

INDISCREET LOGARITHMS IN FINITE FIELDS OF SMALL CHARACTERISTIC

ROBERT GRANGER AND THORSTEN KLEINJUNG

Laboratory for Cryptologic Algorithms
École polytechnique fédérale de Lausanne
Station 14, 1015 Lausanne, Switzerland

JENS ZUMBRÄGEL

Faculty of Computer Science and Mathematics
University of Passau
Innstraße 33, 94032 Passau, Germany

(Communicated by Marcus Greferath)

ABSTRACT. Recently, several striking advances have taken place regarding the discrete logarithm problem (DLP) in finite fields of small characteristic, despite progress having remained essentially static for nearly thirty years, with the best known algorithms being of subexponential complexity. In this expository article we describe the key insights and constructions which culminated in two independent quasi-polynomial algorithms. To put these developments into both a historical and a mathematical context, as well as to provide a comparison with the cases of so-called large and medium characteristic fields, we give an overview of the state-of-the-art algorithms for computing discrete logarithms in all finite fields. Our presentation aims to guide the reader through the algorithms and their complexity analyses *ab initio*.

1. INTRODUCTION

Given a finite cyclic group (G, \cdot) , a generator $g \in G$ and another group element $h \in G$, the discrete logarithm problem (DLP) is the computational problem of finding an integer x such that $g^x = h$. The integer x – denoted by $\log_g h$ – is uniquely determined modulo the group order and is called the *discrete logarithm* of h to the base g . The most natural example is the multiplicative group \mathbb{F}_p^\times of the field \mathbb{F}_p of integers modulo a prime p , which is a cyclic group in which the DLP is believed to be hard in general. Besides this classical case, several other groups have been extensively studied, including the multiplicative group \mathbb{F}_q^\times of any finite field \mathbb{F}_q of prime power order $q = p^n$, the group $E(\mathbb{F}_q)$ of \mathbb{F}_q -rational points on an elliptic curve E , or the group $J_C(\mathbb{F}_q)$ associated with the Jacobian J_C of a (usually hyperelliptic) curve C of genus > 1 .

The study of discrete logarithms can be traced back to at least two centuries ago, when they appeared in Gauß' *Disquisitiones Arithmeticae* and were referred to as *indices* [22, art. 57–60]. Beyond its relevance as a natural computational problem, the study of the DLP really came of age with the advent of public-key cryptography

2010 *Mathematics Subject Classification*: Primary: 11Y16, 11T71.

Key words and phrases: Discrete logarithm problem, finite fields, number field sieve, function field sieve, quasi-polynomial algorithms.

The first author is supported by the Swiss National Science Foundation via grant number 200021-156420.

in 1976, as the hardness of the DLP (originally in \mathbb{F}_p^\times) forms the basis of the famous Diffie-Hellman protocol [18] and other cryptographic primitives.

To describe the hardness of the DLP in a group of order N one usually considers the asymptotic complexity in the input size of the problem, which is proportional to $\log N$. To indicate the complexity, it has become customary to use the notation

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

where $\alpha \in [0, 1]$, $c > 0$ and \log denotes the natural logarithm; we may omit the subscript N when there is no ambiguity and sometimes write $L(\alpha)$ to mean $L(\alpha, c)$ for some $c > 0$. Note that $L(0) = (\log N)^{c+o(1)}$, which corresponds to polynomial time, while $L(1) = N^{c+o(1)}$ corresponds to exponential time. An algorithm with a running time of $L(\alpha)$ for some $0 < \alpha < 1$ is said to be of *subexponential* complexity.

Algorithms for solving the DLP can be broadly classified into two classes. One class consists of the so-called generic algorithms, which do not exploit any representational properties of group elements and may thus be applied to any group. Generic algorithms, like Pollard's rho method [56], necessarily have a running time of $\Omega(\sqrt{N})$ [51, 67], if N is prime; if N is composite then one may reduce the original problem to the DLP in prime order subgroups, using the Pohlig-Hellman algorithm [55]. The other class consists of the so-called index calculus algorithms, which exploit a suitable representation of group elements, which imbues a notion of smoothness. A basic version of index calculus was analysed by Adleman [4] and shown to have complexity $L(1/2)$. Subsequently, more advanced index calculus algorithms with lower complexity have been developed.

In particular, the first $L(1/3)$ algorithm for computing discrete logarithms, published in 1984 by Coppersmith [16], targeted binary finite fields. Later, the number field sieve, originally devised for the integer factoring problem [47], was adapted for the DLP in prime fields [25, 61] and resulted again in an $L(1/3)$ algorithm. Inspired by the number field sieve, the function field sieve was developed for computing discrete logarithms in small (fixed) characteristic [5, 6, 36, 38]. Regarding general finite fields \mathbb{F}_{p^n} , for $n \ll (\log p)^{1/2}$ the number field sieve and for $\log p \ll n^{1/2}$ the function field sieve still have complexity $L(1/3)$. Finally, in 2006 the number field sieve was generalised [39] to work also in the remaining cases, thereby obtaining $L(1/3)$ algorithms for all families of finite field DLPs. Some improvements of the medium number field sieve have been reported lately, see e.g., [8, 10, 11, 43].

Regarding the case of small characteristic, dramatic advances have recently taken place. Indeed, nearly thirty years after Coppersmith's algorithm the $L(1/3)$ barrier was broken in a series of remarkable results, starting with the work of Göloğlu, Granger, McGuire and Zumbrägel [23], and independently Joux [35]. The approach of Joux led to the first heuristic quasi-polynomial algorithm, i.e., one with running time $\exp(O((\log \log N)^2)) = (\log N)^{O(\log \log N)}$, which is in $L(o(1))$, due to Barbulescu, Gaudry, Joux and Thomé [9]. Independently, the approach of Göloğlu *et al.* was later developed into an alternative algorithm by Granger, Kleinjung and Zumbrägel, which also has quasi-polynomial complexity, but is rigorous for a large family of fields [28, 29].

In this article we describe the key insights and constructions underlying the recent algorithms. To put these developments into both a historical and a mathematical context, as well as to provide a comparison with the cases of so-called large and medium characteristic fields, we give an overview of the state-of-the-art algorithms

for computing discrete logarithms in all finite fields. All necessary mathematical prerequisites will be briefly introduced along the way.

As our focus is on DLP algorithms for finite fields, we cover generic algorithms only briefly. We refer to Odlyzko's paper [54] or the recent survey [40] for an overview on these and other general aspects of the DLP, including curve-based groups, which are not covered here.

A remark on terminology. The cardinality of any finite field F is a prime power $|F| = p^n$, where the prime p is called the *characteristic* of F , which is the smallest positive integer m such that $m1_F = 0$. Conversely, given any prime power $q = p^n$ there exists a finite field of size q , which is unique up to isomorphism and is denoted by \mathbb{F}_q .

When considering families of finite fields of order $q = p^n$, where $p = L_q(\alpha)$ and $q \rightarrow \infty$, then different DLP algorithms apply depending on α . Accordingly, in the case $\alpha > \frac{2}{3}$ we speak of *large* characteristic; for $\alpha \in (\frac{1}{3}, \frac{2}{3})$ we speak of *medium* characteristic; and for $\alpha \in (0, \frac{1}{3})$ we say that the characteristic is *medium-small*, while the boundary cases $\alpha = \frac{2}{3}$ and $\alpha = \frac{1}{3}$ are special cases to be treated additionally. Finally, in the case $\alpha = 0$, i.e., if p is of polynomial size in $\log q$, we speak of finite fields of *small* characteristic. In particular, if $q = p$ is a prime or $q = p^n$ with n fixed, then we are in (very) large characteristic $p = L_q(1)$, whereas if the characteristic p is fixed, then we have small characteristic $p = L_q(0)$; note however that $p = L_q(1)$ (or $p = L_q(0)$) does not imply that n (or p) has to be fixed.

Furthermore, for complexity considerations we make use of the notation $f \approx g$ to indicate that $f/g \rightarrow 1$ as the argument goes to infinity.

Outline. We discuss the DLP for general finite cyclic groups in Section 2 and briefly present the most common cryptographic applications and generic algorithms. The index calculus method serves as a framework for all advanced DLP algorithms for finite fields and will be described in Section 3, where we first present it abstractly in a general group and then give some basic concrete instances. Section 4 is devoted to the number field sieve, which is currently the fastest method for DLP in both large and medium characteristic, while Section 5 deals with the function field sieve in finite fields of small (and medium-small) characteristic. Section 6 presents the recent dramatic developments for small characteristic finite fields, as well as the key insights and ideas which led to them. Finally, we conclude in Section 7.

2. THE DLP IN FINITE CYCLIC GROUPS

The discrete logarithm problem can be formulated for any finite cyclic group; indeed, most cryptographic protocols using the DLP can be formulated in this abstract setting. We present in this section the most important cryptographic applications and the common generic algorithms. Since our main focus is the DLP in finite fields to which index calculus algorithms may be applied, we will here be rather brief in our treatment.

Let (G, \cdot) be a finite cyclic group of order N and let $g \in G$ be a generator. We assume that the group operation can be computed efficiently, i.e., in polynomial time, and that the group order is known. The surjective group homomorphism $\mathbb{Z} \rightarrow G$, $x \mapsto g^x$ induces a group isomorphism

$$\psi: \mathbb{Z}/N\mathbb{Z} \rightarrow G, \quad [x] \mapsto g^x,$$

with inverse map $\psi^{-1} = \log_g: G \rightarrow \mathbb{Z}/N\mathbb{Z}$, $h \mapsto \log_g h$. The map ψ can be computed efficiently by using a square-and-multiply method, whereas the computation of ψ^{-1} is believed to be a difficult problem in general; this is of course the DLP. Indeed, for so-called *generic algorithms*, in which one can only perform group operations and check equality of elements, the DLP necessarily has an exponential running time of $\Omega(\sqrt{N})$ if N is prime [51, 67]. However, for particular groups with concrete representations, faster algorithms exist, and it is not proven in any such cases that the DLP is hard.

One should note that the DLP, as well as the integer factorisation problem, can be solved in polynomial time using quantum computers [66], but in this article we confine ourselves to the classical computational model.

Cryptographic applications. The difficulty of the DLP is nowadays widely used, e.g., for secure communication over the Internet. Virtually all public-key cryptosystems in use today are based on either the integer factorisation problem or the DLP. Some common cryptographic protocols using the DLP are briefly presented below. In each case the group (G, \cdot) and a generator $g \in G$ are assumed to be publicly known.

In the Diffie-Hellman key-agreement protocol [18] two parties, usually referred to as Alice and Bob, choose random integers a and b , respectively, and exchange the group elements g^a and g^b over the public channel, hence they both can compute a common session key $(g^b)^a = g^{ab} = (g^a)^b$. Clearly, if the DLP in G is feasible, then the key can be computed from the public information, so it is essential that the DLP is hard.

The Diffie-Hellman protocol can be transformed into a public-key encryption scheme as shown by ElGamal [19]. Indeed, if Bob has announced his public key g^b and Alice has a secret message $m \in G$ for Bob, she chooses a random integer a and sends the pair $(g^a, m(g^b)^a)$ to Bob, who can decrypt by computing $(g^a)^b = (g^b)^a$. Moreover, there are digital signature schemes based on the DLP, e.g., the ElGamal [19] and Schnorr [64] signature schemes.

More recently, the advent of pairing-based cryptography in 2000 provided new cryptographic functionalities, starting with identity-based non-interactive key distribution [59], one-round tripartite key-agreement [33] and identity-based encryption [13], followed by numerous others. It also renewed interest in the underlying DLPs and related security assumptions.

To summarise, the hardness of the DLP in any group in which any of the above protocols is instantiated is essential to their security. It is therefore imperative that these DLPs continue to be studied.

Generic algorithms. We now present examples of generic algorithms for solving the DLP in any finite cyclic group. Suppose that we want to find the discrete logarithm $\log_g h$ for a target element $h \in G$. Recall that $N = |G|$ is the group order.

In the *Baby-Step-Giant-Step* method, attributed to Shanks, let $M = \lceil \sqrt{N} \rceil$. One computes a table $\{(j, g^j) \mid j \in \{0 \dots M-1\}\}$ (baby steps) and sorts it by the second component. Then one computes $k = g^{-M}$, as well as h, hk, hk^2, \dots (giant steps) until a collision $hk^i = g^j$ is detected, in which case the output is $\log_g h = iM + j$. The method requires $O(\sqrt{N})$ database entries and $O(\sqrt{N} \log N)$ group operations (or $O(\sqrt{N})$ if hash tables are used).

Pollard's rho method [56] reduces the memory requirement to a negligible amount while preserving the square-root running time. Therefore, it is the preferred method in practice. However, due to its randomised nature the analysis is more intricate.

The idea is to define pseudorandom sequences (k_i) in G and $(a_i), (b_i)$ in $\mathbb{Z}/N\mathbb{Z}$ such that $g^{a_i} h^{b_i} = k_i$ holds for any i . If $k_j = k_{j+\ell}$ holds for some $j, \ell > 0$, then there exists also i with $k_i = k_{2i}$, and such collisions can be easily detected. In this case, $\log_g h = \frac{a_{2i} - a_i}{b_i - b_{2i}}$ is found, provided that the denominator is invertible modulo N .

Moreover, there is the *Pohlig-Hellman* method [55], the efficiency of which depends on the prime factorisation of $N = p_1^{e_1} \cdots p_r^{e_r}$, which we assume is given. Then one has $\mathbb{Z}/N\mathbb{Z} \cong \mathbb{Z}/p_1^{e_1}\mathbb{Z} \times \cdots \times \mathbb{Z}/p_r^{e_r}\mathbb{Z}$ by the Chinese Remainder Theorem, hence $x = \log_g h$ corresponds to a tuple (x_1, \dots, x_r) and one can consider the DLP for each factor $x_i \in \mathbb{Z}/p_i^{e_i}\mathbb{Z}$ separately. Furthermore, if $e_i > 1$, then $x_i \in \mathbb{Z}/p_i^{e_i}\mathbb{Z}$ can be found by obtaining the “p-ary” digits of x_i , one at a time starting with the least significant digit, by solving a DLP of order p_i . Therefore, the Pohlig-Hellman method essentially reduces the DLP complexity in a group of order N to a group of order the largest prime factor of N , for which either the Baby-Step-Giant-Step or Pollard’s rho method may be applied.

We remark that, although we are mainly interested in the DLP in finite fields – to which more efficient index calculus algorithms apply – to compute the DLP modulo the order of the full multiplicative group it is often preferable to use the Pohlig-Hellman method for computing the solution modulo the small prime power factors.

3. THE INDEX CALCULUS METHOD

The index calculus method can be much more efficient than generic algorithms for certain groups and representations, and is the template for all subexponential DLP algorithms. While we describe the framework for any group, the details will depend on the concrete representation of the group elements.

As before let $\log_g h$ be the target logarithm in a cyclic group G of order N . We choose a subset $F \subseteq G$ such that $\langle F \rangle = G$ (often we have $g \in F$), called the *factor base*. The idea is to first obtain $\log_g f$ for all $f \in F$. Let $A = \mathbb{Z}/N\mathbb{Z}$ and consider the surjective group homomorphism

$$\varphi: A^F \rightarrow G, \quad (e_f)_{f \in F} \mapsto \prod_{f \in F} f^{e_f}.$$

The index calculus method consists of three phases.

1. **Relation generation.** Find vectors $(e_f)_{f \in F}$ in $\ker \varphi$, called *relations*, thus generating a subset $R \subseteq \ker \varphi$.
2. **Linear algebra.** Compute an element $0 \neq (x_f)_{f \in F} \in R^\perp$, i.e., satisfying $\sum_{f \in F} x_f e_f = 0$ for all $(e_f)_{f \in F} \in R$.
3. **Individual logarithm.** Find a preimage $(e_f)_{f \in F} \in \varphi^{-1}(h)$, for then we have found $\log_g h = \sum_{f \in F} e_f \log_g f$, cf. Lemma 3.1.

As the next result shows, provided that enough relations have been found, the discrete logarithms of the factor base elements are determined up to a scalar multiple. In practice, the condition given in the lemma is satisfied, if a few more relations than factor base elements have been obtained.

Lemma 3.1. *Suppose that $\text{span } R = \ker \varphi$ and $(x_f)_{f \in F} \in R^\perp$, then there exists $\lambda \in A$ such that $x_f = \lambda \log_g f$ for all $f \in F$.*

Proof. It holds that $R^\perp = (\text{span } R)^\perp = (\ker \varphi)^\perp$ and $A^F / \ker \varphi \cong \text{im } \varphi \cong A$, while $(\ker \varphi)^\perp \cong A^F / \ker \varphi$ as can be seen by considering the Smith normal form, hence $R^\perp \cong A$. On the other hand, we have $(\log_g f)_{f \in F} \in R^\perp$; indeed, since $\log_g: G \rightarrow A$

is a group homomorphism, it holds $\sum_{f \in F} e_f \log_g f = \log_g(\prod_{f \in F} f^{e_f}) = 0$ for all $(e_f)_{f \in F} \in R \subseteq \ker \varphi$. \square

Sparse linear algebra. After the relation generation step an $r \times m$ -matrix M over $\mathbb{Z}/N\mathbb{Z}$ has been found, where m is the size of the factor base F and $r \approx m$ is the number of relations. In order to obtain the factor base logarithms we need a solution $0 \neq x \in (\mathbb{Z}/N\mathbb{Z})^m$ of $Mx = 0$. Due to the relation generation method the matrix is usually of low average row weight w . For such sparse matrices iterative algorithms are available, most commonly used are the Lanczos method [45] (for an adaption to finite fields see, e.g., [44]) or the Wiedemann algorithm [69]. Their cost is dominated by repeated computations of matrix-vector products Mv , and the running time is $O(m^2w)$ operations in $\mathbb{Z}/N\mathbb{Z}$. Provided that $\log w = o(\log m)$ (and $\log \log N = o(\log m)$), which is usually the case, this is of complexity $m^{2+o(1)}$.

We note that, as $O(m)$ divisions in $\mathbb{Z}/N\mathbb{Z}$ are necessary (when using the Lanczos method), the group order N should avoid small prime factors, therefore the Pohlig-Hellman reduction and Pollard's rho algorithm should be used for the small prime power factors. In practice, for large-scale computations the linear algebra step poses some challenges, as the iterative algorithms are not easily parallelisable. We also remark that so-called structured Gaussian elimination (cf. [44, 37]) can be used to decrease the matrix dimension m , while increasing the weight w only moderately.

3.1. A VARIANT PERMITTING RIGOROUS ANALYSIS. The following variant of the index calculus method, proposed by Enge and Gaudry [20] and subsequently refined by Diem [17], is valuable for theoretical analysis. With this variant one can compute discrete logarithms, provided only that it is feasible to express group elements as products over the factor base, as in the individual logarithm step of the classical index calculus method.

As before, let (G, \cdot) be a cyclic group of order N , let $g \in G$ be a generator and let $h \in G$ be the target element for the DLP. Suppose that $F \subseteq G$ is a factor base of cardinality $|F| = m$. We choose $a, b \in \mathbb{Z}/N\mathbb{Z}$ uniformly and independently at random and try to express $g^a h^b$ as a product $g^a h^b = \prod_{f \in F} f^{e_f}$. Once more than m such expressions have been found, we consider the matrix consisting of the collected rows $(e_f)_{f \in F}$ over $\mathbb{Z}/N\mathbb{Z}$, and compute using invertible row transformations a row echelon form, which contains a vanishing row. Contrary to the classical index calculus method we do not require a rank condition for this matrix. Applying the invertible row transformations also to the numbers a and b , then considering the vanishing row we obtain an identity $g^{a'} h^{b'} = 1_G$. One can show that $b' \in \mathbb{Z}/N\mathbb{Z}$ is uniformly distributed, so that b' is coprime to N with high enough probability, in which case $\log_g h = \frac{a'}{b'} \in \mathbb{Z}/N\mathbb{Z}$ has been found.

Instead of computing a row echelon form by a variant of the Gauß algorithm one may use sparse linear algebra techniques, which have an improved running time, however their analysis is more difficult and the above algorithm has to be modified [20]. In particular, it is then necessary to fulfil rank conditions for the generated matrix following a technique from Pomerance [58].

3.2. BASIC CONCRETE VERSIONS. The class of groups to which the index calculus method is applicable includes the multiplicative groups of prime fields and fields of small fixed characteristic. We describe for these cases a simple index calculus method and provide a running time analysis, which also serves as a basis for the more advanced index calculus algorithms.

Suppose that G is \mathbb{F}_p^\times , the multiplicative group of a prime field \mathbb{F}_p , of order $N = p - 1$, with a given generator $g \in G$. As factor base we choose

$$F = \{f \mid f \leq B, f \text{ prime}\} \subseteq G$$

for some bound B (by slight abuse of notation, for $f \in \mathbb{Z}$ we denote the class $[f] \in \mathbb{F}_p$ also by f). For simplicity, we assume that $g \in F$ (otherwise, we include it into the factor base). To generate relations, for random $e \in \mathbb{Z}/N\mathbb{Z}$ we compute $g^e \in \mathbb{F}_p$, lift it to an element in \mathbb{N} , and check by trial division whether it has only prime divisors $\leq B$. If successful, we obtain a relation $g^e \equiv \prod_{f \in F} f^{e_f} \pmod p$ in G . Once enough relations (more than $|F|$) have been found, we compute $\log_g f$ for all $f \in F$ by solving a linear system over $\mathbb{Z}/N\mathbb{Z}$. Finally, given a target element $h \in G$ we similarly obtain one more relation of the form $hg^e = \prod_{f \in F} f^{e_f}$ to obtain $\log_g h$.

Considering a finite field $\mathbb{F}_q = \mathbb{F}_{p^n}$ of fixed characteristic p , note that \mathbb{F}_{p^n} is usually represented as a quotient ring $\mathbb{F}_p[X]/(I)$, where $I \in \mathbb{F}_p[X]$ is an irreducible polynomial of degree n . For $G = \mathbb{F}_q^\times$, it is then straightforward to adapt the basic index calculus method for \mathbb{F}_p^\times described above to the present situation. In particular, as factor base we choose all irreducible polynomials in $\mathbb{F}_p[X]$ of some bounded degree b , i.e.,

$$F = \{f \mid f \in \mathbb{F}_p[X], \deg f \leq b, f \text{ irreducible}\} \subseteq G$$

(where we employ a similar abuse of notation). It suffices in practice to include only the monic polynomials into the factor base. In fact, one may perform the discrete logarithm computation in $\mathbb{F}_q^\times / \mathbb{F}_p^\times$, i.e., ignoring constants in \mathbb{F}_p^\times , to obtain $\log_g h$ modulo $\frac{N}{p-1}$. Using the Pohlig-Hellman method with the fact that $p - 1$ divides the product of the small prime power divisors of the group order $N = p^n - 1$, the remaining information of $\log_g h$ is deduced easily.

3.3. COMPLEXITY ANALYSIS. A positive integer is called *B-smooth* if all its prime divisors are $\leq B$. The (heuristic) running time analysis for the basic index calculus method in \mathbb{F}_p^\times , as well as for the more advanced algorithms presented in Section 4, is based on the following result on the asymptotic density of smooth numbers among the integers.

Theorem 3.2 (Canfield, Erdős, Pomerance [14]). *A uniformly random integer in $\{1, \dots, M\}$ is B-smooth with probability*

$$P = u^{-u(1+o(1))}, \quad \text{where } u = \frac{\log M}{\log B},$$

provided that $3 \leq u \leq (1 - \varepsilon) \frac{\log M}{\log \log M}$ for some $\varepsilon > 0$.

We recall the notation $L_N(\alpha, c) = \exp((c + o(1))(\log N)^\alpha (\log \log N)^{1-\alpha})$, where $\alpha \in [0, 1]$, $c > 0$, which is originally due to Pomerance [57] for $\alpha = \frac{1}{2}$ and Lenstra and Lenstra [46] for general α .

Corollary 3.3. *Let $M = L_N(2\alpha, \mu)$ and $B = L_N(\alpha, \beta)$, then the expected number of trials until a uniformly random number in $\{1, \dots, M\}$ is B-smooth is $L_N(\alpha, \frac{\alpha\mu}{\beta})$.*

For analysing the basic version of the index calculus method in \mathbb{F}_p^\times , we set the smoothness bound $B = L(\frac{1}{2}, \beta)$ and we have $M = N = L(1, 1)$. As we need about $|F| \approx B/\log B \leq B$ relations, our estimated running time equals

$$L(\frac{1}{2}, \beta) \cdot L(\frac{1}{2}, \frac{1}{2\beta}) = L(\frac{1}{2}, \beta + \frac{1}{2\beta}),$$

and the optimal choice $\beta = \frac{1}{\sqrt{2}}$ results in a running time of $L(\frac{1}{2}, \sqrt{2})$ for the relation generation. The linear algebra running time (using iterative techniques for sparse matrices) is about $B^2 = L(\frac{1}{2}, 2\beta) = L(\frac{1}{2}, \sqrt{2})$ as well, while the individual logarithm phase is of lower complexity.

Similarly, a polynomial is called *b-smooth* if all its irreducible factors are of degree $\leq b$; hence, the 1-smooth polynomials are precisely those that split into linear factors.

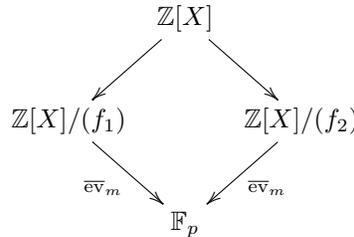
Theorem 3.4 (Odlyzko [53], Lovorn [50]). *A uniformly random polynomial $f \in \mathbb{F}_q[X]$ of degree m is b -smooth with probability $P = u^{-u(1+o(1))}$, where $u = \frac{m}{b}$, provided that $m^{1/100} \leq b \leq m^{99/100}$.*

For the DLP in $G = \mathbb{F}_{p^n}^\times$, where p is fixed, we obtain quite analogously a running time of $L(\frac{1}{2}, \sqrt{2})$.

4. THE NUMBER FIELD SIEVE

The number field sieve is an advanced index calculus method with heuristic $L(1/3)$ -complexity. It was originally devised for the integer factorisation problem, but the method can be adapted to apply for the DLP in prime fields [25, 61] and more generally fields of large or medium characteristic [62, 39]. The principle difference between these algorithms and the basic version of the previous section is that the two sides of each relation are of a considerably smaller “size” than before, so that the smoothness probability is – at least heuristically – greatly increased.

The setup for computing discrete logarithms in \mathbb{F}_p^\times for a prime p is as follows. Let $f_1, f_2 \in \mathbb{Z}[X]$ be coprime irreducible polynomials and $m \in \mathbb{Z}$ such that $f_1(m) \equiv f_2(m) \equiv 0 \pmod p$. To simplify the description we assume that f_1 and f_2 are monic, although this requirement is not essential. We have the following commutative diagram:



Here, $\bar{e}v_m : \mathbb{Z}[X]/(f_i) \rightarrow \mathbb{F}_p$ denotes the map $[h] \mapsto [h(m)]$, for $i = 1, 2$. For applying the index calculus method we need a way to factor elements $[h]$ in the ring $\mathbb{Z}[X]/(f_i)$ over a smoothness base, which is a more intricate issue, requiring some concepts from algebraic number theory (see, e.g., [52]).

For $i = 1, 2$ let d_i be the degree and let x_i be a root of the polynomial f_i . Then $K_i = \mathbb{Q}(x_i)$ is a number field, i.e., a finite extension of \mathbb{Q} , of degree d_i , and its associated ring of integers \mathcal{O}_i is a Dedekind domain, thus every nonzero ideal factors uniquely into a product of prime ideals. We have $\mathbb{Z}[X]/(f_i) \cong \mathbb{Z}[x_i] \subseteq \mathcal{O}_i$ and the quotient group $\mathcal{O}_i/\mathbb{Z}[x_i]$ is finite.

In order to generate relations we make use of the multiplicative norm map $N : K_i \rightarrow \mathbb{Q}$, which satisfies $N(\mathcal{O}_i) \subseteq \mathbb{Z}$. Associating to each nonzero ideal $\mathfrak{a} \subseteq \mathcal{O}_i$ its norm $N(\mathfrak{a}) = \#\mathcal{O}_i/\mathfrak{a}$, we have $|N(\alpha)| = N((\alpha))$ for any nonzero element $\alpha \in \mathcal{O}_i$, and if $(\alpha) = \prod \mathfrak{p}^{e_p}$ is the decomposition of the principal ideal (α) into prime ideals, then $|N(\alpha)| = \prod N(\mathfrak{p})^{e_p}$. Therefore, the prime ideal decomposition of (α) can be

easily obtained from the prime factorisation of the norm $N(\alpha) \in \mathbb{Z}$, provided that every prime ideal \mathfrak{p} containing α has prime norm; this holds for example in the situation $\alpha = h(x_i)$ and $h = h_1X + h_0$ with coprime integers $h_0, h_1 \in \mathbb{Z}$.

Accordingly, as factor base we choose elements, which are represented by prime ideals in \mathcal{O}_i of small norm, and we look for linear polynomials $h \in \mathbb{Z}[X]$ such that both norms $N(h(x_i)) = h_1^{d_i} f_i(-\frac{h_0}{h_1})$, $i = 1, 2$, are smooth. After collecting sufficiently many relations we obtain the “virtual logarithms” of the prime ideals using linear algebra, cf. [37, 63]. We remark that one has to account for the unit groups of \mathcal{O}_1 and \mathcal{O}_2 , for which one usually includes up to $(d_1 + d_2)$ further factor base elements corresponding to certain computable logarithmic maps devised by Schirokauer [61].

Parameter choices and complexity analysis. Classically one chooses $d_2 = 1$, so that there is a rational side $K_2 = \mathbb{Q}$ and $\mathcal{O}_2 = \mathbb{Z}$, as well as $x_2 = m \in \mathbb{Z}$. Denote $f = f_1$, $d = d_1$ and $x = x_1$. Letting $h = h_1X + h_0 \in \mathbb{Z}[X]$, where $|h_i| \leq E$, be the sieving polynomial one wants both $N(h(x))$ and $h(m)$ to be B -smooth. To obtain an $L(1/3)$ algorithm one sets $B = L(\frac{1}{3}, \beta)$ and $E = L(\frac{1}{3}, \varepsilon)$, as well as $d = (\delta + o(1))(\frac{\log N}{\log \log N})^{1/3}$ with parameters $\beta, \delta, \varepsilon > 0$ to be determined.

From the Kalkbrenner bound [42] one gets $\log |N(h(x))| \approx \log \|f\|_\infty + d \log E$, where $\|f\|_\infty$ denotes the maximum absolute value of the coefficients of f , while $\log |h(m)| \approx \log m$. One can choose $\|f\|_\infty \leq m$ and $m \approx N^{1/d}$, so that $\log m \approx \frac{1}{d} \log N$, and this implies $m = L(\frac{2}{3}, \frac{1}{\delta})$. Taking into account $d \log E = \log L(\frac{2}{3}, \delta\varepsilon)$ one thus obtains

$$|N(h(x))h(m)| = L(\frac{2}{3}, \delta\varepsilon + \frac{2}{\delta}).$$

Under the heuristic assumption that the quantity $|N(h(x))h(m)|$ is uniformly distributed for a random polynomial h , we get from Corollary 3.3 for the probability P of this quantity being B -smooth that $\frac{1}{P} = L(\frac{1}{3}, \frac{\delta^2\varepsilon+2}{3\beta\delta})$.

The sieving space size $\approx E^2$ should be equal to the linear algebra complexity $\approx B^2$, therefore $\beta = \varepsilon$, and as about B relations are needed one sets $\frac{1}{P} \approx B$. Therefore, we obtain the condition $\frac{\delta^2\beta+2}{3\beta\delta} = \beta$, or $\delta^2\beta+2 = 3\beta^2\delta$. The optimal choice $\delta^2 = 2/\beta$ then yields $\beta = (\frac{8}{9})^{1/3}$, resulting in a complexity of $L(\frac{1}{3}, (\frac{64}{9})^{1/3} \approx 1.923)$ for the first two phases of the index calculus method.

For prime fields in which the prime is of particular shape the so-called special number field sieve is available [65]. In this case one may have small coefficients of f , namely $\log \|f\|_\infty = o(\log m)$, which leads to a faster algorithm. Indeed, one gets $\delta^2\beta + 1 = 3\beta^2\delta$, and for $\delta^2 = 1/\beta$ one obtains $\beta = (\frac{4}{9})^{1/3}$, resulting in a complexity of $L(\frac{1}{3}, (\frac{32}{9})^{1/3} \approx 1.526)$.

Note that these two complexities are the same as for factoring using the number field sieve and the special number field sieve, respectively.

Individual logarithms. For obtaining an individual logarithm $\log_g h$ the previous approach, i.e., multiplying h by random powers of g until it factors over the factor base, is not viable as the factor base is now much smaller. Instead one uses a recursive strategy commonly referred to as the *descent*. At each step of the descent one *eliminates* a given element which is not in the factor base, i.e., rewrites it as a product of elements of smaller norm. Starting with h , eliminations are carried out recursively until only factor base elements remain. In this way a tree is constructed with h as root and with factor base elements as leaves, from which the logarithm $\log_g h$ can easily be deduced from the factor base logarithms by traversing the tree.

We omit the technical details for this section and simply remark that this phase can be shown to have negligible (heuristic) complexity compared to the other steps (see, e.g., [15]).

4.1. THE MEDIUM NUMBER FIELD SIEVE. This variation of the number field sieve applies to the DLP in finite fields of large or medium characteristic [39].

The setup is as follows. Choose irreducible polynomials $f_1, f_2 \in \mathbb{Z}[X]$ such that $f_1 \bmod p$ and $f_2 \bmod p$ have a common irreducible factor $I \in \mathbb{F}_p[X]$ of degree n . (A simple choice is to let $f_2 = f_1 + p$, if $f_1 \bmod p$ contains an irreducible degree n factor; more advanced selection methods are sketched below.) We have thus the following commutative diagram:

$$\begin{array}{ccc} & \mathbb{Z}[X] & \\ & \swarrow & \searrow \\ \mathbb{Z}[X]/(f_1) & & \mathbb{Z}[X]/(f_2) \\ & \searrow & \swarrow \\ & \mathbb{F}_{p^n} \cong \mathbb{F}_p[X]/(I) & \end{array}$$

As in the number field sieve for prime fields one obtains relations by finding polynomials $h \in \mathbb{Z}[X]$ (of some degree $\leq t$ and $\|h\|_\infty \leq E$) such that both norms $N(h(x_i)) = \text{Res}(h, f_i)$ are B -smooth, where x_i is a root of f_i and Res denotes the resultant. The Kalkbrener bound [42] implies here that

$$\log |N(h(x_1))N(h(x_2))| \approx t(\log \|f_1\|_\infty + \log \|f_2\|_\infty) + (d_1 + d_2) \log E,$$

where d_i is the degree of f_i (if d_i is small enough). Therefore, the running time will crucially depend on the degree and the coefficient size of the selected polynomials f_i in the setup phase of the algorithm. The recent work [8] achieves improvements by a clever polynomial selection.

Indeed, for large characteristic one first chooses a polynomial $f_1 \in \mathbb{Z}[X]$ of degree $d+1$ with $\|f_1\|_\infty$ small such that $f_1 \bmod p$ has an irreducible factor I of degree n . Then a polynomial $f_2 \in \mathbb{Z}[X]$ of degree d is chosen such that $I \mid f_2 \bmod p$ and $\|f_2\|_\infty$ as small as possible; this can be achieved by lattice reduction [48], resulting in the estimation $\log \|f_2\|_\infty \approx \frac{n}{d+1} \log p$, while one has $d_1 = d+1$ and $d_2 = d \geq n$. With suitably chosen parameters this results in a running time of $L(\frac{1}{3}, (\frac{64}{9})^{1/3})$, the same as for the prime field case.

For medium characteristic the so-called Conjugation Method improves upon the original selection method. Here, one lets $\mu = Y^2 + aY + b \in \mathbb{Z}[Y]$ be an irreducible polynomial with small coefficients such that $\mu \bmod p$ has a root $\lambda \in \mathbb{F}_p$. Then one chooses $g_0, g_1 \in \mathbb{Z}[X]$ with $\|g_i\|_\infty$ small and $\deg g_0 < n = \deg g_1$. With this let $I = \lambda g_0 + g_1$, $f_2 = u g_0 + v g_1$, where $\lambda = \frac{u}{v} \bmod p$ with $\log \|f_2\|_\infty \approx \frac{1}{2} \log p$, as well as $f_1 = \text{Res}_Y(\mu, Y g_0 - g_1) = g_1^2 + a g_1 g_0 + b g_0^2$, so that $\|f_1\|_\infty$ is small, while $d_1 = 2n$ and $d_2 = n$. The complexity analysis results in a running time of $L(\frac{1}{3}, (\frac{96}{9})^{1/3} \approx 2.201)$.

The tower number field sieve. An alternative method to deal with the DLP in extension fields stems from work of Schirokauer [62], in which an algorithm for finite fields \mathbb{F}_{p^n} of fixed extension degree n is presented with a heuristic running time of $L(\frac{1}{3}, (\frac{64}{9})^{1/3})$. It uses a field extension on top of a number field F , for which

its ring of integers \mathcal{O}_F satisfies $\mathcal{O}_F/(p) \cong \mathbb{F}_{p^n}$. The setup is then very similar to the number field sieve for prime fields, with \mathbb{Z} replaced by the ring \mathcal{O}_F .

The algorithm has been analysed and extended in a recent work [10], in which the authors coined the term tower number field sieve. In particular, it has been shown that it applies to the whole range of large characteristic fields, achieving the same complexity $L(\frac{1}{3}, (\frac{64}{9})^{1/3})$ as the medium number field sieve of [39], while being practically advantageous for certain cases.

Finally, by combining the tower number field sieve with the medium number field sieve it has been shown lately [11, 60, 43] that for medium characteristic finite fields of the form $\mathbb{F}_{p^{k\ell}}$ with (coprime) $k, \ell > 1$ the DLP has complexity $L(\frac{1}{3}, (\frac{48}{9})^{1/3})$, thus improving upon the conjugation method of [8].

5. THE FUNCTION FIELD SIEVE – CLASSICAL VARIANT

The first (heuristic) $L(1/3)$ algorithm for the DLP in finite fields was devised by Coppersmith in 1984 [16] – thus predating the number field sieve – and as originally described applies to binary finite fields, i.e., finite fields of characteristic 2. As a result of this algorithm, binary field DLPs fell out of favour in cryptography; only much later with the advent of pairing-based cryptography did they become fashionable once again, when the lower security-per-bit they offer was considered tolerable thanks to the functionality afforded by pairings on supersingular curves [21].

The function field sieve was conceived of by Adleman [5] in analogy with the number field sieve, and later refined by Adleman and Huang [6], but can also be seen as a generalisation of Coppersmith’s algorithm. Joux and Lercier later provided a more streamlined version [36], as well as an improved version for medium-size base fields [38]. In this section we briefly describe these algorithms in turn, as not only are they historically relevant to the DLP in small characteristic fields, but they also serve as the benchmark against which the recent developments may be compared.

Before describing the algorithms we remark that it is possible to map between any two representations of a finite extension field efficiently [49]; therefore when solving a DLP one is free to choose a field-defining polynomial to one’s advantage.

5.1. COPPERSMITH’S ALGORITHM. Coppersmith’s algorithm [16] is an index calculus algorithm for solving the DLP in \mathbb{F}_{2^n} , represented as $\mathbb{F}_2[X]/(I)$, where $I \in \mathbb{F}_2[X]$ is an irreducible degree n polynomial of the form $I = X^n + J$, where $J \in \mathbb{F}_2[X]$ is of relatively low degree (less than $n^{2/3}$). Let the factor base consist of the irreducible polynomials up to a degree bound $\leq b$, and choose positive integers h and ℓ such that $h2^\ell \geq n$. For relation generation, consider $f = uX^h + v \in \mathbb{F}_2[X]$, with $u, v \in \mathbb{F}_2[X]$ coprime polynomials of degree $\leq d$, where d is a sieving parameter, and compute

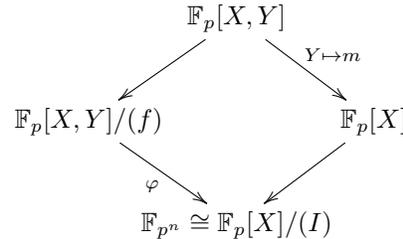
$$f^{2^\ell} = (uX^h + v)^{2^\ell} = u^{2^\ell} X^{h2^\ell} + v^{2^\ell} \equiv u^{2^\ell} X^{h2^\ell - n} J + v^{2^\ell} = g \pmod{I}.$$

A relation is found if the polynomials $f, g \in \mathbb{F}_2[X]$ on both sides are b -smooth. Note that the corresponding degrees, namely $\deg f \leq h + d$ and $\deg g \leq r + 2^\ell d$ with $r = \deg X^{h2^\ell - n} J$, can be made rather small by suitably chosen parameters. Indeed, we let $d = (c + o(1))n^{1/3}(\log n)^{2/3}$ and suppose that $2^\ell \approx \sqrt{\frac{n}{d}}$, as well as $h = \lceil \frac{n}{2^\ell} \rceil$. Then $\deg f$ and $\deg g$ are about $\sqrt{nd} = (\sqrt{c} + o(1))n^{2/3}(\log n)^{1/3}$. Choosing $b = (c + o(1))n^{1/3}(\log n)^{2/3}$, by applying an analogue of Corollary 3.3, we get for the probability P of both polynomials being b -smooth that $\frac{1}{P} = L(\frac{1}{3}, \frac{2}{3\sqrt{c}})$. In order to generate enough relations we set $\frac{2}{3\sqrt{c}} = c$, so that $c = (\frac{4}{9})^{1/3}$, resulting

in an overall complexity of $L(\frac{1}{3}, (\frac{32}{9})^{1/3})$ for relation generation. This matches the linear algebra complexity using sparse matrix techniques, while the individual logarithm phase was shown to have lower complexity.

This analysis supposes that $\sqrt{\frac{n}{d}}$ is close to some power 2^ℓ , which cannot be fulfilled for all n . In general one obtains a complexity between $L(\frac{1}{3}, (\frac{32}{9})^{1/3})$ and $L(\frac{1}{3}, 4^{1/3})$, with $(\frac{32}{9})^{1/3} \approx 1.526$ and $4^{1/3} \approx 1.587$. Observe that the algorithm exploits the basic identity $(u+v)^2 = u^2+v^2$, which holds for any polynomials $u, v \in \mathbb{F}_2[X]$, and can thus be adapted easily to the case of fields of fixed characteristic p , by using the identity $(u+v)^p = u^p+v^p$ for $u, v \in \mathbb{F}_p[X]$, with a correspondingly increased upper bound on the resulting complexity.

5.2. THE FUNCTION FIELD SIEVE. The function field sieve is an adaption of the number field sieve and applies to finite fields \mathbb{F}_{p^n} of small characteristic p . The principle idea is to replace $\mathbb{Z}[X]$ by $\mathbb{F}_p[X, Y]$, and one usually chooses the polynomial f_2 to be of degree 1. More precisely, define \mathbb{F}_{p^n} as $\mathbb{F}_p[X]/(I)$, where the irreducible degree n polynomial $I \in \mathbb{F}_p[X]$ divides $f_d m^d + f_{d-1} m^{d-1} + \dots + f_1 m + f_0 = f(m)$, with an irreducible polynomial $f(Y) = \sum f_i Y^i \in (\mathbb{F}_p[X])[Y]$ and $m \in \mathbb{F}_p[X]$ suitably chosen. Then $f \in \mathbb{F}_p[X, Y]$ defines an algebraic curve C and its associated function field $\mathbb{F}_p(C) = \text{Quot}(\mathbb{F}_p[X, Y]/(f))$. Since $f(m) \equiv 0 \pmod I$ we have the following commutative diagram:

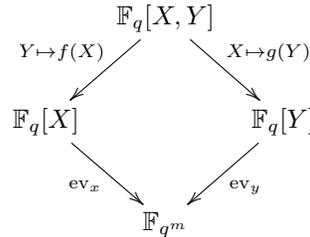


The factor base consists of all irreducible polynomials in $\mathbb{F}_p[X]$ up to degree b . In order to generate relations we search for polynomials $h = h_1 Y + h_0 \in \mathbb{F}_p[X, Y]$ such that both $h_1 m + h_0$ and the norm of h from $\mathbb{F}_p(C)$ to $\mathbb{F}_p(X)$ are b -smooth. In this case, we can express the divisor of $h_1 Y + h_0$ as a sum of places in $\mathbb{F}_p(C)$ over factor base elements, and relate the corresponding so-called surrogates via the map φ with the factors of $h_1 m + h_0$. Then we solve as usual the logarithms of the factor base elements by linear algebra.

As already mentioned there are three variations of the function field sieve, which differ by the choice of the polynomials I, f and m . In Adleman’s original version [5] the field-defining polynomial is written as $I = m^d + f_{d-1} m^{d-1} + \dots + f_1 m + f_0$ using the base- m technique prevalent in the number field sieve. The refined version by Adleman and Huang [6] improves upon this approach by choosing $I = X^n + J$ with $J \in \mathbb{F}_p[X]$ of small degree $\deg J < n^{2/3}$, $f = Y^d + X^{ed-n} J$ and $m = X^e$, which results in a complexity of $L(\frac{1}{3}, (\frac{32}{9})^{1/3})$, the same as for the special number field sieve. It is worth noting that Coppersmith’s algorithm can be seen as a particular case of this variation. Finally, the function field sieve version of Joux and Lercier [36] achieves the same complexity, but has some practical advantages. The idea is to start by choosing f with coefficients of low degree in X and then letting $I = m_2^d f(\frac{m_1}{m_2})$ be the field-defining polynomial, where $m = \frac{m_1}{m_2}$ is now a rational function. In all cases the individual logarithm phase can be shown to have a lower complexity than the other two phases.

5.3. THE MEDIUM FUNCTION FIELD SIEVE. Joux and Lercier in 2006 proposed the following simplified variant of the function field sieve, which employs just the rational function field of a univariate polynomial ring [38]. The algorithm applies to the whole range of finite fields \mathbb{F}_{p^n} of medium-small characteristic, i.e., $p = L_{p^n}(\alpha)$, where $\alpha \leq \frac{1}{3}$. It also applies to extension fields \mathbb{F}_{q^m} , where q is any prime power.

The representation of the field \mathbb{F}_{q^m} is as follows. Let $f, g \in \mathbb{F}_q[X]$ be polynomials such that $g(f(X)) - X$ has an irreducible factor $I \in \mathbb{F}_q[X]$ of degree m . Let x be a root of I in \mathbb{F}_{q^m} and let $y = f(x)$, hence $x = g(y)$. Again we have the following commutative diagram:



Now if $q = L_{q^m}(1/3)$ then for $a, b, c \in \mathbb{F}_q$ we consider $h = XY + aY + bX + c \in \mathbb{F}_q[X, Y]$, which leads in \mathbb{F}_{q^m} to the following identity

$$xf(x) + af(x) + bx + c = g(y)y + ay + bg(y) + c.$$

Let the factor base consist of the linear polynomials in X or in Y . If the corresponding polynomials on both sides, namely $h(X, f(X)) = Xf(X) + af(X) + bX + c$ and $h(g(Y), Y) = g(Y)Y + aY + bg(Y) + c$ are 1-smooth, then a relation has been found. We may choose the polynomials f and g such that $\deg f, \deg g \approx \sqrt{m}$, which leads to an algorithm with complexity $L(\frac{1}{3}, 3^{1/3} \approx 1.442)$. Here, the individual logarithm phase turns out to have the same complexity as the main phase, with the bottleneck being the elimination of degree two polynomials, see [38, Sec. 3].

In the general case, where $q = L_{q^m}(\alpha)$ with $\alpha \leq 1/3$, to obtain an $L(1/3)$ algorithm we set as the degree bound for the factor base $(C + o(1))(\frac{\log q}{\log \log q})^{1/3-\alpha}$ and consider polynomials of the form $h = h_1(X)Y + h_0(X)$ in order to generate relations. Note that for $\alpha = 0$ the case $\log q = o(\log \log q^m)$ has to be treated extra with a slightly modified analysis. If $q = L_{q^m}(0)$, i.e., the case of small characteristic, this results in an algorithm of complexity $L(\frac{1}{3}, (\frac{32}{9})^{1/3})$, with a less costly individual logarithm phase.

6. SMALL CHARACTERISTIC QUASI-POLYNOMIAL TIME ALGORITHMS

In this section we present the two recent quasi-polynomial time algorithms for the small characteristic DLP. The first of these was due to Barbulescu, Gaudry, Joux and Thomé [9], while the second was due to Granger, Kleinjung and Zumbrägel [28, 29]. For the purpose of assisting the reader in comparing the key insights and ideas behind the two algorithms, we present these for relation generation and individual logarithms, in turn.

6.1. RESISTING SMOOTHNESS HEURISTICS. Before presenting the aforementioned developments, we start with some general remarks on how to obviate smoothness heuristics.

As detailed in Section 3, a complexity of $L(1/2)$ could be said to be the “natural complexity” of the DLP when elements are generated uniformly in the multiplicative

group. In order to obtain algorithms of complexity better than $L(1/2)$ – at least for the relation generation phase of the index calculus method – there are (at least) two approaches that one can explore. Firstly, one can try to generate relations between elements of smaller norm, which heuristically would have a higher probability of being sufficiently smooth, which is what the $L(1/3)$ algorithms do. Secondly, one can try to generate relations which have better than expected smoothness properties (or possibly a combination of both of these approaches). 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.

However, in 2013 a series of algorithmic breakthroughs occurred which demonstrated that for small characteristic fields the DLP is, at least heuristically, far easier than originally believed. Central to these developments was the fundamental insight that one can produce families of polynomials that are smooth by construction [23, 35] and use these to generate relations, while uniformly generated polynomials of the same degree are smooth with only an exponentially small probability. These results were the first to succeed in applying the second approach described above, and the families of polynomials used lay the foundation for the two quasi-polynomial algorithms.

6.2. POLYNOMIAL TIME RELATION GENERATION. We now explain the polynomial time relation generation methods due to Göloğlu, Granger, McGuire and Zumbrägel [23] and Joux [35], which while different are essentially isomorphic, and were discovered independently at essentially the same time.

The finite fields to which the methods apply are of the form $\mathbb{F}_Q = \mathbb{F}_{q^{kn}}$, with $k \geq 2$ and $n \approx q$. A given small characteristic finite field \mathbb{F}_{p^n} can be embedded into one of this form by letting $q = p^\ell$ with $\ell = \lceil \log_p n \rceil$, thus increasing the extension degree by a factor of $k \lceil \log_p n \rceil$, which does not impact upon the resulting complexity too much, see Section 6.4.

The GGMZ method. The field setup used in [23] can be seen in the context of the Joux-Lercier function field sieve [38], described in Section 5.3, but for which the degrees of the polynomials f and g are extremely unbalanced. In fact, we consider $f = X^q$ and $g = \frac{h_0}{h_1}$ for some $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ of low degree¹, which leads to the following field representation. We define $\mathbb{F}_{q^{kn}}$ as $\mathbb{F}_{q^k}[X]/(I)$, where $I \in \mathbb{F}_{q^k}[X]$ is an irreducible degree n polynomial dividing $h_1(X^q)X - h_0(X^q)$ for $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ of low degree $\leq d_h$. Letting x be a root of I in \mathbb{F}_Q , we have $y = f(x) = x^q$ and $x = g(y) = \frac{h_0(y)}{h_1(y)}$. The factor base consists of $h_1(x^q)$ and all linear polynomials on the x -side; note that the y -side factor base is not needed since for all $d \in \mathbb{F}_{q^k}$ one has $(y + d) = (x + d^{1/q})^q$.

As in Section 5.3, let a, b, c be in the base field \mathbb{F}_{q^k} , and consider elements of $\mathbb{F}_{q^{kn}}$ of the form $xy + ay + bx + c$. Using the field isomorphisms we have the following identity for such elements:

$$(1) \quad x^{q+1} + ax^q + bx + c = \frac{1}{h_1(y)}(yh_0(y) + ayh_1(y) + bh_0(y) + ch_1(y)).$$

An extremely useful observation is that the l.h.s. of Eq. (1) splits completely over \mathbb{F}_{q^k} with probability $\approx 1/q^3$, which is exponentially higher than the splitting probability

¹In [23] h_1 was not specified (and was thus implicitly 1); in [27] h_1 was introduced in order to exploit other more convenient field representations.

of a uniformly random polynomial of the same degree, which is $\approx 1/(q+1)!$. Indeed, let $k \geq 3$ and consider the polynomial $X^{q+1} + aX^q + bX + c$. If $ab \neq c$ and $a^q \neq b$, this may be transformed (up to a scalar) into

$$F_B(\bar{X}) = \bar{X}^{q+1} + B\bar{X} + B, \quad \text{with } B = \frac{(b - a^q)^{q+1}}{(c - ab)^q},$$

via $X = \frac{c-ab}{b-a^q}\bar{X} - a$. Clearly, the original polynomial splits whenever F_B splits and we have a valid transformation from \bar{X} to X . The following theorem of Blüher counts the number of $B \in \mathbb{F}_{q^k}$ for which F_B splits completely over \mathbb{F}_{q^k} .

Theorem 6.1. (Blüher [12]) *The number of elements $B \in \mathbb{F}_{q^k}^\times$ s.t. the polynomial $F_B(\bar{X}) \in \mathbb{F}_{q^k}[\bar{X}]$ splits completely over \mathbb{F}_{q^k} equals*

$$\frac{q^{k-1} - 1}{q^2 - 1} \quad \text{if } k \text{ is odd,} \quad \frac{q^{k-1} - q}{q^2 - 1} \quad \text{if } k \text{ is even.}$$

Using the expression for B as a function of a, b, c one can easily generate triples (a, b, c) for which the l.h.s. *always* splits. In particular, one first computes the set \mathcal{B} of all $B \in \mathbb{F}_{q^k}$ for which F_B splits over \mathbb{F}_{q^k} . Then for any $a, b \neq a^q$ and $B \in \mathcal{B}$ there is a unique c for which the l.h.s. splits. Furthermore, by Theorem 6.1 there are $\approx q^{3k-3}$ such triples, and for each one whenever the r.h.s. splits as well, one has a relation. For $k = 2$ there are no such F_B , however in this case the set of triples for which the l.h.s. splits non-trivially is easily shown to be

$$\{(a, a^q, c) \mid a \in \mathbb{F}_{q^2} \text{ and } c \in \mathbb{F}_q, c \neq a^{q+1}\}.$$

In all cases, assuming the r.h.s. splits with probability $1/(d_h + 1)!$, in order for there to be sufficiently many relations one requires that $q^{2k-3} > (d_h + 1)!$. For fixed d_h and $q \rightarrow \infty$ the cost of computing the logarithms of all the factor base elements is heuristically $O(q^{2k+1})$ (operations in $\mathbb{Z}/(Q-1)\mathbb{Z}$), using sparse (weight q) linear algebra techniques, which is polynomial in $\log Q = q^{1+o(1)}$ for fixed k .

Joux's method. The method of Joux [35] works for fields of the same shape as those in the GGMZ method (although only $k = 2$ was used for the exposition and initial examples), however the field representation is slightly different. In particular, let $\mathbb{F}_Q = \mathbb{F}_{q^k}(x) = \mathbb{F}_{q^k}[X]/(I)$, where $I \mid h_1(X)X^q - h_0(X)$ for some $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ of low degree $\leq d_h$, and hence $x^q = \frac{h_0(x)}{h_1(x)}$. Fixing some notation, for any $a \in \mathbb{F}_{q^k}[X]$ we have $a(X)^q = \tilde{a}(X^q)$ with $\deg \tilde{a} = \deg a$, where the coefficients of \tilde{a} are the q -th powers of the coefficients of a .

As already mentioned, Joux's method is essentially isomorphic to the GGMZ method and starts with the identity

$$\prod_{\mu \in \mathbb{F}_q} (X - \mu) = X^q - X.$$

Substituting X by $\frac{\alpha X + \beta}{\gamma X + \delta}$ with $\alpha, \beta, \gamma, \delta \in \mathbb{F}_{q^k}$ and $\alpha\delta - \beta\gamma \neq 0$, and multiplying by $(\gamma X + \delta)^{q+1}$ one obtains

$$(2) \quad (\gamma X + \delta) \prod_{\mu \in \mathbb{F}_q} ((\alpha X + \beta) - \mu(\gamma X + \delta)) = (\alpha X + \beta)^q (\gamma X + \delta) - (\alpha X + \beta)(\gamma X + \delta)^q.$$

Observe that the r.h.s. of Eq. (2) has (up to a scalar factor) the same form as the l.h.s. of Eq. (1), and automatically splits over \mathbb{F}_{q^k} by virtue of the l.h.s. of Eq. (2). Using the field equation $x^q = \frac{h_0(x)}{h_1(x)}$, the r.h.s. of Eq. (2) becomes

$$\frac{1}{h_1(x)} (\tilde{\alpha}h_0(x) + \tilde{\beta}h_1(x))(\gamma x + \delta) - (\alpha x + \beta)(\tilde{\gamma}h_0(x) + \tilde{\delta}h_1(x)),$$

and if it also splits over \mathbb{F}_{q^k} then one has a relation between the linear elements and $h_1(x)$. As for the number of different relations one can obtain in this way, observe that the total number of $(\alpha, \beta, \gamma, \delta)$ -transformations is just $|\mathrm{PGL}_2(\mathbb{F}_{q^k})| = q^{3k} - q^k$. However, two transformations will give the same relation (up to multiplication by a scalar in $\mathbb{F}_{q^k}^\times$) if there exists an element of $\mathrm{PGL}_2(\mathbb{F}_q)$ which gives the second transformation when multiplied by the first. Hence the total number of different transformations is $\approx q^{3k-3}$, just as for the GGMZ method. One should therefore compute a set of coset representatives for the quotient $\mathrm{PGL}_2(\mathbb{F}_{q^k})/\mathrm{PGL}_2(\mathbb{F}_q)$ to avoid repetitions; the GGMZ method by contrast already does this implicitly.

6.3. DEGREE TWO ELIMINATION. The methods of the previous subsection demonstrate that the factor base can be extremely small and the logarithms of its elements computed efficiently. Therefore, our attention now falls on the individual logarithm phase. For the $L(1/2)$ and $L(1/3)$ algorithms this phase had those respective complexities, although usually with smaller constants. In contrast, for the new algorithms it becomes the dominant phase.

If one employs the usual descent method for the individual logarithm phase for our setup, then as in [38] the degree two eliminations are the bottleneck. However, the work of GGMZ and Joux shows how to eliminate degree two elements efficiently, and the respective methods were developed into the building blocks of the quasi-polynomial descent algorithms. Thus, due to their importance, we now present them.

The GGMZ degree two elimination method. As before consider $\mathbb{F}_{q^{kn}} = \mathbb{F}_{q^k}(x)$ with $y = x^q$ and $x = \frac{h_0(y)}{h_1(y)}$. Let $P(x)$ be a degree two irreducible element to be eliminated, i.e., written as a product of linear elements, and let $\tilde{P}(y) = P(x)^q$. For $a, b, c \in \mathbb{F}_{q^k}$ consider Eq. (1) from the setup for relation generation. Imposing the condition that $\tilde{P}(y)$ divides the r.h.s. of Eq. (1), we get an \mathbb{F}_{q^k} -linear system in (a, b, c) , whose solution can in general be expressed in terms of a . In particular, there exists $u_0, u_1, v_0, v_1 \in \mathbb{F}_{q^k}$ dependent on $P(x)$ such that $b = u_0 + au_1$ and $c = v_0 + av_1$ and hence

$$x^{q+1} + ax^q + bx + c = \frac{1}{h_1(y)}(h_0y + ah_1y + (u_0 + au_1)h_0 + (v_0 + av_1)h_1).$$

The r.h.s. has degree $d_h + 1$, so the cofactor heuristically splits with probability $\approx 1/(d_h - 1)!$, while the l.h.s. heuristically splits with probability $1/q^3$ for randomly chosen a .

To find such a more directly, we again use the set \mathcal{B} of splitting F_B polynomials. In particular, as a function of a we transform the l.h.s. to obtain an F_B , so that

$$(3) \quad B = \frac{(-a^q + u_0 + au_1)^{q+1}}{(-u_1a^2 + (v_1 - u_0)a + v_0)^q}.$$

Then for each $B \in \mathcal{B}$ we solve Eq. (3) for a , which one can do by taking the greatest common divisor with $a^{q^k} - a$ for instance, which takes time which is polynomial in $\log Q$.

Joux's degree two elimination method. In [35] Joux used the following extension of the degree one relation generation to compute the logarithms of the degree two elements in batches. In particular, let $u \in \mathbb{F}_{q^k}$ and substitute X in $X^q - X$ by

$$\frac{\alpha(X^2 + uX) + \beta}{\gamma(X^2 + uX) + \delta},$$

with $\alpha, \beta, \gamma, \delta \in \mathbb{F}_{q^k}$ and $\alpha\delta - \beta\gamma \neq 0$, and multiply by $(\gamma(X^2 + uX) + \delta)^{q+1}$. This gives

$$(4) \quad (\gamma(X^2 + uX) + \delta) \prod_{\mu \in \mathbb{F}_q} ((\alpha(X^2 + uX) + \beta) - \mu(\gamma(X^2 + uX) + \delta)) \\ = (\alpha(X^2 + uX) + \beta)^q (\gamma(X^2 + uX) + \delta) - (\alpha(X^2 + uX) + \beta)(\gamma(X^2 + uX) + \delta)^q.$$

Observe that in the l.h.s. of Eq. (4) all of the factors are of the form $X^2 + uX + v$ for some $v \in \mathbb{F}_{q^k}$. Replacing all occurrences of X^q by $\frac{h_0(X)}{h_1(X)}$ on the r.h.s. of Eq. (4) gives a low degree expression, which if 1-smooth gives a relation. Once more than $\approx q^k/2$ such relations have been obtained – which is the expected number of irreducibles of this form – one may take logs and solve the resulting system. Doing this for each $u \in \mathbb{F}_{q^k}$ means the logarithms of all degree two elements can be computed in time $O(q^{3k+1})$.

Joux also gave a new elimination method which relies on Gröbner basis computations, whose cost increases with the degree. Balancing the costs of the Gröbner basis descent and the classical descent (whose cost decreases with the degree) results in a heuristic $L(1/4 + o(1))$ algorithm, which was the first algorithm to break the long-standing $L(1/3)$ barrier.

6.4. QUASI-POLYNOMIAL TIME DESCENT. In this subsection we describe the quasi-polynomial descent algorithms of BGJT [9] and GKZ [28, 29].

The BGJT quasi-polynomial algorithm. Joux’s idea for performing degree two elimination in batches can be generalised to polynomials P of any degree, which leads to the quasi-polynomial descent algorithm of Barbulescu *et al.* [9]. As before, let the finite field $\mathbb{F}_Q = \mathbb{F}_{q^{kn}}$ be given as $\mathbb{F}_{q^k}[X]/(I)$, where the degree n irreducible polynomial $I \in \mathbb{F}_{q^k}[X]$ divides $h_1(X)X^q - h_0(X)$ for polynomials $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ of low degree $\leq d_h$. Suppose an element to be eliminated is represented by a polynomial $P \in \mathbb{F}_{q^k}[X]$ of some degree $D < n$. The goal of each elimination step is to find an expression for P in terms of elements of degree less than $D/2$. Substituting X by P in Eq. (2) one obtains

$$(\gamma P + \delta) \prod_{\mu \in \mathbb{F}_q} ((\alpha P + \beta) - \mu(\gamma P + \delta)) = (\alpha P + \beta)^q (\gamma P + \delta) - (\alpha P + \beta)(\gamma P + \delta)^q,$$

for any $\alpha, \beta, \gamma, \delta \in \mathbb{F}_{q^k}$ with $\alpha\delta - \beta\gamma \neq 0$. Observe that all factors in the l.h.s. are (up to a scalar) of the form $P + v$ for some $v \in \mathbb{F}_{q^k}$. Similarly to before, replacing all occurrences of X^q by $\frac{h_0(X)}{h_1(X)}$ on the r.h.s. gives an expression of relatively low degree $(d_h + 1)D$, which if $D/2$ -smooth is viewed as a relation. In order to get a degree $D/2$ -smooth expression for P not involving its \mathbb{F}_{q^k} -translates, one collects more than $\approx q^k$ relations of the above form, so that the associated linear system (in the logarithms of the factors $P + v$, for $v \in \mathbb{F}_{q^k}$) is expected to have full rank, and performs a linear algebra elimination to find an expression for P as a $D/2$ -smooth product, as desired.

By this process the polynomial P is rewritten modulo I as a product $\prod_i R_i^{e_i}$ of polynomials of degree $\leq D/2$, in time polynomial in q^k and n , where the number of factors R_i is in $O(nq^k)$. Iterating these elimination steps down to the factor base gives for fixed d_h and k a heuristic descent running time of

$$(5) \quad q^{O(\log n)} = \exp(O(\log q \log n)),$$

which is, taking into account the polynomial time relation generation, also the heuristic total complexity for solving the DLP in \mathbb{F}_Q .

In fact, the running time of the algorithm can be improved slightly by replacing the degree bound $D/2$ above by $O(D \log \log q / \log q)$, leading to the heuristic complexity $q^{O(\log n / \log \log q)}$. While this quasi-polynomial algorithm is asymptotically the fastest, it has however not yet been used in record computations, see Section 6.5.

Impact on small and medium-small characteristic. Having a quasi-polynomial time algorithm for finite fields of the form $\mathbb{F}_Q = \mathbb{F}_{q^{kn}}$ affects also the DLP in general finite fields of small or medium-small characteristic. Indeed, if a DLP in a finite field \mathbb{F}_{p^n} is considered, we embed this field in \mathbb{F}_Q , where $Q = q^{kn} = (p^n)^{k\ell}$ and $q = p^\ell$.

In the case of small characteristic, i.e., $p = L_{p^n}(0)$, we have $\log p = O(\log n)$ and $\log \log p^n = \log n + \log \log p \approx \log n$. We let $q = p^\ell$ such that $q \geq n - d_h$ and $\log q = O(\log n)$, hence Eq. (5) implies a running time of $\exp(O(\log \log p^n)^2)$, which is quasi-polynomial in the logarithm $\log p^n$ of the group order.

If the characteristic is medium-small, i.e., $p = L_{p^n}(\alpha)$, where $0 < \alpha < 1/3$, we let $q = p$, so that $\log q = \log p = O((\log p^n)^\alpha (\log \log p^n)^{1-\alpha})$. By Eq. (5) and observing that $\log n = O(\log \log p^n)$ we get a complexity of $L_{p^n}(\alpha)^{O(\log \log p^n)} = \exp(O((\log p^n)^\alpha (\log \log p^n)^{2-\alpha})) = L_{p^n}(\alpha + o(1))$, which improves on the function field sieve having $L(1/3)$ -complexity.

Eliminating smoothness heuristics: The GKZ quasi-polynomial algorithm. The prototype algorithm [28] and its rigorous analysis [29] uses a novel descent strategy, which was developed from the GGMZ degree two elimination technique and is thus independent of Joux's degree two elimination and the BGJT algorithm. The algorithm has quasi-polynomial complexity for the DLP in small characteristic (as well the same complexity as the BGJT algorithm for $q = L_Q(\alpha)$), and it features a rigorous complexity analysis, which is based on the absolute irreducibility of a certain algebraic curve, once an appropriate field representation is found. A remarkable property of the descent method is that it does not require any smoothness assumptions about non-uniformly distributed polynomials, in contrast to all previous index calculus algorithms. The new descent method also has practical advantages and has been used in record computations, see Section 6.5.

Consider in $\mathbb{F}_Q = \mathbb{F}_{q^{kn}}$ an element to be eliminated which is represented by an irreducible polynomial $P \in \mathbb{F}_{q^k}[X]$ of degree $2d$, and observe that over the degree d extension $\mathbb{F}_{q^{kd}}$ this polynomial factors into d irreducible quadratics. Applying the degree two elimination method of GGMZ to any of these quadratics enables one to rewrite the quadratic as a product of linear elements over $\mathbb{F}_{q^{kd}}$. Then one simply applies the norm map down to \mathbb{F}_{q^k} , which takes the linear elements to powers of irreducible elements of degree dividing d and the quadratic element back to the original element P , 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. 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 [68, Thm. 5.1] to ensure that any field element is given by an irreducible of degree a power of two only slightly larger than n .

To prove that the degree two elimination step represented in Eq. (3) always succeeds, the following characterisation of the set of $B \in \mathcal{B}$ is used. In particular, by an elementary extension of [32, Theorem 5] the set \mathcal{B} is the image of $\mathbb{F}_{q^{kd}} \setminus \mathbb{F}_{q^2}$

under the map

$$z \mapsto \frac{(z - z^{q^2})^{q+1}}{(z - z^q)^{q^2+1}}.$$

By this and Eq. (3), in order to eliminate P we need to find $(a, z) \in \mathbb{F}_{q^{kd}} \times (\mathbb{F}_{q^{kd}} \setminus \mathbb{F}_{q^2})$ satisfying

$$(z - z^{q^2})^{q+1}(-u_1 a^2 + (v_1 - u_0)a + v_0)^q - (z - z^q)^{q^2+1}(-a^q + u_0 + au_1)^{q+1} = 0.$$

That there are sufficiently many appropriate (a, z) pairs was proven using the classification of subgroups of $\text{PGL}_2(\mathbb{F}_{q^{kd}})$ and the Weil bound for absolutely irreducible curves. One also needs to deal carefully with so-called *descent traps*, i.e., elements that divide $h_1(X)X^{q^{kd+1}} - h_0(X)$ for $d \geq 0$ which can be shown to be ineliminable in the above manner. These considerations led to the following theorem.

Theorem 6.2. *Given a prime power $q > 61$ that is not a power of 4, an integer $k \geq 18$, coprime polynomials $h_0, h_1 \in \mathbb{F}_{q^k}[X]$ of degree at most two and an irreducible degree n factor I of $h_1 X^q - h_0$, the DLP in $\mathbb{F}_{q^{kn}} \cong \mathbb{F}_{q^k}[X]/(I)$ can be solved in expected time*

$$q^{\log_2 n + O(k)}.$$

That the degree of h_0, h_1 is at most two is essential to eliminating smoothness heuristics, since this ensures that the cofactor of the r.h.s. of Eq. (1) has degree at most one, and is thus automatically 1-smooth.

Thanks to Kummer theory, such h_1, h_0 are known to exist when $n = q - 1$, which gives the following easy corollary when $m = ik(p^i - 1)$.

Theorem 6.3. *For every prime p there exist infinitely many explicit extension fields \mathbb{F}_{p^m} in which the DLP can be solved in expected quasi-polynomial time*

$$\exp((1/\log 2 + o(1))(\log m)^2).$$

Proving the existence of h_0, h_1 for general extension degrees as per Theorem 6.2 seems to be a hard problem, even if heuristically it would appear to be almost certain, and in practice it is very easy to find such polynomials.

6.5. PRACTICAL CONSIDERATIONS. While rigorously proving the correctness of a new DLP algorithm is of theoretical interest, a perhaps more immediate measure of its value arises from its practical impact. Furthermore, the challenge of setting new DLP records often leads to theoretical insights that give rise to novel or superior algorithms, which may not have been at all obvious. In addition, such large-scale computations are of cryptologic relevance since they allow one to assess the security of contemporary or future DLP-based cryptosystems. Hence, the value of practical considerations should not be underestimated.

Kummer extensions and automorphisms. Kummer theory provides us with particularly useful polynomials for the field representation, as observed in [34]. In fact, whenever $n \mid q^k - 1$ there exists $c \in \mathbb{F}_{q^k}$ such that $X^n - c \in \mathbb{F}_{q^k}[X]$ is irreducible (one may, for instance, take c to be a multiplicative generator). In particular, for $n = q - 1$ we have $I = X^{q-1} - c \mid X^q - cX = h_1 X^q - h_0$, with $h_1 = 1$, $h_0 = cX$ of degree $\leq d_h = 1$.

Having degree at most one for the polynomials h_0, h_1 in the field representation has some practical advantages for the relation generation and especially for the individual logarithm phase. Furthermore, when defining $\mathbb{F}_{q^{kn}}$ by $\mathbb{F}_{q^k}[X]/(I)$ one can use factor base preserving automorphisms to reduce the complexity of the linear

TABLE 1. Discrete logarithm record computations in finite fields of small or medium characteristic. Details, as well as further announcements, can be found in the number theory mailing list (<https://listserv.nodak.edu/cgi-bin/wa.exe?A0=NMBRTHRY>).

bitlength	charact.	Kummer	who/when	running time
127	2	no	Coppersmith 1984 [16]	$L(1/3, 1.526..1.587)$
401	2	no	Gordon, McCurley 1992 [26]	$L(1/3, 1.526..1.587)$
521	2	no	Joux, Lercier 2001 [36]	$L(1/3, 1.526)$
607	2	no	Thomé 2002	$L(1/3, 1.526..1.587)$
613	2	no	Joux, Lercier 2005	$L(1/3, 1.526)$
556	medium	yes	Joux, Lercier 2006 [38]	$L(1/3, 1.442)$
676	3	no	Hayashi <i>et al.</i> 2010 [31]	$L(1/3, 1.442)$
923	3	no	Hayashi <i>et al.</i> 2012 [30]	$L(1/3, 1.442)$
1175	medium	yes	Joux 24 Dec 2012 [34]	$L(1/3, 1.260)$
1425	medium	yes	Joux 6 Jan 2013 [34]	$L(1/3, 1.260)$
1778	2	yes	Joux 11 Feb 2013 [35]	$L(1/4 + o(1))$
1971	2	yes	GGMZ 19 Feb 2013 [23]	$L(1/3, 0.763)$
4080	2	yes	Joux 22 Mar 2013 [35]	$L(1/4 + o(1))$
809	2	no	CARAMEL 6 Apr 2013 [7]	$L(1/3, 1.526)$
6120	2	yes	GGMZ 11 Apr 2013 [24]	$L(1/4)$
6168	2	yes	Joux 21 May 2013	$L(1/4 + o(1))$
1303	3	no	AMOR 27 Jan 2014 [2]	$L(1/4 + o(1))$
4404	2	no	GKZ 30 Jan 2014 [27]	$L(1/4 + o(1))$
9234	2	yes	GKZ 31 Jan 2014	$L(1/4 + o(1))$
1551	3	no	AMOR 26 Feb 2014 [2]	$L(1/4 + o(1))$
3796	3	no	Joux, Pierrot 15 Sep 2014 [41]	$L(0 + o(1))$
1279	2	no	Kleinjung 17 Oct 2014	$L(0 + o(1))$
4841	3	no	ACCMORR, 18 Jul 2016 [3]	$L(0 + o(1))$

algebra step. In fact, for the q -th power Frobenius one has $(x + a)^q = x^q + a^q = cx + a^q = c(x + \frac{a^q}{c})$. The group generated by the Frobenius automorphism of order kn acts on the factor base, effectively reducing the variables of the linear algebra problem by a factor of about kn . Similar observations hold for so-called twisted Kummer extensions, when $n \mid q + 1$.

New discrete logarithm records. In order to obtain the best performance in practice one typically uses a combination of elimination techniques to develop the descent tree. For currently considered bitlengths these consist of:

- Continued fraction initial split,
- Large degree classical descent,
- quasi-polynomial time descent [28, 29],
- low degree Gröbner Basis descent [35],
- degree two elimination [23, 35].

The new DLP algorithms have lead to a series of records, see Table 1. At the time of writing the largest one was in the field of 2^{9234} elements.

Impact on pairing-based cryptography. The recent advances in solving the DLP in finite fields of small characteristic have had a considerable impact on cryptology research, particularly in the area of identity-based cryptography in which such DLPs are extremely important. This application was investigated by Adj *et al.* [1] and

Granger *et al.* [27]. In particular, in [27] the new methods were improved and extended to make the algorithms more efficient and more widely applicable. At the “128-bit security level” it was shown that a common genus one curve offered only 59 bits of security (which is now much reduced thanks to [28]), while a genus two curve over $\mathbb{F}_{2^{367}}$ with embedding degree 12 was totally broken.

7. SUMMARY AND OUTLOOK

After three decades of relatively little progress in the DLP in small characteristic fields, in the past three years the area has transformed dramatically. Not only are the algorithms much faster, but the ingredients are much simpler too, which raises the question of why these techniques were not discovered earlier? While not historians we proffer two tentative speculations.

Firstly, since Coppersmith’s algorithm was so much faster than the contemporary $L(1/2)$ algorithms, small characteristic fields were avoided in cryptographic applications and so became less important to study as attention moved away to larger characteristic fields (this changed however with the advent of pairing-based cryptography).

Secondly, and perhaps crucially, one might deduce from Theorems 3.2 and 3.4 that for large prime fields and small characteristic fields, DLP algorithms that rely on smoothness heuristics should have the same complexity when index calculus methods are used. As well as being true for the above two algorithms, in which elements are generated uniformly, it is also the case – at least heuristically – for the number field sieve and function field sieve, in which elements are generated non-uniformly. Hence algorithms which exploit the additional structure of small characteristic fields, namely the splitting of $X^q - X$, were simply not considered.

Looking ahead, there are three natural problems that remain open. The first is to provide a rigorous quasi-polynomial algorithm for the small characteristic DLP. The second, perhaps more challenging problem is to find a polynomial time algorithm for small characteristic, either heuristic or rigorous. Finally, the third, probably much harder problem is to develop faster algorithms for medium and large characteristic. Since there is less structure for prime fields in particular, it seems that fundamentally new ideas will be required.

REFERENCES

- [1] G. Adj, A. Menezes, T. Oliveira and F. Rodríguez-Henríquez, Weakness of \mathbb{F}_{3^6-509} for discrete logarithm cryptography, in: *Pairing-Based Cryptography—Pairing 2013*, Springer, LNCS **8365** (2014), 20–44.
- [2] G. Adj, A. Menezes, T. Oliveira and F. Rodríguez-Henríquez, Computing discrete logarithms in \mathbb{F}_{3^6-137} and \mathbb{F}_{3^6-163} using Magma, in: *Arithmetic of Finite Fields*, Springer, LNCS **9061** (2015), 3–22.
- [3] G. Adj, I. Canales-Martínez, N. Cruz-Cortés, A. Menezes, T. Oliveira, L. Rivera-Zamarripa and F. Rodríguez-Henríquez, Computing discrete logarithms in cryptographically-interesting characteristic-three finite fields, *IACR Cryptology ePrint Archive*, (2016), 19 pages, eprint.iacr.org/2016/914.
- [4] L. M. Adleman, A subexponential algorithm for the discrete logarithm problem with applications to cryptography, in: *20th Annual Symposium on Foundations of Computer Science*, (1979), 55–60.
- [5] L. M. Adleman, The function field sieve, in: *Algorithmic Number Theory*, Springer, LNCS **877** (1994), 108–121.
- [6] L. M. Adleman and M.-D. A. Huang, [Function field sieve method for discrete logarithms over finite fields](#), *Inform. and Comput.*, **151** (1999), 5–16.

- [7] R. Barbulescu, C. Bouvier, J. Detrey, P. Gaudry, H. Jeljeli, E. Thomé, M. Videau and P. Zimmermann (the CAMEL group), Discrete logarithm in $\text{GF}(2^{809})$ with FFS, in: *Public-Key Cryptography—PKC 2014*, Springer, LNCS **8383** (2014), 221–238.
- [8] R. Barbulescu, P. Gaudry, A. Guillevic and F. Morain, Improving NFS for the discrete logarithm problem in non-prime finite fields, in: *Advances in Cryptology—EUROCRYPT 2015*, Springer, LNCS **9056** (2015), 129–155.
- [9] R. Barbulescu, P. Gaudry, A. Joux and E. Thomé, A heuristic quasi-polynomial algorithm for discrete logarithm in finite fields of small characteristic, in: *Advances in Cryptology—EUROCRYPT 2014*, Springer, LNCS **8441** (2014), 1–16.
- [10] R. Barbulescu, P. Gaudry and T. Kleinjung, The Tower Number Field Sieve, in: *Advances in Cryptology—ASIACRYPT 2015*, Springer, LNCS **9453** (2015), 31–55.
- [11] R. Barbulescu and T. Kim, Extended tower number field sieve: A new complexity for the medium prime case, in: *Advances in Cryptology—CRYPTO 2016*, Springer, LNCS **9814** (2016), 543–571.
- [12] A. W. Blüher, On $x^{q+1} + ax + b$, *Finite Fields Appl.*, **10** (2004), 285–305.
- [13] D. Boneh and M. Franklin, Identity-based encryption from the Weil pairing, in: *Advances in Cryptology—CRYPTO 2001*, Springer, LNCS **2139** (2001), 213–229.
- [14] E. R. Canfield, P. Erdős and C. Pomerance, On a problem of Oppenheim concerning ‘factorisatio numerorum’, *J. Number Theory*, **17** (1983), 1–28.
- [15] A. Commeine and I. Semaev, An algorithm to solve the discrete logarithm problem with the number field sieve, in: *Public Key Cryptography—PKC 2006*, Springer, LNCS **3958** (2006), 174–190.
- [16] D. Coppersmith, Fast evaluation of logarithms in fields of characteristic two, *IEEE Trans. Inform. Theory*, **30** (1984), 587–594.
- [17] C. Diem, On the discrete logarithm problem in elliptic curves, *Compositio Math.*, **147** (2011), 75–104.
- [18] W. Diffie and M. E. Hellman, New directions in cryptography, *IEEE Trans. Inform. Theory*, **22** (1976), 644–654.
- [19] T. ElGamal, A public key cryptosystem and a signature scheme based on discrete logarithms, in: *Advances in Cryptology—CRYPTO ’84*, Springer, LNCS **196** (1985), 10–18.
- [20] A. Enge and P. Gaudry, A general framework for subexponential discrete logarithm algorithms, *Acta Arithmetica*, **102** (2002), 83–103.
- [21] S. D. Galbraith, Supersingular curves in cryptography, in: *Advances in Cryptology—ASIACRYPT 2001*, Springer, LNCS **2248** (2001), 495–513.
- [22] C. F. Gauß, *Disquisitiones Arithmeticae*, Translated into English by Arthur A. Clarke, S. J. Yale University Press, New Haven, Conn.-London, 1966.
- [23] F. Göloğlu, R. Granger, G. McGuire and J. Zumbrägel, On the function field sieve and the impact of higher splitting probabilities: Application to discrete logarithms in $\mathbb{F}_{2^{1971}}$ and $\mathbb{F}_{2^{3164}}$, in: *Advances in Cryptology—CRYPTO 2013*, Springer, LNCS **8043** (2013), 109–128.
- [24] F. Göloğlu, R. Granger, G. McGuire and J. Zumbrägel, Solving a 6120-bit DLP on a desktop computer, in: *Selected Areas in Cryptography—SAC 2013*, Springer, LNCS **8282** (2014), 136–152.
- [25] D. M. Gordon, Discrete logarithms in $\text{GF}(p)$ using the number field sieve, *SIAM J. Discrete Math.*, **6** (1993), 124–138.
- [26] D. M. Gordon and K. S. McCurley, Massively parallel computation of discrete logarithms, in: *Advances in Cryptology—CRYPTO ’92*, Springer, LNCS **740** (1993), 312–323.
- [27] R. Granger, T. Kleinjung and J. Zumbrägel, Breaking ‘128-bit secure’ supersingular binary curves, in: *Advances in Cryptology—CRYPTO 2014*, Springer, LNCS **8617** (2014), 126–145.
- [28] R. Granger, T. Kleinjung and J. Zumbrägel, On the powers of 2, *IACR Cryptology ePrint Archive*, (2014), 18 pages, eprint.iacr.org/2014/300.
- [29] R. Granger, T. Kleinjung and J. Zumbrägel, On the discrete logarithm problem in finite fields of fixed characteristic, *Trans. Amer. Math. Soc.*, **370** (2018), 3129–3145.
- [30] T. Hayashi, T. Shimoyama, N. Shinohara, T. Takagi, Breaking pairing-based cryptosystems using η_T pairing over $\text{GF}(3^{97})$, in: *Advances in Cryptology—ASIACRYPT 2012*, Springer, LNCS **7658** (2012), 43–60.
- [31] T. Hayashi, N. Shinohara, L. Wang, S. I. Matsuo, M. Shirase and T. Takagi, Solving a 676-bit discrete logarithm problem in $\text{GF}(3^{6n})$, in: *Public Key Cryptography—PKC 2010*, Springer, LNCS **6056** (2010), 351–367.

- [32] T. Helleseth and A. Kholosha, $x^{2^l+1} + x + a$ and related affine polynomials over $GF(2^k)$, *Cryptogr. Commun.*, **2** (2010), 85–109.
- [33] A. Joux, A one round protocol for tripartite Diffie-Hellman, in: *Algorithmic Number Theory*, Springer, LNCS **1838** (2000), 385–393.
- [34] A. Joux, Faster index calculus for the medium prime case; application to 1175-bit and 1425-bit finite fields, in: *Advances in Cryptology—EUROCRYPT 2013*, Springer, LNCS **7881** (2013), 177–193.
- [35] A. Joux, A new index calculus algorithm with complexity $L(1/4 + o(1))$ in small characteristic, in: *Selected Areas in Cryptography—SAC 2013*, Springer, LNCS **8282** (2014), 355–379.
- [36] A. Joux and R. Lercier, *The function field sieve is quite special*, in: *Algorithmic Number Theory*, Springer, LNCS **2369** (2002), 431–445.
- [37] A. Joux and R. Lercier, *Improvements to the general number field sieve for discrete logarithms in prime fields*, *Math. Comp.*, **72** (2003), 953–967.
- [38] A. Joux and R. Lercier, The function field sieve in the medium prime case, in: *Advances in Cryptology—EUROCRYPT 2006*, Springer, LNCS **4004** (2006), 254–270.
- [39] A. Joux, R. Lercier, N. Smart and F. Vercauteren, The number field sieve in the medium prime case, in: *Advances in Cryptology—CRYPTO 2006*, Springer, LNCS **4117** (2006), 326–344.
- [40] A. Joux, A. M. Odlyzko and C. Pierrot, The past, evolving present and future of discrete logarithm, in: *Open Problems in Mathematical and Computational Science*, Springer (2014), 5–36.
- [41] A. Joux and C. Pierrot, Improving the polynomial time precomputation of Frobenius representation discrete logarithm algorithms, in: *Advances in Cryptology—ASIACRYPT 2014*, Springer, LNCS **8873** (2014), 378–397.
- [42] M. Kalkbrener, An upper bound on the number of monomials in determinants of sparse matrices with symbolic entries, *Mathematica Pannonica*, **8** (1997), 73–82.
- [43] T. Kim and J. Jeong, Extended tower number field sieve with application to finite fields of arbitrary composite extension degree, in: *Public-Key Cryptography—PKC 2017*, **10174** (2017), 388–408.
- [44] B. A. LaMacchia and A. M. Odlyzko, Solving large sparse linear systems over finite fields, in: *Advances in Cryptology—CRYPTO '90*, Springer, LNCS **537** (1991), 109–133.
- [45] C. Lanczos, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, *J. Research Nat. Bur. Standards*, **45** (1950), 255–282.
- [46] A. K. Lenstra and H. W. Lenstra, Jr, Algorithms in number theory, in: *Handbook of Theoretical Computer Science (A): Algorithms and Complexity*, Elsevier, (1990), 673–715.
- [47] A. K. Lenstra and H. W. Lenstra, Jr (eds), *The Development of the Number Field Sieve*, Springer, 1993.
- [48] A. K. Lenstra, H. W. Lenstra, Jr and L. Lovász, *Factoring polynomials with rational coefficients*, *Math. Ann.*, **261** (1982), 515–534.
- [49] H. W. Lenstra, Jr, *Finding isomorphisms between finite fields*, *Math. Comp.*, **56** (1991), 329–347.
- [50] R. Lovorn, *Rigorous Subexponential Algorithms for Discrete Logarithms over Finite Fields*, *Ph. D. Thesis*, University of Georgia, 1992.
- [51] V. I. Nechaev, *On the complexity of a deterministic algorithm for a discrete logarithm*, *Mat. Zametki*, **55** (1994), 91–101, 189.
- [52] J. Neukirch, *Algebraic Number Theory*, Translated from the 1992 German original, Springer, 1999.
- [53] A. M. Odlyzko, Discrete logarithms in finite fields and their cryptographic significance, in: *Advances in Cryptology—CRYPTO '84*, Springer, LNCS **209** (1985), 224–314.
- [54] A. M. Odlyzko, *Discrete logarithms: the past and the future*, *Des. Codes Cryptogr.*, **19** (2000), 129–145.
- [55] S. C. Pohlig and M. E. Hellman, *An improved algorithm for computing logarithms over $GF(p)$ and its cryptographic significance*, *IEEE Trans. Inform. Theory*, **24** (1978), 106–110.
- [56] J. M. Pollard, *Monte Carlo methods for index computation (mod p)*, *Math. Comp.*, **32** (1978), 918–924.
- [57] C. Pomerance, Analysis and comparison of some integer factoring algorithms, in: *Computational Methods in Number Theory*, *Math. Centre Tracts*, Math. Centrum, Amsterdam, **154** (1982), 89–139.
- [58] C. Pomerance, Fast, rigorous factorization and discrete logarithm algorithms, in: *Discrete Algorithms and Complexity, Perspect. Comput.*, Academic Press, **15** (1987), 119–143.

- [59] R. Sakai, K. Ohgishi and M. Kasahara, Cryptosystems based on pairing, in: *Symposium on Cryptography and Information Security*, Okinawa, Japan, (2000), 26–28.
- [60] P. Sarkar and S. Singh, A general polynomial selection method and new asymptotic complexities for the tower number field sieve algorithm, in: *Advances in Cryptology—ASIACRYPT 2016*, Springer, LNCS **10031** (2016), 37–62.
- [61] O. Schirokauer, [Discrete logarithms and local units](#), *Philos. Trans. Roy. Soc. London Ser. A*, **345** (1993), 409–423.
- [62] O. Schirokauer, [Using number fields to compute logarithms in finite fields](#), *Math. Comp.*, **69** (2000), 1267–1283.
- [63] O. Schirokauer, [Virtual logarithms](#), *J. Algorithms*, **57** (2005), 140–147.
- [64] C.-P. Schnorr, [Efficient signature generation by smart cards](#), *J. Cryptology*, **4** (1991), 161–174.
- [65] I. A. Semaev, [Special prime numbers and discrete logs in finite prime fields](#), *Math. Comp.*, **71** (2002), 363–377.
- [66] P. W. Shor, [Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer](#), *SIAM J. Computing*, **26** (1997), 1484–1509.
- [67] V. Shoup, Lower bounds for discrete logarithms and related problems, in: *Advances in Cryptology—EUROCRYPT '97*, Springer, LNCS **1223** (1997), 256–266.
- [68] D. Wan, [Generators and irreducible polynomials over finite fields](#), *Math. Comp.*, **66** (1997), 1195–1212.
- [69] D. H. Wiedemann, [Solving sparse linear equations over finite fields](#), *IEEE Trans. Inform. Theory*, **32** (1986), 54–62.

Received April 2016; revised December 2017.

E-mail address: robert.granger@epfl.ch

E-mail address: thorsten.kleinjung@epfl.ch

E-mail address: jens.zumbraegel@uni-passau.de