Polynomial Equivalence Problems: Algorithmic and Theoretical Aspects

Jean-Charles Faugère[1] and Ludovic Perret[2]

[1] LIP6, 8 rue du Capitaine Scott, F-75015, France
Jean-Charles.Faugere@lip6.fr
[2] UCL, Crypto Group, Microelectronic Laboratory, Place du Levant, 3
Louvain-la-Neuve, B 1348, Belgium
ludovic.perret@uclouvain.be

Abstract. The Isomorphism of Polynomials (IP) [28], which is the main concern of this paper, originally corresponds to the problem of recovering the secret key of a C* scheme [26]. Besides, the security of various other schemes (signature, authentication [28], traitor tracing [5], ...) also depends on the practical hardness of IP. Due to its numerous applications, the Isomorphism of Polynomials is thus one of the most fundamental problems in multivariate cryptography. In this paper, we address two complementary aspects of IP, namely its theoretical and practical difficulty. We present an upper bound on the theoretical complexity of "IP-like" problems, i.e. a problem consisting in recovering a particular transformation between two sets of multivariate polynomials. We prove that these problems are not NP-Hard (provided that the polynomial hierarchy does not collapse). Concerning the practical aspect, we present a new algorithm for solving IP. In a nutshell, the idea is to generate a suitable algebraic system of equations whose zeroes correspond to a solution of IP. From a practical point of view, we employed a fast Gröbner basis algorithm, namely F_5 [17], for solving this system. This approach is efficient in practice and obliges to modify the current security criteria for IP. We have indeed broken several challenges proposed in literature [28, 29, 5]. For instance, we solved a challenge proposed by O. Billet and H. Gilbert at Asiacrypt'03 [5] in less than one second.

Keywords : Public-Key Cryptography, Cryptanalysis, Isomorphism of Polynomials (IP), Gröbner bases, F_5 algorithm

1 Introduction

Multivariate cryptography – which can be roughly defined as the cryptography using polynomials in several variables – offers a relatively wide spectrum of problems that can be used in public-key cryptography. The Isomorphism of Polynomials (IP) lies in this family [28]. Briefly, this problem consists in recovering a particular transformation between two sets of multivariate polynomials permitting to obtain one set from the other. It originally corresponds to the problem of recovering the secret key of a C^* scheme [26]. Besides, the security of several other schemes is directly based on the practical difficulty of IP, namely the authentication/signature schemes proposed by J. Patarin at Eurocrypt'96 [28], and the traitor tracing scheme described by O. Billet and H. Gilbert at Asiacrypt'03 [5]. We also mention that IP is in a certain manner related to the security of Sflash [13] – the signature scheme recommended by the European consortium Nessie for low-cost smart cards [27] – and can be alternatively viewed as the problem of detecting affine equivalence between S-Boxes [6]. All in all, one can consider the hardness of IP as one of the major issues in multivariate cryptography. The goal of this paper is to provide new insights on the theoretical and practical complexity of IP and some of its relevant variants.

1.1 Previous Work

To the best of our knowledge, the most significant results concerning IP are presented in [11], where an upper bound on the theoretical complexity of IP is given. Nevertheless, we point out that the proof provided is actually not complete. Anyway, the upper bound presented in that paper is original and general. It is indeed based on a group theoretic approach of IP and actually dedicated to "IP-like" problems. A new algorithm for solving IP, called "To and Fro", is also described in [11]. This algorithm is however devoted to special instances of IP, namely the ones corresponding to a public key of C^* [26]. Thus, it can not be used for solving generic instances of IP. This is not the case for the algorithm presented here. Besides, we present in Section 4 experimental results demonstrating that our algorithm outperforms the "To and Fro" method. Finally, we would like to mention a result due to W. Geiselmann, R. Steinwandt, and T. Beth [23]. In the context of C^*, they showed how to easily recover the affine parts of a solution of IP. A similar property also holds in the context of HFE [20]. Such a kind of result does not exist for generic instances of IP. Nevertheless, it means that in the cryptographic context we can focus our attention on the linear variant of IP, called 2PLE here.

1.2 Organization of the Paper and Main Results

The paper is organized as follows. We begin in Section 2 by introducing our notation and defining essential tools of our algorithm, namely varieties and Gröbner bases. A recent algorithm (i.e. F_5 [17]) for computing these bases is also succinctly described. Finally, we define more formally the Isomorphism of Polynomials (IP) and two of its variants, namely the Isomorphism of Polynomials with one Secret (IP1S) [28], and the linear variant of IP that we name 2PLE. In Section 3, we show that these problems are actually particular instances of a more general problem that we call Polynomial Equivalence (PE). This problem provides a formal definition of an "IP-like" problem. Using classical results of group theory, we conclude this section by providing an upper bound on the theoretical hardness of PE. A new algorithm for solving 2PLE is presented in Section 4. The idea is to generate a suitable polynomial system of equations whose zeroes correspond to a solution of IP. In order to construct this system, we also provide some specific properties of 2PLE. From a practical point of view, we used the most recent (and efficient) Gröbner basis algorithm, namely F_5 [17], for solving this system. It is difficult to obtain a complexity bound really reflecting the practical behavior of the F_5 algorithm. We therefore carried out experimental results illustrating the practical efficiency of our approach. We have indeed broken several challenges proposed in literature [28, 29, 5]. For instance, we solved a challenge proposed by O. Billet and H. Gilbert at Asiacrypt'03 [5] in less than one second.

2 Preliminaries

The notation used throughout this paper is the following. We denote by $\mathbb{F}_q$ the finite field with $q = p^r$ elements (p a prime, and $r \geq 1$), and by $\mathcal{M}_{n,u}(\mathbb{F}_q)$ the set of $n \times u$ matrices whose components are in $\mathbb{F}_q$. As usual, $GL_n(\mathbb{F}_q)$ represents the set of invertible matrices of $\mathcal{M}_{n,n}(\mathbb{F}_q)$, and $AGL_n(\mathbb{F}_q)$ denotes the cartesian product $GL_n(\mathbb{F}_q) \times \mathbb{F}_q^n$. Finally, let $\mathbf{x} = (x_1, \ldots, x_n)$, and $\mathbb{F}_q[\mathbf{x}] = \mathbb{F}_q[x_1, \ldots, x_n]$, be the polynomial ring in the n indeterminates $x_1, \ldots, x_n$ over $\mathbb{F}_q$. By convention, a boldfaced letter will always refer to a row vector.

2.1 Gröbner bases

We define now two essential notions of this paper, namely varieties and Gröbner bases. For a more thorough introduction to these tools, we refer to [1, 15].
Let $\mathbf{p} = (p_1, \ldots, p_s)$ be polynomials in $\mathbb{F}_q[\mathbf{x}]$. We shall call $\mathcal{I} = \langle p_1, \ldots, p_s \rangle = \left\{ \sum_{k=1}^{s} p_k u_k, u_1, \ldots, u_k \in \mathbb{F}_q[\mathbf{x}] \right\} \subset \mathbb{F}_q[\mathbf{x}]$ the *ideal generated* by $p_1, \ldots, p_s$, and denote by $V(\mathcal{I}) = \{\mathbf{z} \in \mathbb{F}_q^n : p_i(\mathbf{z}) = 0, \forall i, 1 \leq i \leq s\}$ the *variety associated* to $\mathcal{I}$. Gröbner bases provide a method for computing this variety. Informally, a Gröbner basis of an ideal $\mathcal{I}$ is a computable generator set of $\mathcal{I}$ with "good" algorithmic properties. These bases are defined with respect to *monomial orders.* Here, we will use the lexicographical (LEX) and degree reverse lexicographical (DRL) orders, which are defined as follows:

Definition 1. *Let $\alpha = (\alpha_1, \ldots, \alpha_n)$ and $\beta = (\beta_1, \ldots, \beta_n) \in \mathbb{N}^n$. Then:*
– $x_1^{\alpha_1} \cdots x_n^{\alpha_n} \prec_{LEX} x_1^{\beta_1} \cdots x_n^{\beta_n}$*, if the left-most nonzero entry of $\alpha - \beta$ is positive.*
– $x_1^{\alpha_1} \cdots x_n^{\alpha_n} \prec_{DRL} x_1^{\beta_1} \cdots x_n^{\beta_n}$*, if $\sum_{i=1}^n \alpha_i > \sum_{i=1}^n \beta_i$, or $\sum_{i=1}^n \alpha_i = \sum_{i=1}^n \beta_i$ and the right-most nonzero entry of $\alpha - \beta$ is negative.*

To define Gröbner bases, we need to introduce the following definitions.

Definition 2. *For any n-tuple $\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathbb{N}^n$, we denote by $\mathbf{x}^\alpha$ the* **monomial** $x_1^{\alpha_1} \cdots x_n^{\alpha_n}$. *We shall define the* total degree *of this monomial by the sum $\sum_{i=1}^n \alpha_i$. The* **leading monomial** *of a polynomial $p \in \mathbb{F}_q[\mathbf{x}]$ is the largest monomial (w.r.t some monomial ordering $\prec$) among the monomials of p. This leading monomial will be denoted by* $\mathrm{LM}(p, \prec)$. *The* **degree** *of p, denoted deg(p), is the total degree of* $\mathrm{LM}(p, \prec)$. *Finally, the* **maximal total degree** *of p is the maximal total degree of the monomials occurring in p.*

We are now in a position to define one of the main objects of this paper.

Definition 3. *A set of polynomials G is a* **Gröbner basis** *– w.r.t. a monomial ordering $\prec$ – of an ideal $\mathcal{I}$ in $\mathbb{F}_q[\mathbf{x}]$ if, for all $p \in \mathcal{I}$, there exists $g \in G$ such that* $\mathrm{LM}(g, \prec)$ *divides* $\mathrm{LM}(p, \prec)$.

Gröbner bases are a fundamental tool to study algebraic systems in theory and practice. They provide an algorithmic solution to several problems related to polynomial systems (see [1] for instance). We pay here particular attention to Gröbner bases computed for a lexicographical ordering. It offers a way of simplifying an algebraic system by giving an equivalent system with a structured shape. A lexicographical Gröbner basis of a zero-dimensional system (i.e. with a finite number of zeroes over the algebraic closure) is indeed always as follows:

$$\{f_1(x_1) = 0, f_2(x_1, x_2) = 0, \ldots, f_{k_2}(x_1, x_2) = 0, f_{k_2+1}(x_1, x_2, x_3) = 0, \ldots, \ldots\}$$

To compute the variety, we simply have to successively eliminate variables by computing zeroes of univariate polynomials and back-substituting results. However, computing a Gröbner basis w.r.t. a lexicographical order is in practice much slower than computing a Gröbner basis w.r.t. another monomial ordering. It is usually for a DRL order that the computation of Gröbner bases is the fastest in practice. Algorithms changing the monomial ordering of a Gröbner basis permit to handle efficiently this problem. The FLGM algorithm [19] allows to transform a Gröbner basis w.r.t. some monomial ordering into a lexicographical Gröbner basis in the zero-dimensional case and is polynomial-time.

The historical method for computing Gröbner bases is Buchberger's algorithm [9, 8]. Recently, more efficient algorithms have been proposed. The F_4 algorithm [16] is based on the intensive use of linear algebra methods. In short, the arbitrary choices – which limit the practical efficiency of Buchberger's algorithm – are replaced by computational strategies related to classical linear algebra problems (mainly the computation of a row echelon form).

In [17], a new criterion (the F_5 criterion) for detecting useless computations has been proposed. We mention that Buchberger's algorithm spends 90% of its time to perform these useless computations. Under some regularity conditions, it has been proved that all useless computations can be avoided. A new algorithm, called F_5, has then been built using this criterion and linear algebra methods. Briefly, it constructs incrementally the following matrices in degree d:

$$A_d = \begin{array}{c} \\ t_1 f_1 \\ t_2 f_2 \\ t_3 f_3 \\ \ldots \end{array} \begin{array}{c} m_1 \succ m_2 \succ m_3 \ldots \\ \begin{bmatrix} \ldots & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots \end{bmatrix} \end{array}$$

where the indices of the columns are monomials sorted for the admissible ordering $\prec$ and the rows are product of some polynomials f_i by some monomials t_j such that $\deg(t_j f_i) \leq d$. For a *regular system* ([17]) the matrices A_d are of full rank. In a second step, row echelon forms of theses matrices are computed, i.e.

$$A'_d = \begin{array}{c} \\ t_1 f_1 \\ t_2 f_2 \\ t_3 f_3 \\ \dots \end{array} \begin{array}{c} m_1\ m_2\ m_3\ \dots \\ \begin{bmatrix} 1 & 0 & 0 & \dots \\ 0 & 1 & 0 & \dots \\ 0 & 0 & 1 & \dots \\ 0 & 0 & 0 & \dots \end{bmatrix} \end{array}$$

For d sufficiently large, A'_d contains a Gröbner basis of the ideal considered. Important parameters to evaluate the complexity of F_5 is the maximal degree d occurring in the computation and the size of the matrix A_d. The overall cost is thus dominated by $(\#A_d)^3$. Very roughly, $(\#A_d)$ can be approximated by $O(n^d)$. A more precise complexity analysis can be found in [3, 4].

From a practical point of view, the gap with other algorithms computing Gröbner basis is consequent. To date, F_5 is the most efficient method for computing Gröbner bases, and hence zero-dimensional varieties. In particular, it has been proved [2] – from both a theoretical and practical point of view – that XL [14] is less efficient than F_5. Due to the range of examples that become computable with F_5, Gröbner basis can be considered as a reasonable computable object in real scale applications. For systems arising in cryptography, F_5 has for instance given impressing results on HFE [18].

2.2 Isomorphism of Polynomials and Related Problems

Before defining formally IP, we briefly come back here to the origin of this problem. To do so, we describe the encryption scheme called C^* [26]. The public key of this system is a set of multivariate quadratic polynomials $\mathbf{b} = \big(b_1(\mathbf{x}), \dots, b_n(\mathbf{x})\big) \in \mathbb{F}_q[\mathbf{x}]^n$. These polynomials are obtained by applying two bijective affine transformations $(S, \mathbf{V})$ and $(U, \mathbf{V})$ of $AGL_n(\mathbb{F}_q)$ to a particular set of polynomials $\mathbf{a} = \big(a_1(\mathbf{x}), \dots, a_n(\mathbf{x})\big) \in \mathbb{F}_q[\mathbf{x}]^n$. That is:

$$\big(b_1(\mathbf{x}), \dots, b_n(\mathbf{x})\big) = \big(a_1(\mathbf{x}S + \mathbf{T}), \dots, a_n(\mathbf{x}S + \mathbf{T})\big)U + \mathbf{V},$$

denoted $\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S + \mathbf{T})U + \mathbf{V}$ in the sequel.
To encrypt, we simply evaluate a message $\mathbf{m} \in \mathbb{F}_q^n$ on $\mathbf{b}$, i.e. $\big(b_1(\mathbf{m}), \dots, b_n(\mathbf{m})\big)$. To recover the correct plaintext, the legitimate recipient uses the bijectivity of the affine transformations combined with the particular structure of the polynomials of $\mathbf{a}$. How these polynomials are constructed is not relevant here. But, due to particular constraints, the polynomials of $\mathbf{a}$ are always considered as a public data. The secret key of C^* is constituted of $(S, \mathbf{T}), (U, \mathbf{V}) \in AGL_n(\mathbb{F}_q)$.

The first approach for attacking this scheme consists in trying to retrieve the message corresponding to a ciphertext $\mathbf{c} \in \mathbb{F}_q^n$, i.e. finding a zero of $\mathbf{b}(\mathbf{x}) = \mathbf{c}$. This corresponds to solving a particular instance of the so-called MQ problem, which is NP-Hard in general [10, 22]. We emphasize that such a kind of result uniquely guarantees the worst-case hardness and does not provide any information concerning the average-case difficulty. For instance, J.-C. Faugère and A. Joux proposed a polynomial-time algorithm for solving instances of MQ corresponding to the public key of HFE [18], which is an extension of C^*.

Another approach for breaking C^* consists in attempting to recover the affine transformations hiding the structure of $\underline{a}$. That is, extracting the secret key from the public key. This problem, introduced by J. Patarin at Eurocrypt'96 [28], is defined as follows:

Isomorphism of Polynomials (IP)
Input: $\mathbf{a} = (a_1, \ldots, a_u)$, and $\mathbf{b} = (b_1, \ldots, b_u\big)$ in $\mathbb{F}_q[\mathbf{x}]^u$.
Question: Find – if any – $(S, \mathbf{V}) \in AGL_n(\mathbb{F}_q)$ and $(U, \mathbf{V}) \in AGL_u(\mathbb{F}_q)$, s. t.:

$$\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S + \mathbf{V})U + \mathbf{V}.$$

More precisely, it is usually the linear variant of IP which is considered in practice [28, 5]. That is, when the vectors $\mathbf{T}$ and $\mathbf{V}$ are both equal to the null vector. This problem, that we call 2PLE is the following:
Input: $\mathbf{a} = (a_1, \ldots, a_u)$, and $\mathbf{b} = (b_1, \ldots, b_u)$ in $\mathbb{F}_q[\mathbf{x}]^u$.
Question: Find – if any – $(S, U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$, such that:

$$\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S)U.$$

However, it is without solving any of the two problems mentioned above that J. Patarin proposed a full cryptanalysis of C^* [30]. This attack uses the very particular structure of the polynomials of $\underline{a}$. This result thus does not then affect at all the practical hardness of IP. The security estimate provided for this problem [29] is based on the complexity of the "To and Fro" (TF) algorithm [11, 12], which is $q^{n/2}$ for quadratic polynomials, and q^n otherwise.

In the rest of this paper, $\big(\mathbf{a} = (a_1, \ldots, a_u), \mathbf{b} = (b_1, \ldots, b_u)\big)$ will always denote an element of $\mathbb{F}_q[\mathbf{x}]^u \times \mathbb{F}_q[\mathbf{x}]^u$. We will always suppose that all the polynomials of $\mathbf{a}$ have the same maximal total degree noted D (in the practical applications, we have $2 \leq D \leq 4$). Note that, if $\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S)U$, for some $(S, U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$, then the polynomials of $\mathbf{b}$ must have the same maximal total degree than the ones of $\mathbf{a}$, i.e. D.

3 A Unified Point of View

The Isomorphism of Polynomials and 2PLE problems have actually a very similar formulation. An input of these problems is formed of two systems of multivariate polynomials and the question consists in recovering a particular transformation permitting to express one system in function of the other. All transformations have the same characteristic: inducing a group action on $\mathbb{F}_q[\mathbf{x}]^u$. Recall that a group $(G, \cdot)$, with identity element e, *acts* on $\mathbb{F}_q[\mathbf{x}]^u$ if there exists a map $\phi : G \times \mathbb{F}_q[\mathbf{x}]^u \to \mathbb{F}_q[\mathbf{x}]^u$ such that $\phi(e, \mathbf{p}) = \mathbf{p}$, for all $\mathbf{p} \in \mathbb{F}_q[\mathbf{x}]^u$, and:

$$\phi\big(g, \phi(g', \mathbf{p})\big) = \phi(g \cdot g', \mathbf{p}), \text{ for all } g, g' \in G, \text{ and for all } \mathbf{p} \in \mathbb{F}_q[\mathbf{x}]^u.$$

Remark 1. In order to simplify the notations, we will write G instead of $(G, \cdot)$.

For 2PLE, one can then easily check that $GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$ acts on $\mathbb{F}_q[\mathbf{x}]^u$ through:

$$\begin{aligned} \phi_{\mathrm{2PLE}} : GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q) \times \mathbb{F}_q[\mathbf{x}]^u &\to \mathbb{F}_q[\mathbf{x}]^u \\ \big((S, U), \mathbf{a}\big) &\mapsto \mathbf{a}(\mathbf{x}S)U \end{aligned}$$

Similarly for IP, $AGL_n(\mathbb{F}_q) \times AGL_u(\mathbb{F}_q)$ acts on $\mathbb{F}_q[\mathbf{x}]^u$ through:

$$\begin{aligned} \phi_{\mathrm{IP}} : AGL_n(\mathbb{F}_q) \times AGL_u(\mathbb{F}_q) \times \mathbb{F}_q[\mathbf{x}]^u &\rightarrow \mathbb{F}_q[\mathbf{x}]^u \\ \big((S,\mathbf{T}),(U,\mathbf{V}),\mathbf{a}\big) \qquad\qquad &\mapsto \mathbf{a}(\mathbf{x}S+\mathbf{T})U+\mathbf{V} \end{aligned}$$

This observation naturally leads to the introduction of the following problem. Let $(G,\cdot)$ be a group, and $\phi : G \times \mathbb{F}_q[\mathbf{x}]^u \rightarrow \mathbb{F}_q[\mathbf{x}]^u$ be an action of G on $\mathbb{F}_q[\mathbf{x}]^u$. Given $(\mathbf{a},\mathbf{b}) \in \mathbb{F}_q[\mathbf{x}]^u \times \mathbb{F}_q[\mathbf{x}]^u$, the problem we call *Polynomial Equivalence*, with respect to $(G,\cdot)$ and ϕ – and denoted by $\mathrm{PE}\big(G,\phi\big)$ – is the one of finding (if any) $g \in G$, verifying:

$$\mathbf{b} = \phi(g,\mathbf{a}),$$

denoted $\mathbf{a} \equiv_{(G,\phi)} \mathbf{b}$ in the sequel. This formulation is very convenient since it procures a unified description of IP and 2PLE. Indeed, $\mathrm{PE}\big(GL_n(\mathbb{F}_q)\times GL_u(\mathbb{F}_q),$ $\phi_{\mathrm{2PLE}}\big)$=2PLE, and $\mathrm{PE}\big(AGL_n(\mathbb{F}_q) \times AGL_u(\mathbb{F}_q),\phi_{\mathrm{IP}}\big)$=IP. More generally, PE provides a unified description of "IP-like" problems. In our mind, such a kind of problems consists in recovering a particular transformation between two sets of multivariate polynomials. For instance, the Isomorphism of Polynomials with one Secret (IP1S) – introduced at Eurocrypt'96 by J. Patarin [28] – falls into this new formalism. This problem, which can be used to design an authentication (resp. signature) scheme [28], is as follows. Given $(\mathbf{a},\mathbf{b}) \in \mathbb{F}_q[\mathbf{x}]^u \times \mathbb{F}_q[\mathbf{x}]^u$, find – if any – $(S,\mathbf{T}) \in AGL_n(\mathbb{F}_q)$, such that $\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S+\mathbf{T})$. Using our formalism, we immediately obtain that $\mathrm{PE}\big(AGL_n(\mathbb{F}_q),\phi_{\mathrm{IP1S}}\big) = \mathrm{IP1S}$, with $\phi_{\mathrm{IP1S}} : AGL_n(\mathbb{F}_q) \times \mathbb{F}_q[\mathbf{x}]^u \rightarrow \mathbb{F}_q[\mathbf{x}]^u$, $\big((S,\mathbf{T}),\mathbf{a}(\mathbf{x})\big) \mapsto \mathbf{a}(\mathbf{x}S+\mathbf{T})$. Finally, the following lemma justifies the use of the word equivalence in PE.

Lemma 1. *Let $(G,\cdot)$ be a group, and $\phi : G \times \mathbb{F}_q[\mathbf{x}]^u \rightarrow \mathbb{F}_q[\mathbf{x}]^u$ be an action of G on $\mathbb{F}_q[\mathbf{x}]^u$. Then, $\equiv_{(G,\phi)}$ is an equivalence relation on $\mathbb{F}_q[\mathbf{x}]^u$.*

3.1 Polynomial Equivalence Problems and Group theory

In the Graph Isomorphism context, the introduction of group theory concepts permitted to achieve significant advances from both a theoretical and algorithmic point of view [24, 21]. The formalism previously given permits to naturally extend these results to Polynomial Equivalence problems.

Definition 4. *Let $(G,\cdot)$ be a group. We shall call $Aut_{(G,\phi)}(\mathbf{a}) = \big\{g \in G : \phi(g,\mathbf{a}) = \mathbf{a}\big\}$, $Aut_{(G,\phi)}(\mathbf{b}) = \big\{g \in G : \phi(g,\mathbf{b}) = \mathbf{b}\big\}$, the* automorphism groups of $\mathbf{a}$ and $\mathbf{b}$ w.r.t. (G,ϕ). *We shall also set $S_{(G,\phi)}(\mathbf{a},\mathbf{b}) = \big\{g \in G : \mathbf{b} = \phi(g,\mathbf{a})\big\}$.*

$Aut_{(G,\phi)}(\mathbf{a})$ and $.Aut_{(G,\phi)}(\mathbf{b})$ are also known as stabilizer of $\mathbf{a}$ (resp. $\mathbf{b}$) w.r.t. (G,ϕ). However, we will rather call these sets automorphism groups. This designation being indeed more usually used in the Graph Isomorphism context [24]. Anyway, the results that we are going to expose are classical results of group theory concerning the stabilizers and orbits, and then given without proofs.

Proposition 1. *Let $(G,\cdot)$ be a group, and $\phi : G \times \mathbb{F}_q[\mathbf{x}]^u \rightarrow \mathbb{F}_q[\mathbf{x}]^u$ be an action of G on $\mathbb{F}_q[\mathbf{x}]^u$. If there exists $g \in G$, such that $\mathbf{b} = \phi(g,\mathbf{a})$, then $S_{(G,\phi)}(\mathbf{a},\mathbf{b})$ is* *a left (resp. right) coset – in G – of the automorphism group $Aut_{(G,\phi)}(\mathbf{a})$ $\big($resp. $Aut_{(G,\phi)}(\mathbf{b})\big)$. That is:*

$$\begin{cases} S_{(G,\phi)}(\mathbf{a},\mathbf{b}) = \{g \cdot h : h \in Aut_{(G,\phi)}(\mathbf{a})\} = g \cdot Aut_{(G,\phi)}(\mathbf{a}), \\ S_{(G,\phi)}(\mathbf{a},\mathbf{b}) = \{h \cdot g : h \in Aut_{(G,\phi)}(\mathbf{b})\} = Aut_{(G,\phi)}(\mathbf{b}) \cdot g. \end{cases}$$

Moreover, the automorphism groups $Aut_{(G,\phi)}(\mathbf{a})$ and $Aut_{(G,\phi)}(\mathbf{b})$ are conjugate, i.e. $Aut_{(G,\phi)}(\mathbf{b}) = g \cdot Aut_{(G,\phi)}(\mathbf{a}) \cdot g^{-1}$, and we have:

$$|S_{(G,\phi)}(\mathbf{a},\mathbf{b})| = |Aut_{(G,\phi)}(\mathbf{b})| = |Aut_{(G,\phi)}(\mathbf{a})|.$$

3.2 A Generic Upper Bound on the Complexity of "IP-like" Problems

Using the Polynomial Equivalence problem previously defined, we give in this part a general upper bound on the theoretical complexity of "IP-like" problems. To do so, Let us fix a group $(G,\cdot)$ acting on $\mathbb{F}_q[\mathbf{x}]^u$ through a map noted ϕ.
For simplicity, we suppose here that G is included in a finite set $\mathcal{E}$. We also suppose that the uniform distribution of the elements of $\mathcal{E}$ can be simulated in polynomial-time. These assumptions allows to facilitate the proofs, and are additionally well adapted to "IP-like" problems. Indeed, $AGL_n(\mathbb{F}_q) \subset \mathcal{M}_{n,n}(\mathbb{F}_q) \times \mathbb{F}_q^n$, $GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q) \subset \mathcal{M}_{n,n}(\mathbb{F}_q) \times \mathcal{M}_{u,u}(\mathbb{F}_q)$, $AGL_n(\mathbb{F}_q) \times AGL_u(\mathbb{F}_q) \subset \mathcal{M}_{n,n}(\mathbb{F}_q) \times \mathbb{F}_q^n \times \mathcal{M}_{u,u}(\mathbb{F}_q) \times \mathbb{F}_q^u$. To obtain our upper bound, we introduce:

Definition 5. *An* **interactive proof** *for a language L $\big($i.e. a subset of $\{0,1\}^*\big)$ is a two party protocol between a verifier $\mathcal{V}$ and a prover $\mathcal{P}$. At the end of the protocol, the verifier has to accept or reject a given input such that the following conditions hold:*
Efficiency. *The verifier strategy is a probabilistic polynomial time procedure.*
Completeness. *For all $x \in L$, $Pr[(\mathcal{V},\mathcal{P})(x)$ accepts$] = 1$.*
Soundness. *For all $x \notin L$, and for any prover $\mathcal{P}^*$, $Pr[(\mathcal{V},\mathcal{P}^*)(x)$ accepts$] \leq \frac{1}{2}$. The probabilities are taken over the random choices of the verifier.*

Let us analyse the following two party protocol:

Input: $(\mathbf{a_0},\mathbf{a_1}) \in \mathbb{F}_q[\mathbf{x}]^u \times \mathbb{F}_q[\mathbf{x}]^u$
Protocol: $\mathrm{PI}(G,\phi)$
The verifier chooses uniformly at random $i \in \{0,1\}$.
He also chooses uniformly at random $g \in \mathcal{E}$ and checks if $g \in G$. If after C trials the verifier does not obtain an element $g \in G$, he accepts directly.
Otherwise, he sends $\mathbf{a}' = \phi(g,\mathbf{a_i})$ to the prover.
The prover replies by sending $j \in \{0,1\}$ to the verifier.
The verifier accepts if $i = j$ and rejects otherwise.

Efficiency. The efficiency of this protocol depends on the cost of computing $\phi(g,\mathbf{a_i})$, for all $g \in G$, and of the number of trials C.

Completeness. If $\mathbf{a_0} \not\equiv_{(G,\phi)} \mathbf{a_1}$, then a prover can always check if $\mathbf{a}' \equiv_{(G,\phi)} \mathbf{a_0}$ or $\mathbf{a}' \equiv_{(G,\phi)} \mathbf{a_1}$. In this situation, the verifier accepts with probability one.

Soundness. If $\mathbf{a_0} \equiv_{(G,\phi)} \mathbf{a_1}$, then by transitivity $\mathbf{a}' \equiv_{(G,\phi)} \underline{a_1}$ and $\mathbf{a}' \equiv_{(G,\phi)} \mathbf{a_0}$. In such a case, we will show that $\mathbf{a}' = \phi(g, \mathbf{a_i})$ yields no information about the bit i chosen by the prover. Let then ψ be a random variable uniformly distributed over $\{0,1\}$, and Σ be a random variable uniformly distributed over G.

Lemma 2. *Let $\mathbf{a_0}, \mathbf{a_1}, \mathbf{a}' \in \mathbb{F}_q[\mathbf{x}]^u$. If $\mathbf{a_0} \equiv_{(G,\phi)} \mathbf{a_1}$ and $\mathbf{a}' \equiv_{(G,\phi)} \mathbf{a_0}$, then:*

$$Pr[\psi = 0 \mid \mathbf{a}_\psi(\mathbf{x}\Sigma) = \mathbf{a}'] = Pr[\psi = 1 \mid \mathbf{a}_\psi(\mathbf{x}\Sigma) = \mathbf{a}'] = \frac{1}{2}.$$

Proof. We have $\Pr[\phi(\Sigma, \mathbf{a}_\psi) = \mathbf{a}' \mid \psi = 0] = \Pr[\phi(\Sigma, \mathbf{a_0}) = \mathbf{a}'\,] = \Pr[\Sigma \in S_{(G,\phi)}(\mathbf{a_0}, \mathbf{a}')\,]$. Moreover, according to Proposition 1:

$$|S_{(G,\phi)}(\mathbf{a_0}, \mathbf{a}')| = |Aut_{(G,\phi)}(\mathbf{a}')| = |S_{(G,\phi)}(\mathbf{a_1}, \mathbf{a}')|.$$

Therefore, $\Pr[\phi(\Sigma, \mathbf{a_0}) = \mathbf{a}'\,] = \Pr[\mathbf{a_1}(\mathbf{x}\Sigma) = \mathbf{a}'\,]$, and thus:

$$\Pr[\phi(\Sigma, \mathbf{a}_\psi) = \mathbf{a}' \mid \psi = 0] = \Pr[\phi(\Sigma, \mathbf{a}_\psi) = \mathbf{a}' \mid \psi = 1].$$

According to the Bayes formula:

$$\begin{aligned} \Pr\psi = 0 \mid \phi(\Sigma, \mathbf{a}_\psi) = \mathbf{a}'] &= \frac{\Pr[\psi=0]\Pr[\phi(\Sigma,\mathbf{a}_\psi)=\mathbf{a}' \mid \psi=0]}{\Pr[\phi(\Sigma,\mathbf{a}_\psi)=\mathbf{a}']} \\ &= \frac{\Pr[\psi=1]\ \Pr[\phi(\Sigma,\mathbf{a}_\psi)=\mathbf{a}' \mid \psi=1]}{\Pr[\phi(\Sigma,\mathbf{a}_\psi)=\mathbf{a}']} \\ &= \Pr[\psi = 1 \mid \phi(\Sigma, \mathbf{a}_\psi) = \mathbf{a}']. \end{aligned}$$

Finally:

$$\begin{aligned} \Pr[\psi = 0 \mid \phi(\Sigma, \mathbf{a}_\psi) = \mathbf{a}'] &= \frac{\Pr[\psi=0]\Pr[\phi(\Sigma,\mathbf{a}_\psi)=\mathbf{a}' \mid \psi=0]}{\Pr[\phi(\Sigma,\mathbf{a}_\psi)=\mathbf{a}']} \\ &= \frac{\Pr[\psi=1]\ \Pr[\phi(\Sigma,\mathbf{a_0})=\mathbf{a}']}{\Pr[\phi(\Sigma,\mathbf{a}_\psi)=\mathbf{a}']} \\ &= \frac{\Pr[\psi=1]\ \Pr[\Sigma\in S_{(G,\psi)}(\mathbf{a}',\mathbf{a_0})\,]}{\Pr[\Sigma\in S_{(G,\phi)}(\mathbf{a}_\psi,\mathbf{a}')]} = \frac{1}{2}. \end{aligned}$$

□

It follows that no prover – no matter what its strategy is – can guess i with probability greater than $\frac{1}{2}$. Finally, using a classical result of R. B. Boppana, J. Hastad, and S. Zachos [7], we get that:

Corollary 1. *If the polynomial hierarchy does not collapse then* IP, 2PLE*, and* IP1S *are not* NP-Hard.

Proof. We sketch the proof for IP1S. Note that for all $g \in AGL_n(\mathbb{F}_q)$, one can compute $\phi_{\text{IP1S}}(g, \mathbf{a}')$ in polynomial-time. Let L_{IP} be the language associated to IP1S (i.e. the set of instances of IP admitting a solution). We study now the number of trials in $\text{PI}\big(AGL_n(\mathbb{F}_q), \phi_{\text{IP1S}}\big)$. Recall that more than 1/4 of the matrices of $\mathcal{M}_{n,n}(\mathbb{F}_q)$ are invertible. Therefore for IP1S, we have $G = AGL_n(\mathbb{F}_q)$,

$\mathcal{E} = \mathcal{M}_{n,n}(\mathbb{F}_q) \times \mathbb{F}_q^n$, and $\Pr[g \in G \mid g \in \mathcal{E}] \geq \frac{1}{4}$. By setting $C = 10$, we get that no prover can guess i with probability greater than

$$\frac{1}{2} + \left(\frac{3}{4}\right)^{10} < \frac{1}{2} + \frac{1}{16} = \frac{9}{16},$$

where $\left(\frac{3}{4}\right)^{10} < \frac{1}{16}$ is the probability of not obtaining an element of $AGL_n(\mathbb{F}_q)$ after ten trials. By repeating the protocol two times, we obtain that no prover can fool the verifier into accepting $\mathbf{a_0} \not\equiv_{(AGL_n(\mathbb{F}_q),\phi_{\mathrm{IP1S}})} \mathbf{a_1}$ with a probability greater than $\left(\frac{9}{16}\right)^2 < \frac{1}{2}$. The protocol $\mathrm{PI}\big(AGL_n(\mathbb{F}_q), \phi_{\mathrm{IP1S}}\big)$ is then an interactive proof for the complementary language of $\mathrm{L_{IP1S}}$ (i.e. $\{0,1\}^* \backslash \mathrm{L_{IP1S}}$), where at most 4 messages are exchanged between the verifier and the prover. We do not detail the proof, but one can easily check that the same result holds for $\mathrm{PI}\big(AGL_n(\mathbb{F}_q) \times AGL_u(\mathbb{F}_q), \phi_{\mathrm{IP}}\big)$ and $\mathrm{PI}\big(GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q), \phi_{\mathrm{2PLE}}\big)$.
The corollary then follows from a result of [7], stating that if the complementary of a language admits a constant round interactive protocol, then this language can not be NP-Complete, unless the polynomial hierarchy collapses. □

The new formalism introduced in this part allows to upper bound the theoretical hardness of IP, 2PLE, and IP1S. More generally, it provides a new insight on the complexity of "IP-like" problems. The previous corollary can be indeed easily adapted to any instance of the Polynomial Equivalence problem. An "IP-like" problem is then intrinsically not NP-Hard. Furthermore, we believe that our formalism is of independent interest. It indeed procures a general framework for studying "IP-like" problems. However, this is out of the scope of this paper. We investigate now another aspect of these problems.

4 An Algorithm for Solving 2PLE

We study here the practical hardness of a particular Polynomial Equivalence problem, namely 2PLE. Precisely, we present a new algorithm for solving this problem. We emphasize that – as explained in 1.1 – it is usually sufficient to consider this problem rather than its affine variant IP. Besides, any algorithm solving 2PLE can be transformed into an algorithm solving IP [11, 12].

4.1 A First Attempt: Evaluation and Linearization

Instead of directly describing the details of our method, we present the different steps that yielded to this algorithm. Anyway, most of the intermediate results that we are going to present will be used in our final algorithm, but differently. Our earliest idea for solving 2PLE was based on the following remark. If $\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S)U$, for $(S, U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$, then:

$$\mathbf{b}(\mathbf{p})U^{-1} = \mathbf{a}(\mathbf{p}S), \text{ for all } \mathbf{p} \in \mathbb{F}_q^n. \tag{1}$$

We hence obtain, for each $\mathbf{p} \in \mathbb{F}_q^n$, u non-linear equations in the $n^2 + u^2$ components of the matrices S and U^{-1}. We point out that the coefficients of U^{-1} only appear linearly in these equations. This is the advantage of considering the inverse of U rather than simply U in (1). The number of equations obtained is then significantly bigger than the number of unknowns. In this situation, one can simply use a linearization method (i.e. associating a new variable to each monomial) for solving the algebraic system. Unfortunately, our experiments rapidly revealed that the equations generated in this way are not all linearly independent. Besides, it also appeared that the number of unknowns is significantly bigger than the number of linearly independent equations. The use of a linearization method is then clearly no longer relevant. Let us explain this phenomenon.

Lemma 3. *Let* $\mathbf{y} = (y_{1,1}, \ldots, y_{1,n}, \ldots, y_{n,1}, \ldots, y_{n,n})$*, and* $\mathbf{z} = (z_{1,1}, \ldots, z_{1,u}, \ldots, z_{u,1}, \ldots, z_{u,u})$*. For each* $i, 1 \leq i \leq u$*, there exists a subset* $S_i \subseteq \mathbb{F}_q^n$ *and polynomials* $p_{\alpha,i} \in \mathbb{F}_q[\mathbf{y}, \mathbf{z}]$*, such that the following equality holds:*

$$\big(\mathbf{b}(\mathbf{x})U^{-1} - \mathbf{a}(\mathbf{x}S)\big)_i = \sum_{\alpha \in S_i} p_{\alpha,i}(S, U^{-1})\mathbf{x}^{\alpha}, \tag{2}$$

$p_{\alpha,i}(S, U^{-1})$ *being the evaluation of* $p_{\alpha,i}$ *on* $S = \{s_{i,j}\}_{1 \leq i,j \leq n}$*,* $U^{-1} = \{u'_{i,j}\}_{1 \leq i,j \leq u}$*.*

Proof. The polynomial $\big(\mathbf{b}(\mathbf{x})U^{-1} - \mathbf{a}(\mathbf{x}S)\big)_i$ can be regarded as an element of:

$$\mathbb{F}_q[s_{1,1}, \ldots, s_{1,n}, \ldots, s_{n,1}, \ldots, s_{n,n}, u'_{1,1}, \ldots, u'_{1,u}, \ldots, u'_{u,u}, \ldots, u'_{u,u}][x_1, \ldots, x_n], \tag{3}$$

i.e. a polynomial with unknowns $x_1, \ldots, x_n$ and whose coefficients are polynomials in the components of S and U^{-1}. In this setting, the polynomials $p_{\alpha,i}$ exactly correspond to the coefficients of the monomials (in $x_1, \ldots, x_n$) occurring in $\big(\mathbf{b}(\mathbf{x})U^{-1} - \mathbf{a}(\mathbf{x}S)\big)_i$. Lastly $S_i = \{\alpha \in \mathbb{F}_q^n : p_{\alpha,i} \neq 0\}$. □

The cost of generating the polynomials $p_{\alpha,i}$ is proportional to the number of monomials occurring in $\big(\mathbf{b}(\mathbf{x})U^{-1} - \mathbf{a}(\mathbf{x}S)\big)_i$ viewed as a polynomial of (3), i.e. $O(n^{2D})$. Note also that each $p_{\alpha,i}$ is by construction the sum of a polynomial in $\mathbf{y}$, plus a linear polynomial in $\mathbf{z}$. Furthermore, the maximal total degree reached by a monomial in the variables $\mathbf{y}$ is equal to D.
From (2), we obtain that for all $i, 1 \leq i \leq u$:

$$\big(\mathbf{b}(\mathbf{p})U^{-1} - \mathbf{a}(\mathbf{p}S)\big)_i = \sum_{\alpha \in S_i} p_{\alpha,i}(S, U^{-1}) p_1^{\alpha_1} \cdots p_n^{\alpha_n}, \text{for all } \mathbf{p} = (p_1, \ldots, p_n) \in \mathbb{F}_q^n.$$

It follows that, for all $\mathbf{p} \in \mathbb{F}_q^n$, the equations procured by (1) are linear combinations of the $p_{\alpha,i}(S, U^{-1})$. The number of polynomials $p_{\alpha,i}$ is limited by the number of monomials occuring in $\big(\mathbf{b}(\mathbf{p})U^{-1} - \mathbf{a}(\mathbf{p}S)\big)_i$. Thus, $u \cdot \mathrm{C}_{n+D}^{D}$ bounds from above the number of linearly independent equations provided by linearizing (1). On the other hand, the number of unknowns in the linearized system is equal to the number of monomials in the variables $\mathbf{y}$ of degree smaller than D, plus the u^2 variables corresponding to $\mathbf{z}$. Using a rough bound, the linearization method yields a linear system of at most $O(u \cdot n^D)$ linearly independent equations with $O(u \cdot n^{2D})$ unknowns.

4.2 The 2PLE algorithm

The linearization can thus not be employed for solving efficiently 2PLE. However, Gröbner basis procures another method for solving the algebraic system given by (1). From a practical point of view, this approach is quite promising. Indeed, the system obtained by evaluating $\mathbf{b}(\mathbf{x})U^{-1} = \mathbf{a}(\mathbf{x}S)$ on several vectors is overdetermined. Nevertheless, all the equations derived from $\mathbf{b}(\mathbf{p})U^{-1} = \mathbf{a}(\mathbf{p}S)$ are according to (2) linear combinations the polynomials $p_{\alpha,i}$. It is hence sufficient to only consider the system formed by these equations. Formally:

Proposition 2. *Let* $\mathcal{I} = \langle p_{\alpha,i} : \text{for all } i, 1 \leq i \leq u, \text{ and for all } \alpha \in S_i \rangle \subset \mathbb{F}_q[\mathbf{y}, \mathbf{z}]$ *be the ideal generated by the polynomials* $p_{\alpha,i}$ *defined as in Lemma 3, and* $V(\mathcal{I})$ *be the following variety:*

$$V(\mathcal{I}) = \left\{ \mathbf{s} \in \mathbb{F}_q^{n^2+u^2} : p_{\alpha,i}(\mathbf{s}) = 0, \text{for all } i, 1 \leq i \leq u, \text{ and for all } \alpha \in S_i \right\}.$$

If $\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S)U$, *for some* $(S, U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$, *then:*

$$\left(\phi_1(S), \phi_2(U^{-1})\right) \in V(\mathcal{I}),$$

with:

$$\phi_1 : \mathcal{M}_{n,n}(\mathbb{F}_q) \to \mathbb{F}_q^{n^2}, S = \{s_{i,j}\}_{1 \leq i,j \leq n} \mapsto (s_{1,1}, \ldots, s_{1,n}, \ldots, s_{n,1}, \ldots, s_{n,n}), \text{ and}$$
$$\phi_2 : \mathcal{M}_{u,u}(\mathbb{F}_q) \to \mathbb{F}_q^{u^2}, U^{-1} = \{u'_{i,j}\}_{1 \leq i,j \leq u} \mapsto (u'_{1,1}, \ldots, u'_{1,u}, \ldots, u'_{u,1}, \ldots, u'_{u,u}).$$

Proof. For all, $i, 1 \leq i \leq u$:

$$\left(\mathbf{b}(\mathbf{x})U^{-1} - \mathbf{a}(\mathbf{x}S)\right)_i = \sum_{\alpha \in S_i} p_{\alpha,i}(S, U^{-1})\mathbf{x}^\alpha = 0.$$

Thus, $p_{\alpha,i}(S, U^{-1}) = 0, \forall i, 1 \leq i \leq u$, and $\forall \alpha \in S_i$, i.e. $\left(\phi_1(S), \phi_2(U^{-1})\right) \in V(\mathcal{I})$. $\square$

In other words, if $\mathbf{b} = \mathbf{a}(\mathbf{x}S)U$, for some $(S, U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$, then the variety $V(\mathcal{I})$ contains the components of the matrices S and U^{-1}. The system associated to $\mathcal{I}$ has n^2+u^2 variables and is of degree D. Once again, we recall that the variables of $\mathbf{z}$ only appear linearly in this system. The number of equations of the system is equal to the number of monomials occurring in the polynomials of $\mathbf{a}$, i.e. $O\left(u \cdot \mathrm{C}_{n+D}^{D}\right)$. The system is then overdetermined.

Remark 2. In order to guarantee that $V(\mathcal{I}) \subseteq \mathbb{F}_q^{2n}$, *we must generally join the field equations to the initial system. The fields considered in our case can be relatively large, leading then to a significant increase of the system's degree. This can artificially render impracticable the computation of a Gröbner basis. Fortunately, our systems are overdetermined and it is not necessary in practice to include the field equations. In our experiments the elements of* $V(\mathcal{I})$ *were indeed – without including these equations – all the times in* $\mathbb{F}_q^{2n}$. *It implies in particular that the hardness of 2PLE is not related to the size of the field. This is an important remark since the current security bound for 2PLE depends on this size.*

The next proposition is fundamental to understand the practical behaviour of our approach. This result permits furthermore to improve the efficiency of our method.

Proposition 3. *Let d be a positive integer, and $\mathcal{I}_d \subset \mathbb{F}_q[\mathbf{y},\mathbf{z}]$ be the ideal generated by the polynomials $p_{\alpha,i}$ of maximal total degree smaller than d. Let also $V(\mathcal{I}_d)$ be the variety associated to $\mathcal{I}_d$. If $\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S)U$, for some $(S,U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$, then:*

$$\big(\phi_1(S), \phi_2(U^{-1})\big) \in V(\mathcal{I}_d), \text{ for all } d, 0 \le d \le D,$$

ϕ_1 and ϕ_2 being defined as in proposition 2.

The proof is obviously deduced from the following result:

Lemma 4. *Let $(S,U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$. We have:*

$$\mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S)U \Longleftrightarrow \mathbf{b}^{(d)}(\mathbf{x}) = \mathbf{a}^{(d)}(\mathbf{x}S)U, \text{ for all } d, 0 \le d \le D,$$

$\mathbf{b}^{(d)}$ $\big($resp. $\mathbf{a}^{(d)}\big)$ being the homogeneous components of degree d (i.e. the sum of the terms of total degree d) of the polynomials of $\mathbf{b}$ (resp. $\mathbf{a}$).

The systems associated to $\mathcal{I}_1$ and $\mathcal{I}_0$ only contain linear equations in the components of S and U^{-1}. Indeed, let $\mathbf{0_n}$ be the null vector of $\mathbb{F}_q^n$, and $A \in \mathcal{M}_{n,u}(\mathbb{F}_q)\big($resp. $B \in \mathcal{M}_{n,u}(\mathbb{F}_q)\big)$ be the matrix representation of $\mathbf{a}^{(1)}$ $\big($resp. $\mathbf{b}^{(1)}\big)$, i.e. $\mathbf{x}A = \mathbf{a}^{(1)}(\mathbf{x})$ $\big($resp. $\mathbf{x}B = \mathbf{b}^{(1)}(\mathbf{x})\big)$. According to Lemma 4:

$$\mathbf{b} = \mathbf{a}(\mathbf{x}S)U, \text{ for } (S,U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q) \Longrightarrow \begin{cases} \mathbf{b}^{(0)}(\mathbf{0_n})U^{-1} = \mathbf{a}^{(0)}(\mathbf{0_n}), \\ BU^{-1} = SA. \end{cases}$$

That is, we get linear dependencies between the components S and U^{-1}. More precisely, we obtain $u(n+1)$ linear equations in the n^2+u^2 components of the matrices solution. Anyway, we can not solve 2PLE just by using these equations. On the other hand, it is not necessary to consider the system formed by all the polynomials $p_{\alpha,i}$. According to Proposition 3, we can actually restrict our attention on $\mathcal{I}_{d_0}$, with d_0 being the smaller integer rendering the system overdetermined. This d_0 can be defined in function of $\mathbf{a}$. Indeed, $d_0 \approx \min\{d > 1 : \mathbf{a}^{(d)} \neq \underline{0_u}\}$. In practice, it is usually sufficient to take $d_0 = 2$. The hardness of an instance of 2PLE is then related to d_0 rather than to the maximal total degree D of this instance. It is also an important remark since the maximal degree of an instance is taken into account in the security estimate of 2PLE given by J. Patarin [28, 29]. Our algorithm for solving this problem is as follows:

Input: $(\mathbf{a},\mathbf{b}) \in \mathbb{F}_q[\mathbf{x}]^u \times \mathbb{F}_q[\mathbf{x}]^u$
Let $d_0 = \min\{d > 1 : \mathbf{a}^{(d)} \neq \mathbf{0_u}\}$
Construct the polynomials $p_{\alpha,i}$ of max. total deg. smaller than d_0
Compute $V(\mathcal{I}_{d_0})$ using the F_5 algorithm
Find an element of $V(\mathcal{I}_{d_0})$ corresponding to a solution of 2PLE
Return this solution

The system associated to $\mathcal{I}_{d_0}$ is overdetermined by its very construction $\big(u^2+n^2$ unknowns, and $O\big(u \cdot \mathrm{C}_{n+d_0}^{d_0}\big)$ equations$\big)$. The variety $V(\mathcal{I}_{d_0})$ is then very likely reduced to a solution of 2PLE (this has been indeed verified in our experiments). The complexity of this algorithm is (theoretically) dominated by the Gröbner basis computation. It is difficult to obtain a complexity bound really reflecting the practical behavior of the F_5 algorithm. We therefore carry out now experimental results illustrating the practical efficiency of our approach.

4.3 Experimental Results

We present in this part experimental results obtained with our algorithm. Before that, we provide the conditions of our experiments.

Generation of the instances
We have only considered instances $(\mathbf{a}, \mathbf{b})$ of 2PLE admitting a solution. We constructed the instances in the following way:
(1) Choose the polynomials of $\mathbf{a}$
(2) Randomly choose $(S, U) \in GL_n(\mathbb{F}_q) \times GL_u(\mathbb{F}_q)$
(3) Return $\big(\mathbf{a}(\mathbf{x}), \mathbf{b}(\mathbf{x}) = \mathbf{a}(\mathbf{x}S)U\big)$

Precisely, we constructed the polynomials of $\mathbf{a}$ in two different ways. The first one simply consists in randomly choosing – w.r.t. a given maximal total degree D – the polynomials of $\mathbf{a}$. Precisely, each polynomial is a random linear combination of all the monomials of total degree smaller (or equal) to D. Note that we obtain in this way dense polynomials. We shall call *random instance*, an instance of 2PLE generated in this manner. In the second method, $\mathbf{a}$ corresponds to the public key of a C^* scheme [26]. An instance of 2PLE generated in this way will be named *C^* instance*.

Programming language – Workstation
The experimental results have been obtained with an Opteron bi-processors 2.4 Ghz, with 8 Gb of Ram. The systems associated to an instance of 2PLE have been generated using the Magma software[25]. We used our own implementation (in language C) of F_5 for computing the Gröbner bases. However, for the sake of comparison, we sometimes used the last version of Magma (i.e. 2.12) for obtaining these bases. This version includes an implementation of the F_4 algorithm.

Table Notations
The following notations are used in the tables below:
– n, the number of variables,
– q, the size of the field,
– deg, the maximal total degree of the considered instance,
– T_{Gen}, the time needed to construct the system,
– T_{F_5}, the time of our algorithm for finding a solution of 2PLE (using the F_5 algorithm for computing the Gröbner bases,
– T, the total time of our algorithm, i.e. $T = T_{F_5} + T_{Gen}$,
– $T_{F_4/Mag}$, the time of our algorithm for recovering a solution of 2PLE, using Magma v. 2.12 for computing Gröbner bases,

– $q^{n/2}$ (resp. q^n), the security bound given in [11, 12] for instances of $deg = 2$ (resp. $deg > 2$).

Practical Results – Random Instances

We present here the results obtained on random instances of 2PLE. We emphasize that this family of instances is the one employed in the authentication and signature schemes based on 2PLE proposed by J. Patarin at Eurocrypt'96 [28, 29]. He suggested to use $u = n$ in practice. Since our main motivation is to study the security of these schemes, we can restrict our attention on the case $u = n$.

n	q	deg	T_{Gen}	T_{F_5}	$T_{F_4/Mag}/T_{F_5}$	T	$q^{n/2}$
8	2^{16}	2	0.35 s.	0.14 s.	6	0.49 s.	2^{64}
10	2^{16}	2	1.66 s.	0.63 s.	10	2.29 s.	2^{80}
12	2^{16}	2	7.33 s.	2.16 s.	16	9.49 s.	2^{96}
15	2^{16}	2	48.01 s.	10.9 s.	23	58.91 s.	2^{120}
17	2^{16}	2	137.21 s.	27.95 s.	31	195.16 s.	2^{136}
20	2^{16}	2	569.14 s.	91.54 s.	41	660.68 s.	2^{160}
10	65521	2	1.21 s.	0.44 s.	10	1.65 s.	$\approx 2^{80}$
15	65521	2	35.58 s.	8.08 s.	23	43.66 s.	$\approx 2^{120}$
20	65521	2	434.96 s.	69.96 s.	41	504.92 s.	$\approx 2^{160}$
23	65521	2	1578.6 s.	235.92 s.		1814 s.	$\approx 2^{184}$

Remark 3. Our implementation of F_5 is faster than the Gröbner basis algorithm available in Magma 2.12. For $n = 20$, F_5 is for instance 41 times faster than Magma. To fix ideas, $u = n = 8$, and $u = n = 16$ were two challenges proposed at Eurocrypt'96 [29]. We obtained exactly the same results as the ones quoted in the previous table for random instances of $deg > 2$. On the other hand, the security estimate for these instances is at least equal to $2^{128}(n = 8)$. The maximal total degree of the systems is indeed the same as for instances of $deg = 2$, i.e. d_0 is equal to 2 independently of D. In other words, increasing the maximal total degree of a random instance will not change its practical hardness. We observe the same behavior for the size of the field, that is increasing q does not really change the hardness of a random instance. This will indeed modify only the cost of the arithmetic operations in the different steps our algorithm.

Interpretation of the results

In all these experiments, the varieties computed were reduced to one element, i.e. the components of the matrices solution of 2PLE. Furthermore, we observe in practice that the complexity of our algorithm is dominated by the time required to construct the system, and not by the Gröbner basis computation. This is surprising, but it clearly highlights that the systems considered here can be easily solved in practice. The generation of the systems being polynomial, we then conclude experimentally that our algorithm solves random instances of 2PLE in polynomial-time. This conclusion is supported by the fact that in all these experiments, the matrices generated by F_5 (see the Appendix) were of size at most equal to n^3. Experimentally, we deduce a complexity of $(n^3)^3 = n^9$ for our algorithm on random instances of 2PLE.

Practical Results – C^* Instances
We now present the results obtained on C^* instances $(\mathbf{a}, \mathbf{b})$ of degree D. We highlight that these instances are used in the traitor tracing scheme described in [5]. In this context, we also have $u = n$. The polynomials of $\mathbf{a}$ correspond to the public-key of a C^* scheme [26]. Precisely, these polynomials are the "multivariate representation" of a univariate monomial (see [5] for details concerning the generation of this multivariate representation). The univariate monomial has the following shape:$m^{1+q^{\theta_1}+q^{\theta_2}+\cdots+q^{\theta_{D-1}}}$, with $\theta_1, \theta_2, \cdots, \theta_{D-1} \in \mathbb{N}^*$.

n	q	deg	T_{Gen}	T_{F_5}	$T_{F_4/Mag}/T_{F_5}$	T	q^n
5	2^{16}	4	0.2 s.	0.13 s.	45	0.33 s.	2^{80}
6	2^{16}	4	0.7 s.	1.03 s.	64	1.73 s.	2^{96}
7	2^{16}	4	1.5 s.	6.15 s.	90	7.65 s.	2^{112}
8	2^{16}	4	3.88 s.	54.34 s.	112	58.22 s.	2^{128}
9	2^{16}	4	5.43 s.	79.85 s.	145	85.28 s.	2^{144}
10	2^{16}	4	12.9 s.	532.33 s.	170	545.23 s.	2^{160}

Remark 4. *$n = 5$, and $deg = 4$ is the first challenge proposed at Asiacrypt'03 [5]. Similarly to random instances, we observed that the size of the field does not really change the practical hardness of the C^* instances. We can conclude that it is a general behaviour of 2PLE instances.*

Interpretation of the results and Future work.
Our algorithm is no longer polynomial for C^* instances. The systems obtained for these instances are indeed harder to solve than the random ones. We believe that it is due to the fact that the systems are here sparser. The equality $\mathbf{b}(\mathbf{0_n}) = \mathbf{a}(\mathbf{0_n})U$ does not provide any information $\big(\mathbf{b}(\mathbf{0_n}) = \mathbf{a}(\mathbf{0_n}) = \mathbf{0_n}$ in the C^* case$\big)$. It is not clear yet but it seems that C^* instances with $n = 19$ (the second challenge proposed in [5]), can not be solved with our approach.
More generally, we think that $d_0 = \min\{d \geq 0 : \mathbf{a}^{(d)} \neq \mathbf{0_u}\}$ provides a relevant measure of the practical hardness of 2PLE instances. It seems actually that this practical difficulty increases in function of d_0. Indeed, for random instances of 2PLE, $d_0 = 0$ and our algorithm solves 2PLE efficiently. For C^* instances, $d_{min} = 1$ and there is a change of complexity class. We also checked that the practical complexity increases for homogeneous instances of degree 2, i.e. $d_0 = 2$. To summarize, for $d_0 = 0$ it is relatively clear that our algorithm solves 2PLE efficiently (likely in polynomial-time). For $d_0 \geq 1$, we conjecture that our algorithm is subexponential in n, and will depend on d_0. This anyway needs further investigations. It is an open problem to precisely determine, as a function of d_0, the asymptotic complexity of our algorithm. It could be possible that techniques presented in [3, 4] provide an answer.

Acknowledgements

We thank Françoise Levy-dit-Vehel and anonymous referees for numerous comments which improved the presentation of the results.

References

1. W.W. Adams and P. Loustaunau. *An Introduction to Gröbner Bases.* Graduate Studies in Mathematics, Vol. 3, AMS, 1994.
2. G. Ars, J.-C. Faugère, H. Imai, M. Kawazoe, and M. Sugita. *Comparison Between XL and Gröbner Basis Algorithms.* Advances in Cryptology – ASIACRYPT 2004, Lecture Notes in Computer Science, vol. 3329, pp. 338-353, 2004.
3. M. Bardet, J-C. Faugère, B. Salvy and B-Y. Yang. *Asymptotic Behaviour of the Degree of Regularity of Semi-Regular Polynomial Systems.* In MEGA 2005, Eighth International Symposium on Effective Methods in Algebraic Geometry, 15 pages, 2005.
4. M. Bardet, J-C. Faugère, and B. Salvy. *On the Complexity of Gröbner Basis Computation of Semi-Regular Overdetermined Algebraic Equations.* In Proc. of International Conference on Polynomial System Solving (ICPSS), pp. 71–75, 2004.
5. O. Billet, and H. Gilbert. *A Traceable Block Cipher.* Advances in Cryptology – ASIACRYPT 2003, Lecture Notes in Computer Science, vol. 2894, Springer–Verlag, pp. 331-346, 2003.
6. A. Biryukov, C. De Cannière, A. Braeken, and B. Preneel. *A Toolbox for Cryptanalysis: Linear and Affine Equivalence Algorithms.* Advances in Cryptology – EUROCRYPT 2003, Lecture Notes in Computer Science, vol. 2656, Springer–Verlag, pp. 33-50, 2003.
7. R. B. Boppana, J. Hastad, and S. Zachos. *Does* co–NP *Have Short Interactive Proofs?* Information Processing Letters, 25(2), pp. 127–132, 1987.
8. B. Buchberger. *Gröbner Bases : an Algorithmic Method in Polynomial Ideal Theory.* Recent trends in multidimensional systems theory. Reider ed. Bose, 1985.
9. B. Buchberger, G.-E. Collins, and R. Loos. *Computer Algebra Symbolic and Algebraic Computation.* Springer-Verlag, second edition, 1982.
10. N. Courtois. *La sécurité des primitives cryptographiques basées sur des problèmes algébriques multivariables: MQ, IP, MinRank, HFE.* Ph.D. Thesis, Paris, 2001.
11. N. Courtois, L. Goubin, and J. Patarin. *Improved Algorithms for Isomorphism of Polynomials.* Advances in Cryptology - EUROCRYPT 1998, Lecture Notes in Computer Science, vol. 1403, Springer-Verlag, pp. 84–200, 1998.
12. N. Courtois, L. Goubin, and J. Patarin. *Improved Algorithms for Isomorphism of Polynomials - Extended Version.* Available from `http://www.minrank.org`.
13. N. Courtois, L. Goubin, and J. Patarin. *SFLASH, a Fast Asymmetric Signature Scheme for low-cost Smartcards – Primitive Specification and Supporting Documentation.* Available at `http://www.minrank.org/sflash-b-v2.pdf`.
14. N. Courtois, A. Klimov, J. Patarin, and A. Shamir. *Efficient Algorithms for Solving Overdefined Systems of Multivariate Polynomial Equations.* Advances in Cryptology – EUROCRYPT 2000, Lecture Notes in Computer Science, vol. 1807, Springer–Verlag, pp. 392-407, 2000.
15. D. A. Cox, J.B. Little and, D. O'Shea. *Ideals, Varieties, and Algorithms: an Introduction to Computational Algebraic Geometry and Commutative Algebra.* Undergraduate Texts in Mathematics. Springer-Verlag. New York, 1992.
16. J.-C. Faugère. *A New Efficient Algorithm for Computing Gröbner Bases (F_4).* Journal of Pure and Applied Algebra, 139(1-3), pp. 61–88, June 1999.

17. J.-C. Faugère. *A New Efficient Algorithm for Computing Gröbner Basis without Reduction to Zero:* F_5. Proceedings of ISSAC, pp. 75–83. ACM press, July 2002.
18. J.-C. Faugère, and A. Joux. *Algebraic Cryptanalysis of Hidden Field Equation (HFE) Cryptosystems using Gröbner bases.* Advances in Cryptology - CRYPTO 2003, Lecture Notes in Computer Science, vol. 2729, Springer-Verlag, pp. 44–60, 2003.
19. J. C. Faugère, P. Gianni, D. Lazard, and T. Mora. *Efficient Computation of Zero-Dimensional Gröbner Bases by Change of Ordering.* Journal of Symbolic Computation, 16(4), pp. 329–344, 1993.
20. P. Felke *On certain Families of HFE-type Cryptosystems.* Proceedings of WCC'05, International Workshop on Coding and Cryptography, March 2005.
21. S. Fortin. *The Graph Isomorphism problem.* Technical Report 96-20, University of Alberta, 1996.
22. M. R. Garey, and D. B. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness.* W. H. Freeman, 1979.
23. W. Geiselmann, R. Steinwandt, and T. Beth. *Attacking the Affine Parts of SFLASH.* Cryptography and Coding, 8th IMA International Conference, vol. 2260, Springer–Verlag, pp. 355-359, 2001.
24. M. Hoffman. *Group-theoretic algorithms and Graph Isomorphism.* Lecture Notes in Computer Science, vol. 136, Springer–Verlag, 1982.
25. http://magma.maths.usyd.edu.au/magma/
26. T. Matsumoto, and H. Imai. *Public Quadratic Polynomial-tuples for efficient signature-verification and message-encryption.* Advances in Cryptology – EUROCRYPT 1988, Lecture Notes in Computer Science, vol. 330, Springer–Verlag, pp. 419–453, 1988.
27. https://www.cosic.esat.kuleuven.be/nessie/deliverables/decision-final.pdf.
28. J. Patarin. *Hidden Fields Equations (HFE) and Isomorphisms of Polynomials (IP): two new families of Asymmetric Algorithms.* Advances in Cryptology – EUROCRYPT 1996, Lecture Notes in Computer Science, vol. 1070, Springer-Verlag, pp. 33–48, 1996.
29. J. Patarin. *Hidden Fields Equations (HFE) and Isomorphisms of Polynomials (IP): two new families of Asymmetric Algorithms – Extended Version.* Available from `http://www.minrank.org/hfe/`.
30. J. Patarin. *Cryptanalysis of the Matsumoto and Imai Public Key Scheme of Eurocrypt'88.* Advances in Cryptology – CRYPTO 1995, Lecture Notes in Computer Science, Springer-Verlag, vol. 963, pp. 248-261, 1995.