

Cover and Decomposition Index Calculus on Elliptic Curves made practical

Application to a previously unreachable curve over \mathbb{F}_{p^6}

Antoine Joux¹ and Vanessa Vitse²

¹ DGA and Université de Versailles Saint-Quentin, Laboratoire PRISM, 45 avenue des États-Unis, F-78035 Versailles cedex, France

`antoine.joux@m4x.org`

² Université de Versailles Saint-Quentin, Laboratoire PRISM, 45 avenue des États-Unis, F-78035 Versailles cedex, France

`vanessa.vitse@prism.uvsq.fr`

Abstract. We present a new “cover and decomposition” attack on the elliptic curve discrete logarithm problem, that combines Weil descent and decomposition-based index calculus into a single discrete logarithm algorithm. This attack applies, at least theoretically, to all composite degree extension fields, and is particularly well-suited for curves defined over \mathbb{F}_{p^6} . We give a real-size example³ of discrete logarithm computations on a curve over a 151-bit degree 6 extension field, which would not have been practically attackable using previously known algorithms.

Key words: elliptic curve, discrete logarithm, index calculus, Weil descent, decomposition attack

1 Introduction

Elliptic curves are used in cryptography to provide groups where the discrete logarithm problem is thought to be difficult. We recall that given a finite group G (written additively) and two elements $P, Q \in G$, the discrete logarithm problem (DLP) consists in computing, when it exists, an integer x such that $Q = xP$. When elliptic curves are used in cryptographic applications, the DLP is usually considered to be as difficult as in a generic group of the same size [31]. As a consequence, for a given security level, the key size is much smaller than for other popular cryptosystems based on factorization or discrete logarithms in finite fields. The first elliptic curves considered in cryptography were defined over either binary or prime fields [20, 24]. But to speed up arithmetic computations, it has been proposed to use various forms of extension fields. In particular, Optimal Extension Fields have been proposed in [4] to offer high performance

³ This work was granted access to the HPC resources of CCRT under the allocation 2010-t201006445 made by GENCI (Grand Equipement National de Calcul Intensif)

in hardware implementations. They are of the form \mathbb{F}_{p^d} where p is a pseudo-Mersenne prime and d is such that there exists an irreducible polynomial of the form $X^d - \omega \in \mathbb{F}_p[X]$. In most examples, the degree d of the extension is rather small. However, when curves defined over extension fields are considered, some non-generic attacks, such as the Weil descent or decomposition attacks, can be applied. The first one aims at transferring the DLP from $E(\mathbb{F}_{q^n})$ to the Jacobian of a curve \mathcal{C} defined over \mathbb{F}_q and then uses index calculus on this Jacobian [2, 12, 15] to compute the logarithm; it works well when the genus of the curve \mathcal{C} is small, ideally equal to n , but this occurs quite infrequently in practice. Many articles have studied the scope of this technique (cf. [7, 10, 11, 14, 16]), but even on vulnerable curves, the Weil descent approach is often just a little more efficient than generic attacks on the DLP. Decomposition-based index calculus, or decomposition attack, is a more recent algorithm (see [9, 13, 18, 26]), which applies equally well to all (hyper-)elliptic curves defined over an extension field. Its asymptotic complexity is promising, but in practice, due to large hidden constants in the complexity, it becomes better than generic attacks for group sizes too large to be threatened anyway.

In this article, we combine both techniques into a cover and decomposition attack, which applies as soon as the extension degree is composite. The idea is to first transfer the DLP to the Jacobian of a curve defined on an intermediate field, then use the decomposition method on this sub-extension instead of the classical index calculus. This new attack is not a mere theoretical possibility: we give concrete examples of curves defined over \mathbb{F}_{p^6} that are practically secure against all other attacks, but for which our method allows to solve the DLP in a reasonable time. In particular, we have been able to compute logarithms on a 149-bit elliptic curve group defined over a degree 6 extension field in about a month real-time, using approximately 110 000 CPU.hours.

The paper is organized as follows: first we briefly recall in Section 2 the principles of Weil descent and of the decomposition method. We then give an explicit description of our attack in Section 3, introducing a useful variant of the decomposition step that can be of independent interest. In particular, we study the case of elliptic curves defined over \mathbb{F}_{p^6} or \mathbb{F}_{p^4} , list all the potentially vulnerable curves and give a complexity analysis and a comparison with previously known attacks. Finally, in Section 6, we describe in details the computations on our 149-bit example.

2 Survey of previous work

2.1 Weil descent and cover attacks

Weil descent has been first introduced in cryptography by Frey [10]; the idea is to view an abelian variety A of dimension d defined over an extension field K/k as an abelian variety $W_{K/k}$ of dimension $d \cdot [K : k]$ over k . If $W_{K/k}$ turns out to be the Jacobian of a curve \mathcal{C}_k or can be mapped into such a Jacobian, then the discrete logarithm in $A(K)$ can be transferred to $\text{Jac}_{\mathcal{C}}(k)$, where it may become much weaker due to the existence of efficient index calculus algorithms.

When the genus of \mathcal{C} is small relative to the cardinality p of k , the complexity is in $O((g^2 \log^3 p)g!p + (g^2 \log p)p^2)$ as p grows to infinity [12]; the first term comes from the relation search and the second from the sparse linear algebra. Following [15], it is possible to rebalance these two terms by using a double large prime variation. In this variant, only a small number p^α of *prime divisors*⁴ are considered as genuine, while the rest of the prime divisors are viewed as “large primes”. The optimal value of α depends of the cost of the two phases; asymptotically the choice that minimizes the total running time is $1 - 1/g$, yielding a complexity in $\tilde{O}(p^{2-2/g})$ for fixed g as p goes to infinity.

The main difficulty of this Weil descent method is to find the curve \mathcal{C} . This problem was first addressed for binary fields by Gaudry, Hess and Smart (GHS [14]) and further generalized by Diem [7] in odd characteristic. To attack an elliptic curve E defined over \mathbb{F}_{p^n} (p a prime power), the GHS algorithm builds a curve \mathcal{C} defined over \mathbb{F}_p such that there exists a cover map $\pi : \mathcal{C} \rightarrow E$ defined over \mathbb{F}_{p^n} . The construction is more easily explained in terms of function fields: the Frobenius automorphism $\sigma_{\mathbb{F}_{p^n}/\mathbb{F}_p}$ can be extended to the composite field $F' = \prod_{i=0}^{n-1} \mathbb{F}_{p^n}(E^{\sigma^i})$, and the function field $F = \mathbb{F}_p(\mathcal{C})$ is defined as the subfield of F' fixed by σ . The GHS algorithm then uses the so-called conorm-norm map $N_{F'/F} \circ \text{Con}_{F'/\mathbb{F}_{p^n}(E)}$ to transfer the discrete logarithm from $E(\mathbb{F}_{p^n})$ to $\text{Jac}_{\mathcal{C}}(\mathbb{F}_p)$. An important condition is that the kernel of this map must not intersect the subgroup in which the discrete logarithm takes place, but as remarked in [7, 16], this is not a problem in most cryptographically relevant situations. This technique is efficient when the genus g of \mathcal{C} is close to n . In particular, for some specific finite fields most elliptic curves are “weak” in the sense that Weil descent algorithms are better, if only by a small margin, than generic attacks [23]. Indeed, when the GHS method does not provide any low genus cover for E , it may be possible to find a sequence of low degree isogenies (a.k.a. an *isogeny walk*) from E to another, more vulnerable elliptic curve E' [11]. Nevertheless, we emphasize that the security loss is quite small for a random curve, and for most curves on most fields \mathbb{F}_{p^n} , g is of the order of 2^n which means that index calculus in the Jacobian of \mathcal{C} is slower than generic attacks on $E(\mathbb{F}_{p^n})$.

2.2 Decomposition attack

Index calculus has become ubiquitous in the last decades for the DLP resolution. However its direct application to elliptic curves faces two major challenges: contrarily to finite fields or hyperelliptic curves, there is no natural choice of factor base and no equivalent of the notion of factorization of group elements. The first main breakthrough was achieved in 2004 by Semaev [30] when he suggested to replace factorization by decomposition into a fixed number of points; for that purpose, he introduced the summation polynomials which give an algebraic expression of the fact that a given point decomposes into a sum of factor base elements. But for a lack of an adequate factor base, this approach fails

⁴ The term *prime divisor* is an abuse of language that denotes the linear irreducible polynomials that are used in the index calculus algorithm on $\text{Jac}_{\mathcal{C}}(k)$

in the general case. Then Gaudry and Diem [9, 13] independently proposed to use Semaev’s idea to attack all curves defined over small degree extension fields $\mathbb{F}_{p^n}/\mathbb{F}_p$. Their method shares the basic outline of index calculus, but to distinguish it from what has been presented in the previous subsection, we follow [26] and call it the decomposition attack. On $E(\mathbb{F}_{p^n})$, a convenient choice of factor base is the set of rational points of the curve having their abscissae in the base field \mathbb{F}_p . By combining Semaev’s summation polynomials and restriction of scalars, the relation search then becomes a resolution of a multivariate polynomial system over \mathbb{F}_p . The complexity of this approach can be estimated using double large prime variation by $\tilde{O}(p^{2-2/n})$ for fixed n as p grows to infinity. Unfortunately, the hidden constants in this complexity become very large as n grows, and the resolution of the systems is intractable as soon as $n \geq 4$ (or $n \geq 5$ with the variant of [18]).

The decomposition attacks can also be applied to higher genus curves. However, Semaev’s polynomials are no longer available in this case and the algebraic expression of the group law is more complicated. In [26], Nagao proposes an elegant way to circumvent this problem, using divisors and Riemann-Roch spaces. For hyperelliptic curves, the decomposition search then amounts to solving a quadratic multivariate polynomial system. This approach is less efficient than Semaev’s in the elliptic case, but is the simplest otherwise. For fixed extension degree n and genus g , the complexity of a decomposition attack is in $\tilde{O}(p^{2-2/ng})$ with a double large prime variation. Again, the resolution of the polynomial system is the main technical difficulty, and is easily feasible for only very few couples (n, g) , namely $(2, 2)$, $(2, 3)$ and $(3, 2)$.

3 Cover and Decomposition attack

Let $\mathbb{F}_{q^d}/\mathbb{F}_p$ be an extension of finite fields, where q is a power of p (in most applications p denotes a large prime but in general, it can be any prime power), and let E be an elliptic curve defined over \mathbb{F}_{q^d} of cryptographic interest, i.e. containing a subgroup G of large prime order. As E is defined over an extension field, it is subject to the attacks presented above. But if the degree $[\mathbb{F}_{q^d} : \mathbb{F}_p]$ of the extension is larger than 5, then we have seen that E is practically immune to decomposition attacks. In the following, we assume that the potential reduction provided by the GHS attack or its variants is not significant enough to threaten the security of the DLP on the chosen curve E .

When q is a strict power of p , we have a tower of extensions given by $\mathbb{F}_{q^d}/\mathbb{F}_q$ and $\mathbb{F}_q/\mathbb{F}_p$. In this context, it becomes possible to combine both cover and decomposition methods and obtain an efficient attack of the DLP on E . The idea is to use Weil descent on the first extension $\mathbb{F}_{q^d}/\mathbb{F}_q$ to get a cover defined over \mathbb{F}_q , with small enough⁵ genus g . Then we can apply a decomposition attack on the curve thus obtained, making use of the second extension $\mathbb{F}_q/\mathbb{F}_p$. As this *cover and*

⁵ Meaning that g should be small relatively to the genus that could be obtained by direct Weil descent, using the extension $\mathbb{F}_{q^d}/\mathbb{F}_p$.

decomposition attack is more efficient when Weil descent provides a hyperelliptic cover over the intermediate field, we focus on this case in the following.

3.1 Description of the attack

We now explicitly detail this cover and decomposition approach. We suppose first that there exists an imaginary hyperelliptic curve \mathcal{H} of small genus g with equation $y^2 = h(x)$, defined over \mathbb{F}_q , together with a covering map $\pi : \mathcal{H} \rightarrow E$ defined over \mathbb{F}_{q^d} . This can be obtained by the GHS attack or its variants, possibly preceded by an isogeny walk. This cover classically allows to transfer the DLP from G to a subgroup $G' \subset \text{Jac}_{\mathcal{H}}(\mathbb{F}_q)$ via the conorm-norm map $N_{\mathbb{F}_{q^d}/\mathbb{F}_q} \circ \pi^* : E(\mathbb{F}_{q^d}) \simeq \text{Jac}_E(\mathbb{F}_{q^d}) \rightarrow \text{Jac}_{\mathcal{H}}(\mathbb{F}_q)$, assuming that $\ker(N_{\mathbb{F}_{q^d}/\mathbb{F}_q} \circ \pi^*) \cap G = \{\mathcal{O}_E\}$.

The decomposition part of the attack is adapted from Gaudry and Nagao; since it is quite recent, we detail the method. As in all index calculus based approaches, there are two time-consuming steps: first we have to collect relations between factor base elements, then we compute discrete logarithms by using linear algebra on the matrix of relations. We consider the same factor base as [13, 26]

$$\mathcal{F} = \{D_Q \in \text{Jac}_{\mathcal{H}}(\mathbb{F}_q) : D_Q \sim (Q) - (\mathcal{O}_{\mathcal{H}}), Q \in \mathcal{H}(\mathbb{F}_q), x(Q) \in \mathbb{F}_p\},$$

which contains approximately p elements. As usual, we can use the hyperelliptic involution ι to reduce the size of \mathcal{F} by a factor 2, so that about $p/2$ relations are needed.

Let n be the extension degree $[\mathbb{F}_q : \mathbb{F}_p]$. In Nagao's decomposition method, one tries to decompose an arbitrary divisor D (typically obtained by considering a large multiple of some element in \mathcal{F}) into a sum of ng divisors in the factor base

$$D \sim \sum_{i=1}^{ng} ((Q_i) - (\mathcal{O}_{\mathcal{H}})). \quad (1)$$

Heuristically, there exist approximately $p^{ng}/(ng)!$ distinct sums of ng elements of \mathcal{F} , so the probability that a given divisor D is decomposable can be estimated by $1/(ng)!$. To check if D can be decomposed, one considers the Riemann-Roch \mathbb{F}_q -vector space

$$\mathcal{L}(ng(\mathcal{O}_{\mathcal{H}}) - D) = \{f \in \mathbb{F}_q(\mathcal{H})^* : \text{div}(f) \geq D - ng(\mathcal{O}_{\mathcal{H}})\} \cup \{0\}.$$

We can assume that the divisor D is reduced and has Mumford representation $(u(x), v(x))$ with $\deg u = g$, so that this \mathbb{F}_q -vector space is spanned by $u(x), u(x)x, \dots, u(x)x^{m_1}, (y - v(x)), x(y - v(x)), \dots, x^{m_2}(y - v(x))$, where $m_1 = \lfloor (n-1)g/2 \rfloor$ and $m_2 = \lfloor ((n-1)g - 1)/2 \rfloor$. A function $f = \lambda_0 u(x) + \lambda_1 u(x)x + \dots + \lambda_{m_1} u(x)x^{m_1} + \mu_0 (y - v(x)) + \mu_1 x(y - v(x)) + \dots + \mu_{m_2} x^{m_2} (y - v(x))$ vanishes on the support of D and exactly ng other points (counted with multiplicity and possibly defined on the algebraic closure of \mathbb{F}_q) if its top-degree coefficient is not zero. We are looking for a condition on $\lambda_0, \dots, \lambda_{m_1}, \mu_0, \dots, \mu_{m_2} \in \mathbb{F}_q$ such that the zeroes Q_1, \dots, Q_{ng} of f disjoint from $\text{Supp}(D)$ have x -coordinate

in \mathbb{F}_p ; this event yields a relation as in (1). Therefore, we consider the polynomial $F(x) = f(x, y)f(x, -y)/u(x)$ where y^2 has been replaced by $h(x)$. Without loss of generality, we can fix either $\lambda_{m_1} = 1$ or $\mu_{m_2} = 1$ in order to have F monic of degree ng . The roots of F are exactly the x -coordinates of the zeroes of f distinct from $\text{Supp}(D)$, thus we are looking for the values of λ and μ for which F splits in linear factors over \mathbb{F}_p . A first necessary condition is that all of its coefficients, which are quadratic polynomials in λ and μ , belong to \mathbb{F}_p ; a scalar restriction on these coefficients then yields a quadratic polynomial system of $(n-1)ng$ equations and variables coming from the components of the variables λ and μ . The corresponding ideal is generically of dimension 0, and the solutions of the system can be found using e.g. a Gröbner basis computation. Since the number of systems to solve is huge (on average $(ng)! \cdot p/2$, or more if a large prime variation is applied), techniques such as the F4 variant of [19] should be preferred. Once the solutions in \mathbb{F}_q are obtained, it remains to check if the resulting polynomial F splits in $\mathbb{F}_p[x]$, and if it is the case, to compute the corresponding decomposition of D .

In this article, we also consider a somewhat different approach to the relation search that offers some similarity with the method used in the number field and function field sieves [1, 22]. More precisely, we no longer have a divisor D to decompose, but instead search for sums of factor base elements equal to 0:

$$\sum_{i=1}^m ((Q_i) - (\mathcal{O}_{\mathcal{H}})) \sim 0. \quad (2)$$

Heuristically, the expected number of relations of the form (2) involving m points of the factor base is approximately $\frac{p^{m-ng}}{m!}$. Since we need to collect at least about $p/2$ relations, we look for sums of $m = ng + 2$ points (assuming that $p \geq (ng+2)!/2$, which will always be the case in practice). As before, we work with the \mathbb{F}_q -vector space $\mathcal{L}(m(\mathcal{O}_{\mathcal{H}}))$, which is spanned by $1, x, \dots, x^{m_1}, y, xy, \dots, x^{m_2}y$, where $m_1 = \lfloor m/2 \rfloor$ and $m_2 = \lfloor (m+1)/2 \rfloor - g$. We consider the function $f = \lambda_0 + \lambda_1x + \dots + \lambda_{m_1}x^{m_1} + \mu_0y + \mu_1xy + \dots + \mu_{m_2}x^{m_2}y$: it vanishes in exactly m points if its top-degree coefficient is not zero, and the abscissae of its zeroes are the roots of

$$F(x) = f(x, y)f(x, -y) = (\lambda_0 + \lambda_1x + \dots + \lambda_{m_1}x^{m_1})^2 - h(x)(\mu_0 + \mu_1x + \dots + \mu_{m_2}x^{m_2})^2.$$

Again, we fix $\lambda_{m_1} = 1$ if m is even or $\mu_{m_2} = 1$ otherwise, so that F is monic. In order to obtain a relation of the form (2), we look for values of λ and μ for which F splits over \mathbb{F}_p . The first condition is that F belongs to $\mathbb{F}_p[x]$; after a scalar restriction on its coefficients, this translates as a quadratic polynomial system of $(n-1)m$ equations and $n(m-g)$ variables. With our choice of $m = ng + 2$, this corresponds to an underdetermined system of $n(n-1)g + 2n - 2$ equations in $n(n-1)g + 2n$ variables. When the parameters n and g are not too large, we remark that it is possible to compute once for all the corresponding lexicographic Gröbner basis. Each specialization of the last two variables should then provide an easy to solve system, namely triangular with low degrees. It remains to check whether the corresponding expression of F is indeed split and to deduce the corresponding relations between the points of \mathcal{F} .

Once enough relations of the form (2) have been collected, and possibly after a structured Gaussian elimination or a large prime variation, we can deduce with linear algebra the logarithms of all elements in \mathcal{F} (up to a multiplicative constant, since we have not specified the logarithm base). To compute the discrete logarithm of an arbitrary divisor D , we proceed to a descent phase: we need to decompose this arbitrary divisor as a sum of factor base elements. This decomposition search can be done using the first method described above. Note that, if D does not decompose as a sum, it suffices to try small multiples $2D$, $3D \dots$ until we find one correct decomposition. Thanks to this descent step, it is possible to compute many discrete logarithms in the same group for negligible additional cost.

When the cover of E is not hyperelliptic, one can still use the Riemann-Roch based approach. It is not difficult to compute a basis of the vector spaces $\mathcal{L}(ng(\mathcal{O}_{\mathcal{H}}) - D)$ or $\mathcal{L}(m(\mathcal{O}_{\mathcal{H}}))$ and to consider a function $f(x, y)$ (depending of parameters λ and μ) in these spaces. Getting rid of the y -variable can be done quite easily by computing the resultant in y of f and the equation of the curve (or multiresultant if the curve is not planar); however, the resulting polynomial $F(x)$ no longer depends quadratically of the parameters λ and μ . Consequently, the system obtained by scalar restriction still has the same number of equations and variables but its degree is greater than 2, so that the resolution is more complicated.

3.2 Sieving for quadratic extensions

This new decomposition technique is already faster than Nagao's when the lexicographic Gröbner basis of the system coming from (2) is efficiently computable, but it can still be further improved. Indeed, checking that F is split has a non-negligible cost, since we need to factor a polynomial of degree $ng + 2$ into linear terms. To avoid this, it is possible to modify the search for relations of the form (2) using a sieving technique when the extension degree $[\mathbb{F}_q : \mathbb{F}_p]$ is equal to 2 in the odd characteristic case. Let t be an element such that $\mathbb{F}_{p^2} = \mathbb{F}_p(t)$; we assume wlog that $t^2 = \omega \in \mathbb{F}_p$. In this case, $f = \lambda_0 + \dots + \lambda_g x^g + \mu y$ and the polynomial F is of the form

$$F(x) = (\lambda_0 x + \dots + \lambda_g x^g + x^{g+1})^2 - \mu^2 h(x).$$

In particular, when the parameter μ equals 0, f is independent of the y variable; the corresponding relation of type (2) is thus necessarily of the form $(P_1) + (\iota(P_1)) + \dots + (P_{g+1}) + (\iota(P_{g+1})) - (2g+2)\mathcal{O}_{\mathcal{H}} \sim 0$, where $\iota(P)$ is the image of P by the hyperelliptic involution. To avoid these trivial relations, we look only for solutions $(\lambda_{0,0}, \dots, \lambda_{g,0}, \lambda_{0,1}, \dots, \lambda_{g,1}, \mu_0, \mu_1) \in \mathbb{V}_{\mathbb{F}_p}(I : (\mu_0, \mu_1)^\infty)$, where I is the ideal corresponding to the $2(g+1)$ quadratic polynomials in $2(g+2)$ variables arising from the scalar restriction on $F \in \mathbb{F}_{p^2}[x]$, setting $\lambda_i = \lambda_{i,0} + t\lambda_{i,1}$ and $\mu^2 = \mu_0 + t\mu_1$. More precisely, with the type of extension considered here, the ideal I is given by the equations corresponding to the vanishing of the coefficients of the $\mathbb{F}_p[x]$ -polynomial

$$2(1 \cdot x^{g+1} + \lambda_{g,0} x^g + \dots + \lambda_{0,0})(\lambda_{g,1} x^g + \dots + \lambda_{0,1}) - \mu_0 h_1(x) - \mu_1 h_0(x),$$

where $h(x) = h_0(x) + th_1(x)$, $h_0, h_1 \in \mathbb{F}_p[x]$. This ideal I is not 2-dimensional, but its saturation is.

An easy but crucial remark is that the ideal is multi-homogeneous, generated by polynomials of bi-degree $(1, 1)$ in the variables $(1 : \lambda_{0,0} : \dots : \lambda_{g,0}), (\lambda_{0,1} : \dots : \lambda_{g,1} : \mu_0 : \mu_1)$. This additional structure has two major consequences. First, the lexicographic Gröbner basis computation is much faster than for a generic quadratic system of the same size. Second, if we denote by π_1 the projection on the first block of variables $(\lambda_{0,0}, \dots, \lambda_{g,0})$, then the image $\pi_1(\overline{\mathbb{V}_{\mathbb{F}_p}(I : (\mu_0, \mu_1)^\infty)}) = \pi_1(\overline{\mathbb{V}_{\mathbb{F}_p}(I) \setminus \mathbb{V}_{\mathbb{F}_p}(\mu_0, \mu_1)})$ is a dimension 1 variety (whose equations are easily deduced from the lex Gröbner basis of I), and each fiber is a 1-dimensional vector space.

From this, we can simplify the relation search. Rather than evaluating a first variable, we choose a point $(\lambda_{0,0}, \dots, \lambda_{g,0}) \in \pi_1(\overline{\mathbb{V}_{\mathbb{F}_p}(I : (\mu_0, \mu_1)^\infty)})$ and express the remaining variables linearly in terms of $\lambda_{0,1}$, so that now F belongs to $\mathbb{F}_p[x, \lambda_{0,1}]$ and has degree $2g + 2$ in x and 2 in $\lambda_{0,1}$. Instead of trying to factor F for many values of $\lambda_{0,1}$, the key idea is to compute for each $x \in \mathbb{F}_p$ the values of $\lambda_{0,1}$ such that $F(x, \lambda_{0,1}) = 0$. Since F has degree 2 in $\lambda_{0,1}$, this can be done very efficiently by computing the square root of the discriminant. In fact, we can speed the process even more by tabulating the square roots of \mathbb{F}_p . Our sieving process consists, for each root $\lambda_{0,1}$, to increment a counter corresponding to this value of $\lambda_{0,1}$; when one of these counters reaches $2g + 2$, then the polynomial F evaluated at the corresponding value of $\lambda_{0,1}$ splits into $2g + 2$ distinct linear terms, yielding a relation. This technique not only allows to skip the factorization of a degree $2g + 2$ polynomial, but is also well-suited to the double large prime variation, as explained in next section.

3.3 Complexity analysis

Constructing the cover $\mathcal{H}_{\mathbb{F}_q}$ of an elliptic curve $E_{\mathbb{F}_{q^d}}$ with the GHS method and transferring the DLP from $G \subset E(\mathbb{F}_{q^d})$ to $G' \subset \text{Jac}_{\mathcal{H}}(\mathbb{F}_q)$ has essentially a unit cost, which is negligible compared to the rest of the attack. The complexity of the decomposition phase is divided between the relation search and the linear algebra steps. In order to collect about $p/2$ relations using Nagao's decomposition method, we need to solve on average $(ng)! \cdot p/2$ quadratic polynomial systems. The resolution cost of this kind of system using e.g. Gröbner bases is hard to estimate precisely, but is at least polynomial in the degree $2^{(n-1)ng}$ of the corresponding zero-dimensional ideal. The linear algebra step then costs $O(ngp^2)$ operations modulo $\#G$, using sparse linear algebra techniques. With the second decomposition method, we need to compute first the lexicographic Gröbner basis of an ideal generated by $n(n-1)g + 2n - 2$ quadratic equations in $n(n-1)g + 2n$ variables. This cost is also at least exponential in n^2g , but the Gröbner basis computation has to be done only once. Afterwards, we have to solve on average $(ng + 2)! \cdot p/2$ "easy" systems. The complexity of the linear algebra step is the same (the cost of the descent is negligible compared to the sieving phase).

When p is large relatively to n and g , the linear algebra becomes the dominating phase. It is nevertheless possible to rebalance the cost of the two steps.

Indeed, collecting extra relations can speed up the logarithm computations; this is the idea behind structured Gaussian elimination [21] and double large prime variation. The analysis of [15] shows that with the latter, the asymptotic complexity of our cover and decomposition attack becomes either $\tilde{O}(p^{2-2/ng})$ or $\tilde{O}(p^{2-2/(ng+2)})$ with the decomposition variant, as p grows to infinity for fixed n and g . Although the complexity of the variant is asymptotically higher, the much smaller hidden constant means that it is actually faster for accessible values of p . Note that it is straightforward to parallelize the relation search phase; this is also possible, but much less efficiently, for the linear algebra step. In particular, the optimal choice of the balance depends not only of the implementation but also of the computing power available.

When $n = 2$, it is possible to improve the double large prime variation by sieving only among the values of x corresponding to the abscissae of points of the “small primes” factor base. As soon as $2g$ values of x are associated to one value of $\lambda_{0,1}$, we obtain a relation involving at most 2 large primes (if the remaining degree 2 factor is split, which occurs with probability close to 1/2). This speeds up the relation search and decreases the overall complexity from $\tilde{O}(p^{2-2/(2g+2)})$ to $\tilde{O}(p^{2-2/(2g+1)})$ as p grows to infinity, thus reducing the asymptotic gap between the two decomposition methods without degrading the practical performances.

Obviously, our approach outperforms generic algorithms only if the genus of the intermediate cover is not too large. Otherwise, it may be possible to transfer the DLP from E to a more vulnerable isogenous curve E' . There exist two “isogeny walk” strategies to find E' (if it exists) [17]: one can sample the isogeny class of E via low-degree isogenies until a weak curve is found, or one can try all the weak curves until a curve isogenous to E is found. The best strategy to use depends on the size of the isogeny class, on the number of weak curves and on the availability of an efficient algorithm for constructing these weak curves. For the cases we have considered, this isogeny walk can become the dominating part in the overall complexity (see below for details).

4 Application to elliptic curves defined over \mathbb{F}_{p^6}

For an elliptic curve E defined over an extension field \mathbb{F}_{p^6} , we can apply our cover and decomposition attack either with the tower $\mathbb{F}_{p^6} - \mathbb{F}_{p^2} - \mathbb{F}_p$ or with the tower $\mathbb{F}_{p^6} - \mathbb{F}_{p^3} - \mathbb{F}_p$. We have seen in Section 2.2 that in practice, we can compute decompositions only for a very limited number of values of (n, g) . In particular, our attack is feasible only if E admits a genus 3 (resp. 2) cover; we give examples of such curves below. Of course, this attack needs to be compared with the classic cover attacks or decomposition attacks using the base field $\mathbb{F}_{p^3}, \mathbb{F}_{p^2}$ or \mathbb{F}_p , as recalled in Section 2.

4.1 Using a genus 3 cover

In the present subsection, we apply our cover and decomposition attack using the first tower $\mathbb{F}_{p^6} - \mathbb{F}_{p^2} - \mathbb{F}_p$. Thanks to the results of [7, 25, 32], in odd char-

acteristic, we know that the only elliptic curves defined over \mathbb{F}_{q^3} (in our case, $q = p^2$) for which the GHS attack yields a cover by a hyperelliptic curve \mathcal{H} of genus 3 defined over \mathbb{F}_q , are of the form

$$y^2 = h(x)(x - \alpha)(x - \sigma(\alpha)) \quad (3)$$

where σ is the Frobenius automorphism of $\mathbb{F}_{q^3}/\mathbb{F}_q$, $\alpha \in \mathbb{F}_{q^3} \setminus \mathbb{F}_q$ and $h \in \mathbb{F}_q[x]$ is of degree 1 or 2. Similar results are also available in characteristic 2 (see [27]), thus our attack is also applicable in characteristic 2; we give details of the construction of the cover in both cases in Appendix A. The number of curves admitting an equation of the form (3) is $\Theta(q^2)$, thus only a small proportion of curves is directly vulnerable to the cover and decomposition attack using this extension tower. However, since this number of weak curves is much larger than the number of isogeny classes (which is about $q^{3/2}$), a rough reasoning would conclude that essentially all curves should be insecure using an isogeny walk strategy. Assuming that the probability for a curve to be weak is independent from its isogeny class, we obtain that the average number of steps before reaching a weak isogenous curve should be about $q = p^2$ steps. It is thus the dominating phase of the algorithm, but is still better than the $\tilde{O}(p^3)$ -generic attacks. Nevertheless, all the curves of the form (3) have a cardinality divisible by 4, so obviously not all curves are vulnerable to this isogeny walk (we recall that two curves are isogenous if and only if they have the same cardinality). Still, we conjecture that all curves with cardinality divisible by 4 are vulnerable to this cover and decomposition attack using an isogeny walk.

We can also consider non-hyperelliptic genus 3 covers. In this case, weak curves have equation

$$y^2 = c(x - \alpha)(x - \sigma(\alpha))(x - \beta)(x - \sigma(\beta)) \quad (4)$$

where $c \in \mathbb{F}_{q^3}$ and either $\alpha, \beta \in \mathbb{F}_{q^3} \setminus \mathbb{F}_q$ or $\alpha \in \mathbb{F}_{q^6} \setminus (\mathbb{F}_{q^2} \cup \mathbb{F}_{q^3})$ and $\beta = \sigma^3(\alpha)$. This targets much more curves: actually, about half of the curves having their full 2-torsion defined over \mathbb{F}_{q^3} admit an equation of this form [25].

For a genus 3 hyperelliptic cover over \mathbb{F}_{p^2} , the quadratic polynomial systems to solve over \mathbb{F}_p are composed of 6 variables and 6 equations, or 8 equations and 10 variables with our variant. Such systems can be solved very quickly by any computational algebra system. Unfortunately, with non-hyperelliptic covers, the systems of equations are much more complicated, and we have not been able to compute decompositions with available Gröbner basis implementations.

4.2 Using a genus 2 cover

We now consider the tower $\mathbb{F}_{p^6} - \mathbb{F}_{p^3} - \mathbb{F}_p$. The existence of genus 2 covers (which are necessarily hyperelliptic) defined over \mathbb{F}_q , where $q = p^3$, has been studied in [3, 29]. In odd characteristic, vulnerable curves admit an equation in so-called Scholten form

$$y^2 = ax^3 + bx^2 + \sigma(b)x + \sigma(a) \quad (5)$$

where $a, b \in \mathbb{F}_{q^2}$ and σ is the Frobenius automorphism of $\mathbb{F}_{q^2}/\mathbb{F}_q$. An elliptic curve E can be transformed into Scholten form as soon as its full 2-torsion is defined over \mathbb{F}_{q^2} [29] or its cardinality is odd and $j(E) \notin \mathbb{F}_q$ [3]. Consequently, a large proportion of curves are vulnerable to our cover and decomposition attack. Moreover, any curve without full 2-torsion but still with a cardinality divisible by 4, is 2-isogenous to a curve with full 2-torsion.

In this setting, the quadratic polynomial systems to solve over \mathbb{F}_p are composed of 12 variables and 12 equations, or 16 equations and 18 variables with the decomposition variant. Solving such systems is still feasible on current personal computers, but is much lower than in the case of hyperelliptic genus 3 cover defined over \mathbb{F}_{p^2} .

4.3 Complexity and comparison with other attacks

Apart from the square-root generic algorithms, the existing ECDLP attacks over sextic extensions are either Gaudry's decomposition method [13] or the GHS attack followed by Gaudry's or Diem's index calculus [15, 8], with base field \mathbb{F}_p or \mathbb{F}_{p^2} (using \mathbb{F}_{p^3} as base field does not provide any advantage in this context). When the base field is \mathbb{F}_{p^2} , the asymptotic complexity is in $\tilde{O}(p^{8/3})$ for both decomposition and GHS (assuming a genus 3 cover), or even in $\tilde{O}(p^2)$ with a degree 4 planar cover. But in all cases, the memory requirement is then very large, in $\tilde{O}(p^2)$. When the base field is \mathbb{F}_p , computing direct decompositions is completely out of reach, and the GHS attack very rarely provides low genus covers: the smallest possible genus is actually 9 (with corresponding degree 10 plane model), but this occurs for at most p^3 curves, see Appendix C. The resulting genus is much higher for most curves [5], implying that this attack is rarely practical.

We give in Table 1 a summary of the performances of the presented approaches. In order to obtain actual (and not just asymptotic) comparisons, we also consider the cryptographically significant example of a curve $E|_{\mathbb{F}_{p^6}}$ where p is a prime close to 2^{27} , whose cardinality is divisible by a 160-bit prime number. The values given are obviously just estimates relying on extrapolations of relation searches done on Magma V2-17-5 [6] with an Intel Core 2 Duo processor; in particular, the two last estimates are greater than what could be expected from the results obtained with optimized implementation presented in Section 6. Details of the computations are provided in Appendix D.

5 Application to elliptic curves defined over \mathbb{F}_{p^4}

If E is an elliptic curve defined over \mathbb{F}_{p^4} , we can use our attack with the tower $\mathbb{F}_{p^4} \text{---} \mathbb{F}_{p^2} \text{---} \mathbb{F}_p$. When p is odd, we have seen that E admits a genus 2 cover over \mathbb{F}_{p^2} as soon as the cardinality of E is either odd (and $j(E) \notin \mathbb{F}_{p^2}$) or divisible by 4. Thus, approximately 3/4 of the elliptic curves defined over \mathbb{F}_{p^4} are directly vulnerable. When the genus 2 cover \mathcal{H} of E has an imaginary model (see Appendix B for a necessary condition), the sieving technique described in Section

Attack	Asymptotic complexity	Memory complexity	Time estimates (years)
Pollard on $E(\mathbb{F}_{p^6})$ [28]	$\tilde{O}(p^3)$	$\tilde{O}(1)$	5×10^{13}
Ind. calc. on $\text{Jac}_{\mathcal{H}}(\mathbb{F}_{p^2})$, $g = 3$ [15]	$\tilde{O}(p^{8/3})$	$\tilde{O}(p^2)$	$6 \times 10^{10} \dagger$
Ind. calc. on $\text{Jac}_{\mathcal{C}}(\mathbb{F}_{p^2})$, $d = 4$ [8]	$\tilde{O}(p^2)$	$\tilde{O}(p^2)$	700 000
Decompositions on $E((\mathbb{F}_{p^2})^3)$ [13]	$\tilde{O}(p^{8/3})$	$\tilde{O}(p^2)$	10^{12}
Ind. calc. on $\text{Jac}_{\mathcal{C}}(\mathbb{F}_p)$, $d = 10$ [8]	$\tilde{O}(p^{7/4})$	$\tilde{O}(p)$	$1500^{(*)}$
Decomp. on $\text{Jac}_{\mathcal{H}}(\mathbb{F}_{p^3})$, $g = 2$ [this work]	$\tilde{O}(p^{5/3})$	$\tilde{O}(p)$	4×10^6
Decomp. on $\text{Jac}_{\mathcal{H}}(\mathbb{F}_{p^2})$, $g = 3$ [this work]	$\tilde{O}(p^{5/3})$	$\tilde{O}(p)$	$750 \dagger$
Sieving on $\text{Jac}_{\mathcal{H}}(\mathbb{F}_{p^2})$, $g = 3$ [this work]	$\tilde{O}(p^{12/7})$	$\tilde{O}(p)$	$300 \dagger$

\dagger : only for $\Theta(p^4)$ curves before isogeny walk $(*)$: only for $O(p^3)$ curves

Table 1. Comparison of the complexity of various attacks on $E(\mathbb{F}_{p^6})$, $\log_2 p \approx 27$.

3.2 can be applied. In particular, we can solve the DLP on E with a complexity of $\tilde{O}(p^{2-2/5})$, where the hidden constant is rather small. With a Nagao-style decomposition, the asymptotic complexity becomes $\tilde{O}(p^{2-2/4})$, but with a larger constant corresponding to the resolution cost of a quadratic system composed of 4 equations in 4 variables. If we directly apply the decomposition attack of Gaudry and Diem to E , the asymptotic complexity is still in $\tilde{O}(p^{2-2/4})$, but the constant is much larger: the systems are also composed of 4 equations in 4 variables, but with total degree 8. For cover attacks, the case of quartic extensions has been studied in [3]; the result is that most elliptic curves defined over \mathbb{F}_{p^4} admit a cover by a non-hyperelliptic genus 9 curve. Using this cover to solve the DLP with an index calculus method yields an asymptotic complexity in $\tilde{O}(p^{2-2/9})$, or potentially $\tilde{O}(p^{2-2/8})$ with the approach of [8]. All these attacks are asymptotically better than generic algorithms, but the improvement is smaller than for elliptic curves defined over \mathbb{F}_{p^6} . Nevertheless, a much larger proportion of curves is directly vulnerable to our attack, and does not necessitate a preliminary isogeny walk. We summarize in Table 2 the asymptotic complexities of the different attacks.

Attack	Asymptotic complexity	Ratio of vulnerable curves
Pollard [28]	$\tilde{O}(p^2)$	1
Ind. calc. on $\mathcal{H}_{ \mathbb{F}_p}$, $g(\mathcal{H}) = 9$ [3]	$\tilde{O}(p^{16/9})$	3/4
Decomp. on $E_{ \mathbb{F}_{p^4}}$ [13]	$\tilde{O}(p^{3/2})$	1
Decomp. on $\mathcal{H}_{ \mathbb{F}_{p^2}}$, $g(\mathcal{H}) = 2$ [this work]	$\tilde{O}(p^{3/2})$	3/4

Table 2. Comparison of the complexity of various attacks on $E(\mathbb{F}_{p^4})$

6 A 149-bit example

In this section, we give a practical example of the cover and decomposition attack for an elliptic curve defined over the Optimal Extension Field \mathbb{F}_{p^6} , where $p = 33\,554\,467 = 2^{25} + 35$ is the smallest 26-bit prime. We define \mathbb{F}_{p^2} as $\mathbb{F}_p[t]$ where $t^2 = 2$ and \mathbb{F}_{p^6} as $\mathbb{F}_{p^2}[\theta]$ where $\theta^3 = t$. The elliptic curve E is given by the following Weierstrass equation:

$$y^2 = x(x - \alpha)(x - \sigma(\alpha))$$

where $\sigma : x \mapsto x^{p^2}$ and $\alpha = 9\,819\,275 + 31\,072\,607\theta + 17\,686\,237\theta^2 + 31\,518\,659\theta^3 + 22\,546\,098\theta^4 + 17\,001\,125\theta^5$. It has a genus 3 cover by the hyperelliptic curve \mathcal{H} defined over \mathbb{F}_{p^2} by $y^2 = \left(x + \phi(x) + \phi^\sigma(x) + \phi^{\sigma^2}(x)\right) N(x)^2$, with $N(x)$ the minimal polynomial of α over \mathbb{F}_{p^2} and $\phi : x \mapsto \frac{(\alpha - \sigma^2(\alpha))(\sigma(\alpha) - \sigma^2(\alpha))}{x - \sigma^2(\alpha)} + \sigma^2(\alpha)$. The cover map π is given by:

$$\pi(x, y) = \left(\frac{x + \phi(x) + \phi^\sigma(x) + \phi^{\sigma^2}(x)}{4}, \frac{y(x - \phi^\sigma(x))(x - \phi^{\sigma^2}(x))}{8N(x)(x - \sigma^2(\alpha))} \right).$$

The common cardinality of E over \mathbb{F}_{p^6} and of the Jacobian of \mathcal{H} over \mathbb{F}_{p^2} is four times the 149-bit prime $\ell = 356814156285346166966901450449051336101786213$, and the number of elements in the factor base is 16 775 441.

For best performances, we use the sieving approach described in Section 3.2. As a first step, we compute a lexicographic Gröbner basis of the system composed of 10 quadratic equations in 8 variables in about 3 s on a 2.6 GHz Intel Core 2 Duo processor with Magma V2.16-12. Instead of the double large prime variation, we execute a structured Gaussian elimination. During the sieving phase, we used 1 024 cores of quadri-core Intel Xeon 5570 processors at 2.93 GHz⁶. After 62 h, we had collected about $1.4 \times 10^{10} \simeq p^2 / (2 \cdot 8!)$ relations, that is all the possible relations of the form (2). For comparison, we also tested Nagao-style decompositions on the same type of processors. We obtained that one decomposition test takes about 22 ms on a single core, thus showing that our decomposition variant is about 960 times faster.

Thanks to the large number of extra relations, structured Gaussian elimination performed quite well and, after 25.5 h on 32 cores, it reduces the number of unknowns to 3 092 914 (a fivefold reduction). For safety, the output system contains 15 000 more equations than this number of unknowns, and each equation involves between 8 and 182 basis elements. The total number of non-zero entries in all the equations is 191 098 665 and all these entries are equal to ± 1 . The most time-consuming step is the iterative linear algebra, which is done with a MPI implementation of the Lanczos algorithm. It took about 28.5 days on 64 cores of the same Intel Xeon processors. A large fraction of this time was taken by the MPI communications, since at each round 200 MB of data had to be broadcast between the 2 involved machines (32 cores/machine). This linear algebra

⁶ This work was granted access to the HPC resources of CCRT under the allocation 2010-t201006445 made by GENCI (Grand Equipement National de Calcul Intensif)

phase produced discrete logarithms for all the basis elements that remained after structured Gaussian elimination. Substituting these values back in the initial linear system, we recovered, in less than 12 h using 32 cores, the discrete logarithms modulo ℓ of all elements in the basis (given by their coordinates on \mathcal{H}):

$$\begin{aligned} \log(5,1\ 646\ 475+19\ 046\ 912\ t) &= 324090233616618447788899446283862317783046006 \\ \log(6,2\ 062\ 691+792\ 228\ t) &= 134424987842262695486611476989769052152832441 \\ \log(7,3\ 868\ 228+21\ 035\ 932\ t) &= 229073181595667785229146623058011111735286961 \\ &\vdots \\ \log(33\ 554\ 465,4\ 471\ 075+14\ 598\ 628\ t) &= 340713467460900419473167933722631654111145151 \end{aligned}$$

With the results of the above precomputation, computing logarithm of arbitrary points on the elliptic curve becomes easy. To demonstrate this, we constructed points on E with the following process and computed their logarithms. First, we let $X_0 = \sum_{j=0}^5 (\lfloor \pi \cdot p^{j+1} \rfloor \bmod p) \theta^j = 4\ 751\ 066 + 748\ 974\ \theta + 8\ 367\ 234\ \theta^2 + 24\ 696\ 290\ \theta^3 + 1\ 372\ 315\ \theta^4 + 7\ 397\ 713\ \theta^5$. We then constructed points on E with abscissa $X_0 + \delta$ for small offsets δ . Let P_1, P_2, P_3, P_4, P_5 and P_6 be the points corresponding to offsets 3, 4, 11, 14, 15 and 16. We lift each of these points to the Jacobian of \mathcal{H} using the conorm-norm map, which takes negligible time in Magma. After that, we apply the descent method of Section 3.1 to small multiples of the lifted element, until we find a multiple that decomposes as a sum of elements from the smoothness basis. Looking up the corresponding logarithms (and dividing back by the small multiples that have been included) yields the logarithm of each point. On average, we expect to try $6! = 720$ multiples before finding a decomposition. To actually decompose the six considered points, we needed 72.5 s. As a consequence, each individual logarithm on E can be performed in less than one minute. We give details in Table 3: the points involved in the decomposition are described by their abscissa together with a + or - sign that indicates whether the “real” part of the ordinate has a positive or negative representative in $(-p/2, p/2)$. Similarly, we indicate the choice of the points on E with a + or a - depending on the representative of the constant term in the ordinate⁷.

Points	Mult. Nagao	Points in decomposition					
$(X_0 + 3)^+$	97	2844007 ⁺	3819744 ⁻	5618276 ⁻	8396644 ⁻	11841629 ⁻	23771773 ⁻
$(X_0 + 4)^-$	36	4673075 ⁻	11272201 ⁺	12937918 ⁻	13869464 ⁻	14428213 ⁺	21399158 ⁻
$(X_0 + 11)^+$	742	4884810 ⁻	6230068 ⁻	8411592 ⁺	12188294 ⁺	20118618 ⁺	20945232 ⁻
$(X_0 + 14)^-$	956	3660673 ⁻	4314732 ⁻	20180301 ⁺	22563519 ⁺	26157093 ⁻	27107773 ⁻
$(X_0 + 15)^-$	682	780652 ⁺	8444164 ⁺	10116987 ⁺	11070139 ⁻	14566563 ⁻	32232816 ⁺
$(X_0 + 16)^-$	19	13089639 ⁻	19783194 ⁻	23921581 ⁻	28500971 ⁺	30393573 ⁺	30478839 ⁺

Table 3. Details of individual logarithm computations.

⁷ We did not really choose the points, but simply took the first point produced by Magma with the specified abscissa.

The group structure of E is $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/(2\ell)\mathbb{Z}$ and all the logarithms are computed mod ℓ . Thus, in order to obtain points of order ℓ , we multiply each of the points P_j by 2. To obtain the discrete logarithms in base P_1 , we simply divide the results by the logarithm of P_1 . Finally, we obtain:

$$\begin{aligned} 2 \cdot P_2 &= 44853429456306375520883685722551142750204929 \cdot 2 \cdot P_1 \\ 2 \cdot P_3 &= 245828744177202642167186866655188704860309093 \cdot 2 \cdot P_1 \\ 2 \cdot P_4 &= 241773698573992897197261454348760406499325884 \cdot 2 \cdot P_1 \\ 2 \cdot P_5 &= 47914434731086497860980273327037833732109767 \cdot 2 \cdot P_1 \\ 2 \cdot P_6 &= 164437442681856563836816034418873153945805017 \cdot 2 \cdot P_1 \end{aligned}$$

7 Conclusion and perspectives

In this paper, we have proposed a new index calculus algorithm to compute discrete logarithms on elliptic curves defined over extension fields of composite degree. In particular, sextic extensions are very well-suited to this method, as we have practically demonstrated on a 149-bit example.

This combination of cover and decomposition techniques raises many questions. For example, it would be interesting to know if elliptic curves of prime cardinality defined over a degree 6 extension field can be efficiently attacked. A related problem is how to target more curves easily: this requires either an improvement of the isogeny walk, or an efficient use of non-hyperelliptic covers. Finally, whether our method applies to different extension degrees is an important issue; clearly, degree 4 extensions are also susceptible, but the advantage over generic methods is then less significant.

Acknowledgements. We acknowledge that the results in this paper have been achieved using the PRACE Research Infrastructure resource Curie based in France at TGCC, Bruyères-le-Chatel.

References

1. L. M. Adleman. The function field sieve. In *Algorithmic number theory (Ithaca, NY, 1994)*, volume 877 of *Lecture Notes in Comput. Sci.*, pages 108–121. Springer, Berlin, 1994.
2. L. M. Adleman, J. DeMarrais, and M.-D. Huang. A subexponential algorithm for discrete logarithms over the rational subgroup of the Jacobians of large genus hyperelliptic curves over finite fields. In *Algorithmic number theory (Ithaca, NY, 1994)*, volume 877 of *Lecture Notes in Comput. Sci.*, pages 28–40. Springer, Berlin, 1994.
3. S. Arita, K. Matsuo, K.-I. Nagao, and M. Shimura. A Weil descent attack against elliptic curve cryptosystems over quartic extension fields. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, E89-A:1246–1254, May 2006.
4. D. V. Bailey and C. Paar. Efficient arithmetic in finite field extensions with application in elliptic curve cryptography. *J. Cryptology*, 14(3):153–176, 2001.

5. I. F. Blake, G. Seroussi, and N. P. Smart, editors. *Advances in elliptic curve cryptography*, volume 317 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2005.
6. W. Bosma, J. Cannon, and C. Playoust. The Magma algebra system. I. The user language. *J. Symbolic Comput.*, 24(3-4):235–265, 1997. Computational algebra and number theory (London, 1993).
7. C. Diem. The GHS attack in odd characteristic. *J. Ramanujan Math. Soc.*, 18(1):1–32, 2003.
8. C. Diem. An index calculus algorithm for plane curves of small degree. In *Algorithmic number theory*, volume 4076 of *Lecture Notes in Comput. Sci.*, pages 543–557. Springer, Berlin, 2006.
9. C. Diem. On the discrete logarithm problem in elliptic curves. *Compos. Math.*, 147(1):75–104, 2011.
10. G. Frey. How to disguise an elliptic curve (Weil descent). Talk at the 2nd Elliptic Curve Cryptography Workshop (ECC), 1998. Preprint, available at <http://www.cacr.math.uwaterloo.ca/conferences/1998/ecc98/frey.ps>.
11. S. D. Galbraith, F. Hess, and N. P. Smart. Extending the GHS Weil descent attack. In *Advances in cryptology—EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Comput. Sci.*, pages 29–44. Springer, Berlin, 2002.
12. P. Gaudry. An algorithm for solving the discrete log problem on hyperelliptic curves. In *Advances in cryptology—EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Comput. Sci.*, pages 19–34. Springer, Berlin, 2000.
13. P. Gaudry. Index calculus for abelian varieties of small dimension and the elliptic curve discrete logarithm problem. *J. Symbolic Comput.*, 44(12):1690–1702, 2008.
14. P. Gaudry, F. Hess, and N. P. Smart. Constructive and destructive facets of Weil descent on elliptic curves. *J. Cryptology*, 15(1):19–46, 2002.
15. P. Gaudry, E. Thomé, N. Thériault, and C. Diem. A double large prime variation for small genus hyperelliptic index calculus. *Math. Comp.*, 76:475–492, 2007.
16. F. Hess. Generalising the GHS attack on the elliptic curve discrete logarithm problem. *LMS J. Comput. Math.*, 7:167–192 (electronic), 2004.
17. F. Hess. Weil descent attacks. In *Advances in elliptic curve cryptography*, volume 317 of *London Math. Soc. Lecture Note Ser.*, pages 151–180. Cambridge Univ. Press, Cambridge, 2005.
18. A. Joux and V. Vitse. Elliptic curve discrete logarithm problem over small degree extension fields. *J. Cryptology*, pages 1–25, 2011. DOI: 10.1007/s00145-011-9116-z.
19. A. Joux and V. Vitse. A variant of the F4 algorithm. In A. Kiayias, editor, *Topics in cryptology—CT-RSA 2011*, volume 6558 of *Lecture Notes in Comput. Sci.*, pages 356–375, Berlin, 2011. Springer.
20. N. Koblitz. Elliptic curve cryptosystems. *Math. Comp.*, 48(177):203–209, 1987.
21. B. A. LaMacchia and A. M. Odlyzko. Computation of discrete logarithms in prime fields. *Des. Codes Cryptogr.*, 1(1):47–62, 1991.
22. A. K. Lenstra and H. W. Lenstra, Jr., editors. *The development of the number field sieve*, volume 1554 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1993.
23. A. Menezes, E. Teske, and A. Weng. Weak fields for ECC. In T. Okamoto, editor, *Topics in cryptology—CT-RSA 2004*, volume 2964 of *Lecture Notes in Comput. Sci.*, pages 366–386. Springer, Berlin, 2004.
24. V. S. Miller. Use of elliptic curves in cryptography. In *Advances in cryptology – CRYPTO 1985*, volume 218 of *Lecture Notes in Comput. Sci.*, pages 417–426. Springer, Berlin, 1986.

25. F. Momose and J. Chao. Scholten forms and elliptic/hyperelliptic curves with weak Weil restrictions. Cryptology ePrint Archive, Report 2005/277, 2005.
26. K. Nagao. Decomposition attack for the Jacobian of a hyperelliptic curve over an extension field. In *Algorithmic Number Theory – ANTS-IX*, volume 6197 of *Lecture Notes in Comput. Sci.*, pages 285–300. Springer, Berlin, 2010.
27. E. Nart and C. Ritzenthaler. Genus 3 curves with many involutions and application to maximal curves in characteristic 2, 2009. To appear in the proceedings of AGCT-12.
28. J. M. Pollard. Monte Carlo methods for index computation (mod p). *Math. Comp.*, 32(143):918–924, 1978.
29. J. Scholten. Weil restriction of an elliptic curve over a quadratic extension. Preprint, available at <http://homes.esat.kuleuven.be/~jscholte/weilres.pdf>, 2003.
30. I. A. Semaev. Summation polynomials and the discrete logarithm problem on elliptic curves. Cryptology ePrint Archive, Report 2004/031, 2004.
31. V. Shoup. Lower bounds for discrete logarithms and related problems. In *Advances in cryptology—EUROCRYPT 1997*, volume 1233 of *Lecture Notes in Comput. Sci.*, pages 256–266. Springer, Berlin, 1997.
32. N. Thériault. Weil descent attack for Kummer extensions. *J. Ramanujan Math. Soc.*, 18(3):281–312, 2003.

A Genus 3 cover

A.1 Odd characteristic

We consider elliptic curves defined over \mathbb{F}_{q^3} of the form

$$y^2 = h(x)(x - \alpha)(x - \sigma(\alpha)) \quad (6)$$

where σ is the Frobenius automorphism of $\mathbb{F}_{q^3}/\mathbb{F}_q$, $\alpha \in \mathbb{F}_{q^3} \setminus \mathbb{F}_q$ and $h \in \mathbb{F}_q[x]$ of degree 1 or 2. Such elliptic curves were studied by [7, 32]; they are the only elliptic curves for which the GHS attack yields a cover by a hyperelliptic curve \mathcal{H} of genus 3 defined over \mathbb{F}_q .

We give now an explicit description of the cover $\pi : \mathcal{H} \rightarrow E$; following [25], we express this cover as a quotient by a bi-elliptic involution, instead of using the GHS approach. For simplicity, we will assume that $h(x) = x$ (this can always be achieved by an appropriate change of coordinates if h has a root in \mathbb{F}_q).

Let $\phi : x \mapsto \frac{D}{x - \sigma^2(\alpha)} + \sigma^2(\alpha)$ be the unique involution of $\mathbb{P}^1(\overline{\mathbb{F}_q})$ sending $\sigma^2(\alpha)$ to ∞ and α to $\sigma(\alpha)$, so that $D = (\alpha - \sigma^2(\alpha))(\sigma(\alpha) - \sigma^2(\alpha))$. If ϕ lifts to an involution of a hyperelliptic curve $\mathcal{H}_{|\mathbb{F}_q}$, then necessarily ϕ^σ and ϕ^{σ^2} will be also involutions of \mathcal{H} . Observing that $\{Id, \phi, \phi^\sigma, \phi^{\sigma^2}\}$ forms a group, this leads us to consider the curve of equation $y^2 = x + \phi(x) + \phi^\sigma(x) + \phi^{\sigma^2}(x)$; a more usual form for this equation is

$$\mathcal{H} : y^2 = F(x)N(x) \quad (7)$$

where $N(x) = (x - \alpha)(x - \sigma(\alpha))(x - \sigma^2(\alpha))$ is the minimal polynomial of α over \mathbb{F}_q and

$F(x) = N(x) \left(x + \phi(x) + \phi^\sigma(x) + \phi^{\sigma^2}(x) \right) \in \mathbb{F}_q[x]$. It is clear that ϕ gives an involution of \mathcal{H} , still denoted by $\phi : (x, y) \mapsto \left(\frac{D}{x - \sigma^2(\alpha)} + \sigma^2(\alpha), y \frac{D^2}{(x - \sigma^2(\alpha))^4} \right)$.

The quotient of this genus 3 hyperelliptic curve \mathcal{H} by ϕ is the elliptic curve

$$E' : y^2 = (x - \alpha - \sigma(\alpha)) (x^2 - 4\alpha\sigma(\alpha))$$

and the quotient map $\pi' : \mathcal{H} \rightarrow E'$ satisfies $\pi'(x, y) = (x + \phi(x), y/(x - \sigma^2(\alpha))^2)$. The curve E' is 2-isogenous to the original curve $E : y^2 = x(x - \alpha)(x - \sigma(\alpha))$ via the map:

$$(x, y) \mapsto \left(\frac{x^2 - 4\alpha\sigma(\alpha)}{4(x - \alpha - \sigma(\alpha))}, y \frac{(x - 2\alpha)(x - 2\sigma(\alpha))}{8(x - \alpha - \sigma(\alpha))^2} \right).$$

Finally, the cover map $\pi : \mathcal{H} \rightarrow E$ has the expression

$$\pi(x, y) = \left(\frac{F(x)}{4N(x)}, \frac{y(x - \phi^\sigma(x))(x - \phi^{\sigma^2}(x))}{8N(x)(x - \sigma^2(\alpha))} \right). \quad (8)$$

In the general case, when E has equation (6), the cover (8) remains the same and the corresponding hyperelliptic curve \mathcal{H} of genus 3 defined over \mathbb{F}_q has the following equation:

$$\mathcal{H} : y^2 = 4N(x)^2 h \left(\frac{F(x)}{4N(x)} \right).$$

A.2 Characteristic 2

Let E be an ordinary curve defined over a binary field \mathbb{F}_{q^3} ; it admits an equation of the form

$$E : y^2 + xy = x^3 + ax^2 + b \quad (9)$$

where $b = 1/j(E)$. As already apparent in [14], the GHS attack produces a genus 3 hyperelliptic cover of E when $\text{Tr}_{\mathbb{F}_{q^3}/\mathbb{F}_q}(b) = 0$, so that $\Theta(q^2)$ curves are directly vulnerable. To describe this cover, we slightly adapt the description of [25, 27], already used in the previous subsection.

Let $\sigma : x \mapsto x^q$ be the Frobenius automorphism and let $v = \sqrt[4]{b}$; by assumption its trace over \mathbb{F}_q is zero. As in the case of odd characteristic, we consider the involution $\phi : x \mapsto \frac{\sigma(v)\sigma^2(v)}{x+v} + v$ of $\mathbb{P}^1(\overline{\mathbb{F}_q})$ sending v to infinity and $\sigma(v)$ to $\sigma^2(v)$. We denote by N the minimal polynomial of v over \mathbb{F}_q and by F the product $N(x) \left(x + \phi(x) + \phi^\sigma(x) + \phi^{\sigma^2}(x) \right) \in \mathbb{F}_q[x]$. Then, ϕ lifts to a bi-elliptic involution of the hyperelliptic curve $\mathcal{H}_{|\mathbb{F}_q}$ defined by

$$\mathcal{H} : y^2 + N(x)y = F(x)N(x) + aN(x)^2. \quad (10)$$

The curve E is up to a change of variable the quotient of \mathcal{H} by ϕ and the cover map from \mathcal{H} to E is given by:

$$\pi : (x, y) \mapsto \left(x + \phi(x) + v, \frac{y(x + \phi(x) + v)}{N(x)} + v^2 \right). \quad (11)$$

B Genus 2 cover

Let E be an elliptic curve defined over \mathbb{F}_{q^2} (where q is odd) in Scholten form:

$$y^2 = ax^3 + bx^2 + \sigma(b)x + \sigma(a)$$

One can observe that the map $x \mapsto 1/\sigma(x)$ permutes the roots of $f(x) = ax^3 + bx^2 + \sigma(b)x + \sigma(a)$. An easy consequence is that f is either irreducible or split in linear factors over \mathbb{F}_{q^2} . In particular, E has either no 2-torsion or its full 2-torsion defined over \mathbb{F}_{q^2} . Note that in the latter case, f admits at least one root with $\mathbb{F}_{q^2}/\mathbb{F}_q$ -norm equal to 1. The elliptic curve E admits a cover by the hyperelliptic genus 2 curve \mathcal{H} defined over \mathbb{F}_q , of equation

$$\mathcal{H} : y^2 = a(x-c)^6 + b(x-c)^4(x-\sigma(c))^2 + \sigma(b)(x-c)^2(x-\sigma(c))^4 + \sigma(a)(x-\sigma(c))^6$$

where $c \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$. The cover map is given by

$$\pi : (x, y) \mapsto \left(\left(\frac{x-c}{x-\sigma(c)} \right)^2, \frac{y}{(x-\sigma(c))^3} \right)$$

The curve \mathcal{H} admits an imaginary model if and only if the polynomial $h(x) = a(x-c)^6 + b(x-c)^4(x-\sigma(c))^2 + \sigma(b)(x-c)^2(x-\sigma(c))^4 + \sigma(a)(x-\sigma(c))^6$ has a \mathbb{F}_q -rational root α . This implies that $\beta = (\alpha-c)^2/(\alpha-\sigma(c))^2$ is a \mathbb{F}_{q^2} -rational root of f . Hence a necessary condition for \mathcal{H} to have an imaginary model is that E has full 2-torsion. Now, $\alpha \in \mathbb{F}_q$ if and only if the $\mathbb{F}_{q^2}/\mathbb{F}_q$ -norm of $(\alpha-c)/(\alpha-\sigma(c))$ is equal to 1. From this, we deduce easily that a sufficient condition for the existence of an imaginary model of \mathcal{H} is that f admits a root β such that $N_{\mathbb{F}_{q^2}/\mathbb{F}_q}(\beta) = 1$ (this implies that β is a square in \mathbb{F}_{q^2}) and $N_{\mathbb{F}_{q^2}/\mathbb{F}_q}(\sqrt{\beta}) = 1$.

C Genus 9 cover

Let E be an elliptic curve defined over a degree 6 extension field \mathbb{F}_{q^6} of odd characteristic; we assume that $j(E) \notin \mathbb{F}_{q^3} \cup \mathbb{F}_{q^2}$. Using techniques as in [7, 32], it is possible to show that the minimum genus of the \mathbb{F}_q -cover obtained by the GHS method is 9, which happens only in one of the following cases:

1. E admits an equation of the form

$$E_\alpha : y^2 = c(x-\alpha)(x-\sigma(\alpha))(x-\sigma^2(\alpha))(x-\sigma^3(\alpha)) \quad (12)$$

where σ is the Frobenius automorphism of $\mathbb{F}_{q^6}/\mathbb{F}_q$, $\alpha \in \mathbb{F}_{q^6} \setminus (\mathbb{F}_{q^2} \cup \mathbb{F}_{q^3})$, and $c \in \mathbb{F}_{q^6}$.

2. E admits an equation of the form

$$E_{\alpha,\beta} : y^2 = c(x-\alpha)(x-\sigma(\alpha))(x-\beta) \quad (13)$$

where σ is the Frobenius automorphism of $\mathbb{F}_{q^6}/\mathbb{F}_q$, $\alpha \in \mathbb{F}_{q^3} \setminus \mathbb{F}_q$, $\beta \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$, and $c \in \mathbb{F}_{q^6}$.

We can give an upper bound on the number of (isomorphism classes of) such curves. It is clear that two curves E_α and $E_{\alpha'}$ are isomorphic (or twists) if α and α' lie in the same $PGL_2(\mathbb{F}_q)$ -orbit; since the number of such orbits in $\mathbb{F}_{q^6} \setminus (\mathbb{F}_{q^2} \cup \mathbb{F}_{q^3})$ is $q^3 + q - 1$, there are at most $O(q^3)$ curves of type (12). For those of type (13), we can use the transitivity of the action of $PGL_2(\mathbb{F}_q)$ on $\mathbb{F}_{q^3} \setminus \mathbb{F}_q$ to fix the value of α ; the number of these curves is thus at most $O(q^2)$. This shows that the proportion of curves defined over \mathbb{F}_{q^6} for which the GHS method yields a genus 9 \mathbb{F}_q -cover is at most $1/q^3$.

D Complexity comparisons of different attacks on $E(\mathbb{F}_{p^6})$ with $\log_2 p \approx 27$

The basis of comparison for all attacks on the ECDLP comes from generic algorithms, i.e. algorithms that do not use any information about the actual group structure and only consider the group law, such as Pollard's Rho [28]. Using Floyd's cycle-finding algorithm, the expected number of iterations is approximately $0.94\sqrt{\ell} \approx 1.14 \times 10^{24}$ where ℓ is the 160-bit prime dividing the cardinality of $E(\mathbb{F}_{p^6})$. Using Magma V2-17-5 on Intel Core 2 Duo 2.6 GHz, it takes 13.91 s to compute 10 000 iterations, corresponding to 5×10^{13} years for the complete DLP resolution.

The main difficulty with the index calculus methods is the estimation of the linear algebra cost, which is needed to find the optimal balance in the large primes variation. We base our extrapolations on the experiment of Section 6, where the resolution of a sparse system of size 3×10^6 took about 44 000 h·CPU. Thus we assume that for a factor base of size n , the linear algebra costs $(n/3\,000\,000)^2 \cdot 44\,000 \cdot 160/148$ h, or $n^2 \cdot 2 \times 10^{-5}$ s. On the other hand, all the relation timings are obtained with Magma as we did not implement optimized versions of all the different attacks.

We first consider index calculus methods for which the size of the factor base is in $p^2/2$. The memory complexity associated to this size is clearly problematic for any real implementation, since the sole storage of the factor base elements requires about 2^{60} bits, which compares to the world's largest databases. In the case where E admits a hyperelliptic genus 3 cover $\mathcal{H}_{|\mathbb{F}_{p^2}}$, we can apply index calculus after transfer to its Jacobian. Our experiment takes 13.27 s to complete 10 000 tests, yielding 1 689 relations; the complete relation search thus requires 2×10^6 years. With our assumption, the linear algebra step (memory issues notwithstanding) takes 5×10^{19} years, a much more larger time. To rebalance the two phases using double large primes, we need to divide the size of factor base by about 40 000; the total computation time then becomes 6×10^{10} years. If E admits a non-hyperelliptic genus 3 cover $\mathcal{C}_{|\mathbb{F}_{p^2}}$, this cover admits a degree 4 plane model on which we can apply Diem's index calculus [8]. It then takes 11.74 s to complete 10 000 tests, yielding 4 972 relations. This means that 700 000 years are necessary to collect $p^2/2$ relations. With the adapted double large prime variation, the optimal small factor base contains about p elements, and the linear algebra cost becomes negligible compared to the relation search. We can finally

apply directly Gaudry’s attack [13] to E with base field \mathbb{F}_{p^2} . Our experiment needs 22.35 s for 100 tests, yielding 36 relations. In other words, finding such a decomposition is 80 times slower than with the genus 3 hyperelliptic cover, and the optimal balances are different. We find that the size of the factor base should be divided by 9 000, for an overall computation time of 10^{12} years.

Now, we consider the index calculus method for which the size of the factor base is in $p/2$. We recall that it is not possible to use Gaudry’s decomposition attack with base field \mathbb{F}_p . In the very rare case where E admits a non-hyperelliptic genus 9 cover (see Appendix C), it is possible to use the attack of [8] on a degree 10 plane model and obtain after 200 000 tests 5 relations in 123 s. An extrapolation gives a time of 50 years for the relation step. With our assumption, the linear algebra costs 3 000 years. With the adapted double large prime variation, the optimal size of the factor base corresponds to a twofold reduction, for an overall computation time of 1 500 years. If E admits a genus 3 hyperelliptic cover $\mathcal{H}_{|\mathbb{F}_{p^2}}$, we can apply the techniques presented in this article and search for decompositions in $\text{Jac}_{\mathcal{H}}(\mathbb{F}_{p^2})$ either with Nagao’s method or our sieving variant. In the first case, it takes 126 sec to run 5 000 tests yielding 9 relations. This means that the relation search would need 30 years, the linear algebra still lasting 3 000 years. The optimal balance corresponds to a reduction by a factor 2.7 of the size of the factor base, for a total computation time of 750 years. In the second case, using the sieving technique we obtained 3 300 relations in 1 800 s, which is 25 times faster than with Nagao’s technique (in practice, we have seen in Section 6 that with optimized implementation, the ratio is rather of the order of 900). With the adapted double large prime variation, the optimal size of the factor base corresponds to a factor 4.4 reduction, for an overall computation time of 300 years. Note that for this sieving method, we have more accurate experimental data obtained with an optimized implementation in C instead of Magma. We detail in Table 4 the timings obtained for curves defined over OEF of sizes 138, 144 and 150 bits; the sieving times are given for the collection of all $p^2/(2 \cdot 8!)$ relations, and the linear algebra is done after a structured Gaussian elimination. Based on these figures, we estimate more accurately that breaking the DLP over a 160-bit elliptic curve group would take about 200 years on a single core.

Size of p	Sieving (CPU.hours)	Sieving (real time)	Lanczos (CPU.hours)	Lanczos (real time)
$\log_2 p \approx 23$	3 600	3.5 hours	4 900	77 hours
$\log_2 p \approx 24$	15 400	15 hours	16 000	250 hours
$\log_2 p \approx 25$	63 500	62 hours	43 800	28.5 days

Table 4. Scaling data for our implementation

Eventually, it is possible to apply our cover and decomposition technique on a hyperelliptic genus 2 cover defined over \mathbb{F}_{p^3} , but without the sieving improve-

ment. On this curve, our experiment takes 3780s for a single decomposition test, which is 150000 times slower than with the same method on a genus 3 cover defined over \mathbb{F}_{p^2} . In particular, no rebalance is needed since the relation search dominates the computation time of about 4×10^6 years.