EFFICIENT INVERSION OF RATIONAL MAPS OVER FINITE FIELDS

ANTONIO CAFURE*, GUILLERMO MATERA†, AND ARIEL WAISSBEIN‡

Abstract. We study the problem of finding the inverse image of a point in the image of a rational map $F : \mathbb{F}_q^n \to \mathbb{F}_q^n$ over a finite field $\mathbb{F}_q$. Our interest mainly stems from the case where F encodes a permutation given by some public–key cryptographic scheme. Given an element $y^{(0)} \in F(\mathbb{F}_q^n)$, we are able to compute the set of values $x^{(0)} \in \mathbb{F}_q^n$ for which $F(x^{(0)}) = y^{(0)}$ holds with $O(\mathsf{T} n^{4.38} D^{2.38} \delta \log^2 q)$ bit operations, up to logarithmic terms. Here T is the cost of the evaluation of $F_1, \ldots, F_n$, D is the degree of F and δ is the degree of the graph of F.

Key words. Finite fields, polynomial system solving, public–key cryptography, matrices of fixed displacement rank.

AMS(MOS) subject classifications. 14G05, 68W30, 11T71, 68Q25, 47B35.

1. Introduction.

Let $\mathbb{F}_q$ be the finite field of q elements, let $\overline{\mathbb{F}}_q$ denote its algebraic closure and let $\mathbb{A}^n$ denote the n–dimensional affine space $\overline{\mathbb{F}}_q^n$ considered as a topological space endowed with the Zariski topology. Let $X := (X_1, \ldots, X_n)$ be a vector of indeterminates and let $F_1, \ldots, F_n$ be elements of the field $\mathbb{F}_q(X) := \mathbb{F}_q(X_1, \ldots, X_n)$ of rational functions in $X_1, \ldots, X_n$ with coefficients in $\mathbb{F}_q$. We consider the rational map $F : \mathbb{A}^n \to \mathbb{A}^n$ defined as $F(x) := (F_1(x), \ldots, F_n(x))$. Assume that the restriction of F to $\mathbb{F}_q^n$ is a partially–defined mapping $F : \mathbb{F}_q^n \to \mathbb{F}_q^n$, i.e., F is well–defined on a nonempty subset of $\mathbb{F}_q^n$. In such a case, we have that $F : \mathbb{F}_q^n \to \mathbb{F}_q^n$ agrees with an $\mathbb{F}_q$–definable polynomial map F^* on its domain (see, e.g., [25]). Unfortunately, the degrees of the polynomials defining F^* may grow exponentially, which prevents us to replace the rational mapping F with the corresponding polynomial map F^*. In this paper we exhibit an algorithm which, given $y^{(0)} \in F(\mathbb{F}_q^n)$, computes $x^{(0)} \in \mathbb{F}_q^n$ such that $F(x^{(0)}) = y^{(0)}$ holds.

A possible approach to this problem consists in computing the inverse mapping of F, provided that F is polynomially or rationally invertible. This is done in [41], where the authors describe an algorithm for inverting a bijective polynomial map $F : \mathbb{A}^n \to \mathbb{A}^n$ defined over $\mathbb{F}_q$, assuming that F

*Depto. de Matemática, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Ciudad Universitaria, Pabellón I, (C1428EHA) Buenos Aires, Argentina. Instituto del Desarrollo Humano, Universidad Nacional de General Sarmiento, J.M. Gutiérrez 1150 (1613) Los Polvorines, Argentina.

†Instituto del Desarrollo Humano, Universidad Nacional de General Sarmiento, J.M. Gutiérrez 1150 (1613) Los Polvorines, Argentina. CONICET, Argentina.

‡CoreLabs, Core Security Technologies, Humboldt 1967 (C1414CTU) Ciudad de Buenos Aires, Argentina. Doctorado en Ingeniería, ITBA: Av. Eduardo Madero 399 (C1106ACD) Cdad. de Buenos Aires, Argentina.

is an automorphism of $\mathbb{F}_q[X]^n$ whose inverse has degree $(dn)^{O(1)}$, where d is the maximum of the degrees of the polynomials $F_1, \ldots, F_n$. The algorithm performs $(\mathsf{T}nd)^{O(1)}$ arithmetic operations in $\mathbb{F}_q$, where T is the number of arithmetic operations required to evaluate F. Nevertheless, the assumption of the existence of a polynomial or rational inverse of F with degree polynomially bounded seems to be too restrictive to be useful in practice.

From the cryptographic point of view, the critical problem is that of computing the inverse image of a given point $y^{(0)} \in F(\mathbb{F}_q^n)$ under the map F, rather than that of inverting F itself. In this sense, our problem may be reduced to that of solving polynomial systems over a finite field. Unfortunately, it is well known that solving polynomial systems over a finite field is an NP–complete problem, even with quadratic equations with coefficients in $\mathbb{F}_2$ [19]. This has led to the construction of several multivariate public–key cryptoschemes whose security is based on this difficulty. In fact, since [24] researchers have tried to construct public–key schemes based on this apparent difficulty (see, e.g., [36]), but proposals are typically proved to be weak through *ad hoc* attacks (see, e.g., [35], [29], [43]). This might be seen as an indication that the polynomial systems used in public–key cryptography are not intrinsically difficult to solve, and calls for the study of parameters to measure such difficulty.

In this article we exhibit a probabilistic algorithm that computes the inverse image of a point $y^{(0)} \in F(\mathbb{F}_q^n)$ with a cost which depends polynomially on two geometric parameters: the degree D of the map F and the degree δ of the graph of F.

1.1. Outline of our approach. Let $Y = (Y_1, \ldots, Y_n)$ be a vector of new indeterminates. We consider the Zariski closure of the graph $Y = F(X)$ of the morphism F as an $\mathbb{F}_q$–variety $V \subset \mathbb{A}^{2n}$, i.e., as the set of common zeros in $\mathbb{A}^{2n}$ of a finite set of polynomials in $\mathbb{F}_q[X, Y]$. We suppose that the projection morphism $\pi : V \to \mathbb{A}^n$ defined as $\pi(x, y) = y$ is a dominant map of degree D and denote by δ the degree of V. It turns out that V is an absolutely irreducible variety of dimension n, and the generic fiber of π is finite. With a slight abuse of notation we shall identity $F^{-1}(y^{(0)})$ with the fiber $\pi^{-1}(y^{(0)})$.

In order to compute all the q–rational points of $F^{-1}(y^{(0)})$, i.e., the points in the set $F^{-1}(y^{(0)}) \cap \mathbb{F}_q^n$, we shall deform the system $F(X) = y^{(0)}$ into a system $F(X) = y^{(1)} := F(x^{(1)})$ with a point $x^{(1)}$ randomly chosen in a suitable finite field extension $\mathbb{K}$ of $\mathbb{F}_q$. The deformation is given by the curve $\mathcal{C} \subset \mathbb{A}^{n+1}$ defined by the equations $F(X) = y^{(1)} + (S-1)(y^{(1)} - y^{(0)})$. For this purpose, we obtain upper bounds on the degree of the generic condition underlying the choice of $x^{(1)}$, which allows us to determine the cardinality of a finite field extension $\mathbb{K}$ of $\mathbb{F}_q$ where such a random choice is possible with good probability of success (see Section 4.2).

The algorithm computing the set $F^{-1}(y^{(0)}) \cap \mathbb{F}_q^n$ may be divided in three main parts. First, we compute a polynomial $m_S(S, T)$ defining a

plane curve birationally equivalent to $\mathcal{C}$. This polynomial is obtained as the solution of a system of linear equations whose matrix is block–Toeplitz. Such a solution is computed by applying an efficient algorithm based on the theory of matrices of fixed displacement rank due to [8] (see Section 5.1). Then, in Section 5.2 we extend the computation of the defining polynomial $m_S(S,T)$ of the plane curve birationally equivalent to $\mathcal{C}$ to the computation of the birational inverse itself. Finally, in Section 5.3 we substitute 0 for S in the birational inverse computed in the previous step and recover the 0–dimensional fiber $\pi^{-1}(y^{(0)})$, from which we obtain all the points of $F^{-1}(y^{(0)})$ with coordinates in $\mathbb{F}_q$.

The cost of our algorithm is roughly $O(\mathsf{T} n^{4.38} D^{2.38} \delta \log^2 q)$ bit operations, up to logarithmic terms, where T is the cost of the evaluation of the rational map F. Therefore, we extend and improve the results of [12], which require F to be a polynomial map defining a bijection from $\mathbb{F}_q^n$ to $\mathbb{F}_q^n$. This extension allows us to deal with cryptographic schemes such as the so–called "Tractable Rational Map" cryptosystem (see [42], [43, Section 6]). On the other hand, we observe that, if the hypotheses of [41] hold, then our algorithm meets also the complexity bound $(\mathsf{T} nd)^{O(1)}$ of [41].

As mentioned before, another alternative approach in order to compute one or all the q–rational points of $F^{-1}(y^{(0)})$ could be to apply a general algorithm for finding one or all the q–rational solutions of a given polynomial system over a finite field. Algorithms directly aimed at solving polynomial systems over a finite field are usually based on Gröbner basis computations ([16], [3], [4]), elimination techniques ([23], [10]) or relinearization ([29], [14]). Unfortunately, all these algorithms have worst–case exponential running time, and only [10] achieves a polynomial cost in the Bézout number of the system (which is nevertheless exponential in worst case). Furthermore, it is not clear how these algorithms could profit from the knowledge that a given map has associated geometric parameters of "low" value, as happens in certain cryptographic situations.

Finally, from the cryptographic point of view, we observe that withstanding the differential cryptanalysis of E. Biham and A. Shamir ([6], [15]) has become a *de facto* requirement for any block cipher. On the other hand, there is no "strong" test which allows one to analyze the security of cryptosystems based on the problem of solving multivariate equations. Our algorithm may be considered as a first step in this direction.

2. Notions and notations. Throughout this paper we shall denote by $|A|$ the number of elements of a given finite set A.

Let $\mathbb{K}$ be a finite field extension of $\mathbb{F}_q$, let $\mathbb{A}^n := \mathbb{A}^n(\overline{\mathbb{F}}_q)$ be the n–dimensional $\overline{\mathbb{F}}_q^n$ endowed with its Zariski topology, and let $V \subset \mathbb{A}^n$ be a $\mathbb{K}$–variety, that is, the set of common zeros in $\mathbb{A}^n$ of a finite set of polynomials of $\mathbb{K}[X]$. We denote by $\mathbf{I}(V) \subset \mathbb{K}[X]$ the ideal of the variety V, by $\mathbb{K}[V] := \mathbb{K}[X]/\mathbf{I}(V)$ its coordinate ring and by $\mathbb{K}(V)$ its field of total fractions.

For a given irreducible $\mathbb{K}$–variety $V \subset \mathbb{A}^n$, we define its *degree* $\deg V$ as the maximum number of points lying in the intersection of V with an affine linear variety $L \subset \mathbb{A}^n$ of codimension $\dim V$ for which $|V \cap L|$ is finite. More generally, if $V = C_1 \cup \cdots \cup C_N$ is the decomposition of V into irreducible $\mathbb{K}$–components, we define the degree of V as $\deg V := \sum_{i=1}^{N} \deg C_i$ (cf. [22]).

We say that a $\mathbb{K}$–variety $V \subset \mathbb{A}^n$ is *absolutely irreducible* if it is an irreducible $\overline{\mathbb{F}}_q$–variety.

Let $V \subset \mathbb{A}^n$ be an irreducible $\mathbb{K}$–variety and let $\pi : V \to \mathbb{A}^n$ be a dominant mapping. Then we have an algebraic field extension $\mathbb{K}(\mathbb{A}^n) \hookrightarrow \mathbb{K}(V)$. The degree $\deg \pi$ of π is defined as the degree of the field extension $\mathbb{K}(\mathbb{A}^n) \hookrightarrow \mathbb{K}(V)$. We called a point $y \in \mathbb{A}^n$ a *lifting point* of π if the number of inverse images of y equals the degree of the morphism π.

2.1. Data structures. Algorithms in elimination theory are usually described using the standard dense (or sparse) complexity model, i.e., encoding multivariate polynomials by means of the vector of all (or of all nonzero) coefficients. Taking into account that a generic n–variate polynomial of degree d has $\binom{d+n}{n} = O(d^n)$ nonzero coefficients, we see that the dense or sparse representation of multivariate polynomials requires an exponential size, and their manipulation usually requires an exponential number of arithmetic operations with respect to the parameters d and n. In order to avoid this exponential behavior, we are going to use an alternative encoding of input and intermediate results of our computations by means of straight–line programs (cf. [9]). A *straight–line program* β in $\mathbb{K}(X) := \mathbb{K}(X_1, \ldots, X_n)$ is a finite sequence of rational functions $(F_1, \ldots, F_k) \in \mathbb{K}(X)^k$ such that for $1 \leq i \leq k$, the function F_i is either an element of the set $\{X_1, \ldots, X_n\}$, or an element of $\mathbb{K}$ (a *parameter*), or there exist $1 \leq i_1, i_2 < i$ such that $F_i = F_{i_1} \circ_i F_{i_2}$ holds, where $\circ_i$ is one of the arithmetic operations $+, -, \times, \div$. The straight–line program β is called *division–free* if $\circ_i$ is different from $\div$ for $1 \leq i \leq k$. A natural measure of the complexity of β is its *time* (cf. [38]) which is the total number of arithmetic operations performed during the evaluation. We say that the straight–line program β *computes* or *represents* a subset S of $\mathbb{K}(X)$ if $S \subset \{F_1, \ldots, F_k\}$ holds.

2.2. The algorithmic model. Our algorithms are of *Monte Carlo* or *BPP* type (see, e.g., [2], [18]), i.e., they return the correct output with a probability of at least a fixed value strictly greater than $1/2$. This in particular implies that the error probability can be made arbitrarily small by iteration of the algorithms. The probabilistic aspect of our algorithms is related to certain random choices of points with coordinates in a given finite field not annihilating certain nonzero polynomials. In order to perform a given random choice with a prescribed probability of success, we must know how many zeros the polynomial under consideration has. For this purpose, we have the following classical result, first shown by Oystein Ore in 1922.

PROPOSITION 2.1 ([32, Theorem 6.13]). *Let $\mathbb{K}$ be a finite field extension of $\mathbb{F}_q$ and let $F \in \overline{\mathbb{F}}_q[X]$ be a polynomial of degree d. The number of zeros of F in $\mathbb{K}$ is at most $d|\mathbb{K}|^{n-1}$.*

For the analysis of our algorithms, we shall interpret the statement of Proposition 2.1 in terms of probabilities. More precisely, assuming a uniform distribution of probability on the elements of the finite field $\mathbb{K}$, we have the following corollary, also known as the Zippel–Schwartz lemma in the computer algebra community (cf. [18, Lemma 6.44]).

COROLLARY 2.1. *Fix $\mu > 0$ and suppose that $|\mathbb{K}| > \mu d$ holds. Then the probability of choosing $x \in \mathbb{K}^n$ with $F(x) \neq 0$ is at least $1 - 1/\mu$.*

2.3. Cost of the basic operations. In this section we briefly review the cost of the basic algorithms we shall employ. The cost of such algorithms will be frequently stated in terms of the quantity

$$\mathsf{M}(m) := m \log^2 m \log\log m.$$

If $\mathbb{K}$ is a finite field, an arithmetic operation in $\mathbb{K}$ requires $O(\mathsf{M}(\log|\mathbb{K}|))$ bit operations. More generally, for a given domain R, the number of arithmetic operations in R necessary to compute the multiplication or division with remainder of two univariate polynomials in $R[T]$ of degree at most m is $O\big(\mathsf{M}(m)/\log(m)\big)$ (cf. [18], [7]).

If R is any field, then we shall use algorithms based on the Extended Euclidean Algorithm (EEA for short) in order to compute the gcd or the resultant of two univariate polynomials in $R[T]$ of degree at most m with $O\big(\mathsf{M}(m)\big)$ arithmetic operations in R (cf. [18], [7]).

Finally, we recall that for the cost $O(n^\omega)$ of the multiplication of two matrices of size $n \times n$ with coefficients in R, we have $\omega < 2.376$ (cf. [7]).

3. Geometric solutions. We shall use a representation of varieties which is well suited for algorithmic purposes (see, e.g., [20], [37], [21]). This representation is called a *geometric solution* or a *rational univariate representation* of the given variety. The notion of a geometric solution of an algebraic variety was first implicitly introduced in the works of Kronecker and König in the last years of the XIXth century. This section is devoted to motivate this notion and to describe certain underlying algorithmic aspects.

Let $\mathbb{K}$ be a perfect field and let $\overline{\mathbb{K}}$ denote its algebraic closure. We start with the definition of a ($\mathbb{K}$–definable) geometric solution of a zero-dimensional $\mathbb{K}$–variety. Let $V = \{P^{(1)}, \ldots, P^{(D)}\}$ be a zero–dimensional $\mathbb{K}$–variety of $\mathbb{A}^n$, and suppose that there exists a linear form $\mathcal{L} \in \mathbb{K}[X]$ which separates the points of V, i.e., which satisfies $\mathcal{L}(P^{(i)}) \neq \mathcal{L}(P^{(k)})$ if $i \neq k$. A geometric solution of V consists of

- a linear form $\mathcal{L} := \lambda \cdot X := \lambda_1 X_1 + \cdots + \lambda_n X_n \in \mathbb{K}[X]$ which separates the points of V,
- the minimal polynomial $m_\lambda := \prod_{1 \leq i \leq D}(T - \mathcal{L}(P^{(i)})) \in \mathbb{K}[T]$ of $\mathcal{L}$ in V (where T is a new variable),

- polynomials $w_1, \ldots, w_n \in \mathbb{K}[T]$ with $\deg w_j < D$ for every $1 \leq j \leq n$ satisfying the identity:

$$V = \{(w_1(\eta), \ldots, w_n(\eta)) \in \overline{\mathbb{K}}^n;\, \eta \in \overline{\mathbb{K}},\, m_\lambda(\eta) = 0\}.$$

Next, we define this notion for irreducible $\mathbb{K}$–varieties of dimension greater than 0. Let $V \subset \mathbb{A}^n$ be an irreducible $\mathbb{K}$–variety of dimension $r > 0$ and degree δ. Suppose that the indeterminates $X_1, \ldots, X_r$ form a separable transcendence basis of the field extension $\mathbb{K} \hookrightarrow \mathbb{K}(V)$, that is, $\mathbb{K}(X_1, \ldots, X_r) \hookrightarrow \mathbb{K}(V)$ is a finite, separable field extension. Denote by D its degree. In particular, we have that the linear projection $\pi : V \to \mathbb{A}^r$ defined by $\pi(x) := (x_1, \ldots, x_r)$ is a dominant morphism of degree D. From the behavior of the degree of a variety under linear maps (see, e.g., [22, Lemma 2]), it follows that $D \leq \delta$ holds.

Let $\Lambda := (\Lambda_1, \ldots, \Lambda_n)$ be a vector of new indeterminates. Observe that the extension $V_{\mathbb{K}(\Lambda)}$ of V to $\mathbb{A}^n(\overline{\mathbb{K}(\Lambda)})$ is an irreducible $\mathbb{K}(\Lambda)$–variety of dimension r and the coordinate ring of $V_{\mathbb{K}(\Lambda)}$ as a $\mathbb{K}(\Lambda)$–variety is isomorphic to $\mathbb{K}(\Lambda) \otimes_{\mathbb{K}} \mathbb{K}[V]$. Consider the generic linear form

$$\mathcal{L}_\Lambda := \Lambda \cdot X := \Lambda_1 X_1 + \cdots + \Lambda_n X_n. \tag{3.1}$$

Let $\xi_1, \ldots, \xi_n \in \mathbb{K}[V]$ be the coordinate functions induced by $X_1, \ldots, X_n$, and set $\widehat{\mathcal{L}} := \mathcal{L}_\Lambda(\xi_1, \ldots, \xi_n) \in \mathbb{K}(\Lambda) \otimes_{\mathbb{K}} \mathbb{K}[V]$. Since $\xi_1, \ldots, \xi_r, \widehat{\mathcal{L}}$ are algebraically dependent over $\mathbb{K}(\Lambda)$, there exists an irreducible polynomial $m_\Lambda \in \mathbb{K}[\Lambda, X_1, \ldots, X_r, T]$, separable with respect to T, such that the following identity holds in $\mathbb{K}(\Lambda) \otimes_{\mathbb{K}} \mathbb{K}[V]$:

$$m_\Lambda(\Lambda, \xi_1, \ldots, \xi_r, \widehat{\mathcal{L}}) = 0. \tag{3.2}$$

From, e.g., [39, Proposition 1], we deduce the following bounds:

- $\deg_T m_\Lambda = D$,
- $\deg_{X_1,\ldots,X_r} m_\Lambda \leq \delta$,
- $\deg_\Lambda m_\Lambda \leq \delta$.

Taking partial derivatives at both sides of equation (3.2) we deduce that for every $j = 1, \ldots, n$ the identity

$$\frac{\partial m_\Lambda}{\partial \Lambda_j}(\Lambda, \xi_1, \ldots, \xi_r, \widehat{\mathcal{L}}) + \xi_j \frac{\partial m_\Lambda}{\partial T}(\Lambda, \xi_1, \ldots, \xi_r, \widehat{\mathcal{L}}) = 0 \tag{3.3}$$

holds in $\mathbb{K}(\Lambda) \otimes_{\mathbb{K}} \mathbb{K}[V]$. As a consequence of the separability of m_Λ with respect to T we see that the polynomial $\partial m_\Lambda / \partial T$ is nonzero.

Assume that there exists $\lambda \in \mathbb{K}^n$ such that the linear form $\mathcal{L} := \lambda \cdot X$ induces a primitive element of the separable field extension $\mathbb{K}(X_1, \ldots, X_r) \hookrightarrow \mathbb{K}(V)$. Let ℓ be the coordinate function of $\mathbb{K}[V]$ defined by $\mathcal{L}$. From the fact that $\deg_T m_\Lambda = D$ holds, it follows that $m_\Lambda(\lambda, X_1, \ldots, X_r, T)$ is the minimal polynomial of ℓ in the field extension $\mathbb{K}(X_1, \ldots, X_r) \hookrightarrow \mathbb{K}(V)$.

Setting

$$m_\lambda(X_1,\ldots,X_r,T) := m_\Lambda(\lambda,X_1,\ldots,X_r,T),$$
$$v_j(X_1,\ldots,X_r,T) := -\frac{\partial m_\Lambda}{\partial \Lambda_j}(\lambda,X_1,\ldots,X_r,T),$$

and substituting λ for Λ in (3.2)–(3.3), we obtain the following identities of $\mathbb{K}[V]$:

$$\begin{gathered} m_\lambda(\xi_1,\ldots,\xi_r,\ell)=0,\\ \frac{\partial m_\lambda}{\partial T}(\xi_1,\ldots,\xi_r,\ell)\xi_j - v_j(\xi_1,\ldots,\xi_r,\ell)=0,\end{gathered} \tag{3.4}$$

which show that the polynomials

$$m_\lambda(X_1,\ldots,X_r,\mathcal{L}), \frac{\partial m_\lambda}{\partial T}(X_1,\ldots,X_r,\mathcal{L})X_j - v_j(X_1,\ldots,X_r,\mathcal{L})\ (1\le j\le n),$$

belong to $\mathbf{I}(V)$. We observe that the polynomials $v_1,\ldots,v_n$ have coefficients in $\mathbb{K}$ and satisfy the conditions $\deg v_j \le \delta$ and $\deg_T v_j \le D$.

Finally, we remark that the polynomials $m_\lambda(X_1,\ldots,X_r,\mathcal{L})$ and $(\partial m_\lambda/\partial T)(X_1,\ldots,X_r,\mathcal{L})X_j - v_j(X_1,\ldots,X_r,\mathcal{L})$ $(r+1\le j\le n)$ constitute a system of equations for the variety V in the Zariski dense open subset $V\cap\{(\partial m_\lambda/\partial T)(X,\mathcal{L})\neq 0\}$ of V. This motivates the definition of a geometric solution of an irreducible $\mathbb{K}$–variety of arbitrary dimension:

Definition 3.1. *With assumptions as above, a* geometric solution *of V consists of the following items:*

- *a linear form $\mathcal{L} := \lambda\cdot X \in \mathbb{K}[X]$ which induces a primitive element ℓ of the field extension $\mathbb{K}(X_1,\ldots,X_r)\hookrightarrow\mathbb{K}(V)$,*
- *the minimal polynomial $m_\lambda\in\mathbb{K}[X_1,\ldots,X_r][T]$ of ℓ,*
- *a generic* parametrization *of the variety V by the zeros of m_λ, of the form $(\partial m_\lambda/\partial T)X_j - v_j$ $(r+1\le j\le n)$, with $v_j\in\mathbb{K}[X_1,\ldots,X_r][T]$, $\deg_T v_j < D$, $\deg_X v_j\le\delta$ and $(\partial m_\lambda/\partial T)(\mathcal{L})X_j - v_j(\mathcal{L})\in\mathbf{I}(V)$.*

We observe that the polynomial $m_\lambda\in\mathbb{K}[X_1,\ldots,X_r][T]$ of the second item of the previous definition can be also defined as follows: consider the linear map $\pi_\lambda : V\to\mathbb{A}^{r+1}$ defined by $\pi_\lambda(x) := (x_1,\ldots,x_r,\lambda\cdot x)$. The Zariski closure of $\pi_\lambda(V)$ is a $\mathbb{K}$–hypersurface H of $\mathbb{A}^{r+1}$ of degree at most δ, which is indeed defined by $m_\lambda(X_1,\ldots,X_r,T)=0$.

3.1. Algorithmic aspects of the computation of a geometric solution. From the algorithmic point of view, the crucial step towards the computation of a geometric solution of a variety V consists in the computation of the minimal polynomial m_Λ of the generic linear form $\mathcal{L}_\Lambda$. In this section we shall show how we can derive an algorithm for computing a geometric solution of an r–dimensional $\mathbb{K}$–variety V from a procedure for computing the minimal polynomial of the generic linear form $\mathcal{L}_\Lambda$ (cf. [1], [21], [39]).

Assume that we have already chosen $\lambda \in \mathbb{K}^n$ such that the linear form $\mathcal{L} := \lambda \cdot X$ induces a primitive element of the separable field extension $\mathbb{K}(X_1, \ldots, X_r) \hookrightarrow \mathbb{K}(V)$. Let $m_\lambda \in \mathbb{K}[X_1, \ldots, X_r, T]$ be its minimal polynomial.

Suppose that we are given an algorithm Φ over $\mathbb{K}(\Lambda)$ for computing the minimal polynomial of the linear form $\mathcal{L}_\Lambda = \Lambda \cdot X$. Suppose further that the vector $(\lambda_1, \ldots, \lambda_n)$ of coefficients of $\mathcal{L}$ does not annihilate any denominator in $\mathbb{K}[\Lambda]$ of any intermediate result of the algorithm Φ. In order to compute the polynomials $v_{r+1}, \ldots, v_n$ of Definition 3.1, we observe that the Taylor expansion of $m_\Lambda(\Lambda, X_1, \ldots, X_r, T)$ in powers of $\Lambda - \lambda := (\Lambda_1 - \lambda_1, \ldots, \Lambda_n - \lambda_n)$ has the following expression:

$$m_\Lambda(\Lambda, X_1, \ldots, X_r, T) = m_\lambda(X_1, \ldots, X_r, T) + \sum_{j=1}^{n} \Big(\frac{\partial m_\lambda}{\partial T}(X_1, \ldots, X_r, T) X_j - v_j(X_1, \ldots, X_r, T) \Big) (\Lambda_j - \lambda_j) \mod (\Lambda - \lambda)^2.$$

We shall compute this first–order Taylor expansion by computing the first–order Taylor expansion of each intermediate result in the algorithm Φ. In this way, each arithmetic operation in $\mathbb{K}(\Lambda)$ arising in the algorithm Φ becomes an arithmetic operation between two polynomials of $\mathbb{K}[\Lambda]$ of degree at most 1, and is truncated up to order $(\Lambda - \lambda)^2$. Since the first–order Taylor expansion of an addition, multiplication or division of two polynomials of $\mathbb{K}[\Lambda]$ of degree at most 1 requires $O(n)$ arithmetic operations in $\mathbb{K}$, we have that the whole step requires $O(n\mathsf{T})$ arithmetic operations in $\mathbb{K}$, where T is the number of arithmetic operations in $\mathbb{K}(\Lambda)$ performed by the algorithm Φ. Summarizing, we have the following result:

LEMMA 3.1. *Suppose that we are given:*

1. *an algorithm Φ in $\mathbb{K}(\Lambda)$ which computes the minimal polynomial $m_\Lambda \in \mathbb{K}[\Lambda, X_1, \ldots, X_r, T]$ of $\mathcal{L}_\Lambda := \Lambda \cdot X$ with T arithmetic operations in $\mathbb{K}(\Lambda)$,*
2. *a separating linear form $\mathcal{L} := \lambda \cdot X \in \mathbb{K}[X]$ such that the vector λ does not annihilate any denominator in $\mathbb{K}[\Lambda]$ of any intermediate result of the algorithm Φ.*

Then a geometric solution of the variety V can be (deterministically) computed with $O\big(n(\mathsf{T} + \mathsf{M}(D))\big)$ arithmetic operations in $\mathbb{K}$.

4. Preparation of the input data. Let $F_1, \ldots, F_n \in \mathbb{F}_q(X)$ be rational functions having a reduced representation $F_i = P_i/Q_i$ with numerator and denominator of degree at most d for $1 \leq i \leq n$. Consider the rational map $F : \mathbb{A}^n \to \mathbb{A}^n$ defined by $F(x) := (F_1(x), \ldots, F_n(x))$. Since the rational functions $F_1, \ldots, F_n$ have coefficients in $\mathbb{F}_q$, we see that the restriction of F to $\mathbb{F}_q^n$ induces a (partially–defined) mapping from $\mathbb{F}_q^n$ to $\mathbb{F}_q^n$ which we shall also denote by F, with a slight abuse of notation.

4.1. The graph of the mapping F. Let $Y := (Y_1, \ldots, Y_n)$ be a vector of new indeterminates. Our algorithm shall deal with the $\mathbb{F}_q$–variety

$V \subset \mathbb{A}^{2n}$ representing the Zariski closure of the graph of the mapping F. More precisely, let $F_i := P_i/Q_i$ be a reduced fraction representing the rational function F_i for $1 \leq i \leq n$ and set $Q := Q_1 \cdots Q_n$. Let $I \subset \mathbb{F}_q[X, Y]$ be the ideal generated by the polynomials $Q_i(X) - Y_i P_i(X)$ $(1 \leq i \leq n)$. Then we define V as

$$V := V(I : Q^\infty),$$

where Q^∞ denotes the multiplicatively closed subset of $\mathbb{F}_q[X, Y]$ generated by 1 and Q and $(I : Q^\infty)$ denotes the saturation of the ideal I by Q^∞, that is, $(I : Q^\infty) := \{P \in \mathbb{F}_q[X, Y]; PQ^s \in I \text{ for some } s \in \mathbb{Z}_{\geq 0}\}$.

Let $\pi : V \to \mathbb{A}^n$ be the projection mapping defined by $\pi(x, y) := y$. In what follows, we shall assume that F and π satisfy the following conditions, which are usually met in the cryptographic situations we are interested in.

- (i) F is partially defined over $\mathbb{F}_q^n$.
- (ii) π is a dominant mapping. In particular, the fiber $V_y := \pi^{-1}(y)$ is a zero–dimensional subvariety of V for a generic $y \in \mathbb{A}^n$.

We observe that (i) is required by most cryptographic schemes based on multivariate equations (see, e.g., [30, Chapter 4], [43]), while (ii) is required for example in cryptographic schemes based on "tractable" rational maps (see, e.g., [42], [43, Section 6]).

Assumption (ii) and the definition of V imply that V is an absolutely irreducible $\mathbb{F}_q$–variety of dimension n. Indeed, it is easy to see that $\mathbb{F}_q(V)$ is isomorphic to $\mathbb{F}_q(X)$, which implies that V is absolutely irreducible and of dimension n. Further consequences of our assumptions are that the set of variables Y is algebraically independent in $\mathbb{F}_q(V)$ and the polynomials $Q_i(X)Y_i - P_i(X)$ $(1 \leq i \leq n)$ generate a radical ideal of the localization $\mathbb{F}_q[X, Y]_{Q^\infty}$.

In the sequel, we shall denote by δ the degree of V and by D the degree of the morphism $\pi : V \to \mathbb{A}^n$.

4.2. Random choices. Let $\mathcal{L} := \lambda \cdot X \in \overline{\mathbb{F}}_q[X]$ be a linear form such that the corresponding coordinate function of $\overline{\mathbb{F}}_q[V]$ is a primitive element of the field extension $\overline{\mathbb{F}}_q(Y) \hookrightarrow \overline{\mathbb{F}}_q(V)$. In particular, the minimal polynomial $m_\lambda \in \overline{\mathbb{F}}_q[Y, T]$ of the coordinate function defined by $\mathcal{L}$ satisfies the degree estimate $\deg_T m_\lambda = D$. By the remark after Definition 3.1 we see that V is birationally equivalent to the hypersurface $H \subset \mathbb{A}^{n+1}$ defined by the polynomial $m_\lambda \in \overline{\mathbb{F}}_q[Y, T]$. We observe that the fact that V is absolutely irreducible implies that H, and thus m_λ, is absolutely irreducible.

For a given point $y^{(0)} \in \mathbb{F}_q^n$, we denote by $V_{y^{(0)}}$ the π–fiber of $y^{(0)}$. In order to compute the points of the set $F^{-1}(y^{(0)}) \cap \mathbb{F}_q^n$, or equivalently, the set $V_{y^{(0)}} \cap \mathbb{F}_q^{2n}$, we shall deform the system $F(X) = y^{(0)}$ into a system $F(X) = F(x^{(1)})$ with a point $x^{(1)}$ randomly chosen in a suitable finite field extension $\mathbb{K}$ of $\mathbb{F}_q$ to be determined. The kind of deformations we shall apply is inspired by the approach of [34]. In our next result we establish

suitable bounds on the degree of the genericity conditions underlying the choice of $x^{(1)}$.

LEMMA 4.1. *There exists a nonzero polynomial $A \in \overline{\mathbb{F}}_q[X]$ of degree at most $3d\delta^4$ such that for any $x \in \mathbb{A}^n$ with $A(x) \neq 0$, the point $y := F(x)$ satisfies the following conditions:*

(i) y is a lifting point of π,

(ii) Let S be a new indeterminate and let $I_{\mathcal{C}} \subset \overline{\mathbb{F}}_q[S,X]$ be the ideal

$$I_{\mathcal{C}} := \Big(P_i(X) - Q_i(X)\big(y + (S-1)(y - y^{(0)})\big); 1 \leq i \leq n\Big).$$

Then the curve $\mathcal{C} := V(I_{\mathcal{C}} : Q^\infty) \subset \mathbb{A}^{n+1}$ is absolutely irreducible.

Proof. Let $m_\lambda \in \overline{\mathbb{F}}_q[Y][T]$ be the minimal (primitive) polynomial of the linear form $\mathcal{L} := \lambda \cdot X \in \overline{\mathbb{F}}_q[X]$. Let $A_1^* \in \overline{\mathbb{F}}_q[Y]$ denote the discriminant of m_λ with respect to the variable T. From the absolutely irreducibility of m_λ we conclude that $A_1^* \neq 0$ holds. Furthermore, for any $y \in \mathbb{A}^n$ with $A_1^*(y) \neq 0$ we have that the fiber V_y consists of D distinct points, that is, y is a lifting point of π. Hence, the nonvanishing of the polynomial $A_1^* \in \overline{\mathbb{F}}_q[Y]$ represents a suitable genericity condition underlying the choice of a lifting point $y \in \mathbb{A}^n$.

In order to obtain a genericity condition underlying the choice of a point $x \in \mathbb{A}^n$ for which $y := \pi(x)$ is a lifting point, consider a reduced representation $A_1^*\big(F(X)\big) = P_1^*/Q_1^*$ of the rational function defined by $A_1^*\big(F(X)\big)$ and set $A_1 := P_1^* Q_1^*$. By definition it follows that $A_1 \in \overline{\mathbb{F}}_q[X]$ has degree bounded by $(2D-1)d\delta$. Since F is a dominant mapping, we have that there exists $x \in \mathbb{A}^n$ such that $A_1(x) \neq 0$ holds (see, e.g., [40, II.6.3, Theorem 4]). This implies that A_1 is a nonzero polynomial.

Next we consider a reduced representation

$$m_\lambda\big(F(X) + (S-1)\big(F(X) - y^{(0)}\big), T\big) = \frac{\mathcal{P}_1(X)\mathcal{P}_2(X,S,T)}{\mathcal{Q}(X)}$$

of the rational function $m_\lambda\big(F(X)+(S-1)\big(F(X)-y^{(0)}\big),T\big) \in \overline{\mathbb{F}}_q(X)[S,T]$, where $\mathcal{P}_2(X,S,T)$ is a primitive polynomial of $\overline{\mathbb{F}}_q[X][S,T]$. Observe that such a representation is unique up to scaling by nonzero elements of $\overline{\mathbb{F}}_q$. Set

$$\widetilde{m}_\lambda(X,S,T) := \frac{\mathcal{Q}(X)}{\mathcal{P}_1(X)}\, m_\lambda\big(F(X)+(S-1)\big(F(X)-y^{(0)}\big),T\big) = \mathcal{P}_2(X,S,T).$$

Let $x \in \mathbb{A}^n$ be any point with $Q(x) \neq 0$. Then the value $F(x)$ is well–defined, and hence $m_\lambda(F(x) + (S-1)(F(x) - y^{(0)}), T)$ and $\widetilde{m}_\lambda(x,S,T)$ are both well–defined nonzero polynomials of $\overline{\mathbb{F}}_q[S,T]$ of degree D. As a consequence, for any lifting point y of π and any $x \in V_y$, the polynomial $\widetilde{m}_\lambda(x,1,T)$ is a nonzero scalar multiple of $m_\lambda(y,T)$, and thus a separable element of $\overline{\mathbb{F}}_q[T]$ of degree D.

Following [28, Theorem 5], in the version of [10, Theorem 3.6], there exists a polynomial $A_2^* \in \overline{\mathbb{F}}_q[Y]$ of degree bounded by $2\delta^4 + \delta$ such that for

any $y \in \mathbb{A}^n$ with $A_2^*(y) \neq 0$ the polynomial $m_\lambda(y + (S-1)(y - y^{(0)}), T)$ is absolutely irreducible. Let $A_2 \in \overline{\mathbb{F}}_q[X]$ be the numerator of a reduced representation of the rational function $A_2^*(F(X)) \in \overline{\mathbb{F}}_q(X)$. It follows that A_2 has degree bounded by $2d\delta^4 + d\delta$ and, for any $x \in \mathbb{A}^n$ with $A_2(x) \neq 0$, the polynomial $\widetilde{m}_\lambda(x, S, T)$ is absolutely irreducible.

Let $A := A_1 A_2$. Observe that $A \in \overline{\mathbb{F}}_q[X]$ and has degree at most $3d\delta^4$. Now, if we consider any point $x \in \mathbb{A}^n$ satisfying $A(x) \neq 0$ and set $y := F(x)$, we claim that conditions (i) and (ii) of the statement of the lemma are satisfied. Indeed, $A_1(x) \neq 0$ implies that $A_1^*(y) \neq 0$, that is, the discriminant of $m_\lambda(y, T)$ with respect to T is nonzero. We deduce that $m_\lambda(y, T)$ has D distinct roots and therefore, y is a lifting point of π. Finally, since y is a lifting point of π and $A_2(x) \neq 0$, the polynomial $\widetilde{m}_\lambda(x, S, T)$ is absolutely irreducible and hence, so is the curve $\mathcal{C}$. □

We remark that in the case of a field of large characteristic, say $char(\mathbb{F}_q) \geq 2\delta^2$ or $char(\mathbb{F}_q) \geq \delta(\delta-1)+1$, the bound of the statement of Lemma 4.1 can be improved applying the approach of [17] or [31] respectively. More precisely, applying [31, Theorem 6] (see also [17, Theorem 5.1] for a slightly worse bound) it follows that there exists a nonzero polynomial $A \in \overline{\mathbb{F}}_q[X]$ of degree at most $4d\delta^2$ for which the conditions of the statement of Lemma 4.1 hold. Nevertheless, taking into account that in cryptographic applications fields of characteristic 2 are very common, we shall not pursue the subject any further.

Suppose that we have already chosen a point $x \in \mathbb{A}^n$ satisfying the conditions of Lemma 4.1 and let $y := F(x)$. Let $\Lambda := (\Lambda_1, \ldots, \Lambda_n)$ be a vector of new indeterminates.

LEMMA 4.2. *There exists a nonzero polynomial $B \in \overline{\mathbb{F}}_q[\Lambda]$ of degree at most $2D^2$ such that for any $\lambda \in \mathbb{A}^n$ with $B(\lambda) \neq 0$, the linear form $\mathcal{L} := \lambda \cdot X$ separates the points of V_y and $V_{y^{(0)}}$.*

Proof. Let $V_y \cup V_{y^{(0)}} := \{P_1, \ldots, P_{D'}\}$. We consider the generic linear form $\mathcal{L}_\Lambda := \Lambda \cdot X$ and define

$$B(\Lambda) := \prod_{1 \leq i < j \leq D'} (\mathcal{L}_\Lambda(P_i) - \mathcal{L}_\Lambda(P_j)).$$

Since $D' \leq 2D$ holds, it follows that $B \in \overline{\mathbb{F}}_q[\Lambda]$ is a nonzero polynomial of degree at most $2D^2$. Any $\lambda \in \mathbb{A}^n$ not annihilating B provides a linear form $\mathcal{L}$ that separates the points of V_y and $V_{y^{(0)}}$. □

Now we can determine the degree of a field extension $\mathbb{K}$ of $\mathbb{F}_q$ for which the existence of points $\lambda, x \in \mathbb{K}^n$ satisfying Lemmas 4.1 and 4.2 can be assured. Our next result states that for a random choice of the coordinates of λ and x in a field extension $\mathbb{K}$ of $\mathbb{F}_q$ of suitable degree the statements of Lemmas 4.1 and 4.2 hold with high probability of success.

COROLLARY 4.1. *With notations as in Lemmas 4.1 and 4.2, fix $\mu > 0$ and let $\mathbb{K}$ be a finite field extension of $\mathbb{F}_q$ such that $|\mathbb{K}| > 4\mu d\delta^4$ holds. Then a random choice of (λ, x) in $\mathbb{K}^{2n}$ satisfies the condition $(AB)(\lambda, x) \neq 0$ with error probability at most $1/\mu$.*

Proof. By Proposition 2.1, the number of zeros in $\mathbb{K}^n$ of the polynomial A is at most $3d\delta^4|\mathbb{K}|^{n-1}$. Then a random choice of $x \in \mathbb{K}^n$ satisfies $A(x) \neq 0$ with probability at least $1 - 3d\delta^4/|\mathbb{K}| \geq 1 - 3/4\mu$. Given such a choice, a random choice of $\lambda \in \mathbb{K}^n$ satisfies $B(\lambda) \neq 0$ with probability at least $1 - 2D^2/|\mathbb{K}| \geq 1 - 1/4\mu$. This shows that a random choice $(\lambda, x) \in \mathbb{K}^{2n}$ satisfies $(AB)(\lambda, x) \neq 0$ with probability at least $(1 - 3/4\mu)(1 - 1/4\mu) \geq 1 - 1/\mu$. ∎

We remark that the polynomial AB of statement of Corollary 4.1 will not be computed during the execution of our algorithm, and therefore our algorithm will proceed with a random choice $(\lambda, x) \in \mathbb{K}$ for which the identity $AB(\lambda, x) = 0$ might hold. In such an unlikely event, certain intermediate values which are expected to be nonzero are equal to zero, and the algorithm must be restarted with another random choice of (λ, x).

A second remark is that, as we do not know in general the values of D and δ *a priori* (although in some cryptosystems such values are known), in order to determine the size of the field $\mathbb{K}$ these values can be estimated by d^n. This will not increase significatively the cost of our algorithm, since the cost depends linearly on the logarithm of $|\mathbb{K}|$.

5. The algorithm. Let $\mathbb{K}$ be a finite field extension of $\mathbb{F}_q$ whose cardinality will be determined later. Let $(\lambda, x^{(1)}) \in \mathbb{K}^{2n}$ be a point randomly chosen. By Corollary 4.1 we have that $(\lambda, x^{(1)})$ satisfies the conditions in the statements of Lemmas 4.1 and 4.2 with error probability at most $4d\delta^4/|\mathbb{K}|$. This means that with such an error probability the following assertions hold:

- $y^{(1)} := F(x^{(1)})$ is a well–defined lifting point of π;
- let $I_{\mathcal{C}} := \big(P_i(X) - Q_i(X)(y_i^{(1)} + (S-1)(y_i^{(1)} - y_i^{(0)})\big)$. Then

$$\mathcal{C} := V(I_{\mathcal{C}} : Q^\infty) \tag{5.1}$$

  is an absolutely irreducible curve of $\mathbb{A}^{n+1}$;
- the linear form $\mathcal{L} := \lambda \cdot X \in \mathbb{K}[X]$ separates the points of the fibers $V_{y^{(1)}}$ and $V_{y^{(0)}}$.

In what follows, we shall assume that all these conditions hold.

We consider the projection $\pi_S : \mathcal{C} \to \mathbb{A}^1$ defined by $\pi_S(s, x) := s$. We have that π_S is a dominant mapping of degree D, that $S = 1$ is a lifting point of π_S and that the identities $\pi_S^{-1}(1) = \{1\} \times \mathcal{C}_1$ and $\pi_S^{-1}(0) = \{0\} \times \mathcal{C}_0$ hold, where $\mathcal{C}_1 := F^{-1}(y^{(1)})$ and $\mathcal{C}_0 := F^{-1}(y^{(0)})$ denote the fibers defined by $y^{(1)}$ and $y^{(0)}$ respectively. Since $\mathcal{L}$ separates the points of $V_{y^{(1)}}$ it follows that $\mathcal{L}$ is a primitive element of the field extension $\mathbb{K}(S) \hookrightarrow \mathbb{K}(\mathcal{C})$.

The algorithm that computes all the points of $V_{y^{(0)}}$ may be divided in three main parts, which will be considered in Sections 5.1, 5.2 and 5.3 below. In the first step, we compute the minimal primitive polynomial $m_S(S, T)$ of $\mathcal{L}$ in the field extension $\mathbb{K}(S) \hookrightarrow \mathbb{K}(\mathcal{C})$. For this purpose, we apply a Newton–Hensel iteration to the rational point $x^{(1)}$ in order to obtain the vector of power series $\Psi \in \mathbb{K}[\![S-1]\!]^n$ which parametrizes the branch

of $\mathcal{C}$ passing through $(x^{(1)}, y^{(1)})$, truncated up to a suitable precision. It turns out that the least–degree nonzero polynomial $m_S(S,T) \in \mathbb{K}[S,T]$ which annihilates the power series $\mathcal{L}(\Psi)$ up to a certain precision equals the minimal polynomial of the coordinate function defined by $\mathcal{L}$ in the field extension $\mathbb{K}(S) \hookrightarrow \mathbb{K}(\mathcal{C})$ (see Lemma 5.1 below).

In the second step we extend the computation of the minimal polynomial $m_S(S,T)$ of $\mathcal{L}$ in the field extension $\mathbb{K}(S) \hookrightarrow \mathbb{K}(\mathcal{C})$ to the computation of a geometric solution of the curve $\mathcal{C}$, applying the algorithm underlying Lemma 3.1. Finally, in the third step we find the coordinates of the q–rational points of $V_{y^{(0)}}$. In order to do this, we first obtain a geometric solution $m_S(1,T), w_1(T), \ldots, w_n(T)$ of the zero–dimensional variety $\mathcal{C}_0 = F^{-1}(y^{(0)})$, substituting 0 for S in the polynomials which form the geometric solution of $\mathcal{C}$ computed in the previous step. Then we easily obtain the q–rational points of $\mathcal{C}_0$ among the points $x := (x_1, \ldots, x_n) \in \mathbb{A}^n$ satisfying the following equations:

$$m_S(1,T) = 0,\ T^{|\mathbb{K}|} - T = 0,\ x_i = w_i(T) \quad (1 \leq i \leq n).$$

The whole algorithm for computing the q–rational points of $F^{-1}(y^{(0)})$ may be briefly sketched as follows:

Algorithm 5.1.

1. *Choose the coefficients of a vector* $(\lambda, x^{(1)}) \in \mathbb{K}^{2n}$ *at random.*
2. *Set* $G(S,X) := F(X) - y^{(1)} - (S-1)(y^{(1)} - y^{(0)})$. *Compute the Newton–Hensel operator* $N_G(X) := X - J_F^{-1}(X)G(S,X)$.
3. *Compute* $\kappa := \lceil \log_2(2D\delta + 1) \rceil$ *iterations of the Newton–Hensel iterator* N_G *applied to* $x^{(1)}$. *Let* Ψ_κ *be the resulting vector of power series truncated up to order* $2D\delta + 1$.
4. *Find the least–degree nonzero polynomial* $m_S \in \mathbb{K}[S,T]$ *such that* $m_S(S, \mathcal{L}(\Psi_\kappa)) \equiv 0 \mod (S-1)^{2D\delta+1}$ *holds. This is the minimal polynomial of* $\mathcal{L}$ *in* $\mathbb{K}(S) \hookrightarrow \mathbb{K}(\mathcal{C})$ *(see Lemma 5.1).*
5. *Obtain a geometric solution of the curve* $\mathcal{C}$ *applying the proof of Lemma 3.1.*
6. *Substitute 0 for* S *in the polynomials which form the geometric solution of* $\mathcal{C}$ *computed in the previous step. The univariate polynomials obtained form a complete description of* $\mathcal{C}_0$ *(eventually including multiplicities).*
7. *Clean multiplicities of the polynomials computed in the previous step to obtain a geometric solution* $m_0, w_1, \ldots, w_n \in \mathbb{K}[T]$ *of the variety* $\mathcal{C}_0$.
8. *Compute* $h := \gcd(m_0, T^{|\mathbb{K}|} - T)$ *and the roots* $\alpha^{(1)}, \ldots, \alpha^{(M)}$ *of* h *in* $\mathbb{K}$.
9. *Compute the q–rational points of* $\mathcal{C}_0 = F^{-1}(y^{(0)})$ *as the intersection* $\big\{\big(w_1(\alpha^{(i)}), \ldots, w_n(\alpha^{(i)})\big); 1 \leq i \leq n\big\} \cap \mathbb{F}_q^n$.

We observe that, in order to determine the size of the field $\mathbb{K}$ and to execute steps (3)–(4), the values D and δ are required. Although in some cases these values are known *a priori*, we cannot in general assume

that they are given. Concerning the determination of the field $\mathbb{K}$, from the complexity point of view we may simply estimate D and δ by d^n and proceed, since the the cost of our procedure depends quasi–linearly on the value $\log_2(D\delta)$. On the other hand, for the execution of steps (3)–(5), the value $N := 2D\delta$ can be found by a process which, roughly speaking, starts with the value $N = 2$, and incrementally doubles the value N until the output of steps (3)–(5) is a geometric solution of the curve $\mathcal{C}$. The efficiency of this process relies on the fact that one can efficiently check if a candidate to be a geometric solution of a given irreducible variety is actually a geometric solution. Such a modification would only contribute with logarithmic factors to the asymptotic cost of our procedure.

5.1. The computation of the polynomial m_S. We consider the factorization of $m_S(S,T)$ in the ring $\mathbb{K}[\![S-1]\!][T]$, where $\mathbb{K}[\![S-1]\!]$ denotes the power series ring in $S-1$. From the fact that $m_S(1,T)$ is separable of degree D we conclude that there exist D distinct power series $\sigma^{(1)},\ldots,\sigma^{(D)} \in \mathbb{K}[\![S-1]\!]$ such that the monic version $\widetilde{m}_S$ of $m_S(S,T)$ can be factored as $\widetilde{m}_S = \prod_{i=1}^{D}(T-\sigma^{(i)})$. Furthermore, $\widetilde{m}_S(1,T)$ can be factored as $\widetilde{m}_S(1,T) = \prod_{i=1}^{D}(T-\sigma^{(i)}(1))$, where $\sigma^{(i)}(1)$ represents the constant term of $\sigma^{(i)}$ for $1 \leq i \leq D$.

Since $\widetilde{m}_S(1,T)$ is the minimal polynomial of the linear form $\mathcal{L}$ in the $\mathbb{K}$–algebra extension $\mathbb{K} \hookrightarrow \mathbb{K}[V_{y^{(1)}}]$, if we write $V_{y^{(1)}} := \{P_1,\ldots,P_D\}$ we have that $\widetilde{m}_S(1,T) = \prod_{i=1}^{D}(T-\mathcal{L}(P_i))$. Given that $(x^{(1)},y^{(1)})$ belongs to the fiber $V_{y^{(1)}}$, there exists a power series $\sigma^{(i)}$ such that $\mathcal{L}(x^{(1)}) = \sigma^{(i)}(1)$. In order to simplifify notations, we shall simply write σ instead of $\sigma^{(i)}$.

The algorithm that computes the polynomial $m_S(S,T)$ starts computing the power series σ truncated up to order $N+1$, where $N := 2D\delta$. Let $\sigma_N \in \overline{\mathbb{F}}_q[S]$ be the polynomial of degree at most N congruent to σ modulo $(S-1)^{N+1}$. Our next result shows that the polynomial $m_S(S,T)$ we want to compute can be obtained as the solution of a suitable congruence equation involving σ_N.

LEMMA 5.1. *Let $g \in \mathbb{K}[S,T]$ be a polynomial with $\deg_S g \leq \delta$ and $\deg_T g \leq D$ satisfying the following congruence*

$$g(S,\sigma_N) \equiv 0 \mod (S-1)^{N+1}. \tag{5.2}$$

Then m_S divides g in $\mathbb{K}[S,T]$.

Proof. Let $g \in \mathbb{K}[S,T]$ be a solution of (5.2) satisfying the conditions on the degree of the statement of the lemma. The resultant $h \in \mathbb{K}[S]$ of g and m_S with respect to T has degree at most N and belongs to the ideal generated by m_S and g. Since $m_S(S,\sigma_N)$ and $g(S,\sigma_N)$ are congruent to 0 modulo $(S-1)^{N+1}$ by hypothesis, we deduce that $h(S) \equiv 0 \mod (S-1)^{N+1}$ holds. Therefore, the fact that $\deg h \leq N$ and $h(S) \equiv 0 \mod (S-1)^{N+1}$ holds imply $h = 0$. In particular, we derive the existence of a common factor of m_S and g in $\mathbb{K}(S)[T]$. Finally, taking into account the

irreducibility of m_S in $\mathbb{K}(S)[T]$ and the Gauss lemma, we easily deduce the statement of the lemma. □

From Lemma 5.1 we conclude that m_S can be characterized as the nonzero solution of (5.2) of minimal (total) degree.

In order to find the least–degree nonzero solution of (5.2), we shall interpret (5.2) as a problem of Hermite–Padé approximation. Indeed, finding a nonzero solution of the congruence equation (5.2) is equivalent to finding $g_0, \ldots, g_D \in \mathbb{K}[S]$ with $\deg g_j \leq \delta$ for $0 \leq j \leq D$ such that the following congruence equation holds:

$$g_0(S) + g_1(S)\sigma_N + \cdots + g_D(S)\sigma_N^D \equiv 0 \mod (S-1)^{N+1}. \tag{5.3}$$

We shall solve (5.3) applying an algorithm due to [8], which is based on fast linear–algebra algorithms for matrices of fixed displacement rank (cf. [7], [33]). This requires the computation of the successive powers $\sigma_N, \ldots, \sigma_N^D$ of the power series σ truncated up to order $N+1$.

The computation of σ_N is based on a multivariate Newton iteration over the power series ring $\mathbb{K}[\![S-1]\!]$, which we now describe. Substituting 1 for S in the polynomials defining the ideal $I_{\mathcal{C}}$ associated to the curve $\mathcal{C}$ of (5.1), we obtain the system $y^{(1)} = F(X)$. Since $y^{(1)}$ is a lifting point of π, it follows that none of the points of $V_{y^{(1)}}$ annihilate the denominator $Q_i(X)$ of the rational function F_i for $1 \leq i \leq n$. Furthermore, from, e.g., [10, Lemma 2.1] we conclude that none of the points of $V_{y^{(1)}}$ annihilate the determinant of the Jacobian matrix $J_F := (\partial F_i)/(\partial X_j)_{1\leq i,j\leq n}$. In particular, $\det J_F(x^{(1)}) \neq 0$ holds.

Observe that the curve $\mathcal{C}$ of (5.1) is locally defined in a neighborhood of each point of $V_{y^{(1)}}$ by the equations $F_i(X) = y_i^{(1)} + (S-1)(y_i^{(1)} - y_i^{(0)})$ $(1 \leq i \leq n)$. Therefore, in order to compute the truncated power series σ_N we consider the Newton–Hensel operator N_G associated to the vector of rational functions $G(S,X) := F(X) - y^{(1)} - (S-1)(y^{(1)} - y^{(0)})$, namely,

$$N_G(X) := X - J_F^{-1}(X)G(S,X).$$

Let $N_G^{(k)}$ denote the k–fold iteration of N_G and define $\Psi_k := N_G^{(k)}(x^{(1)}) \in \mathbb{K}[\![S-1]\!]^n$ for $k \geq 0$. Then it is well known that the following congruence relation holds:

$$G(S, \Psi_k) \equiv 0 \mod (S-1)^{2^k}. \tag{5.4}$$

Since $Q_i(\Psi_k)(1) \neq 0$ holds for $1 \leq i \leq n$, from (5.4) we deduce that

$$P_i(\Psi_k) \equiv Q_i(\Psi_k)\big(y_i^{(1)} + (S-1)(y_i^{(1)} - y_i^{(0)})\big) \mod (S-1)^{2^k}. \tag{5.5}$$

Since the polynomial $m_S(S, \mathcal{L}(X))$ belongs to the ideal of $\mathbb{K}[S,X]_{Q^\infty}$ generated by $P_i(X) - Q_i(X)\big(y_i^{(1)} + (S-1)(y_i^{(1)} - y_i^{(0)})\big)$ $(1 \leq i \leq n)$, from (5.5) we conclude that

$$m_S(S, \mathcal{L}(\Psi_k)) \equiv 0 \mod (S-1)^{2^k}.$$

From the identity $\mathcal{L}(\Psi_k)(1) = \mathcal{L}(x^{(1)})$ we deduce that $\mathcal{L}(\Psi_k) \equiv \sigma \bmod (S-1)^{2^k}$ holds. Hence, we obtain σ_N as the power series $\mathcal{L}(\Psi_\kappa)$ with $\kappa := \lceil \log_2(N+1) \rceil$ truncated up to order $N+1$. From σ_N we easily compute the powers $\sigma_N^2, \ldots, \sigma_N^D$ by successive multiplication and truncation.

We may summarize the algorithm underlying the above considerations as follows:

ALGORITHM 5.2 (Computation of the powers of σ_N).

1. *Set* $\kappa := \lceil \log_2(N+1) \rceil$ *and* $\Psi_0 := x^{(1)}$.
2. *Compute* $\Psi_{k+1} := N_G(\Psi_k) \bmod (S-1)^{2^{k+1}}$ *for* $0 \leq k \leq \kappa - 1$.
3. *Compute* $\sigma_N := \mathcal{L}(\Psi_\kappa) \bmod (S-1)^{N+1}$.
4. *Compute* $\sigma_N^{j+1} := \sigma_N \cdot \sigma_N^j \bmod (S-1)^{N+1}$ *for* $1 \leq j \leq D-1$.

The following proposition provides a complexity estimate of the procedure just described:

PROPOSITION 5.1. *If the rational functions* $F_1, \ldots, F_n$ *are evaluated with* T *operations in* $\mathbb{F}_q$, *the powers* $\sigma_N, \ldots, \sigma_N^D$, *truncated up to order* $N+1$, *can be deterministically computed with* $O\big((\mathsf{T} + n^{1+\omega})\mathsf{M}(D\delta)\big)$ *arithmetic operations in* $\mathbb{K}$.

Proof. First, from the Baur–Strassen theorem [5] it follows that the entries of J_F can be computed with $O(\mathsf{T})$ arithmetic operations. Then, the determinant and adjoint matrix of J_F can be evaluated with $O(\mathsf{T} + n^{1+\omega})$ arithmetic operations (see, e.g., [7]).

In order to compute the $(k+1)$th iteration $\Psi_{k+1} := N_G(\Psi_k)$ from Ψ_k, we compute the inverse matrix $J_F^{-1}(\Psi_k)$ as the multiplication $J_F^{-1}(\Psi_k) = \det J_F(\Psi_k)^{-1} \cdot Adj\big(J_F(\Psi_k)\big)$ of the reciprocal of the (truncated) power series $\det J_F(\Psi_k)$ by each entry of the adjoint matrix $Adj\big(J_F(\Psi_k)\big)$. Using fast power series inversion we can compute $\det J_F(\Psi_k)^{-1}$ with $O\big((\mathsf{T} + n^{1+\omega})\mathsf{M}(2^k)\big)$ arithmetic operations in $\mathbb{K}$ (see, e.g., [18], [7]). With a similar cost we compute the evaluation $Adj\big(J_F(\Psi_k)\big)$ of the adjoint matrix of J_F at Ψ_k and the product $\det J_F(\Psi_k)^{-1} \cdot Adj\big(J_F(\Psi_k)\big)$.

Thus, the computation of Ψ_k for every $2 \leq k \leq \kappa$ requires

$$O\Big((\mathsf{T} + n^{1+\omega}) \sum_{k=0}^{\kappa-1} \mathsf{M}(2^k)\Big) = O\big((\mathsf{T} + n^{1+\omega})\mathsf{M}(D\delta)\big)$$

arithmetic operations in $\mathbb{K}$. The remaining steps do not change the overall asymptotic cost. ∎

Next we discuss how we can solve the Hermite–Padé approximation problem (5.3). This is represented by a linear system with $N+1$ equations and $O(D\delta)$ unknowns, given by the coefficients of the solution $g \in \mathbb{K}[S,T]$ of (5.2). Best general–purpose algorithms solving a system of size $O(D\delta \times D\delta)$ require $O((D\delta)^\omega)$ arithmetic operations [7]. However, in this case we profit from the structure of (5.3): it turns out that for a suitable ordering of the unknowns, the matrix M of the system (5.3) is a block–Toeplitz matrix (see, e.g., [12, Lemma 4.3], [8, Lemma 6]). This allows us to solve (5.3)

using the theory of matrices of fixed displacement rank (cf. [7], [33]). We shall apply the algorithm of [8], which is aimed at solving linear systems defined by matrices of "large" displacement rank.

Further, as we are interested in the least–degree nonzero solution $g \in \mathbb{K}[S,T]$ of (5.3), we combine [8] with a strategy of binary search as in, e.g., [7, Algorithm 8.2]. Fix $\rho \leq \delta$. From [8, Corollary 1] it follows that, if there exist nonzero solutions $g \in \mathbb{K}[S,T]$ of (5.3) with $\deg_S g \leq \rho$, then one such solution can be computed with $O\big(D^{\omega-1}\mathsf{M}(D\rho)\log(D\rho)\big)$ arithmetic operations in $\mathbb{K}$ and error probability at most $2(D\rho)^2/|\mathbb{K}|$. Therefore, applying a binary search we can determine the least–degree solution of (5.3) with at most $\lceil \log \delta \rceil$ such steps. As a consequence, we have the following result:

PROPOSITION 5.2. *Suppose that we are given the dense representation of the powers $\sigma_N, \ldots, \sigma_N^D$, as provided by the algorithm underlying Proposition 5.1. Then the minimal polynomial $m_S \in \mathbb{K}[S,T]$ can be computed with $O\big(D^{\omega-1}\mathsf{M}(D\delta)\log^2\delta\big)$ operations in $\mathbb{K}$ and error probability at most $2(D\delta)^2\log\delta/|\mathbb{K}|$.*

Combining Propositions 5.1 and 5.2 we obtain an algorithm computing the minimal polynomial m_S from the rational functions $F_1, \ldots, F_n$:

PROPOSITION 5.3. *The polynomial $m_S \in \mathbb{K}[S,T]$ can be computed with $O\big((\mathsf{T} + n^{1+\omega} + D^{\omega-1}\log^2\delta)\mathsf{M}(D\delta)\big)$ operations in $\mathbb{K}$ and error probability at most $2(D\delta)^2\log\delta/|\mathbb{K}|$.*

5.2. A geometric solution of $\mathcal{C}$. Our next task consists in extending the algorithm of the previous section to an algorithm computing a geometric solution of the curve $\mathcal{C}$ defined in (5.1). Let $\Lambda := (\Lambda_1, \ldots, \Lambda_n)$ be a vector of new indeterminates and consider the projection map $\pi_\Lambda : \mathbb{A}^n \times \mathcal{C} \to \mathbb{A}^n \times \mathbb{A}^1$ defined by $\pi_\Lambda(\lambda, s, x) := (\lambda, s)$. Since π_S is a dominant morphism, so is π_Λ and $\mathbb{K}(\Lambda, S) \hookrightarrow \mathbb{K}(\Lambda) \otimes_{\mathbb{K}} \mathbb{K}(\mathcal{C})$ is an algebraic field extension. The minimal polynomial $m_\Lambda \in \mathbb{K}[\Lambda, S, T]$ of the linear form $\mathcal{L}_\Lambda := \Lambda \cdot X$ in $\mathbb{K}(\Lambda, S) \hookrightarrow \mathbb{K}(\Lambda) \otimes_{\mathbb{K}} \mathbb{K}(\mathcal{C})$ is a separable element of $\mathbb{K}[\Lambda, S][T]$ satisfying the degree bounds $\deg_T m_\Lambda \leq D$, $\deg_S m_\Lambda \leq \delta$ and $\deg_\Lambda m_\Lambda \leq \delta$ (see, e.g., [11, Proposition 6.1] and [39, Proposition 1]). Notice that substituting λ for Λ we have $m_\Lambda(\lambda, S, T) = m_S(S, T)$.

Applying the algorithm underlying Proposition 5.3 to the linear form $\mathcal{L}_\Lambda$ we compute the minimal polynomial $m_\Lambda(\Lambda, S, T)$ with $O\big((\mathsf{T} + n^{1+\omega} + D^{\omega-1}\log^2\delta)\ \mathsf{M}(D\delta)\big)$ arithmetic operations in $\mathbb{K}(\Lambda)$. Therefore, by Lemma 3.1 we obtain the following result:

PROPOSITION 5.4. *Suppose that the coefficients of the linear form $\mathcal{L}$ are randomly chosen in $\mathbb{K}$. Then we can compute a geometric solution of $\mathcal{C}$ with $O\big((\mathsf{T} + n^{1+\omega} + D^{\omega-1}\log^2\delta)n\mathsf{M}(D\delta)\big)$ operations in $\mathbb{K}$. Furthermore, the algorithm output the right result with error probability at most $9\mathcal{D}\delta\log\delta/|\mathbb{K}|$, where $\mathcal{D} := D(\delta+2)^{\log\delta}$.*

Proof. As explained in the proof of Lemma 3.1, we apply the algorithm for the computation of the minimal polynomial m_S of Proposition 5.3 to

the generic linear form $\mathcal{L}_\Lambda$, truncating each intermediate result up to order $(\Lambda-\lambda)^2$. Therefore, from Lemma 3.1 and Proposition 5.3 we easily deduce the complexity estimate of the proposition.

In order to estimate the error probability of the algorithm we have to estimate the probability of failure of the choice of the vector of coefficients of the linear form $\mathcal{L}$. Recall that the application of Lemma 3.1 requires that the vector of coefficients $\lambda := (\lambda_1, \ldots, \lambda_n)$ of the linear form $\mathcal{L}$ does not annihilate any denominator in $\mathbb{K}[\Lambda]$ of any intermediate result of the algorithm computing the minimal polynomial m_Λ.

The algorithm for obtaining the polynomial m_Λ consists of two steps: the computation of the first D powers of $\sigma_N^{(\Lambda)} := \mathcal{L}_\Lambda(\Psi_k)$, which is considered in Proposition 5.1, and the solution of the Hermite–Padé approximation problem (5.3), which is considered in Proposition 5.2. From Algorithm 5.2 we conclude that the computation of the powers $\sigma_N^{(\Lambda)}, \ldots, (\sigma_N^{(\Lambda)})^D$ does not require any division by a nonconstant polynomial of $\mathbb{K}[\Lambda]$.

Next, we analyze the divisions necessary to solve the Hermite–Padé approximation problem (5.3), which is solved applying an algorithm of [8]. This algorithm is an adaptation of Kaltofen's *Leading Principal Inverse* algorithm ([26], [27]). Kaltofen's algorithm performs a recursive reduction of the computation of the inverse of a "generic–rank–profile" square input matrix

$$A = \begin{pmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{pmatrix}$$

to that of the leading principal submatrix $A_{1,1}$ and its Schur complement $\Delta := A_{2,2} - A_{2,1}A_{1,1}^{-1}A_{1,2}$. The divisions which arise during the execution of this recursive step are related to the computation of $A_{1,1}^{-1}$ and Δ^{-1} and a routine of "compression" (cf. [7, Problem 2.2.11.c]) of the generators of matrices which are obtained as certain products involving $A_{1,1}^{-1}$, Δ^{-1}, $A_{1,2}$ and $A_{2,1}$. The latter in turn requires the computation of the inverses of certain submatrices of the products under consideration.

Each entry of the matrix M of the linear system (5.3) is a coefficient of a power $(\sigma_N^{(\Lambda)})^j$, which is therefore a polynomial in Λ of degree at most $j \leq D$. Since the generic–rank–profile matrix A is obtained by multiplying M with suitable matrices with entries in $\mathbb{K}$, we conclude that the entries of $A_{1,1}$ are polynomials of $\mathbb{K}[\Lambda]$ of degree at most D, while the numerators and denominators of the entries of Δ^{-1} are polynomials of $\mathbb{K}[\Lambda]$ of degree at most $D(\delta+2)$. Therefore, by a simple recursive argument we see that the numerators and denominators of all leading principal submatrices and Schur complements which are inverted during the algorithm have degrees bounded by $\mathcal{D} := D(\delta+2)^{\log\delta}$. This in turn implies that the denominators arising during the compression routine have degrees bounded by $3\mathcal{D}\delta$.

Finally, taking into account that the algorithm of [8] consists of at most $\lfloor \log\delta \rfloor$ recursive steps, and that each recursive step requires the inversion

of at most 4 matrices, we conclude that the product of all the denominators arising during the algorithm has degree bounded by $8\mathcal{D}\delta\log\delta$. By Corollary 2.1, it follows that a random choice of $\lambda := (\lambda_1, \ldots, \lambda_n)$ does not vanish any denominator of the algorithm computing the minimal polynomial m_Λ with error probability at most $8\mathcal{D}\delta\log\delta/|\mathbb{K}|$. Putting together this estimate and the error probability $2(D\delta)^2\log\delta/|\mathbb{K}|$ of the algorithm underlying Proposition 5.3 we deduce the statement of the proposition. □

5.3. Computation of the points of $F^{-1}(y^{(0)})\cap\mathbb{F}_q^n$. In this section we show how to find the solutions in $\mathbb{F}_q^n$ of the system $F(X) = y^{(0)}$.

Assume that we are given a geometric solution defined over $\mathbb{K}$ of the curve $\mathcal{C}$ defined in (5.1), as provided by the algorithm of Proposition 5.4. This geometric solution consists of a linear form $\mathcal{L}\in\mathbb{K}[X]$, the minimal polynomial $m_S\in\mathbb{K}[S,T]$ of $\mathcal{L}$ in the algebraic field extension $\mathbb{K}(S)\hookrightarrow\mathbb{K}(\mathcal{C})$ and the parametrizations $(\partial m_S/\partial T)X_j - v_j(S,T)$ $(1\le j\le n)$ of the variables $X_1,\ldots,X_n$ by the zeros of m_S.

Let $\pi_S^{-1}(0) = \{0\}\times\mathcal{C}_0$, where $\mathcal{C}_0 := F^{-1}(y^{(0)})$. Since $\mathcal{L}$ separates the points of $\pi_S^{-1}(0)$, from a geometric solution of $\mathcal{C}$ we obtain a geometric solution of $\mathcal{C}_0$. Indeed, substituting 0 for S in $m_S, v_1, \ldots, v_n$ we obtain polynomials $m_S(0,T), v_1(0,T), \ldots, v_n(0,T)\in\mathbb{K}[T]$ which represent a complete description of the fiber $\mathcal{C}_0$, i.e., we have the identity

$$\mathcal{C}_0 = \{x\in\mathbb{A}^n; m_S(0,\mathcal{L}(x)) = 0, m_S'(0,\mathcal{L}(x))x_j = v_j(0,\mathcal{L}(x))\ (1\le j\le n)\},$$

where $m_S'(0,T) := \frac{\partial m_S}{\partial T}(0,T)$. Nevertheless, the polynomials $m_S(0,T)$, $v_1(0,T),\ldots,v_n(0,T)\in\mathbb{K}[T]$ do not necessarily form a geometric solution of $\mathcal{C}_0$, because the polynomial $m_S(0,T)$ might have multiple factors. In such a case, it is easy to see that the multiple factors of $m_S(0,T)$ are also factors of $v_1(0,T),\ldots,v_n(0,T)$ (see, e.g., [21, §6.5]). To remove these multiple factors, we proceed in the following way: first, we compute

$$a(T) := \gcd\big(m_S(0,T), m_S'(0,T)\big),\quad m_0(T) := \frac{m_S(0,T)}{a(T)},$$

which yield the square–free representation $m_0(T)$ of $m_S(0,T)$. Next, given that $a(T)$ divides $v_j(0,T)$ for $1\le j\le n$, we obtain polynomials

$$\frac{m_S'(0,T)}{a(T)}X_j - \frac{v_j(0,T)}{a(T)}\quad (1\le j\le n),$$

which vanish on the points of $\mathcal{C}_0$. Finally, since m_0 and $m_S'(0,T)/a(T)$ have no common factors in $\mathbb{K}[T]$, we invert $m_S'(0,T)/a(T)$ modulo $m_0(T)$ and obtain parametrizations $X_j - w_j(T)$ $(1\le j\le n)$ of the coordinates of the points of $\mathcal{C}_0$ by the zeros of $m_0(T)$ which are better suited for our purposes. In the next lemma we state the cost of this procedure:

LEMMA 5.2. *Given a geometric solution of the curve $\mathcal{C}$, as provided by the algorithm underlying Proposition 5.4, we can deterministically compute a geometric solution $m_0(T), X_1 - w_1(T), \ldots, X_n - w_n(T)$ of the zero dimensional variety $\mathcal{C}_0$ with $O(n\delta\mathsf{M}(D))$ operations in $\mathbb{K}$.*

Proof. The dense representation of the polynomials $m_S(0,T), v_1(0,T), \ldots, v_n(0,T)$ can be obtained from the dense representation of the polynomials $m_S(S,T), v_1(S,T), \ldots, v_n(S,T)$ with $O(nD\delta)$ arithmetic operations in $\mathbb{K}$. The remaining computations are $O(n)$ multiplications, greatest common divisors and a modular inversion of univariate polynomials, whose degrees are less than or equal to D, which contribute with $O\big(n\mathsf{M}(D)\big)$ additional arithmetic operations in $\mathbb{K}$. □

Finally, we compute the $\mathbb{K}$–rational points of $\mathcal{C}_0$, which in particular yield the solutions in $\mathbb{F}_q^n$ of $F(X) = y^{(0)}$. For this purpose, set

$$h := \gcd(m_0, T^{|\mathbb{K}|} - T) \in \mathbb{K}[T].$$

Following, e.g., [18, Corollary 14.16], we have that h can be computed with $O\big(\mathsf{M}(D)\log|\mathbb{K}|\big)$ arithmetic operations in $\mathbb{K}$. Since h factors into linear factors, its factorization can be computed with $O\big(\mathsf{M}(D)\log(\mu D|\mathbb{K}|)\big)$ arithmetic operations in $\mathbb{K}$ and error probability at most $1/\mu$ (see, e.g., [18, Theorem 14.9]).

Observe that the roots of h are the values $\mathcal{L}(P)$ resulting from the evaluation of $\mathcal{L}$ in the points $P \in \mathcal{C}_0 \cap \mathbb{K}^n$. In particular, $\mathcal{L}(x) \in \mathbb{K}$ is a root of h for every point $x \in \mathcal{C}_0 \cap \mathbb{K}^n$. Thus, if we substitute the roots α of h for T in the polynomials $w_j(T)$ $(1 \leq j \leq n)$, we obtain all the points of $\mathcal{C}_0 \cap \mathbb{F}_q^n$ as the points $\big(w_1(\alpha), \ldots, w_n(\alpha)\big) \in \mathbb{F}_q^n$ with $h(\alpha) = 0$. Since such substitutions require $O(nD)$ additional arithmetic operations in $\mathbb{K}$, we have the following result:

Proposition 5.5. *Given a geometric solution of the zero–dimensional variety $\mathcal{C}_0 = F^{-1}(y^{(0)})$, as provided by the algorithm underlying Lemma 5.2, we can compute the set $\mathcal{C}_0 \cap \mathbb{F}_q^n$ with $O\big(\mathsf{M}(D)(n + \log(\mu D|\mathbb{K}|))\big)$ arithmetic operations in $\mathbb{K}$ and error probability at most $1/\mu$.*

Putting together Propositions 5.4 and 5.5 we obtain our main result:

Theorem 5.3. *The solutions in $\mathbb{F}_q^n$ of the input system $F(X) = y^{(0)}$ can be computed with*

$$O\Big(\big((\mathsf{T}+n^{1+\omega}+D^{\omega-1}\log^2\delta)\mathsf{M}(D\delta)+\mathsf{M}(D)\mathsf{M}(\log q+\log^2\delta)\big)n\mathsf{M}(\log q+\log^2\delta)\Big)$$

bit operations and error probability at most $1/4$.

Proof. Let $\mathbb{K}$ be a field extension of $\mathbb{F}_q$ of cardinality greater than $128\mathcal{D}\delta\log\delta$, where $\mathcal{D} := D(\delta+2)^{\log\delta}$. Choose randomly a point $(\lambda, x^{(1)}) \in \mathbb{K}^{2n}$ and set $\mathcal{L} := \lambda \cdot X \in \mathbb{K}[X]$. Suppose that $(\lambda, x^{(1)}) \in \mathbb{K}^{2n}$ satisfies the following conditions:

1. $y^{(1)} := F(x^{(1)})$ is a lifting point of $\pi : V \to \mathbb{A}^n$,
2. the curve $\mathcal{C} \subset \mathbb{A}^{n+1}$ defined in (5.1) is absolutely irreducible,
3. the linear form $\mathcal{L}$ separates the fibers $V_{y^{(1)}}$ and $V_{y^{(0)}}$,
4. λ does not annihilate any denominator of the algorithm computing the minimal polynomial m_Λ underlying Proposition 5.3.

Applying the algorithm of Proposition 5.4 we obtain a geometric solution of the curve $\mathcal{C}$ with $O\big((\mathsf{T} + n^{1+\omega} + D^{\omega-1}\log^2\delta)n\mathsf{M}(D\delta)\big)$ operations in $\mathbb{K}$.

Then we apply the algorithm of Proposition 5.5 in order to compute the solutions in $\mathbb{F}_q^n$ of $F(X) = y^{(0)}$ with $O\big(\mathsf{M}(D)(n + \log(D|\mathbb{K}|))\big)$ arithmetic operations in $\mathbb{K}$. Combining both complexity estimates, and taking into account that every arithmetic operation in $\mathbb{K}$ requires $O(\mathsf{M}(\log^2 \delta + \log q))$ bit operations, we deduce the estimate of the statement of the theorem.

By Corollary 4.1 we have that $(\lambda, x^{(1)})$ satisfies conditions (1)–(3) above with error probability at most $4d\delta^4/|\mathbb{K}| \leq 1/16$. Furthermore, from the proof of Proposition 5.3 we conclude that condition (4) is satisfied with error probability at most $1/16$. Finally, assuming that (1)–(4) hold, the algorithms underlying Propositions 5.3 and 5.5 output the right results with error probability at most $2(D\delta)^2 \log \delta/|\mathbb{K}| \leq 1/16$ and $1/16$ (fixing $\mu := 16$ in Proposition 5.5) respectively. This shows that the overall error probability is at most $1/4$ and finishes the proof of the theorem. □

We make a few remarks concerning Theorem 5.3. Observe that our algorithm has a cost of $O(n^{2+\omega} D^\omega \delta \log q + nD \log^2 q)$ bit operations, up to logarithmic terms. This improves and extends the algorithm of [12]. We have further contributed to the latter by providing estimates for the corresponding error probability.

A second remark concerns the behavior of our algorithm under the hypotheses of [41]. Recall that [41] requires $F : \mathbb{A}^n \to \mathbb{A}^n$ to be a polynomial map which is polynomially invertible, with inverse $G := (G_1, \ldots, G_n)$ with degrees polynomially bounded with respect to n and $d := \max_{1 \leq k \leq n} \deg F_k$. Under these hypotheses, the authors exhibit an algorithm which computes the inverse mapping G with a polynomial cost in T, n and d. Under these conditions, we have that the projection mapping $\pi : V \to \mathbb{A}^n$ has degree 1, i.e., the identity $D = 1$ holds. Furthermore, it is easy to see that the minimal polynomial $m_S(S, T)$ has degree bounded by $e := \max_{1 \leq k \leq n} \deg G_k$. Therefore, the algorithms underlying Propositions 5.4 and 5.5 have actually polynomial cost in T, n and e. This shows that the cost of our algorithm meets this polynomial bound assuming the strong hypotheses of [41].

6. Conclusions. Our complexity estimate may be roughly described as polynomial in the cost T of the evaluation of the input rational functions $F_1, \ldots, F_n$, the number of variables n, the logarithm $\log q$ of the cardinality of the field $\mathbb{F}_q$ and two geometric invariants: the degree D of the mapping F and the degree δ of the graph of F. In this sense, we see that the practical convenience of our algorithm, and the subsequent (in)security of cryptosystems based on polynomial or rational mappings over a finite field, essentially relies on these geometric invariants.

In worst case we have $D = \delta = \deg(F_1) \cdots \deg(F_n)$, which implies that our algorithm is exponential in the input size. Furthermore, adapting the arguments of [13] it is possible to prove that any *universal* algorithm solving $F(X) = y^{(0)}$ has necessarily cost $(\deg(F_1) \cdots \deg(F_n))^{\Omega(1)}$, showing thus the security of the corresponding cryptosystem with respect to *universal* decoding algorithms. Since a universal algorithm is one which does not

distinguish input systems according to geometric invariants and represents a model for the standard algorithms based on rewriting techniques, such as Gröbner basis algorithms, such cryptosystems are likely to be secure.

REFERENCES

[1] M. Alonso, E. Becker, M.-F. Roy, and T. Wörmann, *Zeroes, multiplicities and idempotents for zerodimensional systems*, in Proceedings of MEGA'94, Vol. **143** of Progr. Math., Boston, 1996, Birkhäuser, pp. 1–15.

[2] J. Balcázar, J. Díaz, and J. Gabarró, *Structural complexity I*, Vol. **11** of Monogr. Theoret. Comput. Sci. EATCS Ser., Springer, Berlin, 1988.

[3] M. Bardet, *Etude des systèmes algébriques surdétermines. Applications aux codes correcteurs et à la cryptographie*, PhD thesis, Université Paris **6**, 2004.

[4] M. Bardet, J.-C. Faugère, and B. Salvy, *Complexity of Gröbner basis computation for semi–regular overdetermined sequences over $\mathbb{F}_2$ with solutions in $\mathbb{F}_2$*. Rapport de Recherche INRIA RR–5049, `www.inria.fr/rrrt/rr-5049.html`, 2003.

[5] W. Baur and V. Strassen, *The complexity of partial derivatives*, Theoret. Comput. Sci., **22** (1983), pp. 317–330.

[6] E. Biham and A. Shamir, *Differential cryptanalysis of DES-like cryptosystems*, J. Cryptology, **4** (1991), pp. 3–72.

[7] D. Bini and V. Pan, *Polynomial and matrix computations*, Progress in Theoretical Computer Science, Birkhäuser, Boston, 1994.

[8] A. Bostan, C.-P. Jeannerod, and E. Schost, *Solving Toeplitz– and Vandermonde–like linear systems with large displacement rank*. To appear in Proceedings ISSAC'07, `http://www-sop.inria.fr/saga/POL`, 2007.

[9] P. Bürgisser, M. Clausen, and M. Shokrollahi, *Algebraic Complexity Theory*, Vol. **315** of Grundlehren Math. Wiss., Springer, Berlin, 1997.

[10] A. Cafure and G. Matera, *Fast computation of a rational point of a variety over a finite field*, Math. Comp., **75** (2006), pp. 2049–2085.

[11] ———, *Improved explicit estimates on the number of solutions of equations over a finite field*, Finite Fields Appl., **12** (2006), pp. 155–185.

[12] A. Cafure, G. Matera, and A. Waissbein, *Inverting bijective polynomial maps over finite fields*, in Proceedings of the 2006 Information Theory Workshop, ITW2006, G. Seroussi and A. Viola, eds., IEEE Information Theory Society, 2006, pp. 27–31.

[13] D. Castro, M. Giusti, J. Heintz, G. Matera, and L.M. Pardo, *The hardness of polynomial equation solving*, Found. Comput. Math., **3** (2003), pp. 347–420.

[14] N. Courtois, A. Klimov, J. Patarin, and A. Shamir, *Efficient algorithms for solving overdefined systems of multivariate polynomial equations*, in EUROCRYPT 2000, B. Preneel, ed., Vol. 1807 of Lecture Notes in Comput. Sci., Berlin, 2000, Springer, pp. 71–79.

[15] C. de Cannière, A. Biryukov, and B. Preneel, *An introduction to block cipher cryptanalysis*, Proc. IEEE, **94** (2006), pp. 346–356.

[16] J.-C. Faugère, *A new efficient algorithm for computing Gröbner bases without reduction to zero (F5)*, Proceedings ISSAC'02, T. Mora, ed., New York, 2002, ACM Press, pp. 75–83.

[17] S. Gao, *Factoring multivariate polynomials via partial differential equations*, Math. Comp., **72** (2003), pp. 801–822.

[18] J. von zur Gathen and J. Gerhard, *Modern computer algebra*, Cambridge Univ. Press, Cambridge, 1999.

[19] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP–Completeness*, Freeman, San Francisco, 1979.

[20] M. Giusti, K. Hägele, J. Heintz, J.E. Morais, J.L. Montaña, and L.M. Pardo, *Lower bounds for Diophantine approximation*, J. Pure Appl. Algebra, **117, 118** (1997), pp. 277–317.

[21] M. Giusti, G. Lecerf, and B. Salvy, *A Gröbner free alternative for polynomial system solving*, J. Complexity, 17 (2001), pp. 154–211.
[22] J. Heintz, *Definability and fast quantifier elimination in algebraically closed fields*, Theoret. Comput. Sci., **24** (1983), pp. 239–277.
[23] M.-D. Huang and Y.-C. Wong, *Solvability of systems of polynomial congruences modulo a large prime*, Comput. Complexity, **8** (1999), pp. 227–257.
[24] H. Imai and T. Matsumoto, *Public quadratic polynomial-tuples for efficient signature-verification and message-encryption*, in Advances in Cryptology - EUROCRYPT '88, C. Günther, ed., Vol. **330** of Lecture Notes in Comput. Sci., Berlin, 1988, Springer, pp. 419–453.
[25] J.-R. Joly, *Equations et variétés algébriques sur un corps fini*, Enseign. Math., **19** (1973), pp. 1–117.
[26] E. Kaltofen, *Asymptotically fast solution of Toeplitz–like singular linear systems*, in Proceedings ISSAC'94, J. von zur Gathen and M. Giesbrecht, eds., New York, 1994, ACM Press, pp. 297–304.
[27] ———, *Analysis of Coppersmith's block Wiedemann algorithm for the parallel solution of sparse linear systems*, Math. Comp., **64** (1995), pp. 777–806.
[28] ———, *Effective Noether irreducibility forms and applications*, J. Comput. System Sci., **50** (1995), pp. 274–295.
[29] A. Kipnis and A. Shamir, *Cryptanalysis of the HFE Public Key Cryptosystem by relinearization*, in Advances in Cryptology – CRYPTO'99, M. Wiener, ed., Vol. **1666** of Lecture Notes in Comput. Sci., Berlin, 1999, Springer, pp. 19–30.
[30] N. Koblitz, *Algebraic aspects of cryptography*, Vol. **3** of Algorithms Comput. Math., Springer, Berlin Heidelberg New York, corrected 2nd printing ed., 1999.
[31] G. Lecerf, *Improved dense multivariate polynomial factorization algorithms*, J. Symbolic Comput., **42** (2007), pp. 477–494.
[32] R. Lidl and H. Niederreiter, *Finite fields*, Addison–Wesley, Reading, Massachusetts, 1983.
[33] V. Pan, *Structured matrices and polynomials. Unified superfast algorithms*, Birkhäuser, Boston, 2001.
[34] L.M. Pardo and J. San Martín, *Deformation techniques to solve generalized Pham systems*, Theoret. Comput. Sci., **315** (2004), pp. 593–625.
[35] J. Patarin, *Cryptoanalysis of the Matsumoto and Imai Public Key Scheme of Eurocrypt'88*, in Advances in Cryptology - CRYPTO '95, D. Coppersmith, ed., Vol. **963** of Lecture Notes in Comput. Sci., Springer, 1995, pp. 248–261.
[36] ———, *Asymmetric cryptography with a hidden monomial*, in Advances in Cryptology - CRYPTO '96, N. Koblitz, ed., Vol. **1109** of Lecture Notes in Comput. Sci., Springer, 1996, pp. 45–60.
[37] F. Rouillier, *Solving zero–dimensional systems through rational univariate representation*, Appl. Algebra Engrg. Comm. Comput., **9** (1997), pp. 433–461.
[38] J. Savage, *Models of Computation. Exploring the Power of Computing*, Addison Wesley, Reading, Massachussets, 1998.
[39] E. Schost, *Computing parametric geometric resolutions*, Appl. Algebra Engrg. Comm. Comput., **13** (2003), pp. 349–393.
[40] I. Shafarevich, *Basic Algebraic Geometry: Varieties in Projective Space*, Springer, Berlin Heidelberg New York, 1994.
[41] C. Sturtivant and Z.-L. Zhang, *Efficiently inverting bijections given by straight line programs*, in Proceedings of the 31st Annual Symp. Found. Comput. Science, FOCS'90, Vol. **1**, IEEE Computer Society Press, 1990, pp. 327–334.
[42] L.-C. Wang and F.-H. Chang, *Tractable rational map cryptosystem*. Cryptology ePrint Archive, Report 2004/046, `http://eprint.iacr.org/2004/046/`, 2004.
[43] C. Wolf and B. Preneel, *Taxonomy of public key schemes based on the problem of multivariate quadratic equations*. Cryptology ePrint Archive, Report 2005/077, `http://eprint.iacr.org/2005/077/`, 2005.