2^{-52} chance that the curve can be attacked using the extended GHS attack, as opposed to 2^{-122} for the standard GHS attack. Using the methods in Section 4 we searched all isogeny classes to see if there was a curve with $m = 6$ which was isogenous to the IPSEC curve. This search took 31 days on a 500 Mhz Pentium III using the Magma package. Not surprisingly we did not find an isogenous curve and so we can conclude that the IPSEC curve is not susceptible to the extended GHS attack. This shows that further research into Weil descent is required before it can be shown that all elliptic curves over composite extension fields are weak.

6 Acknowledgement

The authors thank Pierrick Gaudry for useful conversations.

References

1. IETF. The Oakley Key Determination Protocol. *IETF RFC 2412*, Nov 1998.
2. D.J. Bernstein. Bounds on $\Psi(x, y)$. **http://cr.yp.to/psibound.html**.
3. I.F. Blake, G. Seroussi and N.P. Smart. *Elliptic Curves in Cryptography.* Cambridge University Press, 1999.
4. J.-M. Couveignes. Computing l-isogenies using the p-torsion. *Algorithmic Number Theory Symposium- ANTS II*, Springer-Verlag LNCS 1122, 59–65, 1996.
5. G. Frey. "How to disguise an elliptic curve". Talk at ECC' 98, Waterloo.
6. G. Frey and H. Rück. A remark concerning m-divisibility and the discrete logarithm in the divisor class group of curves. *Math. Comp.*, **62**, 865–874, 1994.
7. S.D. Galbraith. Constructing isogenies between elliptic curves over finite fields. *LMS J. Comput. Math.*, **2**, 118–138, 1999.
8. S.D. Galbraith and N.P. Smart. A Cryptographic application of Weil descent. *Codes and Cryptography*, Springer-Verlag LNCS 1746, 191–200, 1999.
9. P. Gaudry, F. Hess and N.P. Smart. Constructive and destructive facets of Weil descent on elliptic curves. To appear *J. Cryptology.*
10. M. Jacobson, A. Menezes and A. Stein. Solving elliptic curve discrete logarithm problems using Weil descent. Preprint, 2001.
11. D. Kohel. *Endormorphism rings of elliptic curves over finite fields.* Phd Thesis, Berkeley, 1996.
12. R. Lercier. Computing isogenies in $\mathbb{F}_{2^n}$. *Algorithmic Number Theory Symposium-ANTS II*, Springer-Verlag LNCS 1122, 197–212, 1996.
13. A. Menezes, T. Okamoto and S. Vanstone. Reducing elliptic curve logarithms to logarithms in finite fields. *IEEE Trans. on Infor. Th.*, **39**, 1639–1646, 1993.
14. A. Menezes and M. Qu. Analysis of the Weil descent attack of Gaudry, Hess and Smart. *Topics in Cryptology - CT-RSA 2001*, Springer-Verlag LNCS 2020, 308–318, 2001.
15. S. Pohlig and M. Hellman. An improved algorithm for computing logarithms over $GF(p)$ and its cryptographic significance. *IEEE Trans. on Infor. Th.*, **24**, 106–110, 1978.
16. J. Pollard. Monte Carlo methods for index computations mod p. *Math. Comp.*, **32**, 918–924, 1978.

17. N.P. Smart. How secure are elliptic curves over composite extension fields? *Advances in Cryptology - EUROCRYPT '01*, Springer-Verlag LNCS 2045, 30–39, 2001.
18. J. Vélu. Isogénies entre courbes elliptiques. *Comptes Rendus l'Acad. Sci. Paris, Ser. A*, **273**, 238-241 1971.