On the discrete logarithm problem in finite fields of fixed characteristic

Robert Granger, Thorsten Kleinjung and Jens Zumbrägel

Abstract

For q a prime power, the discrete logarithm problem (DLP) in $\mathbb{F}_q$ consists in finding, for any $g \in \mathbb{F}_q^\times$ and $h \in \langle g \rangle$, an integer x such that $g^x = h$. We present an algorithm for computing discrete logarithms with which we prove that for each prime p there exist infinitely many explicit extension fields $\mathbb{F}_{p^n}$ in which the DLP can be solved in expected quasi-polynomial time. Furthermore, subject to a conjecture on the existence of irreducible polynomials of a certain form, the algorithm solves the DLP in *all* extensions $\mathbb{F}_{p^n}$ in expected quasi-polynomial time.

1. Introduction

In this paper we prove the following result.

Theorem 1. *For every prime p there exist infinitely many explicit extension fields $\mathbb{F}_{p^n}$ in which the DLP can be solved in expected quasi-polynomial time*

$$\exp\big((1/\log 2 + o(1))(\log n)^2\big). \tag{1}$$

Theorem 1 is an easy corollary of the following much stronger result, which we prove by presenting a randomised algorithm for solving any such DLP.

Theorem 2. *Given a prime power $q > 61$ that is not a power of 4, an integer $k \geqslant 18$, coprime polynomials $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ of degree at most two and an irreducible degree l factor I of $h_1 X^q - h_0$, the DLP in $\mathbb{F}_{q^{kl}} \cong \mathbb{F}_{q^k}[X]/(I)$ can be solved in expected time*

$$q^{\log_2 l + O(k)}. \tag{2}$$

To deduce Theorem 1 from Theorem 2, note that thanks to Kummer theory, when $l = q - 1$ such h_0, h_1 are known to exist; indeed, for all k there exists an $a \in \mathbb{F}_{q^k}$ such that $I = X^{q-1} - a \in \mathbb{F}_{q^k}[X]$ is irreducible and therefore $I \mid X^q - aX$. By setting $q = p^i > 61$ for any $i \geqslant 1$ (odd for $p = 2$), $k \geqslant 18$ with $k = o(\log q)$, $l = q - 1 = p^i - 1$ and finally $n = ik(p^i - 1)$, applying (2) proves that the DLP in this representation of $\mathbb{F}_{p^n}$ can be solved in expected time (1). As one can compute an isomorphism between any two representations of $\mathbb{F}_{p^n}$ in polynomial time [Len91],

2010 Mathematics Subject Classification 11Y16, 11T71

Keywords: discrete logarithm problem, finite fields, quasi-polynomial time algorithm

The first author is supported by the Swiss National Science Foundation via grant number 200021-156420. This work was mostly done while the second author was with the Laboratory for Cryptologic Algorithms, EPFL, Switzerland, supported by the Swiss National Science Foundation via grant number 200020-132160, and while the third author was with the Institute of Algebra, TU Dresden, Germany, supported by the Irish Research Council via grant number ELEVATEPD/2013/82.

this completes the proof. Observe that one may replace the prime p in Theorem 1 by a (fixed) prime power p^r by stipulating in the argument above that k is a multiple of r.

In order to apply Theorem 2 to the DLP in $\mathbb{F}_{p^n}$ with p fixed and arbitrary n, one should first embed the DLP into one in an appropriately chosen $\mathbb{F}_{q^{kn}}$. By this we mean that $q = p^i$ should be at least $n-2$ (so that h_0, h_1 may exist) but not too large, and that $18 \leqslant k = o(\log q)$, so that the resulting complexity (2) is given by (1) as $n \to \infty$. Proving that appropriate $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ exist for such q and k would complete our approach and prove the far stronger result that the DLP in $\mathbb{F}_{p^n}$ with p fixed can be solved in expected time (1) for all n. However, this seems to be a very hard problem, even if heuristically it would appear to be almost certain. What is striking about Theorem 2 is that in contrast to all finite field DLP algorithms from the past thirty years, it is rigorous, and our algorithm is therefore guaranteed to work once an appropriate field representation is found.

Note that if one could prove the existence of an infinite sequence of primes p (or more generally prime powers) for which $p-1$ is quasi-polynomially smooth in $\log p$, then the Pohlig-Hellman algorithm [PH78] (discovered independently by Silver) would also give a rigorous – and deterministic – quasi-polynomial time algorithm for solving the DLP in such fields, akin to Theorem 1. However, such a sequence is not known to exist and even if it were, Theorem 1 is arguably more interesting since our algorithm exploits properties of the fields in question rather than just the factorisation of the order of their multiplicative groups. Furthermore, the fields to which our algorithm applies are explicit, whereas it may be very hard to find members of such a sequence of primes (or prime powers), should one exist.

Gauss was probably the first to define discrete logarithms – or indices, as he called them, with respect to a primitive root – noting their usefulness for computing n-th roots modulo primes [Gau65, art. 57–60]. Since he suggested the use of look-up tables for this purpose, the algorithm he used for computing logarithms in the tiny examples to which he applied the technique was almost certainly just tabulation via exponentiation. However, Gauss noted in art. 58 that the table need only consist of indices for primes, implicitly assuming that integers less than the modulus can be factorised efficiently. In the early 1920s Kraitchik developed this observation into what is now called the Index Calculus Method (ICM) [Kra22, Kra24]; evidently a very natural idea, it was also discovered independently by Cunningham at around the same time, see [WM68], and rediscovered by Adleman [Adl79], Merkle [Mer79] and Pollard [Pol78] in the late 1970s. In this context the ICM proceeds by first defining a *factor base* consisting of primes up to some *smoothness bound* B. One then searches for multiplicative relations between elements of the factor base; one can do this for instance by computing random powers of the primitive root g modulo p and storing those which are B-smooth. These relations between factor base elements (and g) each induce a linear equation between their logarithms with respect to g, and once there are sufficiently many relations the logarithms of the factor base elements can be computed via a linear algebra elimination. The second phase of the ICM consists of computing the logarithm of a target element h which is not B-smooth. In this setting one can multiply h by random powers of g until the product is B-smooth, at which point its logarithm is easily determined. Exploiting the distribution of $L_p(1/2)$-smooth integers amongst integers less than p [Dic30, DB51, DB66] gives a heuristic $L_p(1/2)$ algorithm for the DLP in $\mathbb{F}_p$ [Adl79]; here, as is usual for such algorithms, we use the following measure of subexponentiality:

$$L_p(\alpha, c) = \exp((c + o(1))(\log p)^{\alpha}(\log\log p)^{1-\alpha}),$$

where for simplicity we sometimes suppress the subscript, the constant c, or both. The algorithm

just described can be made rigorous for both prime fields and fixed characteristic extension fields [Pom87, EG02].

In 1984 Coppersmith proposed the first heuristic $L(1/3, c)$ algorithm which applies to fields of the form $\mathbb{F}_{2^n}$ [Cop84a, Cop84b], with c being a function of n satisfying $(32/9)^{1/3} \leqslant c \leqslant 4^{1/3}$. Coppersmith's algorithm exhibits similar behaviour for extensions of any fixed base field. In 1994 Adleman proposed the Function Field Sieve (FFS) [Adl94] – an analogue of the famous Number Field Sieve [LL93] – which can also be seen as a generalisation of Coppersmiths algorithm. This was refined by Adleman and Huang in 1999, achieving a heuristic complexity of $L(1/3, (32/9)^{1/3})$ for extension fields of any fixed characteristic [AH99].

For fixed characteristic extension fields, the main difference between the $L(1/2)$ and $L(1/3)$ algorithms is that during relation generation the former generates elements of degree $\approx n$ and searches for sufficiently many which are $\tilde{O}(n^{1/2})$-smooth (where the $\tilde{O}$ indicates suppressed log factors), whereas algorithms of the latter type generate elements of degree $\tilde{O}(n^{2/3})$ and search for sufficiently many which are $\tilde{O}(n^{1/3})$-smooth. In the former case the elements can be generated uniformly and so one can apply smoothness results to obtain a rigorous algorithm. Crucially, for the $L(1/3)$ algorithms the elements generated are not uniformly distributed amongst elements of that degree and hence the complexity analysis is only heuristic. A second difference is that during the individual logarithm phase of the $L(1/3)$ algorithms one needs to recursively express a target element as a product of irreducible elements of lower degrees – with one iteration of this process being known as an *elimination* of that element – which produces a tree with the target element at its root and the elements produced by this process at its nodes. After sufficiently many iterations the elements at the leaves of this tree will be contained entirely in the factor base and so the logarithm of the target element can easily be computed via backtracking. Since this process descends through elements of lower and lower degree, the individual logarithm phase is also known as the *descent*.

In order to obtain algorithms of better complexity – at least for the first phase of the ICM – there are two natural directions that one could explore: firstly, one could attempt to generate relations between elements of lower degree, which heuristically would have a higher probability of being smooth; or secondly, one could attempt to generate relations which have better than expected smoothness properties (or possibly a combination of both). The second idea is perhaps far less obvious and more nuanced than the first; indeed until recently it does not seem to have been appreciated that it was even a possibility, most likely because from an algorithm analysis perspective it is desirable that the expected smoothness properties hold. For nearly three decades there was no progress in either direction; the only development in fixed characteristic being a practical improvement [JL02], while for so-called medium characteristic fields – those for which the base field cardinality satisfies $q = L_{q^n}(1/3)$ – a slight reduction in the constant was achieved, to $c = 3^{1/3} \approx 1.44$ [JL06] and to $c = 2^{1/3} \approx 1.26$ [Jou13a], the latter using a clever method to amplify one relation into many others. Note that we mention the medium characteristic developments because they can be applied to fixed characteristic extensions for appropriate extension degrees. Given the immense importance of the DLP to public key cryptography ever since its inception in 1976 [DH06], this plateau in progress could have been taken as strong evidence of the problem's hardness. However, in 2013 a series of algorithmic breakthroughs occurred which demonstrated that for fixed characteristic fields the DLP is, at least heuristically, far easier than originally believed.

In particular, in February 2013, Göloğlu, Granger, McGuire and Zumbrägel showed that for binary (and more generally fixed characteristic) fields of a certain form, relation generation for

degree one elements runs in heuristic *polynomial time*, as does computing the logarithms of degree two elements using a technique which eliminates them on the fly, i.e., individually and quickly [GGMZ13a, GGMZ13b], which was previously the bottleneck in the descent when using the standard techniques. This was the first example of the second idea alluded to above as it demonstrated how to generate relations which are 1-smooth for arbitrarily large degree, completely contradicting the usual smoothness heuristics. However, the efficient elimination of higher degree elements remained an unresolved problem. For fields of essentially the same form Joux independently gave: a degree one relation generation method which is isomorphic to that of Göloğlu *et al.*; a very different degree two elimination method; and a new small degree element elimination method which resulted in an algorithm with heuristic complexity $L(1/4+o(1))$ [Jou13b, Jou14]. Combinations and variations of these techniques led to several large scale DLP computations and records [Jou13c, GGMZ13c, Jou13d, GGMZ13d, Jou13e, GKZ14c, GKZ14d, GGMZ14, GKZ14a], the largest of which at the time of writing was in the field $\mathbb{F}_{2^{9234}}$.

Then in June 2013, for fields of the same form and of bitlength λ, Barbulescu, Gaudry, Joux and Thomé announced a heuristic *quasi-polynomial time* algorithm (referred to hereafter as the original QPA) for solving the DLP [BGJT14], which has complexity

$$\lambda^{O(\log \lambda)}. \tag{3}$$

Since (3) is smaller than $L(\alpha)$ for any $\alpha > 0$, it is asymptotically the most efficient algorithm known for solving the DLP in finite fields of fixed characteristic. It also results in an immediate $L(\alpha + o(1))$ algorithm when $q = L_{q^n}(\alpha)$ for $0 \leqslant \alpha < 1/3$. The principal idea behind the elimination steps of the original QPA may be viewed as a generalisation of Joux's degree two elimination method [Jou14], which finds the logarithms of all translates of a degree two element simultaneously via the collection of suitable relations and a subsequent linear algebra elimination.

The principal idea[1] behind our new QPA may be viewed as a generalisation of the degree two elimination method of [GGMZ13b]. In particular, for an element of degree $2d$ that we wish to eliminate, observe that over a degree d extension of the base field it factors into a product of d irreducible quadratics. Applying the degree two elimination method of [GGMZ13b] to any one of these quadratics enables one to rewrite the quadratic as a product of linear elements over the degree d extension of the base field. To return to the original base field one simply applies the relevant norm, which takes the linear elements to powers of irreducible elements of degree dividing d and the quadratic element back to the original element which was to be eliminated, thus completing its elimination. If the target element has degree a power of two then this elimination can be applied recursively, halving the degree (or more) of the elements in the descent tree upon each iteration. Central to our proof of Theorem 2 is our demonstration that this recursive step can always be carried out successfully. For the purpose of building a full DLP algorithm which may be applied to any target element, one can use a Dirichlet-type theorem due to Wan [Wan97, Thm. 5.1] to ensure that any field element is equivalent to an irreducible of degree a power of two only slightly larger than the extension degree of the field in question.

A remarkable property of the above descent method is that it does not require any smoothness assumptions about non-uniformly distributed polynomials, in contrast to all previous index calculus algorithms, including the original QPA. So while the polynomial time relation generation techniques of [GGMZ13b, Jou14] in a sense *resisted* smoothness heuristics, our new descent method *completely eliminates* them. We emphasise that our new QPA is radically different from the original QPA of Barbulescu *et al.*, while it is its very algebraic nature that makes our rigorous

[1]This approach was first made public in a preliminary version [GKZ14b] of this article.

analysis possible. Given the essential use of smoothness heuristics in the original QPA, as well as one other heuristic, it seems unlikely that it can be made rigorous, even if the existence of appropriate field representations are assumed or proven. Furthermore, while not of central interest to the results of the present paper, we remark that our elimination steps are extremely practical, even for relatively small fields [PJ14, Kle14], whereas the bitlengths for which the original QPA becomes effective have yet to be determined.

Apart from the existence of suitable $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ in general (cf. Theorem 2 and the ensuing discussion), questions worthy of future consideration include whether or not there exists a polynomial time algorithm (either rigorous or heuristic) for the DLP in fixed characteristic fields, or even harder, what is the true complexity of the DLP in the fixed characteristic case? Note that a result of F. R. K. Chung implies that for fields of our form any element can be represented as a product of a polynomial number of linear elements [Chu89, Thm. 8]. Hence there is no representational barrier to obtaining a polynomial time algorithm, when the factor base consists of linear elements.

The sequel is organised as follows. In Section 2 we describe our algorithm and explain why the steps are sufficient for our purpose. We then give a brief review of the FFS in Section 3 and fix some notation. In Section 4 we provide details of the building block behind our new descent and explain why its successful application implies Theorem 2, and hence Theorem 1. Finally, in Section 5 we complete the proof of these theorems by demonstrating that the descent step is indeed always successful.

2. The algorithm

As per Theorem 2, let $q > 61$ be a prime power that is not a power of 4 and let $k \geqslant 18$ be an integer; the reasons for these bounds are explained in Sections 4 and 5. We also assume there exist $h_0, h_1, I \in \mathbb{F}_{q^k}[X]$ satisfying the conditions of Theorem 2. Finally, let $g \in \mathbb{F}_{q^{kl}}^{\times}$ and let $h \in \langle g \rangle$ be the target element for the DLP to base g.

We now present our algorithm, which differs slightly from the traditional ICM as described in Section 1 in that it does not first compute the logarithms of the factor base elements and then apply a descent strategy. Instead, one computes many descents for elements of the form $g^{\alpha}h^{\beta}$ (just one more than the number of factor base elements suffices) and then applies a linear algebra elimination. This approach and its analysis was first used by Enge and Gaudry [EG02], however the algorithm and argument we present follows very closely those used by Diem in the context of the elliptic curve DLP [Die11]. A small but important difference between our algorithm and Diem's is that we cannot assume that we know the factorisation of the order of the relevant group, since the fastest proven factorisation algorithms have complexity $L(1/2)$ [Pom87, Val91, LP92] and are therefore insufficient for our purpose.

Input: A prime power $q > 61$ that is not a power of 4; an integer $k \geqslant 18$; a positive integer l; polynomials $h_0, h_1, I \in \mathbb{F}_{q^k}[X]$ with h_0, h_1 being coprime, $\deg(h_0), \deg(h_1) \leqslant 2$ and I a degree l irreducible factor of $h_1 X^q - h_0$; $g \in \mathbb{F}_{q^{kl}}^{\times}$ and $h \in \langle g \rangle$.

Output: An integer x such that $g^x = h$.

1. Let $N = q^{kl} - 1$, let $\mathcal{F} = \{F \in \mathbb{F}_{q^k}[X] \mid \deg F \leqslant 1,\ F \neq 0\} \cup \{h_1\}$ and denote its elements by $F_1, \ldots, F_m$, where $m = |\mathcal{F}| = q^{2k}$ (or $q^{2k} - 1$ if $\deg h_1 \leqslant 1$).

2. Construct a matrix $R = (r_{i,j}) \in (\mathbb{Z}/N\mathbb{Z})^{(m+1)\times m}$ and column vectors $\alpha, \beta \in (\mathbb{Z}/N\mathbb{Z})^{m+1}$ as follows. For each i with $1 \leqslant i \leqslant m+1$ choose $\alpha_i, \beta_i \in \mathbb{Z}/N\mathbb{Z}$ uniformly and independently at random and apply the (randomised) descent algorithm of Section 4 to $g^{\alpha_i} h^{\beta_i}$ to express this as
$$g^{\alpha_i} h^{\beta_i} = \prod_{j=1}^{m} F_j^{r_{i,j}}.$$
3. Compute a lower row echelon form R' of R by using invertible row transformations; apply these row transformations also to α and β, and denote the results by α' and β'.
4. If $\gcd(\beta'_1, N) > 1$, go to Step 2.
5. Return an integer x such that $\alpha'_1 + x\beta'_1 \equiv 0 \pmod{N}$.

We now explain why the algorithm is correct and discuss the running time, treating the descent in Step 2 as a black box algorithm for now. Henceforth, we assume that any random choices used in the descent executions are independent from each other and of the randomness of α and β. For the correctness, note that $g^{\alpha'_1} h^{\beta'_1} = 1$ holds after Step 3, since the first row of R' vanishes. Thus for any integer x such that $\alpha'_1 + x\beta'_1 \equiv 0 \pmod{N}$ we have $g^x = h$, provided that β'_1 is invertible in $\mathbb{Z}/N\mathbb{Z}$.

Lemma 3. *After Step 3 of the algorithm the element $\beta'_1 \in \mathbb{Z}/N\mathbb{Z}$ is uniformly distributed. Therefore, the algorithm succeeds with probability $\varphi(N)/N$, where φ denotes Euler's phi function.*

Proof. We follow the argument from [EG02, Sec. 5] and [Die11, Sec. 2.3]. As $h \in \langle g \rangle$, for any fixed value $\beta_i = b \in \mathbb{Z}/N\mathbb{Z}$ the element $g^{\alpha_i} h^b$ is uniformly distributed over the group $\langle g \rangle$, therefore the element $g^{\alpha_i} h^{\beta_i}$ is independent of β_i. As the executions of the descent algorithm are assumed to be independent, we have that the row $(r_{i,1}, \ldots, r_{i,m})$ is also independent of β_i. It follows that the matrix R is independent of the vector β. Then the (invertible) transformation matrix $U \in (\mathbb{Z}/N\mathbb{Z})^{(m+1)\times(m+1)}$ is also independent of β, so that $\beta' = U\beta$ is uniformly distributed over $(\mathbb{Z}/N\mathbb{Z})^{m+1}$, since β is. From this the lemma follows. □

Regarding the running time, for Step 3 we note that a lower row echelon form of R can be obtained using invertible row transformations as for the Smith normal form, which along with the corresponding transformation matrices can be computed in polynomial time [KB79], so that Step 3 takes time polynomial in m and $\log N$. Furthermore, from [RS62] we obtain $N/\varphi(N) \in O(\log \log N)$. Altogether this implies that the DLP algorithm has quasi-polynomial expected running time (in $\log N$), provided the descent is quasi-polynomial. We defer a detailed complexity analysis of the descent to Section 4.

Observe that the algorithm does not require g to be a generator of $\mathbb{F}_{q^{kl}}^{\times}$, which is in practice hard to test without factorising N. In fact, the algorithm gives rise to a Monte Carlo method for deciding group membership $h \in \langle g \rangle$. Indeed, if a discrete logarithm $\log_g h$ has been computed, then obviously $h \in \langle g \rangle$; thus if $h \notin \langle g \rangle$, we always must have $\gcd(\beta'_1, N) > 1$ in Step 4.

Practitioners may have noticed inefficiencies in the algorithm. In particular, in the usual index calculus method one precomputes the logarithms of all factor base elements and then applies a single descent to the target element to obtain its logarithm. Moreover, one usually first computes the logarithm in $\mathbb{F}_{q^{kl}}^{\times}/\mathbb{F}_{q^k}^{\times}$, i.e., one ignores multiplicative constants and therefore includes only monic polynomials in the factor base, obtaining the remaining information by solving an additional DLP in $\mathbb{F}_{q^k}^{\times}$. However, the setup as presented simplifies and facilitates our rigorous analysis.

3. Overview of the Function Field Sieve

In this section we briefly review the classical FFS and describe some of the recent techniques. The knowledgeable reader may omit this section, having familiarised themself with the notation via a brief look at Fig. 1.

Given the embedding of $\mathbb{F}_{p^n}$ into $\mathbb{F}_{q^{kl}}$ as described in the introduction, we focus purely on the latter. A relation in $\mathbb{F}_{q^{kl}}$ is an equality of products of elements in $\mathbb{F}_{q^{kl}}^\times$, or, equivalently, a linear combination of logarithms of elements in $\mathbb{F}_{q^{kl}}^\times$ whose sum is zero. All variants of the FFS rely on the following basic method for obtaining relations. Let $R = \mathbb{F}_{q^k}[X, Y]$ and let $f_1, f_2 \in R$ be two irreducible polynomials such that $R_{12} = R/(f_1, f_2)$ is a finite ring surjecting onto the target field $\mathbb{F}_{q^{kl}}$. Furthermore, for $i = 1, 2$, let $R_i = \mathbb{F}_{q^k}[X, Y]/(f_i)$ and $Z_i \in R$ such that the quotient field $\mathrm{Quot}(R_i)$ is a finite extension of the rational function field $\mathrm{Quot}(Q_i)$ where $Q_i = \mathbb{F}_{q^k}[Z_i]$. This is summarised in Fig. 1.

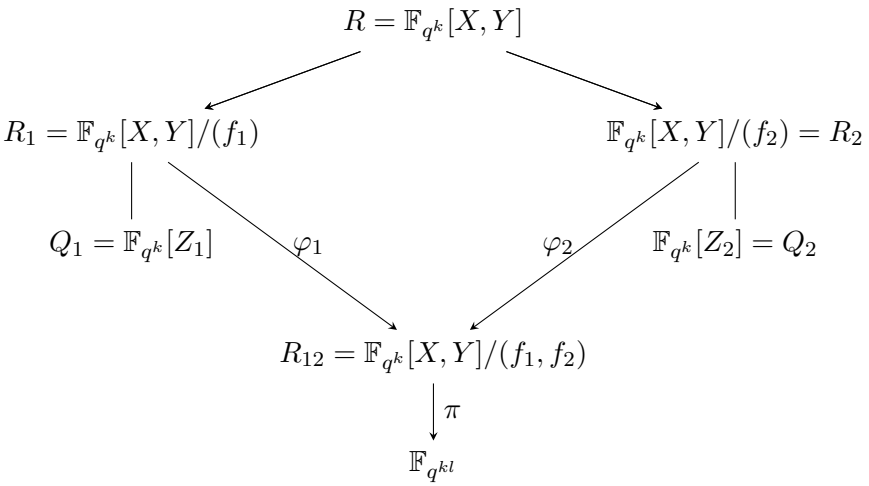

Figure 1. Setup for the FFS

Via the maps π, φ_1 and φ_2, logarithms in $\mathbb{F}_{q^{kl}}^\times$ can be extended to a notion of logarithms in $R_i \setminus (\pi \circ \varphi_i)^{-1}(0)$, $i = 1, 2$. Therefore, relations can also be viewed as linear combinations of logarithms of elements in R_1 and in R_2 whose sum is zero. It is always implicitly assumed that all logarithms are defined, i.e., that the sets $(\pi \circ \varphi_i)^{-1}(0)$, $i = 1, 2$, are avoided.

A polynomial $P \in R$ gives rise to a relation by decomposing $P \bmod f_i$ in R_i for $i = 1, 2$ (and mapping down to R_{12} or $\mathbb{F}_{q^{kl}}$ if desired). Sufficiently many non-trivial relations amongst elements of a set of bounded size allow one to compute logarithms in this set. If the multiplicative closure of such a set is $\mathbb{F}_{q^{kl}}^\times$, arbitrary logarithms can be computed by expressing an element as a product of elements of this set. As was described in Section 1, this is done by following a descent strategy in which elements, also called special-Q, are recursively rewritten as 'easier' elements using relations as above.

In the classical FFS the polynomials f_1, f_2 are chosen such that their degrees are as low as possible, typically of the form $f_1 = Y - a(X)$, $f_2 = \sum_{j=0}^{d} b_j(X)Y^j$ with $\deg_X(a) = e$, $\deg_X(b_j) < e$ and $de > l$, and $Z_1 = Z_2 = X$ so that the extensions $\mathrm{Quot}(R_i)/\mathrm{Quot}(Q_i)$, $i = 1, 2$, are of degree 1 and degree d, respectively. By choosing P as a low-degree polynomial, the degrees

of the norms $N_{\mathrm{Quot}(R_i)/\mathrm{Quot}(Q_i)}(P \bmod f_i)$, $i = 1, 2$, are not too big and therefore the chance of both norms splitting into low-degree polynomials is sufficiently high. With judiciously selected parameters this gives a heuristic running time of $L(1/3)$.

The main difference between the classical FFS and the recent variations [GGMZ13b, Jou14, BGJT14] is where the relation generation begins. In the recent variations a product of low-degree polynomials $\tilde{P} = \prod \tilde{P}_j$ in R_1 is constructed in such a way that it can be lifted to a low-degree polynomial $P \in R$ and such that its reduction $P \bmod f_2$ is of sufficiently low degree, where by low degree we mean that the norm has low degree. This can be achieved by choosing q to be of the order of l, $f_1 = Y - X^q$ and f_2 of low degree.[2] Then $R_1 = \mathbb{F}_{q^k}[X]$ and low-degree polynomials $F, G \in R_1$ give rise to relations via

$$\tilde{P} = F^q G - F G^q = G \prod_{\alpha \in \mathbb{F}_q} (F - \alpha G) = \prod \tilde{P}_j, \tag{4}$$

since F^q (resp. G^q) can be expressed as a degree $\deg F$ (resp. $\deg G$) polynomial in Y, and thus $\tilde{P}$ can be lifted to a low-degree polynomial P. This yields a heuristic polynomial time algorithm for finding relations between elements of $\mathbb{F}_{q^{kl}}$ that are, via π, φ_1 and φ_2, images of polynomials of bounded degree.

In the descent phase it is advantageous to choose f_2 such that its degree in X or in Y is one (cf. [GKZ14a] and [Jou14] respectively), which implies that $\mathrm{Quot}(R_2) = \mathrm{Quot}(Q_2)$ with $Z_2 = Y$ or $Z_2 = X$, respectively. More precisely, writing $f_2 = h_1 X - h_0$ or $f_2 = h_1 Y - h_0$ respectively, with $h_i \in Q_2$, $i = 0, 1$, implies $R_2 = \mathbb{F}_{q^k}[Z_2][\frac{1}{h_1}]$. Up to the logarithm of h_1, the logarithm of a polynomial of R_1 can be related to the logarithm of a corresponding polynomial in R_2 (the same polynomial for $Z_2 = X$ and a Frobenius twist for $Z_2 = Y$) which allows one to view a special-Q (the element to be eliminated) as coming from R_1 or from R_2. In the latter case, the condition that a polynomial $Q \in R_2$, a lift of the special-Q element, divides $P \bmod f_2$ for a P arising via (4), can be expressed as a bilinear quadratic system which gives, for appropriate parameter choices, an algorithm with heuristic running time $L(1/4 + o(1))$.

In the other case, namely the special-Q element being lifted to $Q \in R_1$, a certain set of polynomials in R_1 containing Q is chosen in such a way that pairs F, G from this set generate via (4) sufficiently many relations with $P \bmod f_2$ splitting into polynomials of sufficiently low degree. Solving a linear system of equations then expresses the logarithm of the special-Q element as a linear combination of logarithms of polynomials in R_2 of sufficiently low degree (and h_1), resulting in the original QPA.

Actually, the relations in the original QPA (and in [Jou14]) are generated in a slightly different manner by applying linear fractional transformations to the polynomial $A = X^q - X = \prod_{\alpha \in \mathbb{F}_q}(X - \alpha)$. The subgroup $\mathrm{PGL}_2(\mathbb{F}_q) \subset \mathrm{PGL}_2(\mathbb{F}_{q^k})$ is the largest subgroup fixing this polynomial, so that the action of $\mathrm{PGL}_2(\mathbb{F}_{q^k})/\mathrm{PGL}_2(\mathbb{F}_q)$ on A produces $\frac{q^{3k}-q^k}{q^3-q}$ polynomials, each splitting into linear polynomials and whose only non-zero terms correspond to the monomials X^{q+1}, X^q, X and 1.

[2] An interesting historical aside is that this specialisation was also proposed by Shinohara *et al* in January 2012 in order to half the size of the factor base when q is a power of the characteristic [SSHT12, Sec. 4.1], but its impact on relation generation was not considered. Furthermore, in December 2012 Joux used $f_1 = Y - X^d$ for medium characteristic fields with prime base fields [Jou13a], which does not help in finding a relation, but does allow one to generate many relations once one has been found, via transformations of the roots. Viewed in this context the selection of $f_1 = Y - X^q$ in [GGMZ13a] and [Jou13b] is a very natural (and indeed fertile) one, even if the ensuing approaches diverge in terms of field representation, relation generation and small degree elements elimination.

4. The descent

In this section we detail the building block behind our new descent and explain why its successful application implies Theorem 2. In the terminology of the previous section, the setup for $\mathbb{F}_{q^{kl}}$ has irreducible polynomials $f_1 = Y - X^q$ and $f_2 = h_1 Y - h_0$ with $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ coprime of degree at most two and $h_1 X^q - h_0$ having an irreducible factor I of degree l, i.e., $R_{12} = \mathbb{F}_{q^k}[X,Y]/(f_1, f_2)$ surjects onto $\mathbb{F}_{q^{kl}}$.[3] This implies $R_1 = \mathbb{F}_{q^k}[X]$ and $R_2 = \mathbb{F}_{q^k}[X][\frac{1}{h_1}]$. By the phrase "rewriting a polynomial Q (in R_1 or R_2) in terms of polynomials P_i (in R_1 or R_2)" we henceforth mean that in the target field the image of Q equals a product of powers of images of P_i. Since h_1 appears in almost every relation, we adjoin it to the factor basis $\mathcal{F}$, and for the sake of simplicity it is suppressed in the following description.

4.1 On-the-fly degree two elimination

In this subsection we review the on-the-fly degree two elimination method from [GGMZ13b], adjusted for the present framework. In [Blu04] the affine portion of the set of polynomials obtained as linear fractional transformations of $X^q - X$ is parameterised as follows. Let $\mathcal{B}$ be the set of $B \in \mathbb{F}_{q^k}$ such that the polynomial $X^{q+1} - BX + B$ splits completely over $\mathbb{F}_{q^k}$, the cardinality of which is approximately q^{k-3} [Blu04, Lemma 4.4]. Scaling and translating these polynomials means that all the polynomials $X^{q+1} + aX^q + bX + c$ with $c \neq ab$, $b \neq a^q$ and $B = \frac{(b-a^q)^{q+1}}{(c-ab)^q}$ split completely over $\mathbb{F}_{q^k}$ whenever $B \in \mathcal{B}$.

Let Q (viewed as a polynomial in R_2) be an irreducible quadratic polynomial to be eliminated. We let $L_Q \subset \mathbb{F}_{q^k}[X]^2$ be the lattice defined by

$$L_Q = \{(w_0, w_1) \in \mathbb{F}_{q^k}[X]^2 \mid w_0 h_0 + w_1 h_1 \equiv 0 \pmod{Q}\}. \tag{5}$$

In the case that Q divides $w_0 h_0 + w_1 h_1 \neq 0$ for some $w_0, w_1 \in \mathbb{F}_{q^k}$, then $Q = w(w_0 h_0 + w_1 h_1)$ for some $w \in \mathbb{F}_{q^k}^\times$, since the degree on the right hand side is at most two. Therefore, the relation generated from $P = w_0 Y + w_1 \in R$ relates Q with $w_0 X^q + w_1 = (w_0^{1/q} X + w_1^{1/q})^q \in R_1$ (and h_1). We will say in this case that the lattice is degenerate.

In the other (non-degenerate) case, L_Q has a basis of the form $(1, u_0 X + u_1), (X, v_0 X + v_1)$ with $u_i, v_i \in \mathbb{F}_{q^k}$. Since the polynomial $P = XY + aY + bX + c$ maps to $\frac{1}{h_1}((X+a)h_0 + (bX+c)h_1)$ in R_2, Q divides $P \bmod f_2$ if and only if $(X + a, bX + c) \in L_Q$. Note that the numerator of $P \bmod f_2$ is of degree at most three, thus it can at worst contain a linear factor besides Q. If the triple (a, b, c) also satisfies $c \neq ab$, $b \neq a^q$ and $\frac{(b-a^q)^{q+1}}{(c-ab)^q} \in \mathcal{B}$, then $P \bmod f_1$ splits into linear factors and thus Q has been rewritten in terms of linear polynomials.

Algorithmically, a triple (a, b, c) satisfying all conditions can be found in several ways. Choosing a $B \in \mathcal{B}$, considering $(X + a, bX + c) = a(1, u_0 X + u_1) + (X, v_0 X + v_1)$ and rewriting $b = u_0 a + v_0$ and $c = u_1 a + v_1$ gives the condition

$$B = \frac{(-a^q + u_0 a + v_0)^{q+1}}{(-u_0 a^2 + (-v_0 + u_1)a + v_1)^q}. \tag{6}$$

By expressing a in an $\mathbb{F}_{q^k}/\mathbb{F}_q$ basis, (6) results in a quadratic system in k variables [GGMZ14]. Using a Gröbner basis algorithm the running time is exponential in k. Alternatively, and this is one of the key observations for the present work, equation (6) can be considered as a polynomial of degree $q^2 + q$ in a whose roots can be found in polynomial time in q and in k by taking a GCD

[3] One can equally well work with $f_2 = h_1 X - h_0$ with $h_i \in \mathbb{F}_{q^k}[Y]$ of degree at most two, where $h_1(X^q)X - h_0(X^q)$ has a degree l irreducible factor, as proposed in [GKZ14a], with all subsequent arguments holding *mutatis mutandis*.

with $a^{q^k} - a$ in $\mathbb{F}_{q^k}[a]$ [GGMZ13b]. One can also check for random (a, b, c) such that the lattice condition holds, whether $X^{q+1} + aX^q + bX + c$ splits into linear polynomials, which happens with probability q^{-3}. Each such instance is also polynomial time in q and in k.

These degree 2 elimination methods will fail when Q divides $h_1X^q - h_0$, because this would imply that the polynomial $P \bmod f_1 = X^{q+1} + aX^q + bX + c$ is divisible by Q whenever $P \bmod f_2$ is, a problem first discussed in [CWZ14]. Such polynomials Q or their roots will be called traps of level 0. Similarly, these degree 2 elimination methods might also fail when Q divides $h_1X^{q^{k+1}} - h_0$, in which case such polynomials Q or their roots will be called traps of level k.

Note that for Kummer extensions, i.e., when $h_1 = 1$ and $h_0 = aX$ for some $a \in \mathbb{F}_{q^k}$, there are no traps and hence much of the following treatment is not required for proving only Theorem 1. However, it is essential to consider traps for proving the far more general Theorem 2.

4.2 Elimination requirements

As briefly explained in the introduction, the on-the-fly degree two elimination method can be transformed into an elimination method for irreducible even degree polynomials. We now present a theorem which states that under some assumptions this degree two elimination is guaranteed to succeed, and subsequently demonstrate that it implies Theorem 2.

An element $\tau \in \overline{\mathbb{F}}_{q^k}$ for which $[\mathbb{F}_{q^k}(\tau) : \mathbb{F}_{q^k}] = 2d$ is even and $h_1(\tau) \neq 0$, is called a *trap root* if it is a root of $h_1X^q - h_0$ or $h_1X^{q^{kd+1}} - h_0$, or if $\frac{h_0}{h_1}(\tau) \in \mathbb{F}_{q^{kd}}$. Note that the sets of trap roots is invariant under the absolute Galois group of $\mathbb{F}_{q^k}$. A polynomial in R_1 or R_2 is said to be *good* if it has no trap roots; the same definitions are used when the base field of R_1 and R_2 is extended. This definition encompasses traps of level 0, of level kd, and the case where for $Q \neq h_1$ the lattice L_Q is degenerate.

Theorem 4. *Let $q > 61$ be a prime power that is not a power of 4, let $k \geqslant 18$ be an integer and let $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ be coprime polynomials of degree at most two with $h_1X^q - h_0$ having an irreducible degree l factor. Moreover, let $d \geqslant 1$ be an integer, let $Q \in \mathbb{F}_{q^{kd}}[X]$, $Q \neq h_1$ be an irreducible quadratic good polynomial, and let $(1, u_0X + u_1), (X, v_0X + v_1)$ be a basis of the lattice L_Q in (5). Then the number of solutions $(a, B) \in \mathbb{F}_{q^{kd}} \times \mathcal{B}$ of (6) resulting in good descendents is at least q^{kd-5}.*

This theorem is of central importance for our rigorous analysis and is proven in Section 5.

4.3 Degree $2d$ elimination and descent complexity

Now we demonstrate how the on-the-fly degree two elimination gives rise to a method for eliminating irreducible even degree polynomials, which is the crucial building block for our descent algorithm. As per Theorem 4, let $q > 61$ be a prime power that is not a power of 4, let $k \geqslant 18$, and let h_0, h_1, I as before.

Propsition 5. *Let $Q \in R_2$, $Q \neq h_1$, be an irreducible good polynomial of degree $2d$. Then Q can be expressed in terms of at most $q + 2$ irreducible good polynomials of degrees dividing d, in an expected running time polynomial in q and in d.*

Proof. Over the extension $\mathbb{F}_{q^{kd}}$ the polynomial Q splits into d irreducible good quadratic polynomials; let Q' be one of them. Since $Q' \neq h_1$ is good it does not divide $w_0h_0 + w_1h_1 \neq 0$ for some $w_0, w_1 \in \mathbb{F}_{q^{kd}}$. By Theorem 4, with an expected polynomial number of trials, the on-the-fly degree two elimination method for $Q' \in \mathbb{F}_{q^{kd}}[X]$ produces a polynomial $P' \in \mathbb{F}_{q^{kd}}[X, Y]$ such

that P' mod f_1 splits into a product of at most $q+1$ good polynomials of degree one over $\mathbb{F}_{q^{kd}}$ and such that $(P' \bmod f_2)h_1$ is a product of Q' and a good polynomial of degree at most one. Let P be the product of all conjugates of P' under $\mathrm{Gal}(\mathbb{F}_{q^{kd}}/\mathbb{F}_{q^k})$. Since the product of all conjugates of a linear polynomial under $\mathrm{Gal}(\mathbb{F}_{q^{kd}}/\mathbb{F}_{q^k})$ is the d_1-th power of an irreducible degree d_2 polynomial for d_1 and d_2 satisfying $d_1 d_2 = d$, the rewriting assertion of the proposition follows.

The three steps of this method – computing Q', the on-the-fly degree two elimination (when the second or third approach listed above for solving (6) is used), and the computation of the polynomial norms – all have running time polynomial in q and in d, which proves the running time assertion. □

By recursively applying Proposition 5 we can express a good irreducible polynomial of degree 2^e, $e \geqslant 1$, in terms of at most $(q+2)^e$ linear polynomials. The final step of this recursion, namely eliminating up to $(q+2)^{e-1}$ quadratic polynomials, dominates the running time, which is thus upper bounded by $(q+2)^e$ times a polynomial in q.

Lemma 6. *Any nonzero element in $\mathbb{F}_{q^{kl}}$ can be lifted to an irreducible good polynomial of degree 2^e, provided that $2^e > 4l$.*

Proof. By the effective Dirichlet-type theorem on irreducibles in arithmetic progressions [Wan97, Thm. 5.1], for $2^e > 4l$ the probability of irreducibility for a random lift is lower bounded by 2^{-e-1}. One may actually find an irreducible polynomial of degree 2^e which is good, since the number of possible trap roots ($< q^{k2^{e-1}+2}$) is much smaller than the number ($> q^{k(2^e-l)}2^{-e-1}$) of irreducibles produced by this Dirichlet-type theorem. □

Putting everything together, this proves the quasi-polynomial expected running time of the descent and therefore the running time of our algorithm in Section 2, establishing Theorem 2.

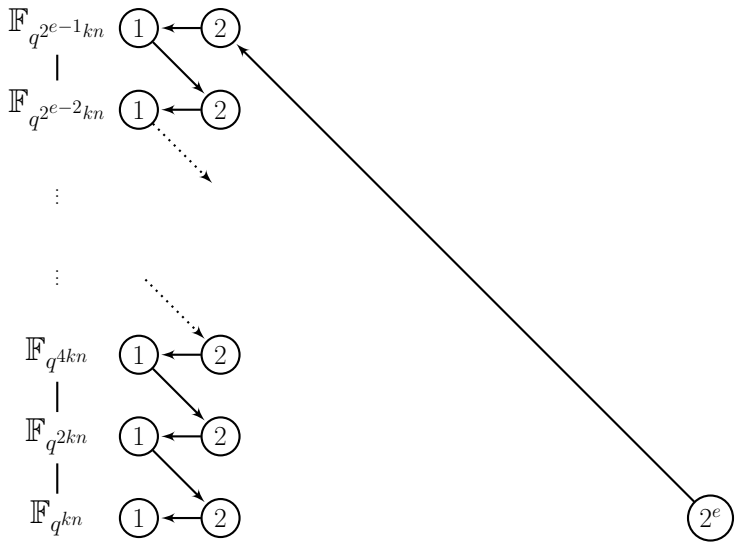

Figure 2. Elimination of irreducible polynomials of degree a power of 2 when considered as elements of $\mathbb{F}_{q^k}[X]$. The arrow directions $\nwarrow$, $\leftarrow$ and $\searrow$ indicate factorisation, degree 2 elimination and taking a norm with respect to the indicated subfield, respectively. (We have suppressed the rare cases, where linear polynomials are already in a subfield of index 2.)

Note that for $q = L_{q^{kl}}(\alpha)$, just as in [BGJT14], the complexity stated in Theorem 2 is $L(\alpha + o(1))$, which is therefore better than the classical FFS for $\alpha < \frac{1}{3}$.

Finally note that during an elimination step, one need not use the basic building block as stated, which takes the norms of the linear polynomials produced back down to $\mathbb{F}_{q^k}$. Instead, one need only take their norms to a subfield of index 2, thus becoming quadratic polynomials, and then recurse, as depicted in Fig. 2.

5. Proof of Theorem 4

In this section we prove Theorem 4, which by the arguments of the previous section demonstrates the correctness of our algorithm and our main theorems.

5.1 Notation and statement of supporting results

Let $K = \mathbb{F}_{q^{kd}}$ with $kd \geqslant 18$, let $L = \mathbb{F}_{q^{2kd}}$ be its quadratic extension, and let Q be an irreducible quadratic polynomial in $K[X]$ such that $(1, u_0X + u_1), (X, v_0X + v_1)$ is a basis of its associated lattice L_Q in (5). Then Q is a scalar multiple of $-u_0X^2 + (-u_1 + v_0)X + v_1$.

Let $\mathcal{B}$ be the set of $B \in K$ such that the polynomial $X^{q+1} - BX + B$ splits completely over K. Using an elementary extension of [HK10, Theorem 5] the set $\mathcal{B}$ can be characterised as the image of $K \setminus \mathbb{F}_{q^2}$ under the map

$$u \mapsto \frac{(u - u^{q^2})^{q+1}}{(u - u^q)^{q^2+1}}. \tag{7}$$

By this and (6), in order to eliminate Q we need to find $(a, u) \in K \times (K \setminus \mathbb{F}_{q^2})$ satisfying

$$(u - u^{q^2})^{q+1}(-u_0a^2 + (-v_0 + u_1)a + v_1)^q - (u - u^q)^{q^2+1}(-a^q + u_0a + v_0)^{q+1} = 0.$$

The two terms have a common factor $(u - u^q)^{q+1}$ which motivates the following definitions. Let $\alpha = -u_0$, $\beta = u_1 - v_0$, $\gamma = v_1$ and $\delta = -v_0$ with $\alpha, \beta, \gamma, \delta \in K$, as well as

$$D = \frac{U^{q^2} - U}{U^q - U} = \prod_{\epsilon \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q} (U - \epsilon),$$

$$E = U^q - U = \prod_{\epsilon \in \mathbb{F}_q} (U - \epsilon),$$

$$F = \alpha A^2 + \beta A + \gamma = \alpha(A - \rho_1)(A - \rho_2) \quad \text{with} \quad \rho_1, \rho_2 \in L,$$

$$G = A^q + \alpha A + \delta \quad \text{and}$$

$$P = D^{q+1}F^q - E^{q^2-q}G^{q+1} \in K[A, U].$$

Note that F equals $Q(-A)$ (up to a scalar), so that $\deg(F) = 2$, F is irreducible and $\rho_1, \rho_2 \notin K$. We consider the curve C defined by $P = 0$ and are interested in the number of (affine) points $(a, u) \in C(K)$ with $u \notin \mathbb{F}_{q^2}$. More precisely, we want to prove the following.

Theorem 7. *Let $q > 61$ be a prime power that is not a power of 4. If the conditions*

$$(*) \qquad \rho_1^q + \alpha\rho_2 + \delta \neq 0$$

$$(**) \qquad \rho_1^q + \alpha\rho_1 + \delta \neq 0$$

hold then there are at least q^{kd-1} pairs $(a, u) \in K \times (K \setminus \mathbb{F}_{q^2})$ satisfying $P(a, u) = 0$.

The relation of the two conditions to the quadratic polynomial Q as well as properties of traps are described in the following propositions.

Propsition 8. *If condition* $(*)$ *is not satisfied, then* Q *divides* $h_1X^q - h_0$*, i.e.,* Q *is a trap of level* 0*. If condition* $(**)$ *is not satisfied, then* Q *divides* $h_1X^{q^{kd+1}} - h_0$*, i.e.,* Q *is a trap of level* kd*. In particular, if* Q *is a good polynomial then conditions* $(*)$ *and* $(**)$ *are satisfied.*

Propsition 9. *Let* $(a, u), (a', u') \in K \times (K \setminus \mathbb{F}_{q^2})$ *be two solutions of* $P = 0$ *with* $a \neq a'$*, corresponding to the polynomials* $\mathcal{P}_a = XY + aY + bX + c$ *and* $\mathcal{P}_{a'} = XY + a'Y + b'X + c'$*, respectively. Then* $\mathcal{P}_a \bmod f_1$ *and* $\mathcal{P}_{a'} \bmod f_1$ *have no common roots. Furthermore, the common roots of* $\mathcal{P}_a \bmod f_2$ *and* $\mathcal{P}_{a'} \bmod f_2$ *are precisely the roots of* Q*.*

Now we explain how (for $q > 61$ not a power of 4) Theorem 4 follows from the above theorem and the propositions. Since the irreducible quadratic polynomial Q is good, the lattice L_Q is non-degenerate so that a basis as above exists, and by Proposition 8 the two conditions of Theorem 7 are satisfied. The map (7) is $q^3 - q : 1$ on $K \setminus \mathbb{F}_{q^2}$, hence there are at least q^{kd-4} solutions $(a, B) \in K \times \mathcal{B}$ of (6), which contain at least q^{kd-4} different values $a \in K$. Observe that a trap root τ that may occur in this situation is a root of $h_1X^q - h_0$, or of $h_1X^{q^{kd'}+1} - h_0$ for $d' \mid \frac{d}{2}$, or it satisfies $\frac{h_0}{h_1}(\tau) \in \mathbb{F}_{q^{kd/2}}$. The cardinalities of these trap roots is at most $q^{\frac{kd}{2}+3}$. By Proposition 9 a trap root can appear in $\mathcal{P}_a \bmod f_j$ for at most two values a, at most once for $j = 1$ and at most once for $j = 2$. Hence there are at most $q^{\frac{kd}{2}+4} < q^{kd-5}$ values a for which a trap root appears in $\mathcal{P}_a \bmod f_j$, $j = 1, 2$. Thus there are at least q^{kd-5} different values a for which a solution (a, B) leads to an elimination into good polynomials. This finishes the proof of Theorem 4, hence we focus on proving the theorem and the two propositions above.

5.2 Outline of the proof method

The main step of the proof of the theorem consists in showing that, subject to conditions $(*)$ and $(**)$, there exists an absolutely irreducible factor P_1 of P that lies already in $K[A, U]$. Since the (total) degree of P_1 is at most $q^3 + q$, restricting to the component of the curve defined by P_1 and using the Weil bound for possibly singular plane curves gives a lower bound on the cardinality of $C(K)$ which is large enough to prove the theorem after accounting for projective points and points with second coordinate in $\mathbb{F}_{q^2}$. This argument is given in the next subsection before dealing with the more involved main step.

For proving the main step the action of $\mathrm{PGL}_2(\mathbb{F}_q)$ on the variable U is considered. An absolutely irreducible factor P_1 of P is stabilised by a subgroup $S_1 \subset \mathrm{PGL}_2(\mathbb{F}_q)$ satisfying some conditions. The first step is to show that, after possibly switching to another absolutely irreducible factor, there are only a few cases for the subgroup. Then for each case it is shown that the factor is defined over $K[A, U]$ or that one of the conditions on the parameters is not satisfied.

The propositions are proven in the final subsection.

5.3 Weil bound

Let C_1 be the absolutely irreducible plane curve defined by P_1 of degree $d_1 \leqslant q^3+q^2$. Corollary 2.5 of [AP96] shows that

$$|\#C_1(K) - q^{kd} - 1| \leqslant (d_1 - 1)(d_1 - 2)q^{\frac{kd}{2}}.$$

Since $\deg_A(P_1) \leqslant q^2 + q$ there are at most $q^4 + q^3$ affine points with $u \in \mathbb{F}_{q^2}$. The number of points at infinity is at most $d_1 \leqslant q^3 + q^2 < q^4$. Denoting by $C_1(K)^{\sim}$ the set of affine points in

$C_1(K)$ with second coordinate $u \notin \mathbb{F}_{q^2}$ one obtains

$$|\#C_1(K)^\frown| > q^{kd} - (q^4 + q^3) - d_1 - (d_1 - 1)(d_1 - 2)q^{\frac{kd}{2}} > q^{kd} - q^{\frac{kd}{2}+8} \geqslant q^{kd-1},$$

since $kd \geqslant 18$, thus proving the theorem if there exists an absolutely irreducible factor P_1 defined over $K[A, U]$.

5.4 PGL_2 action

Here the following convention for the action of $\mathrm{PGL}_2(\mathbb{F}_q)$ on $\mathbb{P}^1$ and on polynomials is used. A matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{PGL}_2(\mathbb{F}_q)$ acts on $\mathbb{P}^1(M)$, where M is an arbitrary field containing $\mathbb{F}_q$, by $(x_0 : x_1) \mapsto \begin{pmatrix} a & b \\ c & d \end{pmatrix}(x_0 : x_1) = (ax_0 + bx_1 : cx_0 + dx_1)$ or, via $\mathbb{P}^1(M) = M \cup \{\infty\}$, by $x \mapsto \frac{ax+b}{cx+d}$. This is an action on the left, i.e., for $\sigma, \tau \in \mathrm{PGL}_2(\mathbb{F}_q)$ and $x \in \mathbb{P}^1(M)$ the following holds: $\sigma(\tau(x)) = (\sigma\tau)(x)$. On a homogeneous polynomial H in the variables $(X_0 : X_1)$ the action of $\sigma = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is given by $H^\sigma(X_0 : X_1) = H(aX_0 + bX_1 : cX_0 + dX_1)$. This is an action on the right, satisfying $H^{(\sigma\tau)} = (H^\sigma)^\tau$. In the following we will usually use this action on the dehomogenised polynomials given by $H^\sigma(X) = H(\frac{aX+b}{cX+d})$, clearing denominators in the appropriate way.

The polynomial $P \in (K[A])[U]$ is invariant under $\mathrm{PGL}_2(\mathbb{F}_q)$ acting on the variable U; this can be checked by considering the actions of $\begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix}$, $\begin{pmatrix} a & 0 \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, and noticing that $\mathrm{PGL}_2(\mathbb{F}_q)$ is generated by these matrices. Let

$$P = s \prod_{i=1}^{g} P_i, \qquad P_i \in (\overline{K}[A])[U],\ s \in \overline{K}[A],$$

be the decomposition of P in $(\overline{K}[A])[U]$ into irreducible factors P_i and possibly reducible s. Notice that s must divide F^q and G^{q+1}, hence it divides a power of $\gcd(F, G)$. As F is irreducible, $\gcd(F, G)$ is either constant or of degree two. In the latter case ρ_1 is a root of G contradicting condition $(**)$. Therefore one can assume that $s \in \overline{K}$ is a constant.

Let

$$P = F^q \prod_{i=1}^{q^3-q} (U - r_i), \qquad r_i \in \overline{K(A)},$$

be the decomposition of P in $\overline{K(A)}[U]$. Then $\mathrm{PGL}_2(\mathbb{F}_q)$ permutes the set $\{r_i\}$ and, since fixed points of $\mathrm{PGL}_2(\mathbb{F}_q)$ lie in $\mathbb{F}_{q^2}$ but $r_i \notin \mathbb{F}_{q^2}$, the action is free. Since $\#\,\mathrm{PGL}_2(\mathbb{F}_q) = q^3 - q$ the action is transitive.

Therefore the action on the decomposition over $\overline{K}[A, U]$ is also transitive (adjusting the P_i by scalars in $\overline{K}[A]$ if necessary). Denoting by $S_i \subset \mathrm{PGL}_2(\mathbb{F}_q)$ the stabiliser of P_i it follows that all S_i are conjugates of each other, thus they have the same cardinality and hence $q^3 - q = g \cdot \#S_i$. Moreover the degree of P_i in U is constant, namely $\#S_i$, and also the degree of P_i in A is constant, thus $g \mid q^2 + q = \deg_A(P)$. In particular, $q - 1 \mid \#S_i$.

5.5 Subgroups of PGL_2

The classification of subgroups of $\mathrm{PSL}_2(\mathbb{F}_q)$ is well known [Dic01] and allows to determine all subgroups of $\mathrm{PGL}_2(\mathbb{F}_q)$ [COTR06]. Since $\#S_i$ is divisible by $q - 1$ (in particular $\#S_i > 60$), only

the following subgroups are of interest (per conjugation class only one subgroup is listed):

1. the cyclic group $\begin{pmatrix} * & 0 \\ 0 & 1 \end{pmatrix}$ of order $q-1$,
2. the dihedral group $\begin{pmatrix} * & 0 \\ 0 & 1 \end{pmatrix} \cup \begin{pmatrix} 0 & 1 \\ * & 0 \end{pmatrix}$ of order $2(q-1)$ and, if q is odd, its two dihedral subgroups
$$\left\{\begin{pmatrix} a & 0 \\ 0 & 1 \end{pmatrix} \mid a \neq 0 \text{ a square}\right\} \cup \left\{\begin{pmatrix} 0 & 1 \\ c & 0 \end{pmatrix} \mid c \neq 0 \text{ a square}\right\} \quad \text{and}$$
$$\left\{\begin{pmatrix} a & 0 \\ 0 & 1 \end{pmatrix} \mid a \neq 0 \text{ a square}\right\} \cup \left\{\begin{pmatrix} 0 & 1 \\ c & 0 \end{pmatrix} \mid c \text{ not a square}\right\},$$
both of order $q-1$,
3. the Borel subgroup $\begin{pmatrix} * & * \\ 0 & 1 \end{pmatrix}$ of order q^2-q,
4. if q is odd, $\mathrm{PSL}_2(\mathbb{F}_q)$ of index 2,
5. if $q = q'^2$ is a square, $\mathrm{PGL}_2(\mathbb{F}_{q'})$ of order $q'^3 - q' = q'(q-1)$, and
6. $\mathrm{PGL}_2(\mathbb{F}_q)$.

In the last case P is absolutely irreducible, thus it remains to investigate the first five cases which are treated in the next subsection.

Remark: The condition $q > 61$ rules out some small subgroups as A_4, S_4, and A_5. In many of the finitely many cases $q \leqslant 61$ the proof of the theorem also works (e.g., q not a square and $q - 1 \nmid 120$). The condition of q not being a power of even exponent of 2 eliminates the fifth case in characteristic 2; removing this condition would be of some interest.

5.6 The individual cases

Since the stabilisers S_i are conjugates of each other, one can assume without loss of generality that S_1 is one of the explicit subgroups given in the previous subsection. Then the polynomial P_1 is invariant under certain transformations of U, so that P_1 and P can be rewritten in terms of another variable as stated in the following.

If a polynomial (in the variable U) is invariant under $U \mapsto aU$, $a \in \mathbb{F}_q^\times$, it can be considered as a polynomial in the variable $V = U^{q-1}$. For the polynomials D and E^{q-1} one obtains
$$D = \frac{V^{q+1} - 1}{V - 1} \qquad \text{and} \qquad E^{q-1} = V(V-1)^{q-1}.$$

Similarly, in the case of odd q, if a polynomial is invariant under $U \mapsto aU$ for all squares $a \in \mathbb{F}_q^\times$, it can be rewritten in the variable $V' = U^{\frac{q-1}{2}}$. For D and E^{q-1} this gives
$$D = \frac{V'^{2q+2} - 1}{V'^2 - 1} \qquad \text{and} \qquad E^{q-1} = V'^2(V'^2 - 1)^{q-1}.$$

If a polynomial is invariant under $U \mapsto U + b$, $b \in \mathbb{F}_q$, it can be considered as a polynomial in $\tilde{V} = U^q - U$ which gives
$$D = \tilde{V}^{q-1} + 1 \qquad \text{and} \qquad E^{q-1} = \tilde{V}^{q-1}.$$

Combining the above yields that a polynomial which is invariant under both $U \mapsto aU$, $a \in \mathbb{F}_q^\times$, and $U \mapsto U + b$, $b \in \mathbb{F}_q$, can be considered as a polynomial in $W = \tilde{V}^{q-1} = (U^q - U)^{q-1}$. For D

and E^{q-1} one obtains

$$D = W + 1 \qquad \text{and} \qquad E^{q-1} = W.$$

This is now applied to the various cases for S_1.

5.6.1 *The cyclic case* Rewriting P and P_1 in terms of $V = U^{q-1}$ one obtains

$$P = \Big(\frac{V^{q+1}-1}{V-1}\Big)^{q+1} F^q - V^q(V-1)^{q^2-q} G^{q+1}$$

and $\deg_V(P_1) = 1$, i.e., $P_1 = p_1 V - p_0$ with $p_i \in \overline{K}[A]$, $\gcd(p_0, p_1) = 1$, $\max(\deg(p_0), \deg(p_1)) = 1$ and it can be assumed that p_0 is monic.

The divisibility $P_1 \mid P$ transforms into the following polynomial identity in $\overline{K}[A]$:

$$\Big(\frac{p_0^{q+1} - p_1^{q+1}}{p_0 - p_1}\Big)^{q+1} F^q = p_1^q p_0^q (p_0 - p_1)^{q^2-q} G^{q+1}.$$

The degree of the first factor on the left hand side is either $q^2 + q$ or $q^2 - 1$ (if $p_0 - \zeta p_1$ is constant for some $\zeta \in \mu_{q+1}(\mathbb{F}_{q^2}) \setminus \{1\}$). Since the degrees of the other factors are all divisible by q, the latter case is impossible. Since $\deg(F) = 2$ one gets $\deg(F^q) = 2q$. Furthermore, $\deg((p_0 p_1)^q) \in \{q, 2q\}$, $\deg((p_0 - p_1)^{q^2-q}) \in \{0, q^2 - q\}$ and $\deg(G^{q+1}) = q^2 + q$ which implies $\deg(p_0 - p_1) = 0$, $\deg(p_0) = \deg(p_1) = 1$ since $q > 2$.

Let $p_0 - p_1 = c_1 \in \overline{K}$; in the following c_i will be some constants in $\overline{K}$. Since the first factor on the left hand side is coprime to $p_0 p_1$, it follows

$$\frac{p_0^{q+1} - p_1^{q+1}}{p_0 - p_1} = c_2 G, \quad F = c_3 p_0 p_1 \quad \text{and} \quad c_2^{q+1} c_3^q = c_1^{q^2-q}.$$

Exchanging ρ_1 and ρ_2, if needed, one obtains

$$p_0 = A - \rho_1, \quad p_1 = A - \rho_2, \quad c_3 = \alpha \quad \text{and} \quad c_1 = \rho_2 - \rho_1.$$

Considering the coefficient of A^q in the equation for G gives $c_2 = 1$ and evaluating this equation at $A = \rho_2$ gives

$$\rho_1^q + \alpha\rho_2 + \delta = 0.$$

This means that condition $(*)$ does not hold.

5.6.2 *The dihedral cases* The case of the dihedral group of order $2(q-1)$ is considered first. Then, as above, P and P_1 can be expressed in terms of V, and, since P and P_1 are also invariant under $V \mapsto \frac{1}{V}$, they can be expressed in terms of $W_+ = V + \frac{1}{V}$. This gives $\deg_{W_+}(P_1) = 1$ and with $\mathcal{Z} = \mu_{q+1}(\mathbb{F}_{q^2}) \setminus \{1\}$

$$D^{q+1} V^{-\frac{q^2+q}{2}} = \prod_{\zeta \in \mathcal{Z}} (W_+ - (\zeta + \zeta^q))^{\frac{q+1}{2}} \qquad \text{and}$$

$$P V^{-\frac{q^2+q}{2}} = \Big(\prod_{\zeta \in \mathcal{Z}} (W_+ - (\zeta + \zeta^q))^{\frac{q+1}{2}} \Big) F^q - (W_+ - 2)^{\frac{q^2-q}{2}} G^{q+1}.$$

In characteristic 2 each factor of the product over $\mathcal{Z}$ appears twice, thus justifying their exponent $\frac{q+1}{2}$.

By writing $P_1 = p_1 W_+ - p_0$, with $p_i \in \overline{K}[A]$, $\gcd(p_0, p_1) = 1$, $\max(\deg(p_0), \deg(p_1)) = 2$ and p_0 being monic, the divisibility $P_1 \mid P$ transforms into the following polynomial identity

in $\overline{K}[A]$:

$$\Big(\prod_{\zeta\in\mathcal{Z}}(p_0-(\zeta+\zeta^q)p_1)^{\frac{q+1}{2}}\Big)F^q=p_1^q(p_0-2p_1)^{\frac{q^2-q}{2}}G^{q+1}.$$

Again the degree of the first factor on the left hand side must be divisible by q (respectively, $\frac{q}{2}$ in characteristic 2), and since $p_0-(\zeta+\zeta^q)p_1$ can be constant or linear for at most one sum $\zeta+\zeta^q$, the degree of the first factor must be q^2+q for $q>4$. Also the degree of p_0-2p_1 must be zero since $q>3$ and thus the degree of p_1 is 2.

In even characteristic $p_0-2p_1=p_0$ is a constant, thus $p_0=1$ (p_0 is monic). The involution $\zeta\mapsto\zeta^q=\zeta^{-1}$ on $\mathcal{Z}$ has no fixed points, and, denoting by $\mathcal{Z}_2$ a set of representatives of $\mathcal{Z}$ modulo the involution, one obtains

$$\prod_{\zeta\in\mathcal{Z}_2}(1-(\zeta+\zeta^q)p_1)=c_1G,\quad F=c_2p_1\quad\text{and}\quad c_1^{q+1}c_2^q=1.$$

Modulo F one gets $F\mid c_1G-1$ which implies $c_1\in K$. Thus $c_2\in K$, $p_1\in K[A]$ and therefore $P_1\in K[A,U]$.

In odd characteristic the factor corresponding to $\zeta=-1$, namely $(p_0+2p_1)^{\frac{q+1}{2}}$, is coprime to the other factors in the product and coprime to $p_1(p_0-2p_1)$. Hence p_0+2p_1 must be a square and its square root must divide G. Moreover, one gets $F=c_1p_1$. Since $p_0-2p_1=c_2$ is a constant and p_0 is monic, one gets $c_1=2\alpha$, implying $p_1\in K[A]$. Since $p_0+2p_1=4p_1+c_2$ is a square, its discriminant is zero, thus $c_2\in K$ and hence $P_1\in K[A,U]$.

If S_1 is one of the two dihedral subgroups of order $q-1$ (which implies that q is odd), the argumentation is similar. The polynomials P and P_1 are expressed in terms of $V'=U^{\frac{q-1}{2}}$ and then, since $U\mapsto\frac{1}{cU}$ becomes $V'\mapsto c^{-\frac{q-1}{2}}\frac{1}{V'}$ with $c^{-\frac{q-1}{2}}=\pm1$, in terms of $W'_+=V'+\frac{1}{V'}$ or $W'_-=V'-\frac{1}{V'}$, respectively. In the first case P is rewritten as

$$PV'^{-(q^2+q)}=\Big(\prod_{\zeta\in\mathcal{Z}'}(W'_+-(\zeta+\zeta^{-1}))^{\frac{q+1}{2}}\Big)F^q-(W'_+-2)^{\frac{q^2-q}{2}}(W'_++2)^{\frac{q^2-q}{2}}G^{q+1}$$

where $\mathcal{Z}'=\mu_{2(q+1)}(\mathbb{F}_{q^2})\setminus\{\pm1\}$. By setting $P_1=p_1W'_+-p_0$ with $p_i\in\overline{K}[A]$, $\gcd(p_0,p_1)=1$, $\max(\deg(p_0),\deg(p_1))=1$ and p_0 being monic, one obtains

$$\Big(\prod_{\zeta\in\mathcal{Z}'}(p_0-(\zeta+\zeta^{-1})p_1)^{\frac{q+1}{2}}\Big)F^q=p_1^{2q}(p_0-2p_1)^{\frac{q^2-q}{2}}(p_0+2p_1)^{\frac{q^2-q}{2}}G^{q+1}.$$

Since one of $p_0\pm2p_1$ is not constant, the degree of the right hand side exceeds the degree of the left hand side for $q>5$ which is a contradiction.

In the second case P is rewritten as

$$PV'^{-(q^2+q)}=\Big(\prod_{\zeta\in\mathcal{Z}'}(W'_--(\zeta-\zeta^{-1}))^{\frac{q+1}{2}}\Big)F^q-W'^{q^2-q}_-G^{q+1}$$

and by setting $P_1=p_1W'_--p_0$ with $p_i\in\overline{K}[A]$, $\gcd(p_0,p_1)=1$, $\max(\deg(p_0),\deg(p_1))=1$ and p_0 being monic, one obtains

$$\Big(\prod_{\zeta\in\mathcal{Z}'}(p_0-(\zeta-\zeta^{-1})p_1)^{\frac{q+1}{2}}\Big)F^q=p_1^{2q}p_0^{q^2-q}G^{q+1}.$$

Considering the degrees for $q>3$ it follows that p_0 must be constant and hence p_1 is of degree one. Since p_1 is coprime to the first factor on the left hand side, it must divide F^q which implies $\rho_1=\rho_2\in K$, contradicting the irreducibility of F.

5.6.3 *The Borel case* In this case, rewriting P and P_1 in terms of $W = (U^q - U)^{q-1}$ gives

$$P = (W+1)^{q+1}F^q - W^q G^{q+1}$$

and $\deg_W(P_1) = 1$, $P_1 = p_1 W - p_0$, with $p_i \in \overline{K}[A]$, $\gcd(p_0, p_1) = 1$, $\max(\deg(p_0), \deg(p_1)) = q$ and p_1 being monic. Then the divisibility $P_1 \mid P$ transforms into the following polynomial identity in $\overline{K}[A]$:

$$(p_0 + p_1)^{q+1}F^q = p_1 p_0^q G^{q+1}.$$

From $\deg(G^{q+1}) = q^2 + q$, $\deg(p_1 p_0^q) \geqslant q$ and $\deg(F^q) = 2q$ it follows that the degree of $p_0 + p_1$ must be q. This implies $\deg(F^q) = \deg(p_1 p_0^q)$, thus $\deg(p_0) \leqslant 2$ and therefore $\deg(p_1) = q$, since $q > 2$, and $\deg(p_0) = 1$.

Since $p_0 + p_1$ is coprime to $p_0 p_1$, it follows

$$p_0 + p_1 = c_1 G, \quad p_1 = \tilde{p}^q, \quad F = c_2 \tilde{p} p_0 \quad \text{and} \quad c_1^{q+1} c_2^q = 1$$

for a monic linear polynomial $\tilde{p} \in \overline{K}[A]$.

Exchanging ρ_1 and ρ_2, if needed, one obtains

$$\tilde{p} = A - \rho_1, \quad p_0 = c_3(A - \rho_2), \quad c_1 = 1, \quad c_2 = 1 \quad \text{and} \quad c_3 = \alpha.$$

Evaluating $p_0 + p_1 = G$ at $A = 0$ gives

$$\rho_1^q + \alpha\rho_2 + \delta = 0.$$

This means that condition $(*)$ does not hold.

5.6.4 *The* PSL_2 *case* This case can only occur for odd q, and then P splits as $P = sP_1P_2$ with a scalar $s \in \overline{K}$. The map $U \mapsto aU$ for a non-square $a \in \mathbb{F}_q$ exchanges P_1 and P_2. Since $\mathrm{PSL}_2(\mathbb{F}_q)$ is a normal subgroup of $\mathrm{PGL}_2(\mathbb{F}_q)$, P_2 is invariant under $\mathrm{PSL}_2(\mathbb{F}_q)$ as well. By rewriting P in terms of $W' = (U^q - U)^{\frac{q-1}{2}}$ one obtains

$$P = (W'^2 + 1)^{q+1}F^q - W'^{2q}G^{q+1} = sP_1(W')P_1(-W').$$

Denoting by $p_0 \in \overline{K}[A]$ the constant coefficient of $P_1 \in (\overline{K}[A])[W']$ this becomes modulo W'

$$F^q = sp_0^2$$

which implies $\rho_1 = \rho_2 \in K$, contradicting the irreducibility of F.

5.6.5 *The case* $\mathrm{PGL}_2(\mathbb{F}_{q'})$ Since $\mathrm{PGL}_2(\mathbb{F}_{q'}) \subset \mathrm{PSL}_2(\mathbb{F}_q)$ in odd characteristic, one can reduce this case to the previous case as follows.

Let $I_1 \subset \{1, \ldots, g\}$ be the subset of i such that S_i is a conjugate of S_1 by an element in $\mathrm{PSL}_2(\mathbb{F}_q)$, and let $I_2 = \{1, \ldots, g\} \setminus I_1$. These two sets correspond to the two orbits of the action of $\mathrm{PSL}_2(\mathbb{F}_q)$ on the S_i (or P_i). Both orbits contain $\#I_1 = \#I_2 = \frac{g}{2}$ elements and an element in $\mathrm{PGL}_2(\mathbb{F}_q) \setminus \mathrm{PSL}_2(\mathbb{F}_q)$ transfers one orbit into the other.

Let $\tilde{P}_j = \prod_{i \in I_j} P_i$, $j = 1, 2$, then P splits as $P = s\tilde{P}_1\tilde{P}_2$, $s \in \overline{K}$, and both $\tilde{P}_j$, $j = 1, 2$, are invariant under $\mathrm{PSL}_2(\mathbb{F}_q)$. Notice that the absolute irreducibility of P_1 and P_2 was not used in the argument in the PSL_2 case.

This completes the proof of Theorem 7.

5.7 Traps

In the following Proposition 8 and Proposition 9 are proven.

Let Q be an irreducible quadratic polynomial in $K[X]$ such that $(1, u_0X + u_1), (X, v_0X + v_1)$ is a basis of the lattice L_Q, so that Q is a scalar multiple of $-u_0X^2 + (-u_1 + v_0)X + v_1 = F(-X)$ and has roots $-\rho_1$ and $-\rho_2$. By definition of L_Q the pair (h_0, h_1) must be in the dual lattice (scaled by Q), given by the basis $(u_0X + u_1, -1), (v_0X + v_1, -X)$.

For the assertions concerning conditions $(*)$ and $(**)$, assume that $\rho_1, \rho_2 \in L \setminus K$ and that

$$\rho_1^q + \alpha\rho_j + \delta = 0$$

holds for $j = 1$ or $j = 2$.

First consider the case $j = 2$, i.e., condition $(*)$. To show that $-\rho_i$, $i = 1, 2$, are roots of $h_1X^q - h_0$ it is sufficient to show this for the basis of the dual lattice of L_Q given above. For $(u_0X + u_1, -1)$ one computes

$$-(-\rho_1^q) - u_0(-\rho_1) - u_1 = \rho_1^q - \alpha\rho_1 - \beta + \delta = -\alpha\rho_2 - \alpha\rho_1 - \beta = 0,$$

and for $(v_0X + v_1, -X)$ one obtains

$$-(-\rho_1)(-\rho_1^q) - v_0(-\rho_1) - v_1 = (-\rho_1^q - \delta)\rho_1 - \gamma = \alpha\rho_1\rho_2 - \gamma = 0.$$

Therefore $h_1X^q - h_0$ is divisible by Q, which is then a trap of level 0.

In the case $j = 1$ an analogous calculation shows that $-\rho_i$, $i = 1, 2$, are roots of $h_1X^{q^{kd+1}} - h_0$, namely for $(u_0X + u_1, -1)$ one has

$$-(-\rho_2^{q^{kd+1}}) - u_0(-\rho_2) - u_1 = \rho_1^q - \alpha\rho_2 - \beta + \delta = -\alpha\rho_1 - \alpha\rho_2 - \beta = 0$$

and for $(v_0X + v_1, -X)$ one gets

$$-(-\rho_2)(-\rho_2^{q^{kd+1}}) - v_0(-\rho_2) - v_1 = (-\rho_1^q - \delta)\rho_2 - \gamma = \alpha\rho_1\rho_2 - \gamma = 0$$

Therefore $h_1X^{q^{kd+1}} - h_0$ is divisible by Q, which is then a trap of level kd. This finishes the proof of Proposition 8.

Regarding Proposition 9, note that a solution (a, B) gives rise to the polynomial $\mathcal{P}_a = a(u_0X + (Y + u_1)) + ((Y + v_0)X + v_1)$. If, for $j = 1$ or $j = 2$, ρ is a root of $\mathcal{P}_a \bmod f_j$ for two different values of a, then ρ is a root of $u_0X + (Y + u_1) \bmod f_j$ and of $(Y + v_0)X + v_1 \bmod f_j$. Since

$$-X(u_0X + (Y + u_1)) + (Y + v_0)X + v_1 = -u_0X^2 + (-u_1 + v_0)X + v_1 = F(-X),$$

which equals Q up to a scalar, it follows that ρ is also a root of Q. Furthermore, in the case $j = 1$ the polynomial $\mathcal{P}_a \bmod f_1$ splits completely, so that $\rho \in K$, contradicting the irreducibility of Q, finishing the proof of Proposition 9.

This completes the proof of Theorem 4.

Acknowledgements

The authors are indebted to Claus Diem for explaining how one can obviate the need to compute the logarithms of the factor base elements, and wish to thank him also for some enlightening discussions.

References

Adl79 Leonard M. Adleman. A subexponential algorithm for the discrete logarithm problem with applications to cryptography. In *Proceedings of the 20th Annual Symposium on Foundations*

of Computer Science, SFCS '79, pages 55–60, Washington, DC, USA, 1979. IEEE Computer Society.

Adl94 Leonard M. Adleman. The function field sieve. In Leonard M. Adleman and Ming-Deh Huang, editors, *Algorithmic Number Theory*, volume 877 of *Lecture Notes in Computer Science*, pages 108–121. Springer Berlin Heidelberg, 1994.

AH99 Leonard M. Adleman and Ming-Deh A. Huang. Function field sieve method for discrete logarithms over finite fields. *Inform. and Comput.*, 151(1-2):5–16, 1999.

AP96 Yves Aubry and Marc Perret. A Weil theorem for singular curves. In *Arithmetic, geometry and coding theory (Luminy, 1993)*, pages 1–7. de Gruyter, Berlin, 1996.

BGJT14 Razvan Barbulescu, Pierrick Gaudry, Antoine Joux, and Emmanuel Thomé. A heuristic quasi-polynomial algorithm for discrete logarithm in finite fields of small characteristic. In *Advances in Cryptology—EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 1–16. Springer, 2014.

Blu04 Antonia W. Bluher. On $x^{q+1} + ax + b$. *Finite Fields and Their Applications*, 10(3):285–305, 2004.

Chu89 Fan-Rong K. Chung. Diameters and eigenvalues. *J. Amer. Math. Soc.*, 2(2):187–196, 1989.

Cop84a Don Coppersmith. Evaluating logarithms in $\mathrm{GF}(2^n)$. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, STOC '84, pages 201–207, New York, NY, USA, 1984. ACM.

Cop84b Don Coppersmith. Fast evaluation of logarithms in fields of characteristic two. *IEEE Trans. Inf. Theor.*, 30(4):587–594, 1984.

COTR06 Peter J. Cameron, Gholam R. Omidi, and Behruz Tayfeh-Rezaie. 3-designs from $\mathrm{PGL}(2, q)$. *Electron. J. Combin.*, 13(1):Research Paper 50, 11, 2006.

CWZ14 Qi Cheng, Daqing Wan, and Jincheng Zhuang. Traps to the bgjt-algorithm for discrete logarithms. *LMS Journal of Computation and Mathematics*, 17:218–229, 2014.

DB51 Nicolaas G. De Bruijn. On the number of positive integers $\leqslant x$ and free of prime factors $> y$. *Indagationes Mathematicae*, 13:50–60, 1951.

DB66 Nicolaas G. De Bruijn. On the number of positive integers $\leqslant x$ and free of prime factors $> y$, II. *Indagationes Mathematicae*, 28:239–247, 1966.

DH06 Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theor.*, 22(6):644–654, September 2006.

Dic01 Leonard E. Dickson. *Linear groups: With an exposition of the Galois field theory.* Teubner, Leipzig, 1901.

Dic30 Karl Dickman. On the frequency of numbers containing prime factors of a certain relative magnitude. *Arkiv för Matematik, Astonomi och Fysik*, 22A (10):1–14, 1930.

Die11 Claus Diem. On the discrete logarithm problem in elliptic curves. *Compositio Mathematica*, 147:75–104, 1 2011.

EG02 Andreas Enge and Pierrick Gaudry. A general framework for subexponential discrete logarithm algorithms. *Acta Arithmetica*, 102:83–103, 2002.

Gau65 Carl F. Gauss. *Disquisitiones Arithmeticae.* Translated by Arthur A. Clarke. Yale University Press, 1965.

GGMZ13a Faruk Göloğlu, Robert Granger, Gary McGuire, and Jens Zumbrägel. On the function field sieve and the impact of higher splitting probabilities. Available from `eprint.iacr.org/2013/074`, 15th Feb 2013.

GGMZ13b Faruk Göloğlu, Robert Granger, Gary McGuire, and Jens Zumbrägel. On the function field sieve and the impact of higher splitting probabilities. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology—CRYPTO 2013*, volume 8043 of *LNCS*, pages 109–128. Springer, 2013.

GGMZ13c Faruk Göloğlu, Robert Granger, Gary McGuire, and Jens Zumbrägel. Discrete Logarithms in $GF(2^{1971})$. NMBRTHRY list, 19/2/2013.

GGMZ13d Faruk Göloğlu, Robert Granger, Gary McGuire, and Jens Zumbrägel. Discrete Logarithms in $GF(2^{6120})$. NMBRTHRY list, 11/4/2013.

GGMZ14 Faruk Göloğlu, Robert Granger, Gary McGuire, and Jens Zumbrägel. Solving a 6120-bit DLP on a desktop computer. In *Selected Areas in Cryptography—SAC 2013*, volume 8282 of *LNCS*, pages 136–152. Springer, 2014.

GKZ14a Robert Granger, Thorsten Kleinjung, and Jens Zumbrägel. Breaking '128-bit secure' supersingular binary curves - (or how to solve discrete logarithms in $\mathbb{F}_{2^{4\cdot 1223}}$ and $\mathbb{F}_{2^{12\cdot 367}}$). In *Advances in Cryptology—CRYPTO 2014*, volume 8617 of *LNCS*, pages 126–145. Springer, 2014.

GKZ14b Robert Granger, Thorsten Kleinjung, and Jens Zumbrägel. On the powers of 2. Available from `eprint.iacr.org/2014/300`, 29th Apr 2014.

GKZ14c Robert Granger, Thorsten Kleinjung, and Jens Zumbrägel. Discrete logarithms in the Jacobian of a genus 2 supersingular curve over $GF(2^{367})$. NMBRTHRY list, 30/1/2014.

GKZ14d Robert Granger, Thorsten Kleinjung, and Jens Zumbrägel. Discrete Logarithms in $GF(2^{9234})$. NMBRTHRY list, 31/1/2014.

HK10 Tor Helleseth and Alexander Kholosha. $x^{2^l+1} + x + a$ and related affine polynomials over GF(2^k). *Cryptogr. Commun.*, 2(1):85–109, 2010.

JL02 Antoine Joux and Reynald Lercier. The function field sieve is quite special. In Claus Fieker and David R. Kohel, editors, *Algorithmic number theory (Sydney, 2002)*, volume 2369 of *LNCS*, pages 431–445. Springer, 2002.

JL06 Antoine Joux and Reynald Lercier. The function field sieve in the medium prime case. In Serge Vaudenay, editor, *Advances in Cryptology—EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 254–270. Springer, 2006.

Jou13a Antoine Joux. Faster index calculus for the medium prime case. application to 1175-bit and 1425-bit finite fields. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology—EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 177–193. Springer, 2013.

Jou13b Antoine Joux. A new index calculus algorithm with complexity $L(1/4 + o(1))$ in very small characteristic. Available from `eprint.iacr.org/2013/095`, 20th Feb 2013.

Jou13c Antoine Joux. Discrete Logarithms in $GF(2^{1778})$. NMBRTHRY list, 11/2/2013.

Jou13d Antoine Joux. Discrete Logarithms in $GF(2^{4080})$. NMBRTHRY list, 22/3/2013.

Jou13e Antoine Joux. Discrete Logarithms in $GF(2^{6168})$. NMBRTHRY list, 21/5/2013.

Jou14 Antoine Joux. A new index calculus algorithm with complexity $L(1/4 + o(1))$ in small characteristic. In Tanja Lange, Kristin Lauter, and Petr Lisoněk, editors, *Selected Areas in Cryptography—SAC 2013*, volume 8282 of *LNCS*, pages 355–379. Springer, 2014.

KB79 Ravindran Kannan and Achim Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. *SIAM J. Comput.*, 8(4):499–507, 1979.

Kle14 Thorsten Kleinjung. Discrete logarithms in GF(2^{1279}). NMBRTHRY list, 17/10/2014.

Kra22 Maurice Kraitchik. *Théorie des nombres*, volume 1. Paris: Gauthier-Villars, 1922.

Kra24 Maurice Kraitchik. *Recherches sur la théorie des nombres*, volume 1. Paris: Gauthier-Villars, 1924.

Len91 Hendrik W. Lenstra, Jr. Finding isomorphisms between finite fields. *Math. Comp.*, 56(193):329–347, 1991.

LL93 Arjen K. Lenstra and Hendrik W. Lenstra, Jr., editors. *The development of the number field sieve*, volume 1554 of *Lecture Notes in Mathematics*. Springer, Heidelberg, 1993.

LP92 Hendrik W. Lenstra, Jr. and Carl Pomerance. A rigorous time bound for factoring integers. *J. Amer. Math. Soc.*, 5(3):483–516, 1992.

Mer79 Ralph C. Merkle. *Secrecy, Authentication, and Public Key Systems.* PhD thesis, Stanford University, Stanford, CA, USA, 1979.

PH78 Stephen C. Pohlig and Martin E. Hellman. An improved algorithm for computing logarithms over gf(p) and its cryptographic significance (corresp.). *IEEE Trans. Inf. Theory*, 24(1):106–110, 1978.

PJ14 Cecile Pierrot and Antoine Joux. Discrete logarithm record in characteristic 3, GF($3^{5\cdot 479}$) a 3796-bit field. NMBRTHRY list, 15/9/2014.

Pol78 John M. Pollard. Monte Carlo Methods for Index Computation (mod p). *Mathematics of Computation*, 32:918–924, 1978.

Pom87 Carl Pomerance. Fast, rigorous factorization and discrete logarithm algorithms. In *Discrete algorithms and complexity (Kyoto, 1986)*, volume 15 of *Perspect. Comput.*, pages 119–143. Academic Press, Boston, MA, 1987.

RS62 J. Barkley Rosser and Lowell Schoenfeld. Approximate formulas for some functions of prime numbers. *Illinois J. Math.*, 6:64–94, 1962.

SSHT12 Naoyuki Shinohara, Takeshi Shimoyama, Takuya Hayashi, and Tsuyoshi Takagi. Key length estimation of pairing-based cryptosystems using η_t pairing. In Mark D. Ryan, Ben Smyth, and Guilin Wang, editors, *Information Security Practice and Experience*, volume 7232 of *Lecture Notes in Computer Science*, pages 228–244. Springer Berlin Heidelberg, 2012.

Val91 Brigitte Vallée. Generation of elements with small modular squares and provably fast integer factoring algorithms. *Math. Comp.*, 56(194):823–849, 1991.

Wan97 Daqing Wan. Generators and irreducible polynomials over finite fields. *Mathematics of Computation*, 66:1195–1212, 1997.

WM68 A. E. Western and Jefferey C. P. Miller. *Tables of indices and primitive roots.* Royal Society Mathematical Tables, vol. 9, Cambridge University Press, 1968.

Robert Granger robert.granger@epfl.ch
Laboratory for Cryptologic Algorithms, School of Computer and Communication Sciences, École polytechnique fédérale de Lausanne, Switzerland

Thorsten Kleinjung thorsten.kleinjung@epfl.ch
Institute of Mathematics, Universität Leipzig, Germany

Jens Zumbrägel jens.zumbragel@epfl.ch
Laboratory for Cryptologic Algorithms, School of Computer and Communication Sciences, École polytechnique fédérale de Lausanne, Switzerland