

A NEW PERSPECTIVE ON THE POWERS OF TWO DESCENT FOR DISCRETE LOGARITHMS IN FINITE FIELDS

THORSTEN KLEINJUNG AND BENJAMIN WESOŁOWSKI

ABSTRACT. A new proof is given for the correctness of the powers of two descent method for computing discrete logarithms. The result is slightly stronger than the original work, but more importantly we provide a unified geometric argument, eliminating the need to analyse all possible subgroups of $\mathrm{PGL}_2(\mathbf{F}_q)$. Our approach sheds new light on the role of PGL_2 , in the hope to eventually lead to a complete proof that discrete logarithms can be computed in quasi-polynomial time in finite fields of fixed characteristic.

1. INTRODUCTION

In this paper we prove the following result.

Theorem 1.1. *Given a prime power q , a positive integer d , coprime polynomials h_0 and h_1 in $\mathbf{F}_{q^d}[x]$ of degree at most two, and an irreducible degree ℓ factor I of $h_1x^q - h_0$, the discrete logarithm problem in $\mathbf{F}_{q^{d\ell}} \cong \mathbf{F}_{q^d}[x]/(I)$ can be solved in expected time $q^{\log_2 \ell + O(d)}$.*

It was originally proven in [GKZ18] when $q > 61$, q is not a power of 4, and $d \geq 18$. Even though we eliminate these technical conditions, the main contribution is the new approach to the proof. The theorem represents the state of the art of provable quasi-polynomial time algorithms for the discrete logarithm problem (or DLP) in finite fields of fixed characteristic. The obstacle separating Theorem 1.1 from a full provable algorithm for DLP is the question of the existence of a good field representation: polynomials h_0, h_1 and I for a small d . A direction towards a full provable algorithm would be to find analogues of this theorem for other field representations, but this may require in the first place a good understanding of why Theorem 1.1 is true.

The integers q, d and ℓ , and the polynomials h_0, h_1 and I are defined as in the above theorem for the rest of the paper. The core of that result is Proposition 1.3 below, which essentially states that elements of $\mathbf{F}_{q^{d\ell}}$ represented by a *good* irreducible polynomial in $\mathbf{F}_{q^d}[x]$ of degree $2m$ can be rewritten as a product of *good* irreducible polynomials of degrees dividing m — a process called *degree two elimination*, first introduced for $m = 1$ in [GGMZ13].

Definition 1.2 (Traps and good polynomials). An element $\tau \in \overline{\mathbf{F}}_q$ for which $[\mathbf{F}_{q^d}(\tau) : \mathbf{F}_{q^d}]$ is an even number $2m$ and $h_1(\tau) \neq 0$ is called

- (1) a *degenerate trap root* if $\frac{h_0}{h_1}(\tau) \in \mathbf{F}_{q^{dm}}$,
- (2) a *trap root of level 0* if it is a root of $h_1x^q - h_0$, or
- (3) a *trap root of level dm* if it is a root of $h_1x^{q^{dm+1}} - h_0$.

Analogously, a polynomial in $\overline{\mathbf{F}}_q[x]$ that has a trap root is called a *trap*. A polynomial is *good* if it is not a trap.

Proposition 1.3 (Degree two elimination). *Given an extension k/\mathbf{F}_{q^d} of degree m such that $dm \geq 23$, and a good irreducible quadratic polynomial $Q \in k[x]$, there is an algorithm which finds a list of good linear polynomials (L_0, \dots, L_n) in $k[x]$ such that $n \leq q + 1$ and*

$$Q \equiv h_1 L_0^{-1} \cdot \prod_{i=1}^n L_i \pmod{I},$$

and runs in expected polynomial time in q , d and m .

The difficulty of proving Theorem 1.1 lies mostly in Proposition 1.3. We recall briefly in Section 1.2 how the proposition implies the theorem. The main contribution of the present paper is a new proof of Proposition 1.3, which hopefully provides a better understanding of the degree two elimination method, the underlying geometry, and the role of traps. The action of PGL_2 on the polynomial $x^q - x$ became a crucial ingredient in the recent progress on the discrete logarithm problem for fields of small characteristic, since [Jou13] (and implicitly in [GGMZ13]). While the proof in [GKZ18] resorted to an intricate case by case analysis enumerating through all possible subgroups of $\mathrm{PGL}_2(\mathbf{F}_q)$, we provide a unified geometric argument, shedding new light on the role of PGL_2 .

1.1. Degree two elimination algorithm. The key observation allowing degree two elimination is that a polynomial of the form $\alpha x^{q+1} + \beta x^q + \gamma x + \delta$ has a high chance to split completely over its field of definition. Furthermore, we have the congruence

$$(1.1) \quad \alpha x^{q+1} + \beta x^q + \gamma x + \delta \equiv h_1^{-1}(\alpha x h_0 + \beta h_0 + \gamma x h_1 + \delta h_1) \pmod{I},$$

and the numerator of the right-hand side has degree at most 3. Consider the $\overline{\mathbf{F}}_q$ -vector space V spanned by x^{q+1}, x^q, x and 1 in $\overline{\mathbf{F}}_q[x]$, and the linear subspace

$$V_Q = \{\alpha x^{q+1} + \beta x^q + \gamma x + \delta \in V \mid \alpha x h_0 + \beta h_0 + \gamma x h_1 + \delta h_1 \equiv 0 \pmod{Q}\}.$$

As long as Q is a good irreducible polynomial, V_Q is of dimension two. The algorithm simply consists in sampling uniformly at random elements $f \in V_Q(k)$ (or equivalently in its projectivisation $\mathbf{P}_Q^1(k)$) until f splits completely over k into good linear polynomials $(L_1, \dots, L_{\deg f})$. Since $f \in V_Q$, the polynomial Q divides the numerator of the right-hand side of (1.1), and the quotient is a polynomial L_0 of degree at most 1. The algorithm returns $(L_0, \dots, L_{\deg f})$.

To prove that the algorithm terminates in expected polynomial time, we need to show that a random polynomial in $V_Q(k)$ has good chances to split into good linear polynomials over k . In this paper, we prove this by constructing a morphism $C \rightarrow \mathbf{P}_Q^1$ where C is an absolutely irreducible curve defined over k , such that the image of any k -rational point of C is a polynomial that splits completely over k . This construction is the object of Section 4. The absolute irreducibility implies that C has a lot of k -rational points, allowing us to deduce that a lot of polynomials in $\mathbf{P}_Q^1(k)$ split over k . This is done in Section 5.

1.2. Proof of Theorem 1.1. We briefly explain in this section how Proposition 1.3 implies Theorem 1.1. Consider the factor base

$$\mathfrak{F} = \{f \in \mathbf{F}_{q^d}[x] \mid \deg f \leq 1, f \neq 0\} \cup \{h_1\}.$$

First, the following proposition extends the degree two elimination to a full descent algorithm from any polynomial down to the factor base.

Proposition 1.4. *Suppose $d \geq 23$. Given a polynomial $F \in \mathbf{F}_{q^d}[x]$, there is an algorithm that finds integers $(\alpha_f)_{f \in \mathfrak{F}}$ such that*

$$F \equiv \prod_{f \in \mathfrak{F}} f^{\alpha_f} \pmod{I},$$

and runs in expected time $q^{\log_2 \ell + O(d)}$.

Proof. This is essentially the *zigzag* descent presented in [GKZ18]. We recall the main idea for the convenience of the reader. First, one finds a good irreducible polynomial $G \in \mathbf{F}_{q^d}[x]$ of degree 2^e such that $F \equiv G \pmod{I}$ (this can be done for $e = \lceil \log_2(4\ell + 1) \rceil$, see [Wan97, Th. 5.1] and [GKZ18, Lem. 2]). Over the extension $\mathbf{F}_{q^{d2^{e-1}}}$, the polynomial G splits into 2^{e-1} good irreducible quadratic polynomials, all conjugate under $\text{Gal}(\mathbf{F}_{q^{d2^{e-1}}}/\mathbf{F}_{q^d})$. Let Q be one of them, and apply the algorithm of Proposition 1.3 to rewrite Q in terms of linear polynomials (L_0, \dots, L_n) in $\mathbf{F}_{q^{d2^{e-1}}}[x]$ and h_1 . For any index i , let L'_i be the product of all the conjugates of L_i in the extension $\mathbf{F}_{q^{d2^{e-1}}}/\mathbf{F}_{q^d}$. Then,

$$F \equiv h_1^{2^{e-1}} L_0'^{-1} \cdot \prod_{i=1}^n L'_i \pmod{I},$$

and each L'_i factors into good irreducible polynomials of degree a power of 2 at most 2^{e-1} . The descent proceeds by iteratively applying this method to each L'_i until all the factors are in the factor base \mathfrak{F} . \square

Then, as in [GKZ18, Sec. 2], the descent algorithm of Proposition 1.4 can be used to compute discrete logarithms, following ideas from [EG02] and [Die11]. To compute the discrete logarithm of an element h in base g , the idea is to collect relations between g , h , and elements of the factor base by applying the descent algorithm on $g^\alpha h^\beta$ for a few uniformly random exponents α and β . That proves Theorem 1.1 for $d \geq 23$. To remove the condition on d , suppose that $d \leq 22$, and let $d' \leq 44$ be the smallest multiple of d larger than 22. Let I' be an irreducible factor of I in $\mathbf{F}_{q^{d'}}[x]$. The DLP can be solved in expected time $q^{\log_2(\deg I') + O(d')} = q^{\log_2 \ell + O(1)}$ in $\mathbf{F}_{q^{d'}}[x]/(I')$, and therefore also in the subfield $\mathbf{F}_{q^d}[x]/(I)$.

2. THE ACTION OF PGL_2 ON $x^q - x$

As already mentioned, a crucial fact behind degree two elimination is that a polynomial of the form $\alpha x^{q+1} + \beta x^q + \gamma x + \delta$ has a high chance to split completely over its field of definition. This fact is closely related to the action of 2×2 matrices on such polynomials.

Definition 2.1. We denote by \star the action of invertible 2×2 matrices on univariate polynomials defined as follows:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \star f(x) = (cx + d)^{\deg f} f\left(\frac{ax + b}{cx + d}\right).$$

Consider the $\overline{\mathbf{F}}_q$ -vector subspace V spanned by x^{q+1}, x^q, x , and 1 in $\overline{\mathbf{F}}_q[x]$. The above action induces an action of the group PGL_2 on the projective space $\mathbf{P}(V)$, which we also write \star . Parameterizing the polynomials in $\mathbf{P}(V)$ as $\alpha x^{q+1} + \beta x^q +$

$\gamma x + \delta$, let S be the quadratic surface in $\mathbf{P}(V)$ defined by the equation $\alpha\delta = \beta\gamma$. This surface is the image of the morphism

$$\psi : \mathbf{P}^1 \times \mathbf{P}^1 \longrightarrow \mathbf{P}(V) : (a, b) \longmapsto (x - a)(x - b)^q.$$

Note that to avoid heavy notation, everything is written affinely, but we naturally have $\psi(\infty, b) = (x - b)^q$, $\psi(a, \infty) = x - a$ and $\psi(\infty, \infty) = 1$. More generally, we say that $f(x) \in V$ has a root of degree n at infinity if f is of degree $q + 1 - n$. Now, the following lemma shows that apart from the surface S , the polynomials of $\mathbf{P}(V)$ form exactly one orbit for PGL_2 .

Lemma 2.2. *We have $\mathbf{P}(V) \setminus S = \mathrm{PGL}_2 \star (x^q - x)$.*

Proof. First notice that both S and $\mathbf{P}(V) \setminus S$ are closed under the action of PGL_2 . In particular, $\mathrm{PGL}_2 \star (x^q - x) \subseteq \mathbf{P}(V) \setminus S$. Let $f(x) \in \mathbf{P}(V) \setminus S$. Suppose by contradiction that $f(x)$ has a double root $r \in \mathbf{P}^1$, and let $g \in \mathrm{PGL}_2$ be a linear transformation sending 0 to r . The polynomial $g \star f(x)$ has a double root at 0, so has no constant or linear term, and must be of the form $\alpha x^{q+1} + \beta x^q$, so it is in S , a contradiction. Therefore $f(x)$ has $q + 1$ distinct roots. Let $g \in \mathrm{PGL}_2$ send 0, 1 and ∞ to three of these roots. Then, $g \star f(x)$ has a root at 0 and at ∞ so is of the form $\beta x^q + \gamma x$, and since it also has a root at 1, it can only be $x^q - x$. \square

This result implies that most polynomials of $\mathbf{P}(V)$ are of the form $g \star (x^q - x)$, which splits completely over the field of definition of the matrix g .

3. THE ROLE OF TRAPS

Consider a finite field extension k/\mathbf{F}_{q^d} of degree m . Let Q be an irreducible quadratic polynomial in $k[x]$ coprime to h_1 . Let a_1 and a_2 be the roots of Q in $\overline{\mathbf{F}}_q$. The degree two elimination aims at expressing Q modulo $h_1 x^q - h_0$ as a product of linear polynomials. To do so, we study a variety $\mathbf{P}_Q^1 \subset \mathbf{P}(V)$ parameterizing polynomials that can possibly lead to an elimination of Q (i.e., such that Q divides the right hand side of (1.1)). In this section, we define \mathbf{P}_Q^1 and show how the notion of traps and good polynomials determine how it intersects the surface S from Lemma 2.2.

Recall that V is the $\overline{\mathbf{F}}_q$ -vector subspace V spanned by x^{q+1}, x^q, x , and 1 in $\overline{\mathbf{F}}_q[x]$. Consider the linear map

$$(3.1) \quad \varphi : V \longrightarrow \overline{\mathbf{F}}_q[x][h_1^{-1}] : \begin{cases} 1 & \longmapsto 1, \\ x & \longmapsto x, \\ x^q & \longmapsto h_0/h_1, \\ x^{q+1} & \longmapsto xh_0/h_1. \end{cases}$$

We want \mathbf{P}_Q^1 to parameterise the polynomials $f \in V$ such that $\varphi(f)$ is divisible by Q . For any $P \in \overline{\mathbf{F}}_q[x]$ coprime with h_1 , write $\varphi_P = \pi_P \circ \varphi$ where $\pi_P : \overline{\mathbf{F}}_q[x][h_1^{-1}] \rightarrow \overline{\mathbf{F}}_q[x]/P$ is the canonical projection. We can now define \mathbf{P}_Q^1 as

$$(3.2) \quad \mathbf{P}_Q^1 = \mathbf{P}(\ker \varphi_Q).$$

The variety \mathbf{P}_Q^1 is the intersection of the two planes $\mathbf{P}(\ker \varphi_{x-a_1})$ and $\mathbf{P}(\ker \varphi_{x-a_2})$.

Lemma 3.1. *If Q is not a degenerate trap, then $|\mathbf{P}_Q^1 \cap S|(\overline{\mathbf{F}}_q)| = 2$, and these two points are of the form $\psi(a_1, b_1)$ and $\psi(a_2, b_2)$, with $a_1 \neq a_2$ and $b_1 \neq b_2$.*

Proof. For $a \in \{a_1, a_2\}$, we have

$$\mathbf{P}(\ker \varphi_{x-a}) \cap S = \psi(\{a\} \times \mathbf{P}^1) \cup \psi\left(\mathbf{P}^1 \times \left\{\frac{h_0}{h_1}(a)^{1/q}\right\}\right).$$

Since the polynomial Q is irreducible, we have $a_1 \neq a_2$. Furthermore, assuming that Q is not a degenerate trap, we have $\frac{h_0}{h_1}(a_1) \notin k$, and thereby $\frac{h_0}{h_1}(a_1) \neq \frac{h_0}{h_1}(a_2)$. Therefore $\mathbf{P}_Q^1 \cap S$ is equal to

$$\mathbf{P}(\ker \varphi_{x-a_1}) \cap \mathbf{P}(\ker \varphi_{x-a_2}) \cap S = \left\{ \psi\left(a_1, \frac{h_0}{h_1}(a_2)^{1/q}\right), \psi\left(a_2, \frac{h_0}{h_1}(a_1)^{1/q}\right) \right\}.$$

□

In particular, when Q is not a degenerate trap, \mathbf{P}_Q^1 is exactly the line passing through the two points $s_1 = \psi(a_1, b_1)$ and $s_2 = \psi(a_2, b_2)$. We get a k -isomorphism $\mathbf{P}^1 \rightarrow \mathbf{P}_Q^1 : \alpha \mapsto s_1 - \alpha s_2$. For this reason the two points s_1 and s_2 play a central role in the rest of the analysis, and the following proposition shows that they behave nicely when Q is a good polynomial.

Proposition 3.2. *Suppose Q is a good polynomial. Then, $(\mathbf{P}_Q^1 \cap S)(\overline{\mathbf{F}}_q) = \{s_1, s_2\}$, where $s_1 = (x - a_1)(x - b_1)^q$, and $s_2 = (x - a_2)(x - b_2)^q$, and the roots a_1, a_2, b_1 and b_2 are all distinct.*

Proof. From Lemma 3.1, we can write $(\mathbf{P}_Q^1 \cap S)(\overline{\mathbf{F}}_q) = \{s_1, s_2\}$ with $a_1 \neq a_2$ and $b_1 \neq b_2$. If $a_1 = b_2$ or $a_2 = b_1$, then Q divides $x^q h_1 - h_0$, a trap of level 0. Now, suppose $a_1 = b_1$ (the case $a_2 = b_2$ is similar). Since a_1 and a_2 are the two roots of Q , and Q divides $(x - a_1)(h_0 - a_1^q h_1)$, then a_2 is a root of $h_0 - a_1^q h_1$. We get that $h_0(a_2) = a_1^q h_1(a_2)$, so a_2 is a root of $h_1 x^{q^{dm+1}} - h_0$, a trap of level dm . □

4. IRREDUCIBLE COVERS OF \mathbf{P}_Q^1

In this section we suppose that Q is a good polynomial, and we consider the polynomials $s_1 = (x - a_1)(x - b_1)^q$ and $s_2 = (x - a_2)(x - b_2)^q$ as defined in Proposition 3.2, where a_1, a_2, b_1 and b_2 are all distinct. Consider the variety \mathbf{P}_Q^1 from (3.2).

Recall that our goal is to prove that a significant proportion of the polynomials of $\mathbf{P}_Q^1(k)$ splits completely over k . As mentioned in Section 1.1, our method consists in constructing a morphism $C \rightarrow \mathbf{P}_Q^1$ where C is an absolutely irreducible curve defined over k , such that the image of any k -rational point of C is a polynomial that splits completely over k . The absolute irreducibility is crucial as it implies that C has a lot of k -rational points. The idea is to consider the algebraic set

$$C = \{(u, r_1, r_2, r_3) \mid \text{the } r_i\text{'s are three distinct roots of } u\} \subset \mathbf{P}_Q^1 \times \mathbf{P}^1 \times \mathbf{P}^1 \times \mathbf{P}^1,$$

and the canonical projection $C \rightarrow \mathbf{P}_Q^1$.

Proposition 4.1. *If $(u, r_1, r_2, r_3) \in C(k)$, then u splits completely over k .*

Proof. Suppose that (u, r_1, r_2, r_3) is a k -rational point of C . From Lemma 2.2, we get $u = g \star (x^q - x)$ where g is the matrix $g \in \text{PGL}_2(k)$ sending the three points r_1, r_2 and r_3 to 0, 1 and ∞ . In particular, the set of roots of u is $g^{-1}(\mathbf{P}^1(\mathbf{F}_q))$ which are all in $\mathbf{P}^1(k)$. □

In the rest of this section, we prove that C is absolutely irreducible (Proposition 4.6). The strategy is the following. Instead of considering directly C , which encodes three roots for each polynomial of \mathbf{P}_Q^1 , we start with the variety

$$X = \{(u, r) \mid u(r) = 0\} \subset \mathbf{P}_Q^1 \times \mathbf{P}^1,$$

which considers a single root for each polynomial. We can then “add” roots by considering fibre products. Recall that given two covers $\nu : Z \rightarrow Y$ and $\mu : Z' \rightarrow Y$, the geometric points of the fibre product $Z \times_Y Z'$ are pairs (z, z') such that $\nu(z) = \mu(z')$. In particular, the fibre product over the projection $X \rightarrow \mathbf{P}_Q^1$ is

$$\begin{aligned} X \times_{\mathbf{P}_Q^1} X &= \{((u_1, r_1), (u_2, r_2)) \mid u_1(r_1) = 0, u_2(r_2) = 0, u_1 = u_2\} \\ &\cong \{(u, r_1, r_2) \mid u(r_1) = 0, u(r_2) = 0\}. \end{aligned}$$

This product $X \times_{\mathbf{P}_Q^1} X$ contains a trivial component, the diagonal, corresponding to triples (u, r, r) . The rest is referred to as the non-trivial part, and we prove that it is an absolutely irreducible curve (Corollary 4.3). Iterating this construction, the fibre product $(X \times_{\mathbf{P}_Q^1} X) \times_X (X \times_{\mathbf{P}_Q^1} X)$ (over the projection $X \times_{\mathbf{P}_Q^1} X \rightarrow X$ to the first component) encodes quadruples (u, r_1, r_2, r_3) . Therefore the curve C naturally embeds into the non-trivial part of this product. We prove that this non-trivial part is itself an absolutely irreducible curve (Lemma 4.5).

Instead of the projection $X \rightarrow \mathbf{P}_Q^1$, we work with an isomorphic cover θ . It is easy to see that the canonical projection $X \rightarrow \mathbf{P}^1$ is an isomorphism, with inverse $r \mapsto (s_2(r)s_1 - s_1(r)s_2, r)$. Through the isomorphisms $X \cong \mathbf{P}^1$ and $\mathbf{P}_Q^1 \cong \mathbf{P}^1$, this projection is isomorphic to the cover θ in the following commutative diagram (where, again, the morphisms are written affinely for convenience):

$$\begin{array}{ccccc} & & (u, r) \longmapsto & u & \\ & & & & \\ (u, r) & & X & \longrightarrow & \mathbf{P}_Q^1 & & s_1 - \alpha s_2 \\ & \downarrow & \downarrow \wr & & \downarrow \wr & & \downarrow \alpha \\ & r & \mathbf{P}^1 & \xrightarrow{\theta} & \mathbf{P}^1 & & \\ & & r \longmapsto & s_1(r)/s_2(r) & & & \end{array}$$

For convenience, consider θ as a cover $X_1 \rightarrow X_0$ where $X_0 = X_1 = \mathbf{P}^1$. As a first step, we study the induced fibre product $X_1 \times_{X_0} X_1$. It contains the diagonal Δ_1 , isomorphic to X_1 . We wish to show that $Y_2 = X_1 \times_{X_0} X_1 \setminus \Delta_1$ is absolutely irreducible. The second step consists in showing that $X_2 \times_{X_1} X_2 \setminus \Delta_2$ is also absolutely irreducible, where X_2 is a desingularisation of Y_2 and Δ_2 is the diagonal. The following lemma provides a general method used in both steps.

Lemma 4.2. *Let Y and Z be two absolutely irreducible, smooth, complete curves over k , and consider a cover $\eta : Z \rightarrow Y$. If there exists a point $a \in Z$ such that η is not ramified at a and $\#(\eta^{-1}(\eta(a))) = 2$, then $Z \times_Y Z \setminus \Delta$ is absolutely irreducible, where Δ is the diagonal component.*

Proof. By contradiction, suppose that $Z \times_Y Z \setminus \Delta$ is not absolutely irreducible, and can be decomposed as two components $A \cup B$. Let $\text{pr} : Z \times_Y Z \rightarrow Z$ be the projection on the first factor. Since $Z \times_Y Z$ is complete, both A and B are

complete, so we have $\text{pr}(A) = \text{pr}(B) = \text{pr}(\Delta) = Z$. Observe that $\text{pr}^{-1}(a)$ consists of $\#(\eta^{-1}(\eta(a))) = 2$ points, so one of them must belong to two of the components A , B and Δ . That point must therefore be singular in $Z \times_Y Z$, contradicting the fact that η is not ramified at a (recall that a point $(z_1, z_2) \in Z \times_Y Z$ is singular if and only if η is ramified at both z_1 and z_2). \square

Corollary 4.3. *The curve $Y_2 = X_1 \times_{X_0} X_1 \setminus \Delta_1$ is absolutely irreducible.*

Proof. First observe that θ is ramified only at b_1 and b_2 (as can be verified from the explicit formula $\theta(r) = s_1(r)/s_2(r)$). In particular, it is not ramified at a_1 . Since $\#(\theta^{-1}(\theta(a_1))) = \#\{a_1, b_1\} = 2$, we apply Lemma 4.2. \square

Lemma 4.4. *The desingularisation morphism $\nu : X_2 \rightarrow Y_2$ is a bijection between the geometric points.*

Proof. It is sufficient to prove that for any singular point P on Y_2 , and $\varphi : \tilde{Y}_2 \rightarrow Y_2$ the blowing-up at P , the preimage $\varphi^{-1}(P)$ consists of a single smooth point. Up to a linear transformation of $X_1 = \mathbf{P}^1$, we can assume that s_1 and s_2 are of the form $s_1(x) = (x-1)x^q$ and $s_2(x) = x-a$, for some $a \neq 0, 1$. The intersection A of the curve Y_2 with the affine patch $\mathbf{A}^2 \subset \mathbf{P}^1 \times \mathbf{P}^1$ is then defined by the polynomial

$$f(x, y) = \frac{s_1(x)s_2(y) - s_1(y)s_2(x)}{x-y} = \frac{x^q(x-1)(y-a) - y^q(y-1)(x-a)}{x-y}.$$

It remains to blow up A at the singularity $(0, 0)$ (which corresponds to (b_1, b_1) through the linear transformation), and check the required properties. This is easily done following [Har77, Ex. 4.9.1], and we include details for the benefit of the reader. Let $\psi : Z \rightarrow \mathbf{A}^2$ be the blowing-up of \mathbf{A}^2 at $(0, 0)$. The inverse image of A in Z is defined in $\mathbf{A}^2 \times \mathbf{P}^1$ by the equations $f(x, y) = 0$ and $ty = xu$ (where t and u parameterize the factor \mathbf{P}^1). It consists of two irreducible components: the blowing-up \tilde{A} of A at $(0, 0)$ and the exceptional curve $\psi^{-1}(0, 0)$. Suppose $t \neq 0$, so we can set $t = 1$ and use u as an affine parameter (since f is symmetric, the case $u \neq 0$ is similar). We have the affine equations $f(x, y) = 0$ and $y = xu$, and substituting we get $f(x, xu) = 0$, which factors as

$$f(x, xu) = x^{q-1} \frac{(x-1)(xu-a) - u^q(xu-1)(x-a)}{1-u}.$$

The blowing-up \tilde{A} is defined on $t = 1$ by the equations $g(x, u) = f(x, xu)/x^{q-1} = 0$ and $y = xu$. It meets the exceptional line only at the point $u = 1$, which is non-singular. \square

The projection $X_1 \times_{X_0} X_1 \rightarrow X_1$ on the first component induces another cover $\theta_2 : X_2 \rightarrow X_1$, through which we build the fibre product $X_2 \times_{X_1} X_2$. As above, it contains a diagonal component Δ_2 isomorphic to X_2 .

Lemma 4.5. *The curve $Y_3 = X_2 \times_{X_1} X_2 \setminus \Delta_2$ is absolutely irreducible.*

Proof. Let $\nu : X_2 \rightarrow Y_2$ be the bijective morphism from Lemma 4.4. Since θ_1 is only ramified at b_1 and b_2 , the cover θ_2 is ramified at most at the points $\nu^{-1}(b_i, b_i)$ and $\nu^{-1}(a_i, b_i)$ (for $i \in \{1, 2\}$). In particular, it is not ramified at $\nu^{-1}(b_1, a_1)$. Since $\#(\theta_2^{-1}(\theta_2(\nu^{-1}(b_1, a_1)))) = \#\{\nu^{-1}(b_1, a_1), \nu^{-1}(b_1, b_1)\} = 2$, we apply Lemma 4.2. \square

Proposition 4.6. *The curve C is absolutely irreducible.*

Proof. Let $\nu : X_2 \rightarrow Y_2$ be the morphism from Lemma 4.4. It is an isomorphism away from the singularities of Y_2 , so

$$C \longrightarrow Y_3 : (u, r_1, r_2, r_3) \longmapsto (\nu^{-1}(r_1, r_2), \nu^{-1}(r_1, r_3))$$

is a morphism. It is an embedding, and the result follows from Lemma 4.5. \square

5. COUNTING SPLIT POLYNOMIALS IN \mathbf{P}_Q^1

Recall that we wish to prove Proposition 1.3 by showing that $\mathbf{P}_Q^1(k)$ contains a lot of polynomials that split into good polynomials over k . The results of Section 4 allow us to prove in Theorem 5.1 that a lot of polynomials in $\mathbf{P}_Q^1(k)$ do split. We then show in Proposition 5.2 that all these polynomials are coprime, which implies that bad polynomials cannot appear too often.

Theorem 5.1. *Let k/\mathbf{F}_{q^a} be a field extension of degree m , and Q be a good irreducible quadratic polynomial in $k[x]$ coprime to h_1 . If $dm \geq 23$, there are at least $\#k/2q^3$ polynomials in \mathbf{P}_Q^1 that split completely over the field k .*

Proof. Let $\Theta : Y_3 \rightarrow \mathbf{P}_Q^1$ be the cover resulting from the composition of the successive covers of Section 4. Let $S_3 = \Theta^{-1}(\mathbf{P}_Q^1 \cap S)$. The embedding $C \rightarrow Y_3$ from Proposition 4.6 has image $Y_3 \setminus S_3$. The morphism

$$\mu : Y_3 \rightarrow \mathbf{P}^1 \times \mathbf{P}^1 \times \mathbf{P}^1 : (\nu^{-1}(r_1, r_2), \nu^{-1}(r_1, r_3)) \mapsto (r_1, r_2, r_3)$$

restricts to an embedding of $Y_3 \setminus S_3$. Let A be the intersection of $\mu(Y_3)$ with the affine patch \mathbf{A}^3 . The curve A is a component of the (reducible) curve defined by the equations $\theta(r_1) = \theta(r_2)$ and $\theta(r_1) = \theta(r_3)$. Therefore A is of degree at most $4(q+1)^2$. If B is the closure of A in \mathbf{P}^3 , then [Bac96, Th. 3.1] shows that

$$|\#B(k) - \#k - 1| \leq 16(q+1)^4 \sqrt{\#k}.$$

Since Y_3 is complete, $\mu(Y_3)$ is closed, so all the points of $B \setminus A$ are at infinity, and there are at most $\deg(B) \leq 4(q+1)^2$ of them. Also, at most $2(q^3 - q)$ points of B are in $\mu(S_3)$ (because $\#S = 2$ and Θ is of degree $q^3 - q$). Therefore

$$\#C(k) = \#(Y_3 \setminus S_3)(k) \geq \#k + 1 - 16(q+1)^4 \sqrt{\#k} - 4(q+1)^2 - 2(q^3 - q).$$

Since $q \geq 2$ and $dm \geq 23$, we get $\#C(k) \geq \#k/2$. From Proposition 4.1, and the fact that the map Θ is $q^3 - q$ to one, we get that at least $\#k/2q^3$ polynomials in \mathbf{P}_Q^1 split completely over k . \square

Let φ be the morphism defined in (3.1).

Proposition 5.2. *Suppose Q is a good polynomial. For any two distinct polynomials f and g in $\mathbf{P}_Q^1(\overline{\mathbf{F}}_q)$, we have $\gcd(f, g) = 1$ and $\gcd(h_1\varphi(f), h_1\varphi(g)) = Q$.*

Proof. Let s_1 and s_2 be as in Proposition 3.2. They have no common root. Since f and g are distinct, all the polynomials of \mathbf{P}_Q^1 are of the form $\alpha f + \beta g$ for $(\alpha : \beta) \in \mathbf{P}^1$. Then, if r is a root of f and g , r is a root of all the polynomials of \mathbf{P}_Q^1 . In particular, it is a root of both s_1 and s_2 , a contradiction. This shows that $\gcd(f, g) = 1$.

Similarly, if a polynomial h divides $h_1\varphi(f)$ and $h_1\varphi(g)$, it must also divide both $h_1\varphi(s_1) = (x - a_1)(h_0 - b_1^q h_1)$, and $h_1\varphi(s_2) = (x - a_2)(h_0 - b_2^q h_1)$. Since $h_0 - b_1^q h_1$ and $h_0 - b_2^q h_1$ are coprime, h must divide Q . \square

Proof of Proposition 1.3. As discussed in Section 1.1, it is sufficient to prove that a uniformly random element of $\mathbf{P}_Q^1(k)$ has a good probability to lead to an elimination into good polynomials. A polynomial $f \in \mathbf{P}_Q^1(k)$ leads to an elimination into good polynomials if f splits completely over k into good linear polynomials, and $\varphi(f)$ is itself a good polynomial.

Let A be the set of polynomials of $\mathbf{P}_Q^1(k)$ that split completely over k . From Theorem 5.1, A contains at least $q^{dm-3}/2$ elements. Trap roots τ occurring in A or $\varphi(A)$ must be roots of $h_1x^q - h_0$, or of $h_1x^{q^{dn+1}} - h_0$ for $n \mid m/2$, or satisfy $\frac{h_0}{h_1}(\tau) \in \mathbf{F}_{q^{dm/2}}$. There are at most $q^{\frac{dm}{2}+3}$ such trap roots. From Proposition 5.2, any trap root can only occur once in A and in $\varphi(A)$. So there are at most $2q^{\frac{dm}{2}+3}$ polynomials in A for which trap roots appear. Therefore the number of elements in A leading to a good reduction is at least

$$\frac{1}{2}q^{dm-3} - 2q^{\frac{dm}{2}+3} \geq \frac{1}{2}(q^{dm-3} - 4q^{dm-8}) \geq \frac{1}{4}q^{dm-3},$$

using $dm \geq 23$. Since $\mathbf{P}_Q^1(k)$ contains $q^{dm} + 1$ elements, the probability of a random element to lead to a good elimination is $1/O(q^3)$. \square

ACKNOWLEDGEMENTS

Part of this work was supported by the Swiss National Science Foundation under grant number 200021-156420.

REFERENCES

- [Bac96] Eric Bach, *Weil bounds for singular curves*, *Applicable Algebra in Engineering, Communication and Computing* **7** (1996), no. 4, 289–298.
- [Die11] Claus Diem, *On the discrete logarithm problem in elliptic curves*, *Compositio Mathematica* **147** (2011), no. 1, 75–104.
- [EG02] Andreas Enge and Pierrick Gaudry, *A general framework for subexponential discrete logarithm algorithms*, *Acta Arithmetica* **102** (2002), 83–103.
- [GGMZ13] Faruk Göloğlu, Robert Granger, Gary McGuire, and Jens Zumbrägel, *On the function field sieve and the impact of higher splitting probabilities*, *Advances in Cryptology – CRYPTO 2013*, Springer Berlin Heidelberg, 2013, pp. 109–128.
- [GKZ18] Robert Granger, Thorsten Kleinjung, and Jens Zumbrägel, *On the discrete logarithm problem in finite fields of fixed characteristic*, *Trans. Amer. Math. Soc.* **270** (2018), no. 5, 3129–3145.
- [Har77] R. Hartshorne, *Algebraic Geometry*, Springer-Verlag, New York, 1977, Graduate Texts in Mathematics, No. 52.
- [Jou13] Antoine Joux, *A new index calculus algorithm with complexity $L(1/4 + o(1))$ in small characteristic*, *Selected Areas in Cryptography - SAC 2013*, Lecture Notes in Computer Science, vol. 8282, Springer, 2013, pp. 355–379.
- [Wan97] Daqing Wan, *Generators and irreducible polynomials over finite fields*, *Mathematics of Computation* **66** (1997), no. 219, 1195–1212.