Chapter 5

Factoring Multivariate Polynomials over Finite Fields

Summary:

We consider the deterministic complexity of the problem of polynomial factorization over finite fields - given a finite field $\mathbb{F}_q$ and a polynomial $h(x, y) \in \mathbb{F}_q[x, y]$ compute the unique factorization of $h(x, y)$ as a product of irreducible polynomials. This problem admits a randomized polynomial-time algorithm and no deterministic polynomial-time algorithm is known. In this chapter, we give a deterministic polynomial-time algorithm that *partially* factors the input polynomial $h(x, y)$. The algorithm can be generalized to partially factor multivariate polynomials in an arbitary number of variables.

We now describe precisely the output of our partial factoring algorithm. Associated with every $\mathbb{F}_q$-irreducible factor $f(x, y)$ of $h(x, y)$ are two objects - its total degree n and the smallest extension field $\mathbb{F}_{q^d}$ of $\mathbb{F}_q$ over which $f(x, y)$ splits into absolutely irreducible factors. Collecting all the $\mathbb{F}_q$-irreducible factors of $h(x, y)$ which have the same degree and the same splitting field, we get a unique factorization of $h(x, y)$ into a product of "uniform polynomials" - polynomials whose component $\mathbb{F}_q$-irreducible factors all have the same degree and the same splitting field. It is this unique representation of $h(x, y)$ as a product of uniform polynomials that is outputted by our algorithm.

5.1 Introduction

A fundamental theorem of algebra states that polynomials over any field $\mathbb{F}$ admit a unique factorization into a product of (a finite number of) $\mathbb{F}$-irreducible factors. Computing this factorization for polynomials over various fields is a very well-studied problem in algorithmic number theory. For densely represented polynomials (that is, polynomials of degree n in m variables that are specified by giving all the $\binom{n+m}{m}$-possible coefficients of monomials), the problem of factoring multivariate polynomials is known to reduce to the problem of factoring univariate polynomials [Kal82]. For univariate polynomials over $\mathbb{Q}$, the field of rational numbers, Lenstra, Lenstra and Lovasz [LLL82] gave a deterministic polynomial-time algorithm.

Over finite fields, the problem admits random polynomial time algorithms [Ber67, Ber70, CZ81] but no deterministic polynomial-time algorithm is known. In a very interesting development, Kaltofen devised an algorithm that given an algebraic circuit computing a moderate degree polynomial in a large number of variables, computes its factorization in random polynomial time. Kaltofen's algorithm has been widely used in theoretical computer science with applications in list decoding of codes [GS99, Gur01] and hardness-randomness tradeoffs for arithmetic circuits [KI04].

The deterministic complexity of factoring polynomials over finite fields has also made partial progress. Berlekamp gave a deterministic algorithm for computing the distinct-degree factorization of univariate polynomials. This was subsequently generalized by Gao, Kaltofen and Lauder [GKL04] for deterministic distinct degree factorization of multivariate polynomials over finite fields. Motivated by the solvability problem to be tackled in the next chapter, we continue this line of work and develop a deterministic algorithm for partially factoring multivariate polynomials over finite fields. Moreover our algorithm can be parallelized so that the parallel time complexity is polylogarithmic in the degree of the input polynomial to be factored.

In order to describe the output of our algorithm we need to introduce some terms.

Definition 5.1.1. A bivariate polynomial $h(x, y) \in \mathbb{F}_q[x, y]$ is said to be *absolutely irreducible* if it is irreducible over $\mathbb{F}_q$ and remains irreducible over the algebraic closure $\overline{\mathbb{F}}_q$ of $\mathbb{F}_q$.

Example: For example, $(y^2 - x^3) \in \mathbb{F}_7[x, y]$ is absolutely irreducible whereas $(y^2 + x^2) \in \mathbb{F}_7[x, y]$ is irreducible over $\mathbb{F}_7$ but factors into $(y + \sqrt{-1}x)(y - \sqrt{-1}x)$ over the extension

$\mathbb{F}_{7^2} = \mathbb{F}_7(\sqrt{-1})$ and hence is not absolutely irreducible over $\mathbb{F}_7$.

Remark. Note that a univariate polynomial $f(x) \in \mathbb{F}_q[x]$ is absolutely irreducible if and only if it is a linear polynomial. To see this, observe that if $f(x) \in \mathbb{F}_q[x]$ is a univariate irreducible polynomial of degree $d \geq 2$ then then it splits properly over $\mathbb{F}_{q^d}$, and therefore cannot be absolutely irreducible.

The polynomial $h(x, y)$ has a unique factorization over the algebraic closure $\overline{\mathbb{F}}_q$ of $\mathbb{F}_q$. Now collect all the elements of $\overline{\mathbb{F}}_q$ that occur as the coefficient of some monomial $x^i y^j$ in some absolutely irreducible factor $g(x, y)$ of $h(x, y)$ over $\overline{\mathbb{F}}$. Since this is a finite set, all these coefficients lie in some finite extension $\mathbb{K}$ of $\mathbb{F}_q$. We will call the smallest such extension field $\mathbb{K}$ the splitting field of $h(x, y)$. We will denote by $dim_{\mathbb{F}_q}(h(x, y))$ the dimension of the splitting field $\mathbb{K}$ of $h(x, y)$ over $\mathbb{F}_q$. That is, $dim_{\mathbb{F}_q}(h(x, y)) \overset{\text{def}}{=} [\mathbb{K} : \mathbb{F}_q]$.

We will call a polynomial $f(x, y) \in \mathbb{F}_q[x, y]$ *a uniform polynomial* if any two of its $\mathbb{F}_q$-irreducible factors have the same total degree and the same splitting field $\mathbb{K}$. In this chapter, we build upon the distinct degree factorization algorithm of Gao, Kaltofen and Lauder [GKL04] to split a given polynomial $h(x, y) \in \mathbb{F}_q[x, y]$ into a product of uniform polynomials. We summarize our main result as a theorem:

Theorem 5.1.2. *[Uniform factoring] There exists a deterministic algorithm that on input a polynomial $h(x, y) \in \mathbb{F}_q[x, y]$ of degree n outputs*

$$\langle (h_1(x, y), n_1, d_1), \ldots, (h_k(x, y), n_k, d_k) \rangle$$

such that

$$h(x, y) = h_1(x, y) \cdot \ldots \cdot h_k(x, y)$$

where each $h_i(x, y)$ is a uniform polynomial consisting of $\mathbb{F}_q$-irreducible factors of degree n_i and splitting field $\mathbb{F}_{q^{d_i}}$.

The algorithm has a time complexity of poly$(n \cdot \log q)$. Moreover, the algorithm can be implemented parallely to get a family of P -uniform circuits of depth poly$(\log n \cdot \log q)$ and size poly$(n \cdot \log q)$.

Note that the output of the algorithm of Theorem 5.1.2 is a refinement of the distinct degree factorization of $h(x, y)$ over $\mathbb{F}_q$.

We now give the overall idea behind our algorithm.

5.1.1 Basic Idea

The starting point of our algorithm is the procedure (due to Kaltofen [Kal82]) for reducing the problem of factoring bivariate polynomials to the problem of factoring univariate polynomials. Let $\mathbb{F}_q$ be a finite field and $h(x, y) \in \mathbb{F}_q[x, y]$ be a square-free bivariate polynomial of degree n that we wish to factor. By applying a suitable linear transformation if necessary, we can assume without loss of generality that $w(z) = h(z, 0)$ is square-free (cf. Kaltofen [Kal82]). Suppose we know an $\alpha \in \overline{\mathbb{F}}_q$ which is a root of some $\mathbb{F}_q$-irreducible factor $t(z)$ of $w(z)$. Let $R_t \overset{\text{def}}{=} \mathbb{F}_q[z]/\langle t(z) \rangle = \mathbb{F}_q(\alpha)$ be the splitting field of $t(z)$. Then by the squarefree-ness of $h(z, 0)$ there exists a unique (upto constant factors) minimal degree factor $h_{t(z)}(x, y) \in R_t[x, y]$ of $h(x, y)$ such that α is a root of $h_{t(z)}(z, 0)$. With this background in mind, Kaltofen's algorithm can be viewed as follows: using the root α having minimal polynomial $t(z)$ over $\mathbb{F}_q$, it simply writes down a system $\mathcal{R}_{t(z),m}$ of homogeneous linear equations over R_t whose solutions correspond to polynomials in $R_t[x, y]$ of degree at most m and which are multiples of $h_{t(z)}(x, y)$. Setting $m = n - 1$ and taking the gcd of all the polynomials corresponding to a basis of the solution space of $\mathcal{R}_{t(z),m}$ gives us the factor $h_{t(z)}(x, y) \in R_t[x, y]$ of $h(x, y)$.

Unfortunately the absence of a deterministic algorithm for univariate factoring over finite fields prevents us from obtaining irreducible factors of $w(z)$. Suppose that $v(z) \in \mathbb{F}_q[z]$ is *any* (not necessarily irreducible) factor of $w(z)$. As before, we construct the ring $R_v \overset{\text{def}}{=} \mathbb{F}_q[z]/\langle v(z) \rangle$ (note that R_v is no longer a field). We then view the element $\alpha' \in R_v$, $\alpha' \overset{\text{def}}{=} z \pmod{v(z)}$ as a *'pseudo-root'* of the polynomial $w(x) = h(x, 0) \in R_v[x]$. Proceeding as before, we write down a system $\mathcal{R}_{v(z),m}$ of homogeneous linear equations over R_v. We then ask the question - what do the solutions of $\mathcal{R}_{v(z),m}$ correspond to 'in reality'? Examining this question minutely, we deduce that by setting $v(z) = w(z)$ and varying m, the solutions of $\mathcal{R}_{w(z),m}$ can be used to factor out divisors of $h(x, y)$ having distinct degree or distinct splitting fields over $\mathbb{F}_q$.

Remark. Subsequently, Kaltofen [Kal85] essentially observed that $\mathcal{R}_{w(z),(n-1)}$ does **not** have a nontrivial solution if and only if $h(x, y)$ is absolutely irreducible. Combining this with efficient parallel algorithms for linear algebraic computations, he obtained a fast parallel deterministic algorithm for absolute irreducibility testing.

5.2 Mathematical machinery.

This section forms the core of this chapter. Its organized as follows - following tradition, we first introduce *nice* bivariate polynomials. We then examine how an $\mathbb{F}_q$-irreducible bivariate polynomial factors over various possible field extensions of $\mathbb{F}_q$. Next, we define some systems of linear equations $\mathcal{R}_{v(z),m}$, $\mathcal{F}_{v(z),m}$ and $\mathcal{B}_{v(z),m}$ and prove the basic properties of their solution spaces. Finally we show how these solution spaces can be used to obtain factors of $h(x, y)$. In this and the next section, we will use $h(x, y)$ for the reducible input polynomial to be factored and $f(x, y)$ for an $\mathbb{F}_q$-irreducible factor of $h(x, y)$.

5.2.1 Nice bivariate polynomials

Definition 5.2.1. A bivariate polynomial $f(x, y) \in \mathbb{F}_q[x, y]$ of total degree n is *nice* if $f(x, 0)$ is squarefree and of degree n.

Note that the coefficient of x^i of a nice polynomial $f(x, y)$ as a polynomial in y has degree no more than $n - i$, in particular the leading coefficient of $f(x, y)$ with respect to x is in $\mathbb{F}_q$.

Also observe that a nice polynomial $f(x, y) \in \mathbb{F}_q[x, y]$ remains nice over any extension field $\mathbb{K}$ of $\mathbb{F}_q$ and that any factor of a nice polynomial is also a nice polynomial. By doing a square-free factorization of the input polynomial followed by a suitable linear transformation of the variables, the problem of general bivariate factoring can be reduced to factoring a nice bivariate polynomial (cf. Kaltofen [Kal82] for details).

Throughout the rest of this chapter we will use $\mathbb{F}_q$ to denote the input field and unless mentioned otherwise, all the algebras that we come across in this chapter will be over $\mathbb{F}_q$. Also we shall throughout use $h(x, y) \in \mathbb{F}_q[x, y]$ to denote the input polynomial to be factored.

5.2.2 How $\mathbb{F}_q$-irreducible bivariate polynomials behave over extensions of $\mathbb{F}_q$.

We will now examine how an $\mathbb{F}_q$-irreducible factor $f(x, y)$ of $h(x, y)$ factors over an extension field $\mathbb{F}_{q^d}$ of $\mathbb{F}_q$. We will show that over any extension field $f(x, y)$ splits into a product of *conjugate* factors and if the extension field happens to be isomorphic to $\mathbb{F}_q[z]/\langle v(z)\rangle$ where $v(z)$ is an irreducible factor of $f(z, 0)$ then $f(x, y)$ splits into absolutely irreducible factors over it.

▮ *Conjugacy - an equivalence relation.*

Let $\mathbb{K}$ be a field extension of the finite field $\mathbb{F}_q$. Let $\phi \in Gal_{\mathbb{K}/\mathbb{F}_q}$ be an automorphism of $\mathbb{K}$. We extend ϕ to an automorphism of the ring $\mathbb{K}[x,y]$ in the natural way:

Definition 5.2.2. Let $\phi \in Gal_{\mathbb{K}/\mathbb{F}_q}$ be an automorphism of $\mathbb{K}$. Define the map $\phi : \mathbb{K}[x,y] \mapsto \mathbb{K}[x,y]$ as

$$\phi(f(x,y)) = \sum_{1 \leq k,l \leq n} \phi(a_{kl}) x^k y^l$$

where

$$f(x,y) = \sum_{1 \leq k,l \leq n} a_{kl} x^k y^l$$

Observe that the map $\phi : \mathbb{K}[x,y] \mapsto \mathbb{K}[x,y]$ is an automorphism of the ring $\mathbb{K}[x,y]$ that fixes the subring $\mathbb{F}_q[x,y]$. In particular,

- $\phi(f(x,y) + g(x,y)) = \phi(f(x,y)) + \phi(g(x,y))$
- $\phi(f(x,y) \cdot g(x,y)) = \phi(f(x,y)) \cdot \phi(g(x,y))$

We now define an equivalence relation on $\mathbb{K}[x,y]$ induced by such automorphisms of $\mathbb{K}[x,y]$.

Definition 5.2.3. Let $f(x,y), g(x,y) \in \mathbb{K}[x,y]$ be two bivariate polynomials. $g(x,y)$ is said to be a conjugate of $f(x,y)$ over $\mathbb{F}_q$, or an $\mathbb{F}_q$-conjugate of $f(x,y)$, if there exists an automorphism $\phi \in Gal_{\mathbb{K}/\mathbb{F}_q}$ such that $g(x,y) = \phi(f(x,y))$.

Observe that conjugacy is an equivalence relation on $\mathbb{K}[x,y]$.

▮ *Factorization of $\mathbb{F}_q$-irreducible polynomials over extension fields.*

Now consider a nice $\mathbb{F}_q$-irreducible polynomial $f(x,y) \in \mathbb{F}_q[x,y]$. Let $\mathbb{K} \supseteq \mathbb{F}_q$ be a finite field extension of $\mathbb{F}_q$. How does $f(x,y)$ factor over $\mathbb{K}$? We claim that all the $\mathbb{K}$-irreducible factors of $f(x,y)$ in $\mathbb{K}$ are in fact $\mathbb{F}_q$-conjugates of each other. In particular, all the $\mathbb{K}$-irreducible factors of $f(x,y)$ in $\mathbb{K}[x,y]$ are of equal degree.

Lemma 5.2.4. *Let $f(x,y) \in \mathbb{F}_q[x,y]$ be a nice $\mathbb{F}_q$-irreducible polynomial of total degree n. Let $\mathbb{K}$ be any finite field extension of $\mathbb{F}_q$. If $f_1(x,y) \in \mathbb{K}[x,y]$ and $f_2(x,y) \in \mathbb{K}[x,y]$ are any two $\mathbb{K}$-irreducible factors of $f(x,y)$ then $f_1(x,y)$ and $f_2(x,y)$ are $\mathbb{F}_q$-conjugates.*

Proof. For a polynomial $g(x,y) \in \mathbb{K}[x,y]$, define $H_g \leq Gal_{\mathbb{K}/\mathbb{F}_q}$ to be the subgroup of $Gal_{\mathbb{K}/\mathbb{F}_q}$ consisting of automorphisms in $Gal_{\mathbb{K}/\mathbb{F}_q}$ that fix $g(x,y)$. Since the galois groups of finite extensions of finite fields are cyclic groups, H_g must be a normal subgroup of $Gal_{\mathbb{K}/\mathbb{F}_q}$.

Let $g(x,y) \in \mathbb{K}[x,y]$ be a $\mathbb{K}$-irreducible factor of $f(x,y)$. Let the set of distinct cosets of H_g in $Gal_{\mathbb{K}/\mathbb{F}_q}$ be

$$Gal_{\mathbb{K}/\mathbb{F}_q}/H_g = \{H_g\phi_1, H_g\phi_2, \cdots H_g\phi_t\}$$

Then $\phi_1(g(x,y)), \phi_2(g(x,y)), \cdots \phi_t(g(x,y))$ are all the distinct conjugates of $g(x,y)$. We claim that the unique factorization of $f(x,y)$ into $\mathbb{K}$-irreducible polynomials over $\mathbb{K}$ is simply the product of all these distinct conjugates of $g(x,y)$. That is,

$$f(x,y) = \prod_{H_g\phi \in Gal_{\mathbb{K}/\mathbb{F}_q}/H_g} \phi(g(x,y)) \tag{5.1}$$

We first observe that any $\mathbb{F}_q$-conjugate of $g(x,y)$ is also a $\mathbb{K}$-irreducible factor of $f(x,y)$.

Claim 5.2.4.1. *Every conjugate of $g(x,y)$ is a $\mathbb{K}$-irreducible factor of $f(x,y)$.*

Proof. Since $g(x,y)|f(x,y)$, therefore $\exists g'(x,y) \in \mathbb{K}[x,y]$ such that $f(x,y) = g(x,y) \cdot g'(x,y)$. Suppose that ϕ is any automorphism in $Gal_{\mathbb{K}/\mathbb{F}_q}$. Applying ϕ to both sides we get:

$$\begin{aligned}
&\phi(f(x,y)) = \phi(g(x,y)) \cdot \phi(g'(x,y)) \\
\Rightarrow &f(x,y) = \phi(g(x,y)) \cdot \phi(g'(x,y)) \\
\Rightarrow &\phi(g(x,y))|f(x,y)
\end{aligned}$$

By the same reasoning $\phi(g(x,y)) \in \mathbb{K}[x,y]$ is $\mathbb{K}$-irreducible for if any $\widehat{g}(x,y) \in \mathbb{K}[x,y]$ is a proper divisor of $\phi(g(x,y))$ then $\phi^{-1}(\widehat{g}(x,y))$ is a a proper divisor $g(x,y)$, contradicting the $\mathbb{K}$-irreducibility of $g(x,y)$. Thus any conjugate of $g(x,y)$ is also an $\mathbb{K}$-irreducible factor of $f(x,y)$. □

Now $g(x,y)$ being $\mathbb{K}$-irreducible, is coprime to all $\mathbb{F}_q$-conjugates distinct from itself. Thus the rhs of equation (5.1) divides $f(x,y)$. Moreover the rhs of equation (5.1) is fixed by all the automorphisms in $Gal_{\mathbb{K}/\mathbb{F}_q}$. Since finite extensions of finite fields are normal extensions, so any polynomial in $\mathbb{K}[x,y]$ that is fixed by all the automorphisms in $Gal_{\mathbb{K}/\mathbb{F}_q}$ is in fact a polynomial in $\mathbb{F}_q[x,y]$. Hence the rhs of equation (5.1) is in fact a polynomial in $\mathbb{F}_q[x,y]$ that divides $f(x,y)$. By the $\mathbb{F}_q$-irreducibility of $f(x,y)$, we deduce that equation

(5.1) is indeed the unique factorization of $f(x,y)$. Thus all the $\mathbb{K}$-irreducible factors of $f(x,y)$ over $\mathbb{K}$ are precisely all the distinct conjugates of $g(x,y)$. $\square$

Now consider an $\mathbb{F}_q$-irreducible polynomial $f(x,y) \in \mathbb{F}_q[x,y]$ that factors in the algebraic closure $\overline{\mathbb{F}}_q$ of $\mathbb{F}_q$. What is the splitting field of $f(x,y)$? Can we put a bound on the dimension of the splitting field over $\mathbb{F}_q$? Assuming that $f(x,y)$ is a nice polynomial, the following proposition shows that if $t(z)$ is an $\mathbb{F}_q$-irreducible factor of $f(z,0)$, then the splitting field of $f(x,y)$ is a subfield of the finite field $\mathbb{F}_q[z]/\langle t(z)\rangle$. In particular, if $f(z,0)$ has a root $\alpha \in \mathbb{F}_q$, then $f(x,y)$ must be absolutely irreducible.

Proposition 5.2.5. *Let $f(x,y) \in \mathbb{F}_q[x,y]$ be a nice $\mathbb{F}_q$-irreducible polynomial of total degree n whose splitting field is $\mathbb{F}_{q^d}$. Also let $t(z) \in \mathbb{F}_q[z]$ be an $\mathbb{F}_q$-irreducible factor of $f(z,0)$. Then $d|deg(t(z))$ and $f(x,y)$ breaks into absolutely irreducible factors over $\mathbb{K} := \mathbb{F}_q[z]/\langle t(z)\rangle$, each absolutely irreducible factor being of degree $m = \frac{n}{d}$.*

Proof. Let $g(x,y) \in \mathbb{K}[x,y]$ be a $\mathbb{K}$-irreducible factor of $f(x,y)$ in $\mathbb{K}[x,y]$. Suppose if possible that $g(x,y)$ is not absolutely irreducible but breaks further over some finite extension $\mathbb{L} \supset \mathbb{K}$.

Let H_g be as in lemma 5.2.4. By lemma 5.2.4

$$f(x,y) = \prod_{H_g\phi \in G/H_g} \phi(g(x,y)) \tag{5.2}$$

Let $\alpha \in \mathbb{K}$ be a root of the polynomial $t(z)$. We start with the observation that some $\mathbb{F}_q$-conjugate of α must be a root of $g(z,0)$. Since α is a root of $f(z,0)$ we have $(z-\alpha)|(f(z,0) = \prod_{H_g\phi\in G/H_g} \phi(g(z,0)))$. Being irreducible, $(z-\alpha)$ must divide one of the factors on the rhs. That is, $\exists\phi \in Gal_{\mathbb{K}/\mathbb{F}_q}$ such that $(z-\alpha)|\phi(g(z,0))$. Applying ϕ^{-1} to both sides, we get $(z-\beta)|g(z,0)$, where $\beta = \phi^{-1}(\alpha)$. This $\beta = \phi^{-1}(\alpha) \in \mathbb{K}$ is the required $\mathbb{F}_q$-conjugate of α that is a $\mathbb{K}$-root of the polynomial $g(z,0)$.

By lemma 5.2.4 the $\mathbb{L}$-irreducible factors of $g(x,y)$ in $\mathbb{L}[x,y]$ are all $\mathbb{K}$-conjugates. Let $g_1(x,y) \in \mathbb{L}[x,y]$ be such an $\mathbb{L}$-irreducible factor of $g(x,y)$ with $(z-\beta)$ dividing $g_1(z,0)$. Let $\psi \in Gal_{\mathbb{L}/\mathbb{K}}$ be such that $\psi(g_1(x,y)) \in \mathbb{L}[x,y]$ is another $\mathbb{L}$-irreducible factor of $g(x,y)$ distinct from $g_1(x,y)$. Now since $(z-\beta)|g_1(z,0)$, applying ψ on both sides we get that $(z-\psi(\beta))|\psi(g_1(z,0))$. But $\psi(\beta) = \beta$ and therefore $(z-\beta)$ divides two distinct coprime factors $g_1(z,0)$ and $\psi(g_1(z,0))$ of $g(z,0)$. This implies that $(z-\beta)^2$ divides $g(z,0)$ which is a contradiction since $f(z,0)$ and hence $g(z,0)$ are squarefree.

Thus the $\mathbb{K}$-irreducible factors of $f(x,y)$ are in fact absolutely irreducible. Hence there exists a subfield $\mathbb{F} \subseteq \mathbb{K}$ which is the splitting field of $f(x,y)$. Therefore $d = [\mathbb{F} : \mathbb{F}_q]$ divides $deg(t(z)) = [\mathbb{K} : \mathbb{F}_q] = [\mathbb{K} : \mathbb{F}][\mathbb{F} : \mathbb{F}_q]$.

By the definition of the splitting field of $f(x,y)$, the coefficients occuring in $g(x,y)$ lie in the field $\mathbb{F}$ and do not all lie in any proper subfield of $\mathbb{F}$. Hence $\mathbb{F}$ is precisely the subfield of $\mathbb{K}$ which is fixed by every automorphism in H_g. So

$$d = [\mathbb{F} : \mathbb{F}_q] = ord(Gal_{\mathbb{K}/\mathbb{F}_q}/H_g).$$

Further $ord(Gal_{\mathbb{K}/\mathbb{F}_q}/H_g)$ is the number of distinct absolutely irreducible factors of $f(x,y)$. Since all the absolutely irreducible factors of $f(x,y)$ are of the same degree, say m, we have

$$\begin{aligned} m.ord(Gal_{\mathbb{K}/\mathbb{F}_q}/H_g) &= deg(f(x,y)) \\ \Rightarrow \ m.d &= n \\ \Rightarrow \ m &= \frac{n}{d} \end{aligned}$$

□

This proposition means that if $f(x,y) \in \mathbb{F}_q[x,y]$ is a nice $\mathbb{F}_q$-irreducible polynomial and $t_1(z), t_2(z) \in \mathbb{F}_q[z]$ are any two $\mathbb{F}_q$-irreducible factors of $f(z,0)$ then the degree of an irreducible factor of $f(x,y)$ over $\mathbb{K}_1 \stackrel{\text{def}}{=} \mathbb{F}_q[z]/\langle t_1(z)\rangle$ is the same as the degree of an irreducible factor of $f(x,y)$ over $\mathbb{K}_2 \stackrel{\text{def}}{=} \mathbb{F}_q[z]/\langle t_2(z)\rangle$. This observation will be the key to our uniform-factoring algorithm.

5.2.3 Defining the linear systems.

We will now define some linear systems over R_v whose solutions capture different factors of $h(x,y)$. To be able to specify how these factors relate to a "seed polynomial" $v(z)$ we need to make the following definition.

Definition 5.2.6. Let R be any ring and let $v(z) \in R[z]$ be a univariate polynomial $f(x,y) \in R[x,y]$ be a bivariate polynomial. We will say that $f(x,y)$ *sits above* $v(z)$ if $v(z)$ divides $f(z,0)$.

We also extend the usual notion of squarefreeness of polynomials over fields to polynomials over arbitary rings.

Definition 5.2.7. Let R be any ring and $v(z) \in R[z]$ be a univariate polynomial over R. Let $v'(z) \in R[z]$ be the formal derivative of $v(z)$. We say that $v(z)$ is *squarefree* if $v(z)$ is coprime (see (2.1.11) for definition of coprimality) to $v'(z)$.

▮ *Fixing Some notation.*

We recall some of the quantities from the previous section and define and fix some other quantities that will be be used through the rest of this chapter.

As before, $h(x, y) \in \mathbb{F}_q[x, y]$ is a nice bivariate polynomial of degree n that we wish to factor. $w(z) \overset{\text{def}}{=} h(z, 0) \in \mathbb{F}_q[z]$ and $v(z) \in \mathbb{F}_q[z]$ is any factor of $w(z)$. Let the $\mathbb{F}_q$-irreducible factors of $v(z)$ be $v_j(z),\ 1 \leq j \leq r$.

$$R_v \overset{\text{def}}{=} \mathbb{F}_q[z]/\langle v(z)\rangle \cong \bigotimes_{j=1}^{r} \mathbb{F}_q[z]/\langle v_j(z)\rangle.$$

We will denote by π_{v_j} the projection of R_v onto the j-th component field,

$$R_{v_j} \overset{\text{def}}{=} \mathbb{F}_q[z]/\langle v_j(z)\rangle.$$

That is, for any $u \in R_v$,

$$\pi_{v_j}(u) \overset{\text{def}}{=} u \pmod{v_j}.$$

Note that every π_{v_j} extends in a natural manner to a homomorphism from polynomial rings over R_v to corresponding polynomial rings over R_{v_j}. We shall denote by B_v the Berlekamp subalgebra of R_v, defined as the subalgebra of R_v fixed by the automorphism $\phi : \zeta \mapsto \zeta^q$ of R_v.

The element $\alpha \in R_v$ is defined as $\alpha \overset{\text{def}}{=} z \pmod{v(z)}$ and it is an R_v-root of $h(x, 0) \in R_v[x]$.

We will now define three linear systems $\mathcal{R}_{v(z),m}$, $\mathcal{B}_{v(z),m}$ and $\mathcal{F}_{v(z),m}$. The solutions of each of these linear systems correspond to factors of $h(x, y) \in R_v[x, y]$ of degree at most m which sit above the polynomial $(x - \alpha)$. The difference is in which subring of $R_v[x, y]$ are these factors allowed to lie in, that is which subring of R_v do the coefficients come from.

The solutions of $\mathcal{F}_{v(z),m}$ are intended to capture factors (of degree at most m) which lie in the subring $\mathbb{F}_q[x, y]$ of $R_v[x, y]$. The solutions of $\mathcal{B}_{v(z),m}$ are intended to capture factors which lie in the subring $B_v[x, y]$ of $R_v[x, y]$. Finally, the solutions of $\mathcal{R}_{v(z),m}$ are intended to capture factors which lie in $R_v[x, y]$ itself.

Notational convention: In the rest of this chapter we will use $\mathtt{r}(x, y)$ to denote polynomials in $R_v[x, y]$, $\mathtt{b}(x, y)$ to denote polynomials in $B_v(x, y)$ and $\mathtt{f}(x, y)$ to denote polynomials in $\mathbb{F}_q[x, y]$.

Moreover for $m = \deg(h(x, y))$, the solutions of $\mathcal{R}_{v(z),m}$ are all going to be multiples of some particular well-defined polynomial $\mathtt{r}_{v(z)}(x, y) \in R_v[x, y]$. Similar thing is true for the linear systems $\mathcal{F}_{v(z),m}$ and $\mathcal{B}_{v(z),m}$. We will shortly define this factor $\mathtt{r}_{v(z)}(x, y) \in R_v[x, y]$ and its analogues. We prove a lemma first.

Lemma 5.2.8. *In the component field R_{v_j} of R_v, there exists a unique (upto constant multiples from R_{v_j}) minimal degree factor $\mathtt{r}_j(x, y) \in R_{v_j}[x, y]$ of $h(x, y)$ in $R_{v_j}[x, y]$ which sits above $(x - \pi_{v_j}(\alpha))$.*

Proof. **Existence.** Clearly $h(x, y) \in R_{v_j}[x, y]$ is itself a factor of $h(x, y)$ which sits above $\pi_{v_j}(x - \alpha)$ and therefore there does exist a minimal degree factor $\mathtt{r}_j(x, y)$ of $h(x, y)$ in $R_{v_j}[x, y]$ sitting above $\pi_{v_j}((x - \alpha))$.

Uniqueness. Note that $\mathtt{r}_j(x, y)$ must be an R_{v_j}-irreducible polynomial for if

$$\mathtt{r}_j(x, y) = \mathtt{r}_1(x, y) \cdot \mathtt{r}_2(x, y)$$

then either $\mathtt{r}_1(x, y)$ or $\mathtt{r}_2(x, y)$ sits above $(x - \pi_{v_j}(\alpha))$ and they are both factors of $h(x, y)$ in $R_{v_j}[x, y]$ of degree smaller than $\mathtt{r}_j(x, y)$, contradicting the assumption of the minimality of the degree of $\mathtt{r}_j(x, y)$. Suppose $\mathtt{r}'_j(x, y)$ is another factor of $h(x, y)$ sitting above $(x - \pi_{v_j}(\alpha))$, having the same degree as $\mathtt{r}_j(x, y)$. Then arguing as above, $\mathtt{r}'_j(x, y)$ is also irreducible in $R_{v_j}[x, y]$. Then their product $\mathtt{r}_j(x, y) \cdot \mathtt{r}'_j(x, y) \in R_{v_j}[x, y]$ must be factor of $h(x, y)$. But then $(x - \pi_{v_j}(\alpha))^2$ divides $\mathtt{r}_j(x, 0) \cdot \mathtt{r}'_j(x, 0)$ contradicting the squarefreeness of $h(x, 0)$.

□

In a similar manner, by the $\mathbb{F}_q$-irreducibility of $v_j(z)$, there exists a unique $\mathbb{F}_q$-irreducible factor $\mathtt{b}_j(x, y) \in \mathbb{F}_q[x, y]$ such that $v_j(z)$ divides $\mathtt{b}_j(z, 0)$. We will denote by $\mathtt{r}_{v(z)}(x, y)$ the unique element of $R_v[x, y]$ such that

$$\pi_{v_j}(\mathtt{r}_{v(z)}(x, y)) = \mathtt{r}_j(x, y) \forall 1 \leq j \leq r.$$

Analogously, we will denote by $\mathtt{b}_{v(z)}(x, y) \in B_v[x, y]$ the unique element of $B_v[x, y]$ such that

$$\pi_{v_j}(\mathtt{b}_{v(z)}(x, y)) = \mathtt{b}_j(x, y) \forall 1 \leq j \leq r.$$

Finally we define $\mathtt{f}_{v(z)}(x,y) \in \mathbb{F}_q[x,y]$ to be the polynomial

$$\mathtt{f}_{v(z)}(x,y) \overset{\text{def}}{=} \prod_{1 \le j \le r} \mathtt{b}_{v_j(z)}(x,y)$$

The linear systems $\mathcal{R}_{v(z),m}$, $\mathcal{F}_{v(z),m}$ and $\mathcal{B}_{v(z),m}$.

The polynomials $(x-\alpha)$ and $\frac{w(x)}{(x-\alpha)}$ in $R_v[x]$ are formally coprime (since they are coprime in each of the projected fields). That is α is an ordinary root of $w(x) = h(x,0)$ in $R_v = R_v[y]/\langle y \rangle$. We fix $k \in \mathbb{Z}_{>0}$ to be $k \overset{\text{def}}{=} 2n(n-1)$. By the well-known Hensel lifting lemma 2.1.12, there exists a unique $\alpha(y) = \alpha + \alpha_1 y + \alpha_2 y^2 + \ldots + \alpha_k y^k \in R_v[y]/\langle y^{k+1} \rangle$ such that

$$h(\alpha(y), y) = 0 \ (\text{mod} \ \ y^{k+1}).$$

Moreover, $\alpha(y)$ is easily computed by iteratively solving linear equations over R_v.

Definition 5.2.9. The linear system $\mathcal{R}_{v(z),m}$ over R_v is defined to be the system

$$\sum_{i=0}^{m} u_i(y)\alpha(y)^i = 0 \ \ (\text{mod } y^{k+1}) \tag{5.3}$$

with unknowns

$$u_i(y) \in R_v[y], \ deg(u_i(y)) \le (m-i).$$

The definitions of the linear systems $\mathcal{B}_{v(z),m}$ and $\mathcal{F}_{v(z),m}$ are very similar except that the unknown polynomials are restricted to lie in the respective subrings of R_v.

Definition 5.2.10. The linear system $\mathcal{B}_{v(z),m}$ over B_v is defined to be the system

$$\sum_{i=0}^{m} u_i(y)\alpha(y)^i = 0 \ \ (\text{mod } y^{k+1}) \tag{5.4}$$

with unknowns

$$u_i(y) \in B_v[y], \ deg(u_i(y)) \le (m-i).$$

Definition 5.2.11. The linear system $\mathcal{F}_{v(z),m}$ over $\mathbb{F}_q$ is defined to be the system

$$\sum_{i=0}^{m} u_i(y)\alpha(y)^i = 0 \ \ (\text{mod } y^{k+1}) \tag{5.5}$$

with unknowns

$$u_i(y) \in \mathbb{F}_q[y], \ deg(u_i(y)) \le (m-i).$$

By *a solution $\mathtt{r}(x,y)$ of $\mathcal{R}_{v(z),m}$ in $R_v[x,y]$* we will a mean a solution vector

$$(u_0(y), u_1(y), \ldots, u_m(y))$$

of the linear system $\mathcal{R}_{v(z),m}$, with $\mathtt{r}(x,y) \in R_v[x,y]$ being

$$\mathtt{r}(x,y) = \sum_{i=0}^{m} u_i(y)x^i \in R_v[x,y].$$

In an analogous manner we will identify solutions of $\mathcal{B}_{v(z),m}$ and $\mathcal{F}_{v(z),m}$ with bivariate polynomials $\mathtt{b}(x,y) \in B_v[x,y]$ and $\mathtt{f}(x,y) \in \mathbb{F}_q[x,y]$ respectively.

▮ *Properties of the linear systems for irreducible factors of $v(z)$.*

We will use $\mathcal{R}_{v_j(z),m}$ to denote the projection linear system $\mathcal{R}_{v(z),m}$ onto the j-th component:

$$\mathcal{R}_{v_j(z),m} \overset{\text{def}}{=} \pi_{v_j}(\mathcal{R}_{v(z),m}).$$

The projected linear systems $\mathcal{F}_{v_j(z),m}$ and $\mathcal{B}_{v_j(z),m}$ are defined analogously. We are now all set to prove the fundamental property of the solution space of these linear systems.

Proposition 5.2.12. *For all $1 \le j \le r$:*

1. *The projected linear system $\mathcal{R}_{v_j(z),m}$ has a non-zero solution if and only if*

   $$\textsc{Deg}(\mathtt{r}_{v_j(z)}(x,y)) \le m.$$

   Moreover, the gcd of all the polynomials in $R_{v_j}[x,y]$ corresponding to a basis of the solution space of $\mathcal{R}_{v_j(z),m}$ is precisely the polynomial $\mathtt{r}_{v_j(z)}(x,y) \in R_{v_j}[x,y]$.

2. *The projected linear system $\mathcal{B}_{v_j(z),m}$ has a non-zero solution if and only if*

   $$\textsc{Deg}(\mathtt{b}_{v_j(z)}(x,y)) \le m.$$

   Moreover, the gcd of all the polynomials in $B_{v_j}[x,y]$ corresponding to a basis of the solution space of $\mathcal{B}_{v_j(z),m}$ is precisely the polynomial $\mathtt{b}_{v_j(z)}(x,y) \in B_{v_j}[x,y]$.

3. *The projected linear system $\mathcal{F}_{v_j(z),m}$ has a non-zero solution if and only if*

   $$\textsc{Deg}(\mathtt{f}_{v_j(z)}(x,y)) \le m.$$

   Moreover, the gcd of all the polynomials in $\mathbb{F}_q[x,y]$ corresponding to a basis of the solution space of $\mathcal{F}_{v_j(z),m}$ is precisely the polynomial $\mathtt{f}_{v_j(z)}(x,y) \in \mathbb{F}_q[x,y]$.

Proof. The proofs of parts (ii) and (iii) are analogous to that of part (i) and we omit them for the sake of brevity. To emphasize that R_{v_j} is a field we will let $\mathbb{K}$ stand for it in the rest of this proof.

Existence of solution. Let $\mathtt{r}_{v_j(z)}(x,y) = v_0(y) + v_1(y)x + \ldots v_d(y)x^d$ where $d = \textsc{Deg}(\mathtt{r}_{v_j(z)}(x,y))$. Moreover $\mathtt{r}_{v_j}(x,y)$, being a factor of a nice polynomial $h(x,y)$ is itself a nice polynomial and so $\textsc{Deg}(v_i(y)) \leq (d-i)$. Now if $d \leq m$ then

$$(v_0(y), v_1(y), \cdots, v_d(y), 0, \ldots, 0)$$

is clearly a non-zero solution of the linear system $\mathcal{R}_{v_j(z),m}$. Conversely suppose that the system $\mathcal{R}_{v_j(z),m}$ has a nontrivial solution $g(x,y)$ with

$$g(x,y) := \sum_{i=0}^{m} u_i(y)x^i \in \mathbb{K}[x,y]$$

We claim that $\mathtt{r}_{v_j(z)}(x,y)$ must divide $g(x,y)$ thereby implying that $m \geq d$. Let

$$\rho(y) := \textsc{Resultant}_x(\mathtt{r}_{v_j(z)}(x,y), g(x,y)) \in \mathbb{K}[y]$$

Then $deg(\rho(y)) \leq (2n-1)n = k$. Then there exist polynomials $a(x,y), b(x,y) \in \mathbb{K}[x,y]$ such that

$$\rho(y) = a(x,y)\mathtt{r}_{v_j(z)}(x,y) + b(x,y)g(x,y) \tag{5.6}$$

Substituting $x := \alpha(y)$ in equation (5.6), we have

$$\rho(y) = 0 \pmod{y^{k+1}}.$$

But $deg(\rho(y)) \leq k$ and hence we must have that $\rho(y)$ is identically zero. Thus $gcd_x(\mathtt{r}_{v_j(z)}(x,y), g(x,y))$ is nontrivial whence by the irreducibility of $\mathtt{r}_{v_j(z)}(x,y)$ we deduce that $g(x,y)$ is a multiple of $\mathtt{r}_{v_j(z)}(x,y)$ as claimed. Thus we have shown that $\mathcal{R}_{v_j(z),m}$ has a non-zero solution if and only if

$$\textsc{Deg}(\mathtt{r}_{v_j(z}(x,y)) \leq m$$

and moreover $\mathtt{r}_{v_j(z)}(x,y)$ divides the bivariate polynomial in $\mathbb{K}[x,y]$ corresponding to any solution of $\mathcal{R}_{v_j(z),m}$.

The gcd of the basis vectors. Every solution of $\mathcal{R}_{v_j(z),m}$ corresponds to a bivariate polynomial over $\mathbb{K}$ in the natural way and let $g(x,y)$ be the gcd of all the basis polynomials which are solutions of $\mathcal{R}_{v_j(z),m}$. We must have that

$FactorR_{v_j(z)}(x,y)$ divides $g(x,y)$ because it divides every polynomial in the basis of $\mathcal{R}_{v_j(z),m}$. In the converse direction, observe that by definition, any solution of $\mathcal{R}_{v_j(z),m}$ is a $\mathbb{K}$-linear combination of the basis polynomials and therefore $g(x,y)$ divides any polynomial in the solution space. Since $\mathtt{r}_{v_j(z)}(x,y)$ is a solution of $\mathcal{R}_{v_j(z),m}$, we must have that $g(x,y)$ divides $\mathtt{r}_{v_j(z)}(x,y)$. Thus $\mathtt{r}_{v_j(z)}(x,y) = g(x,y)$ as was to be shown.

□

▮ *Using $\mathcal{F}_{v(z),n}$ to compute a factor of $h(x,y)$.*

Recall that n is the degree of $h(x,y)$ and now we set $m = n$ and look at solutions of $\mathcal{F}_{v(z),n}$. Note that the linear system $\mathcal{F}_{v(z),n}$ lies over the field $\mathbb{F}_q \subset R_v$ which is common to all the components R_{v_j}. Since $\textsc{Deg}(\mathtt{f}_{v_j(z)}(x,y)) \leq n$, by Proposition 5.2.12 all the projected linear systems $\mathcal{F}_{v_j(z),n}$ have a solution. In fact, among all factors $f(x,y)$ of $h(x,y)$ sitting above $v(z)$, $\mathtt{f}_{v(z)}(x,y)$ is the unique one with the minimal possible degree. By the above proposition, we can compute it efficiently by taking the gcd of all the basis polynomials in the solution space of $\mathcal{F}_{v(z),n}$. We record this discussion as a corollary.

Corollary 5.2.13. *Given a factor $v(z)$ of $h(z,0)$ we can compute in deterministic polynomial time the unique minimal degree factor $f(x,y)$ of $h(x,y)$ such that $v(z)$ divides $f(z,0)$.*

The linear systems $\mathcal{L}$ such as $\mathcal{R}_{v(z),m}$ and $\mathcal{B}_{v(z),m}$ will have nontrivial solutions in a projected component field R_{v_j} depending on whether the projected linear system $\mathcal{R}_{v_j(z),m}$ has a nontrivial solution there or not. The next proposition shows that if $\mathcal{L}$ has a nontrivial solution for some but not all the $v_j(z)$'s, then we can use the solutions of $\mathcal{L}$ to factor $v(z)$.

5.2.4 Factoring $v(z)$ using linear systems over R_v.

Recall that $v(z)$ is the product of r irreducible polynomials $v_j(z)$s.

Proposition 5.2.14. *Let $S \subseteq \{1, 2, \cdots r\}$ with the following property: the dimension over $\mathbb{F}_q$ of the solution space of the projected system $\mathcal{L}_{v_j}$ is non-zero if and only if $j \in S$. Then we can compute in deterministic polynomial time the nontrivial factor $(\prod_{j \in S} v_j(z))$ of $v(z)$.*

Proof. (We reproduce the following proof from Gao-Kaltofen-Lauder [GKL04].) Certainly any solution $\mathtt{r}(x,y)$ of $\mathcal{R}_{v(z),m}$ will be sent under the map π_{v_j} to a solution of $\mathcal{R}_{v_j(z),m}$ with entries in R_{v_j}. Moreover this solution will be non-zero if and only if $v_j(z)$ does not divide all of the coefficients in $\mathtt{r}(x,y)$ thought of as polynomials in $\mathbb{F}_q[z]$. Conversely any solution of $\mathcal{R}_{v_j(z),m}$ with entries in R_{v_j} can be lifted using the Chinese Remainder Theorem to a solution for $\mathcal{R}_{v(z),m}$ with entries in R_v.

Now compute a basis over $\mathbb{F}_q$ for the space of solutions in R_v of the linear system $\mathcal{R}_{v(z),m}$. We claim that the greatest commong divisor $g(z)$ say of $v(z)$ and the polynomials that occur as entries in the basis vectors (viewed as polynomials in z) is exactly $\prod_{j\notin S} v_j$.

To see this, suppose $j \in S$. Then there exists some non-zero solution $\mathtt{r}_j(x,y)$ of the linear system $\mathcal{R}_{v_j(z),m}$ which can be lifted to a non-zero solution $\mathtt{r}(x,y)$ of the linear system $\mathcal{R}_{v(z),m}$ as previously described. This solution $\mathtt{r}(x,y)$ has the property that at least one of the entries is not divisible by $v_j(z)$. This solution $\mathtt{r}(x,y)$ of $\mathcal{R}_{v_j(z),m}$ must lie in the $\mathbb{F}_q$-span of the basis vectors of the solution space of $\mathcal{R}_{v(z),m}$. Now if $v_j(z)$ divided all the entries in the basis vectors we would have that $v_j(z)$ divides all the entries of of all vectors in the solution space of $\mathcal{R}_{v(z),m}$ - a contradiction. Hence $v_j(z)$ does not divide $g(z)$. Now suppose, if possible, that $j \notin S$ and also that $v_j(z)$ does not divide $g(z)$. Then $v_j(z)$ does not divide all the entries in the basis vectors of the solution space of $\mathcal{R}_{v_j(z),m}$. Thus there exists at least one basis element $\mathtt{r}(x,y)$ which projects down to a non-zero solution of $\mathcal{R}_{v_j(z),m}$ under π_{v_j} - a contradiction. Thus $g(z)$ is as claimed.

Now one may compute the factor $g(z)$ in deterministic polynomial time using only a deterministic algorithm for computing the solution space over $\mathbb{F}_q$ of the linear system $\mathcal{R}_{v(z),m}$ and the euclidean algorithm for greatest common divisors of univariate polynomials. Moreover, this can be done efficiently in parallel.

□

5.3 The Algorithm.

Proposition 5.3.1. *Let $m \geq 1$ be a natural number and $h(x,y) \in \mathbb{F}_q[x,y]$ a nice polynomial. There is a deterministic polynomial-time algorithm that given $\langle \mathbb{F}_q, h(x,y), m\rangle$ obtains the product of all the $\mathbb{F}_q$-irreducible factors of $h(x,y)$ having degree at most m.*

Proof. Let $f(x,y) \in \mathbb{F}_q[x,y]$ be the product of all $\mathbb{F}_q$-irreducible factors of $h(x,y)$ having degree at most m. Set $v(z)$ to be $h(z,0)$. We claim that the projected linear system $\mathcal{B}_{v_j(z),m}$ has a solution in R_{v_j} if and only if $v_j(z)$ divides $f(z,0)$.

($\Rightarrow$) By Proposition 5.2.12, $\mathtt{b}_{v_j(z)}(x,y) \in \mathbb{F}_q[x,y]$ which is an $\mathbb{F}_q$-irreducible factor of $h(x,y)$ is a solution of the projected system $\mathcal{B}_{v_j(z),m}$. Moreover from the definition of the linear system $\mathcal{B}_{v_j(z),m}$, $\mathtt{b}_{v_j(z)}(x,y)$ has degree at most m. Therefore $\mathtt{b}_{v_j(z)}(x,y)|f(x,y)$. But $v_j(z)|\mathtt{b}_{v_j(z)}(z,0)$ and therefore $v_j(z)|f(z,0)$ as required.

($\Leftarrow$) Since $v_j(z)|f(z,0)$, by the squarefree-ness of $f(z,0)$, there exists a unique $\mathbb{F}_q$-irreducible factor $g(x,y)$ of $f(x,y)$ of degree at most m such that $v_j(z)|g(z,0)$. From the definition of the linear system $\mathcal{B}_{v_j(z),m}$ this polynomial $g(x,y)$ is clearly a solution of $\mathcal{B}_{v_j(z),m}$.

By Proposition 5.2.14, we can recover $f(z,0)$ and using this seed factor of $h(z,0)$ as input to the algorithm of Proposition 5.2.13, we can compute $f(x,y)$ in deterministic polynomial time.

□

Given any polynomial $h(x,y)$ of degree we obtain by the above algorithm a factor $f(x,y)$ consisting of $\mathbb{F}_q$-irreducible factors of degree at most $m := \frac{n}{2}$. Recursively repeating this process (in parallel) on the polynomials $f(x,y)$ and $\frac{h(x,y)}{f(x,y)}$, we obtain a distinct degree factorization of $h(x,y)$ in deterministic time poly$(n \cdot \log q)$. Moreover implementing all the fundamental linear-algebraic operations over $\mathbb{F}_q$ in parallel we can do this in parallel time poly$(\log n \cdot \log q)$.

Proposition 5.3.2. *Let $m, d \geq 1$ be natural numbers and $h(x,y) \in \mathbb{F}_q[x,y]$ a nice polynomial, each of whose $\mathbb{F}_q$-irreducible factors has degree at most m. There is a deterministic polynomial-time algorithm given $\langle \mathbb{F}_q, h(x,y), m, d\rangle$ obtains the product of all the $\mathbb{F}_q$-irreducible factors of $h(x,y)$ having a splitting field of size at least q^d.*

Proof. Let $f(x,y) \in \mathbb{F}_q[x,y]$ be the product of all $\mathbb{F}_q$-irreducible factors of $h(x,y)$ having a splitting field of size at least q^d. Set $v(z)$ to be $h(z,0)$ and $k = \frac{m}{d}$. We claim that the projected linear system $\mathcal{R}_{v_j(z),k}$ has a solution in R_{v_j} if and only if $v_j(z)$ divides $f(z,0)$.

($\Rightarrow$) By Proposition 5.2.12, $\mathtt{r}_{v_j(z)}(x,y) \in R_{v_j}[x,y]$ which is an absolutely irreducible factor of $h(x,y)$ is a solution of the projected system $\mathcal{R}_{v_j(z),k}$. Moreover from the definition of the linear system $\mathcal{R}_{v_j(z),k}$, $\mathtt{r}_{v_j(z)}(x,y)$ has degree at most k. Also $\mathtt{b}_{v_j(z)}(x,y) \in \mathbb{F}_q[x,y]$ is an $\mathbb{F}_q$-irreducible factor of $h(x,y)$ and since all $\mathbb{F}_q$-irreducible factors of $h(x,y)$ have degree m, therefore $\mathtt{b}_{v_j(z)}(x,y)$ also has degree m. Now, $\mathtt{r}_{v_j}(x,y)$ divides $\mathtt{b}_{v_j(z)}(x,y) \in \mathbb{F}_q[x,y]$, an $\mathbb{F}_q$-irreducible factor of $h(x,y)$. By Proposition 5.2.5,

$$\begin{aligned}\text{dimension of } \mathtt{b}_{\mathrm{v_j(z)}}(\mathrm{x,y}) &= \textsc{Deg}(\mathtt{b}_{v_j(z)}(x,y))/\textsc{Deg}(h_{v_j(z)}(x,y))\\ &= m/\textsc{Deg}(h_{v_j(z)}(x,y))\\ &\geq d\end{aligned}$$

Therefore $\mathtt{b}_{v_j(z)}(x,y)|f(x,y)$. But $v_j(z)|\mathtt{b}_{v_j(z)}(z,0)$ and therefore $v_j(z)|f(z,0)$ as required.

($\Leftarrow$) Since $v_j(z)|f(z,0)$, by the squarefree-ness of $f(z,0)$, there exists a unique $\mathbb{F}_q$-irreducible factor $g(x,y)$ of $f(x,y)$ of degree at most m such that $v_j(z)|g(z,0)$. From the definition of the linear system $\mathcal{B}_{v_j(z),m}$ this polynomial $g(x,y)$ is clearly a solution of $\mathcal{B}_{v_j(z),m}$.

By Proposition 5.2.14, we can recover $f(z,0)$ and using this seed factor of $h(z,0)$ as input to the algorithm of Proposition 5.2.13, we can compute $f(x,y)$ in deterministic polynomial time.

□

Given a $h(x,y)$ and m as in the statement of this proposition and setting $d = \frac{m}{2}$, we obtain by the above algorithm a factor $f(x,y)$ consisting of $\mathbb{F}_q$-irreducible factors of dimension at most $d := \frac{m}{2}$. Recursively repeating this process (in parallel) on the polynomials $f(x,y)$ and $\frac{h(x,y)}{f(x,y)}$, we obtain a uniform factorization of $h(x,y)$ in deterministic time $\mathrm{poly}(n \cdot \log q)$, and in parallel time $\mathrm{poly}(\log n \cdot \log q)$.

This completes the proof of Theorem 5.1.2.

5.4 Discussion

The presentation here was complicated by the fact that we also wanted an algorithm that was parallelizable. An easier description for a *sequential* deterministic algorithm achieving the same task can be found in [Kay05]. Finally, we note that in general, the deterministic complexity of factoring polynomials over finite fields remains an open problem and hope that some of the ideas here can also be used to tackle that.

Chapter 6

Solvability of Polynomial Equations over Finite Fields

Summary:
We investigate the complexity of the following polynomial solvability problem: given a finite field $\mathbb{F}_q$ and a set of polynomials $f_1, f_2, \cdots, f_m \in \mathbb{F}_q[x_1, x_2, \cdots, x_n]$ of total degree at most d determine the $\mathbb{F}_q$-solvability of the system $f_1 = f_2 = \cdots = f_m = 0$. This problem is easily seen to be NP-complete even when the field size q is as small as 2 and the degree of each polynomial is bounded by $d = 2$. Here we investigate the deterministic complexity of this problem when the number of variables n in the input is bounded. We show that there is a *deterministic* algorithm for this problem whose running time, for any fixed n, is bounded by a polynomial in d, m and $\log q$.

6.1 Introduction

6.1.1 Motivation

Studying the solution set of a system of polynomial equations is one of the main preoccupations of mathematics. Indeed, three of the most celebrated results of the twentieth century pertain to the solutions of polynomial equations:

- **Weil's Theorem**, also known as the Riemann Hypothesis for curves over finite fields, which gives bounds on the number of rational points on smooth projective curves over finite fields.

- **Falting's Theorem** which states that any curve over $\mathbb{Q}$, the field of rational numbers, of genus greater than 1 has only a finite number of rational points.
- **Wiles' Theorem** which states that the curve $x^n + y^n = 1$ has no nontrivial $(xy \neq 0)$ solution over the field of rational numbers for $n \geq 3$.

This motivates the study of the corresponding computational problems - given a set of polynomials over a field $\mathbb{F}$:

- **Solvability:** Determine whether there exists a common zero of the polynomials.
- **Counting solutions:** Determine the number of common zeroes.
- **Computing a solution:** Compute a common zero, if it exists.

One gets different computational problems depending on whether one is looking for common zeroes in $\mathbb{F}$ itself (i.e. $\mathbb{F}$-rational points) or in the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$. The decidability of the solvability problem for rational points over $\mathbb{Q}$ is an intensively investigated open problem (Poonen [Poo02] gives a survey). In this chapter we consider the solvability problem for rational points over finite fields. We give a deterministic polynomial-time algorithm for the solvability problem over finite fields when the number n of variables in the system is bounded. Our results can be viewed as the natural algorithmic outcome of Weil's theorem. Indeed using Weil's bounds, we get an algorithm with similar complexity for the approximate counting version of the problem. We remark here that given a set of polynomial equations over $\mathbb{Q}$, the field of rational numbers, one can deduce certain properties of the solution set by looking at the reduction of the system of equations modulo p for various primes p and use this information to deduce global values of those properties of the solution set over $\mathbb{Q}$. For certain particularly special sets of polynomial equations over $\mathbb{Q}$, it might be sufficient to verify solvability modulo lots of primes p in order to deduce the existence of a solution over $\mathbb{Q}$. We make one such conjecture in the chapter on open problems. In general, however, there exist polynomials which have lots of $\mathbb{F}_p$-solutions for all primes p but no solution over the rational numbers. Nevertheless, given such a set of equations over $\mathbb{Q}$, one can determine almost all the *geometric properties* such as the numer of $\mathbb{C}$-irreducible components, their dimension and degree of the solution set by looking at the solution set modulo p (see Huang and Wong, [HW00] for details).

Our basic solvability algorithm can be extended in two ways to give more information about the algebraic set $\mathtt{X}$ defined by the given set of polynomials over $\mathbb{F}_q$. We also get efficient deterministic algorithms for:

- Approximating the number of $\mathbb{F}_q$-points on $\mathtt{X}$.
- Computing the number of irreducible components of $\mathtt{X}$ together with the degree and dimension of each such irreducible component.

6.1.2 Problem Definition

Here we are interested in the computational complexity of the solvability problem over the domain of finite fields. The most general version of the polynomial system problem is:

Problem - Existence of solution to a polynomial system (Solvability)

Input. The input is $\langle \mathbb{F}_q, f_1, f_2, \cdots f_m \rangle$ where : (i) $\mathbb{F}_q$ is a finite field with $q = p^r$ being a prime power. The finite field can be specified in the usual way by giving a prime p and an irreducible polynomial of degree r over $\mathbb{F}_p$. (ii) $f_1, f_2, \cdots, f_m \in \mathbb{F}_q[x_1, x_2, \cdots, x_n]$ are m polynomials in the n variables $x_1, x_2, \cdots, x_n$ with coefficients coming from the field $\mathbb{F}_q$. The polynomials are specified using the *dense* representation. That is, a polynomial of degree d in n variables over $\mathbb{F}_q$ has input size $\binom{d+n}{d} \cdot \log q$.

Question. Does there exist a point $(a_1, a_2, \cdots, a_n) \in \mathbb{F}_q^n$ such that

$$f_i(a_1, a_2, \cdots, a_n) = 0 \text{ for all } 1 \leq i \leq m$$

The general polynomial system problem is easily seen to be NP-complete even over a field as small as $\mathbb{F}_2$ and even when all the polynomials in the specified system are of total degree at most 2. This suggests that the problem becomes intractable when the number of variables is large. We examine the complexity of this problem when the number n of variables in the input system is bounded. Huang and Wong [HW96] give a randomized polynomial time (**Z**PP) algorithm for the bounded-variable version of this problem leaving the determintistic complexity unresolved. Our contribution to this problem is to give a *deterministic* polynomial-time algorithm. Moreover, our algorithm works for arbitary finite fields and not just prime fields.

Remark. Consider the slightly more general problem - given a finite field $\mathbb{F}_q$ and polynomials $f_1, f_2, \ldots, f_m$ and $g_1, g_2, \ldots, g_l \in \mathbb{F}_q[\bar{\mathbf{x}}]$ in n variables over $\mathbb{F}_q$, determine if there exists a point $\bar{\mathbf{a}} \in \mathbb{F}_q^n$ such that

$$f_1(\bar{\mathbf{a}}) = \ldots = f_m(\bar{\mathbf{a}}) = 0 \text{ and } g_1(\bar{\mathbf{a}}) \neq 0, \ g_2(\bar{\mathbf{a}}) \neq 0 \ \ldots, \ g_l(\bar{\mathbf{a}}) \neq 0$$

Such an apparently more general problem, involving both equations and 'inequations' over a field is easily seen to reduce to the solvability problem via what is known as the "'Rabinovich trick"' - introduce a new variable y and determine the $\mathbb{F}_q$-solvability of the following system of equations instead:

$$f_1(\bar{\mathbf{x}}) = f_2(\bar{\mathbf{x}}) = \ldots = f_m(\bar{\mathbf{x}}) = 0, \;\; y \cdot g_1(\bar{\mathbf{x}}) \cdot \ldots \cdot g_l(\bar{\mathbf{x}}) = 1$$

Remark. Let $f(x) = \frac{g(x)}{h(x)} \in \mathbb{F}_q(x)$ be a rational function over $\mathbb{F}_q$ with $\gcd(g(x), h(x)) = 1$. Then $f(x)$ induces a partial mapping $\mathbb{F}_q \mapsto \mathbb{F}_q$ via the map $a \mapsto f(a)$ for $a \in \mathbb{F}_q$. If $f(x)$ is total and bijective then $f(x)$ is called a *permutation function* over $\mathbb{F}_q$. In the special case that $h(x) = 1$, so that $f(x) = g(x) \in \mathbb{F}_q[x]$, it is called a *permutation polynomial* over $\mathbb{F}_q$. Permutation functions have been investigated theoretically [Wil68, Mac67, DL63, BD66, Hay67, Coh70], applied in cryptography [LM83] and the complexity of recognizing them dealth with [Shp92, Gat91, Gat89, MG95]. Shparlinski [Shp92] gave a deterministic **superpolynomial**-time algorithm for this problem while Ma and Gathen [Gat91, MG95] gave an efficient **randomized** algorithm. The existence of an efficient deterministic algorithm was open.

Now note that $f(x) = \frac{g(x)}{h(x)}$ is a permutation function if and only if $f(x)$ is total ($h(x) = 0$ has no $\mathbb{F}_q$-solution) and

$$g(x)h(y) - g(y)h(x) = 0, \quad x \neq y$$

has no $\mathbb{F}_q$-solution. Thus, by the remark above, recognizing permutation functions boils down to the solvability problem in 3 variables. Our deterministic solvability algorithm now implies an efficient deterministic algorithm for recognizing permutation functions and thus resolves the deterministic complexity of this problem as well.

6.1.3 Our results

We summarize our main result as a theorem:

Theorem 6.1.1. *There exists a deterministic algorithm which solves the decision version of the Solvability problem on an input consisting of a finite field $\mathbb{F}_q$ and polynomials $f_1, f_2, \cdots, f_m \in \mathbb{F}_q[x_1, x_2, \cdots, x_n]$ of total degree bounded by d in time poly($d^{c_n} \cdot m \cdot \log q$), where c_n is a constant that depends on n alone and is of size $n^{O(n)}$. Moreover, the algorithm can be implemented parallelly to get a family of P -uniform circuits of depth poly($c_n \cdot \log d \cdot \log m \cdot \log q$) and size poly($d^{c_n} \cdot m \cdot \log q$) for the solvability problem.*

The basic algorithm for solvability can be easily extended to get an approximation algorithm of the same complexity for the counting version of the problem. More precisely, the algorithm calculates two non-negative integers N and D, such that $|\#V - Nq^D|$ is bounded by $d^{c_n} q^{D-1/2}$ for some constant c_n that depends on n alone, where $\#V$ denotes the number of common $\mathbb{F}_q$-solutions of the given set of polynomials.

6.1.4 The Idea

The input polynomials with coefficients from $\mathbb{F}_q$ describe an algebraic closed set $\mathtt{X}$. Our aim is to determine if the given closed set $\mathtt{X}$ over the given field $\mathbb{F}_q$ has any $\mathbb{F}_q$-rational point or not. The basic idea is to decompose the given closed set $\mathtt{X}$ into a union of (possibly reducible) closed sets $\mathtt{X}_i$, each $\mathtt{X}_i$ being birational to a hypersurface $\mathtt{Y}_i$. Now Weil's theorem and its generalizations [Sch74, CM03, CM04] imply the abundance of $\mathbb{F}_q$-rational points on any absolutely irreducible $\mathbb{F}_q$-hypersurface. We use the partial factoring algorithm developed in the previous chapter to determine, for each i, if any of the component $\mathbb{F}_q$-irreducible hypersurfaces of $\mathtt{Y}_i$ is absolutely irreducible or not. If $\mathtt{Y}_i$ happens to have an absolutely irreducible $\mathbb{F}_q$-factor, we use Weil's theorem to deduce an abundance of rational points on $\mathtt{Y}_i$ and, via the birational correspondence, on $\mathtt{X}_i$ as well. Otherwise a rational point on $\mathtt{X}_i$, if it exists, must lie on a closed proper subset of $\mathtt{X}_i$. We compute this subset of $\mathtt{X}_i$ and determine the existence of a rational point on it recursively.

Comparison with previous algorithms. Our approach parallels that of Huang and Wong ([HW99]) and it can be viewed as a deterministic modification of their algorithm. Indeed, [HW99] remark that their method actually gives a deterministic reduction to univariate factorization so that the only point that prevents their algorithm from being deterministic is the lack of a deterministic polynomial time algorithm for factoring univariate polynomials over finite fields. *The key contribution of our work on this problem is to observe that as far as the decision version of the problem of solvability is concerned, we do not need to completely factor the multivariate polynomials that arise during this computation process.* In both the works, the algorithm consists of two phases: we first decompose the algebraic closed set corresponding to the given set of equations and reduce the problem to the case of hypersurfaces and then determine the existence of a rational point on the hypersurface by testing for absolutely irreducibility. The difference is that in the first phase, while their algorithm decomposes the set into $\mathbb{F}_q$-irreducible components, the output components of the first phase in our case need not be $\mathbb{F}_q$-irreducible. Our contribution here is to observe that the operations involved and the proofs which hold for irreducible components and their corresponding fields go through with minor modifications when we are working with reducible algebraic sets and their corresponding rings. In the second phase, instead of testing the absolute irreducbility of an $\mathbb{F}_q$-irreducible polynomial, our algorithm uses the output of the partial factoring algorithm developed in the previous chapter. Moreover, they use efficiently parallelizable subroutines developed earlier by Grigoriev, Chistov, et al in order to ensure that the algorithm is efficiently parallelizable with respect to d and m. *We give a self-contained treatment here which preserves this parallelism while eliminating randomness.* Finally, the algorithm in [HW99] works only over prime fields while our algorithm works over all finite fields $\mathbb{F}_q$, even those with a small characteristic p. The difficulty in going from prime fields (q is prime) to general finite fields (q is prime power) is the existence of polynomials $f(x_1, \ldots, x_n) \in \mathbb{F}_q[x_1, \ldots, x_n]$ of degree $d \ll q$ which are squarefree and yet not separable. For example, $f(x_1, x_2) = x_2^p - x_1$ viewed as a univariate polynomial in x_2 over the function field $\mathbb{F} \overset{\mathrm{def}}{=} \mathbb{F}_q(x_1)$ is squarefree and yet has repeated roots in the algebraic closure of $\mathbb{F}$. We overcome this difficulty by observing that a random linear transformation $\sigma \in \mathbb{F}_q^{n \times n}$ on the variables transforms a square-free non-separable polynomial $f(\bar{\mathbf{x}})$ to a separable polynomial in x_n. We then replace $f(\bar{\mathbf{x}})$ by $\sigma(f(\bar{\mathbf{x}}))$ in our computations and work with this transformed polynomial instead.

We flesh out this basic idea in more detail in a later section, after introducing the appropriate terminology and proving some basic facts.

6.2 Basic Algebraic Geometry with Examples

In this section we give a very quick overview of some basic facts from algebraic geometry and introduce the terminology to be used. For proofs see any basic text in algebraic geometry such as Shafarevich [Sha94]. We then give some representative examples.

Algebraic Closed Sets. Let $\mathbb{F}$ be a field. The algebraic closure of $\mathbb{F}$ will be denoted by $\overline{\mathbb{F}}$. A *closed algebraic set over* $\mathbb{F}$ is a subset $\mathtt{X}$ of $\overline{\mathbb{F}}^n$ consisting of all common zeroes of a finite number of polynomials in n variables with coefficients in $\mathbb{F}$. When the field $\mathbb{F}$ is understood from context we will simply refer to $\mathtt{X}$ as a closed algebraic set or just a closed set. A $\mathbb{F}$-*rational point of* $\mathtt{X}$ is a point $P \in \mathtt{X}$ all of whose coordinates are in $\mathbb{F}$.

We shall write $f(\bar{\mathbf{x}})$ to denote a polynomial in n variables, allowing $\bar{\mathbf{x}}$ to stand for the n-tuple of variables $(x_1, x_2, \ldots, x_n)$. If a closed set $\mathtt{X}$ consists of all common zeroes of polynomials $f_1(\bar{\mathbf{x}}), \ldots, f_m(\bar{\mathbf{x}})$, then we refer to $f_1(\bar{\mathbf{x}}) = \cdots = f_m(\bar{\mathbf{x}}) = 0$ as the equations of the set $\mathtt{X}$. We say that $\mathtt{X}$ is a *hypersurface* when it is specified by a single equation ($m = 1$). Observe that a point $P = (a_1, \ldots, a_n) \in \overline{\mathbb{F}}^n$ belongs to the closed algebraic set $\mathtt{X}$ if and only if for all $i \in [m]$, $f_i(x_1 + a_1, \ldots, x_n + a_n)$ has no constant term. $P \in \mathtt{X}$ is said to be a *singular point* of $\mathtt{X}$ iff for all $i \in [m]$, $f_i(x_1 + a_1, \ldots, x_n + a_n)$ has no constant as well as no linear terms. We will say that a closed set $\mathtt{Y}$ is a *singular closed subset of* $\mathtt{X}$ iff every point $P \in \mathtt{Y}$ is a singular point.

A closed algebraic set $\mathtt{X}$ is said to be *reducible* if there exist proper closed subsets $\mathtt{X}_1, \mathtt{X}_2 \subsetneq \mathtt{X}$ such that $\mathtt{X} = \mathtt{X}_1 \cup \mathtt{X}_2$. Otherwise $\mathtt{X}$ is *irreducible*. An irreducible algebraic closed set $\mathtt{X}$ is also referred to as a *variety*.

It is a fundamental theorem in algebraic geometry that any closed agebraic set $\mathtt{X}$ is a finite union of irreducible algebraically closed sets. Now if $\mathtt{X} = \bigcup \mathtt{X}_i$ is an expression of $\mathtt{X}$ as a finite union of irreducible closed sets, and if $\mathtt{X}_i \subseteq \mathtt{X}_j$ then we can delete $\mathtt{X}_i$ from the representation. Repeating this several times, we arrive at a representation $\mathtt{X} = \bigcup \mathtt{X}_i$ in which no $\mathtt{X}_i$ is a subset of any $\mathtt{X}_j$. We say that such a representation is *irredundant*, and the $\mathtt{X}_i$ are the *irreducible components* of $\mathtt{X}$. Such a representation of $\mathtt{X}$ as an irredundant union of a finite number of irreducible algebraic sets is unique.

Let $\mathtt{X} \subseteq \overline{\mathbb{F}}^n$ be an irreducible algebraic closed set (variety) residing in an ambient space of dimension n. Suppose that the minimum possible number of equations required to completely describe $\mathtt{X}$ is m. Then the *dimension* of $\mathtt{X}$, denoted $\ell_{\mathtt{X}}$, is the number $(n - m)$. The varieties contained in an arbitary algebraic closed set are in general of varying dimensions. When all the varieties in a closed set have the same dimension, we will refer to it as a *uniform-dimensional* algebraic closed set.

Correspondence between rings and algebraic sets. Corresponding to the given closed set $\mathtt{X}$ there is a ring $R_{\mathtt{X}}$ obtained by quotienting the polynomial ring $\mathbb{F}_q[\bar{\mathbf{x}}]$ with the ideal generated by the polynomials which are equations of $\mathtt{X}$. That is, if $\mathtt{X}$ is the set of common zeroes of the polynomials $f_1(\bar{\mathbf{x}}), \ldots, f_m(\bar{\mathbf{x}}) \in \mathbb{F}_q[\bar{\mathbf{x}}]$ *the ring* $R_{\mathtt{X}}$ *corresponding to* $\mathtt{X}$ is

$$R_{\mathtt{X}} \overset{\text{def}}{=} \mathbb{F}_q[\bar{\mathbf{x}}]/\langle f_1(\bar{\mathbf{x}}), \cdots, f_m(\bar{\mathbf{x}})\rangle.$$

The elements of $R_{\mathtt{X}}$ can be thought of as functions from $\mathtt{X}$ to $\mathbb{F}_q$, this set of functions itself being endowed with a ring structure. The homomorphisms from $R_{\mathtt{X}}$ to $\mathbb{F}_q$ then correspond to the $\mathbb{F}_q$-rational points on $\mathtt{X}$. Indeed, $\bar{\mathbf{a}} = (a_1, \ldots, a_n) \in \mathbb{F}_q^n$ is an $\mathbb{F}_q$-rational point on $\mathtt{X}$ if and only if the map

$$\phi : R_{\mathtt{X}} \mapsto \mathbb{F}_q, \quad \phi : x_i \mapsto a_i \ \ \forall 1 \leq i \leq n$$

is a homomorphism from $R_{\mathtt{X}}$ to $\mathbb{F}_q$.

In this way the ring $R_{\mathtt{X}}$ captures the algebraic set $\mathtt{X}$ and the structure of the ring $R_{\mathtt{X}}$ corresponds to the structure of $\mathtt{X}$. In particular, $\mathtt{X}$ is $\mathbb{F}_q$-irreducible if and only if $R_{\mathtt{X}}$ is indecomposable. $\mathtt{X}$ is absolutely irreducible (or $\overline{\mathbb{F}}_q$-irreducible) if and only if the ring $\overline{R}_{\mathtt{X}} \overset{\text{def}}{=} \overline{\mathbb{F}}_q[\bar{\mathbf{x}}]/\langle f_1(\bar{\mathbf{x}}), \cdots, f_m(\bar{\mathbf{x}})\rangle$ is indecomposable. In this chapter all the rings R that we will come across will be of the above form (a polynomial ring over $\mathbb{F}_q$ quotiented by some ideal $\mathcal{I}$). We will refer to the closed algebraic set corresponding to the ideal $\mathcal{I}$ as *the closed set of* R. We will denote by $R_{\mathtt{X}}$ the ring corresponding to the closed set $\mathtt{X}$ and by $\mathtt{X}_R$ the algebraic set corresponding to the ring R. We will denote by $R_{\mathtt{X}}^{fr}$ the ring of fractions of $R_{\mathtt{X}}$.

Rational maps between algebraic sets. A map of the form

$$\begin{aligned} y_1 &= \psi_1(x_1, x_2, \ldots, x_n) \\ y_2 &= \psi_2(x_1, x_2, \ldots, x_n) \\ &\vdots \\ y_m &= \psi_m(x_1, x_2, \ldots, x_n), \end{aligned}$$

where the $\psi_i = \frac{G(x_1,\ldots,x_n)}{H(x_1,\ldots,x_n)}$ are ratios of polynomials in the x_j is referred to as a *rational map*. In general, a rational map may be thought of as a function that transforms some set of points $\mathtt{X}$ in $[x_1, \ldots, x_n]$-space to a set of points $\mathtt{Y}$ in $[y_1, \ldots, y_m]$-space. Note that the denominators are polynomials and can have zeroes. Thus the map may not be defined at all points. We denote this map by $\psi : \mathtt{X} \mapsto \mathtt{Y}$. Note that for algebraic sets $\mathtt{X}$ and $\mathtt{Y}$, ψ maps points on $\mathtt{X}$ to points on $\mathtt{Y}$ if and only if the map $y_i \mapsto \psi_i(x_1, \ldots, x_n) \ \ \forall i \in [m]$ is a homomorphism from $R_{\mathtt{Y}}^{fr}$ to $R_{\mathtt{X}}^{fr}$. We will denote this ring homomorphism also by ψ itself.

A rational map $\psi : \mathtt{X} \mapsto \mathtt{Y}$ is called *birational* if it admits an inverse. That is, there exists a rational map $\phi : \mathtt{Y} \mapsto \mathtt{X}$ such that $\psi(\mathtt{X})$ has the same dimension as $\mathtt{Y}$, $\phi(\mathtt{Y})$ has the same dimension as $\mathtt{X}$, $\psi \cdot \phi = 1$ almost everywhere, and $\phi \cdot \psi = 1$ almost everywhere. In terms of the corresponding rings, it means that $(\phi \cdot \psi) : R_{\mathtt{Y}}^{fr} \mapsto R_{\mathtt{Y}}^{fr}$ is the identity map on $R_{\mathtt{Y}}^{fr}$ and $(\psi \cdot \phi) : R_{\mathtt{X}}^{fr} \mapsto R_{\mathtt{X}}^{fr}$ is the identity map on $R_{\mathtt{X}}^{fr}$.

Two algebraic closed sets $\mathtt{X}$ and $\mathtt{Y}$ are said to be *birationally equivalent* or *birational* if there exists a birational map between $\mathtt{X}$ and $\mathtt{Y}$.

A classical theorem from algebraic geometry states that *'Any algebraic variety* $\mathtt{X}$ *is birational to a hypersurface* $\mathtt{Y}$ *of the appropriate dimension'.* This theorem is a direct consequence of the well-known theorem in algebra that every finite-dimensional field extension $\mathbb{K}$ of some base field $\mathbb{F}$ is generated by some element $\gamma \in \mathbb{K}$ (i.e. $\mathbb{K} = \mathbb{F}(\gamma)$) . Moreover it can be arranged that the map $\psi : \mathtt{X} \mapsto \mathtt{Y}$ is just a linear map. That is, each unknown y_j of $\mathtt{Y}$ is expressed as a linear combination of the variables x_i of $\mathtt{X}$. The *degree* of the variety $\mathtt{X}$ is then defined to be the degree of the hypersurface $\mathtt{Y}$ birationally equivalent to $\mathtt{X}$.

6.2.1 Examples

Example: The algebraic set $\mathtt{X}$ defined by the polynomials

$$\begin{aligned} f_1(x,y,z) &= (x+y+z)(x+2y+z) \\ \text{and } f_2(x,y,z) &= (x-y)(x+y-z) \end{aligned}$$

is the irredundant union of four lines -

$$\begin{aligned} \text{line } L_1 :& (x+y+z) = (x-y) = 0, \\ \text{line } L_2 :& (x+y+z) = (x+y-z) = 0, \\ \text{line } L_3 :& (x+2y+z) = (x-y) = 0, \\ \text{and line } L_4 :& (x+2y+z) = (x+y-z) = 0. \end{aligned}$$

Generalization. In general, for polynomials $f_1(\bar{\mathbf{x}}), \ldots, f_m(\bar{\mathbf{x}}) \in \mathbb{F}[x_1, \ldots, x_n]$ where each polynomial $f_i(\bar{\mathbf{x}})$ is the product of d_i linear polynomials *in general position*, the corresponding algebraic set defined by these polynomials is the irredundant union of $(\prod_{i=1}^{m} d_i)$ hyperlines of dimension $(n-m)$.

Example: The algebraic set $\mathtt{X}$ defined by the polynomials

$$\begin{aligned} f_1(x,y,z) &= (x-y)(x+y+z)(x+2y+z) \\ \text{and } f_2(x,y,z) &= (x-y)(x+y-z) \end{aligned}$$

is the irredundant union of a plane

$$\text{plane } P_1 :(x-y) = 0$$

and two lines

$$\begin{aligned} \text{line } L_1 :& (x+y+z) = (x+y-z) = 0 \\ \text{and line } L_2 :& (x+2y+z) = (x+y-z) = 0. \end{aligned}$$

Generalization. We can generalize this example a little. Suppose that $\mathtt{X}$ is an algebraic set defined by the polynomials

$$f_1(x, y, z) = f_2(x, y, z) = 0,$$

where both f_1 and f_2 are products of linear polynomials. Moreover, suppose that $\textsc{Deg}(f_1) = d_1$, $\textsc{Deg}(f_2) = d_2$ and $\textsc{Deg}(\gcd(f_1, f_2)) = d$. Then the closed set $\mathtt{X}$ is the irredundant union of d planes and $(d_1 - d) \cdot (d_2 - d)$ lines.

Example: The algebraic closed set $\mathtt{X}$ in 3 variables x_1, x_2, x_3 defined by the equations

$$x_1^2 - x_3 = x_2^2 - (x_3 + 1) = 0$$

is an irreducible one-dimensional closed set birational to the planar curve $\mathtt{Y}$

$$y^4 - 2(2x + 1)y^2 + 1 = 0$$

via the map

$$\psi : \mathtt{X} \mapsto \mathtt{Y},\ \psi : (x_1, x_2, x_3) \mapsto (x_3, x_1 + x_2)$$

The inverse map ϕ is given by

$$\phi : \mathtt{Y} \mapsto \mathtt{X},\ \phi : (x, y) \mapsto ((\frac{1}{2})(y^3 - (4x + 1)y), \quad (-\frac{1}{2})(y^3 - (4x + 3)y), \quad x).$$

In this example both ψ and ϕ happen to be well-defined *everywhere.*

Generalization. More generally: Suppose that $\mathtt{X}$ is an algebraic closed set in $(n + 1)$ variables $x_1, x_2, \ldots, x_{n+1}$ with defining equations

$$x_1^2 - (x_{n+1} + a_1) = x_2^2 - (x_{n+1} + a_2) = \ldots = x_n^2 - (x_{n+1} + a_n) = 0.$$

Suppose further that the a_i's are all distinct. Then the closed set $\mathtt{X}$ is irreducible and birational to a planar curve of degree 2^n.

Further Generalization. Now suppose that $\mathtt{X}$ is an algebraic closed set in $(n + 1)$ variables $x_1, x_2, \ldots, x_{n+1}$ with defining equations $f_1(\bar{\mathbf{x}}) = \ldots = f_n(\bar{\mathbf{x}}) = 0$ where each $f_i(\bar{\mathbf{x}})$, $1 \le i \le n$ is of the form:

$$f_i(x_1, \ldots, x_n, x_{n+1}) = \prod_{j=1}^{d} (x_i^2 - (x_{n+1} + a_{ij})).$$

Suppose further that the a_{ij}'s are all distinct. Then the closed set $\mathtt{X}$ is a union of d^n irreducible closed sets, each irreducible component being birational to a planar curve of degree 2^n.

Example: We now give an example of a *reducible* one-dimensioal closed set $\mathtt{X}$ being birational to a (reducible) planar curve. Suppose that $f_1(y), f_2(y), g_1(y), g_2(y)$ are univariate polynomials. The algebraic closed set $\mathtt{X}$ in $[x_1, x_2, y]$-space defined by the equations:

$$(x_1 - f_1(y))(x_1 - f_2(y)) = (x_2 - g_1(y))(x_2 - g_2(y)) = 0$$

is reducible and is the union of four irreducible one-dimensional closed sets. $\mathtt{X}$ is birational to the planar curve $\mathtt{Y}$ in $[z, y]$-space defined by the equation:

$$(z - f_1(y) - g_1(y))(z - f_1(y) - g_2(y))(z - f_2(y) - g_1(y))(z - f_2(y) - g_2(y)) = 0$$

via the map

$$\psi : \mathtt{X} \mapsto \mathtt{Y},\ \psi : (x_1, x_2, y) \mapsto (x_1 + x_2, y).$$

The inverse map ϕ is given by

$$\begin{aligned} &\phi : \mathtt{Y} \mapsto \mathtt{X},\ \phi : (z, y) \mapsto (B_1(z, y), B_2(z, y), y) \\ &\qquad \text{with}\ \ B_1(z, y) \overset{\text{def}}{=} A_{11} f_1(y) + A_{12} f_1(y) + A_{21} f_2(y) + A_{22} f_2(y) \\ &\qquad \text{and}\ \ B_2(z, y) \overset{\text{def}}{=} A_{11} g_1(y) + A_{12} g_2(y) + A_{21} g_1(y) + A_{22} g_2(y), \end{aligned}$$

where the coefficient polynomial A_{ij}'s are defined as follows. For $1 \leq i, j \leq 2$ define the polynomial $h_{ij}(u, y)$ as

$$h_{ij}(u, y) \overset{\text{def}}{=} \frac{g(u, y)}{(u - f_i(y) - g_j(y))}.$$

Then for $1 \leq i, j \leq 2$ the coefficent polynomial A_{ij} is

$$A_{ij} \overset{\text{def}}{=} \frac{h_{ij}(z, y)}{h_{ij}(f_i(y) + g_j(y), y)}.$$

6.2.2 Notation

- For an ideal $\mathcal{I} \subseteq \mathbb{F}_q[\bar{\mathbf{x}}]$, we will denote by $\textsc{Rad}(\mathcal{I})$ the radical (square-free part) of the ideal $\mathcal{I}$ defined as

$$\textsc{Rad}(\mathcal{I}) \overset{\text{def}}{=} \{f(\bar{\mathbf{x}}) \in \mathbb{F}_q[\bar{\mathbf{x}}] \mid f(\bar{\mathbf{x}})^m \in \mathcal{I}\ \text{ for some }\ m \geq 1\}.$$

- By the term *total degree* of a rational function $\psi(\bar{\mathbf{x}}) = \frac{F(\bar{\mathbf{x}})}{G(\bar{\mathbf{x}})} \in \mathbb{F}(\bar{\mathbf{x}})$, we will mean the sum of the total degrees of the numerator and the denominator. We denote it by $\textsc{Deg}(\psi)$. That is,

$$\textsc{Deg}(\psi) \overset{\text{def}}{=} \textsc{Deg}(F(\bar{\mathbf{x}})) + \textsc{Deg}(G(\bar{\mathbf{x}}))$$

6.3 Algorithm Description

6.3.1 Overview

In this section we describe in words the proposed algorithm. Our aim is to determine if a given algebraic closed set $\mathtt{X}$ over a given field $\mathbb{F}_q$ has any $\mathbb{F}_q$-rational point or not. (The set $\mathtt{X}$ is specified to us by means of polynomial equations with coefficients from $\mathbb{F}_q$). The basic idea is to decompose the given closed set $\mathtt{X}$ into a union of (possibly reducible) closed sets $\mathtt{X}_i$, each $\mathtt{X}_i$ being birational to a hypersurface $\mathtt{Y}_i$ of the appropriate dimension. We then use the partial factoring algorithm developed in the previous chapter to determine, for each i, the existence of an $\mathbb{F}_q$-rational point on the set $\mathtt{X}_i$.

We flesh out this basic idea in more detail through the rest of this section. We first describe precisely the output of the (deterministic) decomposition algorithm and show how to use our partial factoring algorithm for determining the existence of a rational point on the components of the decomposition. We then describe the decomposition algorithm itself in more detail. Finally, we remark how to improve the parallel time-complexity of the algorithm.

6.3.2 The output of the decomposition and rational points on hypersurfaces

```
Input: A finite field $\mathbb{F}_q$ and a set of polynomials $f_1, \ldots, f_m \in \mathbb{F}_q[x_1, \ldots, x_n]$ of
       total degree at most $d$.
Output: TRUE if $\exists$ an $\mathbb{F}_q$-solution to the system $f_1 = \ldots = f_m = 0$, FALSE
        otherwise.
begin
    let $c_n := 2^n$
    if $q \leq 10^5 n^3 d^{10c_n}$ then
        Check if any of the $q^n$ points in $\mathbb{F}_q^n$ is a common solution to the given
        equations and return accordingly.
    Let $\mathtt{Y}$ be the hypersurface defined by $g(y_1, \ldots, y_n) := \mathrm{RAD}(f_1(y_1, \ldots, y_n))$, $\psi$ be
    the trivial map $\forall i \in [n], \psi : y_i \mapsto x_i$ and $\phi$ be its inverse. Let $\mathtt{X} \subset \overline{\mathbb{F}}_q^n$ be
    $\mathtt{X} := \langle (n-1), \mathtt{Y}, \psi, \phi \rangle$
    return Solvability($\mathtt{X}$, $f_2(\bar{\mathrm{x}}), \ldots, f_m(\bar{\mathrm{x}})$).
end
```

Algorithm 1: **SolvabilityMain :** Determine the existence of an $\mathbb{F}_q$-rational point.

```
Input: A finite field $\mathbb{F}_q$, a component $\mathtt{X} \subset \overline{\mathbb{F}}_q^n$ and a set of polynomials
       $f_1, \ldots, f_m \in \mathbb{F}_q[x_1, \ldots, x_n]$.
Output: TRUE if $\exists$ an $\mathbb{F}_q$-rational point $\bar{\mathbf{a}} \in \mathbb{F}_q^n$ which satisfies $\bar{\mathbf{a}} \in \mathtt{X}$ and
        $f_1(\bar{\mathbf{a}}) = \ldots = f_m(\bar{\mathbf{a}}) = 0$. FALSE otherwise.
begin
    Call Decompose($\mathtt{X}$, $f_1(\bar{\mathbf{x}}), \ldots, f_m(\bar{\mathbf{x}})$ ) to obtain a list $(\mathtt{X}_1, \ldots, \mathtt{X}_t)$ of
    subcomponents of $\mathtt{X}$.
    foreach component $\mathtt{X}_i := \langle \ell, \mathtt{Y}_i, \psi, \phi \rangle$ do
        Let the equation of $\mathtt{Y}_i$ be $g(y_1, \ldots, y_{\ell+1}) = 0$
        if $g(\bar{\mathbf{y}})$ has any absolutely irreducible $\mathbb{F}_q$-factor then return TRUE
        else
            for $j \leftarrow 1$ to $(\ell + 1)$ do
                Compute $h_j(\bar{\mathbf{x}}) := \psi(\frac{\partial g(\bar{\mathbf{y}})}{\partial y_j}) \in \mathbb{F}_q[x_1, \ldots, x_n]$.
            Then the closed set $\mathtt{X}'_i \subsetneq \mathtt{X}_i$,
                $$\mathtt{X}'_i \stackrel{\mathrm{def}}{=} \mathtt{X}_i \cap \left( \bigcap_{j=1}^{\ell+1} \{\bar{\mathbf{a}} \in \overline{\mathbb{F}}_q^n | h_j(\bar{\mathbf{a}}) = 0\} \right),$$
            consists of points $\mathrm{P} \in \mathtt{X}_i$ such that $\psi(\mathrm{P}) \in \mathtt{Y}_i$ is a singular point.
            Recursively determine existence of $\mathbb{F}_q$-rational point on $\mathtt{X}'_i$ by calling
            Solvability($\mathtt{X}_i$, $h_1(\bar{\mathbf{x}}), \ldots, h_{\ell+1}(\bar{\mathbf{x}})$ )
            if $\mathtt{X}'_i$ contains a rational point then return TRUE
    return FALSE
end
```

Algorithm 2: **Solvability :** Determine the existence of an $\mathbb{F}_q$-rational point.

The number of variables is n. We will denote by c_n a constant that depends on n alone and is of size $2^{O(n)}$. Our algorithm is interesting only for large values of q; for if the size q of the given field is *small* ($q = \mathrm{O}(\mathrm{poly}(d^{c_n}))$), we simply do a brute force search over all possible $\mathbb{F}_q$-rational points (q^n many of them) and check if any of them belongs to $\mathtt{X}$. In what follows we shall assume that q is *large* ($q \gg d^{c_n}$).

We break the given algebraic set $\mathtt{X}$ into a union of uniform-dimensional algebraic sets $\mathtt{X}_i$: $\mathtt{X} = \bigcup \mathtt{X}_i$. These $\mathtt{X}_i$'s we call *the components of* $\mathtt{X}$. We represent a component $\mathtt{X}_i$ of $\mathtt{X}$ by a four-tuple $\langle \ell, \mathtt{Y}_i, \psi, \phi \rangle$. where:

- ℓ is the dimension of $\mathtt{X}_i$ and of $\mathtt{Y}_i$,
- $\mathtt{Y}_i$ is a hypersurface with equation $g(y_1, \ldots, y_{\ell+1}) = 0$ for some squarefree $g(\bar{\mathbf{y}}) \in \mathbb{F}_q[\bar{\mathbf{y}}]$.
- $\psi : \mathtt{X}_i \mapsto \mathtt{Y}_i$ is a rational map,
- and $\phi : \mathtt{Y}_i \mapsto \mathtt{X}_i$ is the inverse rational map of ψ.

Note that now $\mathtt{X}$ contains a $\mathbb{F}_q$-rational point if and only if some $\mathtt{X}_i$ contains a $\mathbb{F}_q$-rational point. This computation of the decomposition of $\mathtt{X}$ satisfies the following properties:

P-i). Neither $\mathtt{X}_i$ nor $\mathtt{Y}_i$ contains any singular (repeated) varieties.

P-ii). The map $\psi : \mathtt{X}_i \mapsto \mathtt{Y}_i$ is an $\mathbb{F}_q$-rational map and so is $\phi : \mathtt{Y}_i \mapsto \mathtt{X}_i$. That is the coefficients of all the rational functions occuring in ψ and ϕ are from $\mathbb{F}_q$. In particular this means that $\mathbb{F}_q$-rational points on $\mathtt{X}_i$ get mapped to $\mathbb{F}_q$-rational points on $\mathtt{Y}_i$ and vice-versa.

P-iii). The map $\psi : \mathtt{X}_i \mapsto \mathtt{Y}_i$ is well-defined on all points of $\mathtt{X}_i$. This happens because the corresponding ring homomorphism $\psi : R^{fr}_{\mathtt{Y}_i} \mapsto R^{fr}_{\mathtt{X}_i}$ is actually a linear map, mapping each generator y_i of $R_{\mathtt{Y}_i}$ to a linear combination of the generators x_j's in $R_{\mathtt{X}_i}$.

P-iv). On the other hand, the map $\phi : \mathtt{Y}_i \mapsto \mathtt{X}_i$ is well-defined everywhere except possibly at the singular points of $\mathtt{Y}_i$.

These properties ensure that if there is a $\mathbb{F}_q$-rational point on $\mathtt{X}_i$ then there is one on $\mathtt{Y}_i$ as well. In the other direction, if there is no $\mathbb{F}_q$-rational point on $\mathtt{Y}_i$ **and** there is also no singular point on $\mathtt{Y}_i$ then $\mathtt{X}_i$ does not contain any $\mathbb{F}_q$-rational point as well.

Now consider one such algebraic closed set $\mathtt{Y}_i$ of dimension ℓ. Let the equation of $\mathtt{Y}_i$ be

$$g(y_1, y_2, \ldots, y_l, y_{\ell+1}) = 0.$$

We first handle the zero-dimensional case ($\ell = 0$). In this case the components of $\mathtt{Y}_i$ are simply individual points. Moreover, by the second property, there are no singular points on $\mathtt{Y}_i$. Thus, in this case $\mathtt{X}_i$ has a rational point if and only if the univariate $g(y_1) = 0$ has an $\mathbb{F}_q$-root, or equivalently, if and only if $g(y_1)$ has an absolutely irreducible $\mathbb{F}_q$-factor (see the remark at the end of the defintion of absolute irreducibility 5.1.1).

Now consider the case when $\ell \geq 1$. We use the partial factoring algorithm described in the previous chapter to determine if $g(\bar{\mathbf{y}}) \in \mathbb{F}_q[\bar{\mathbf{y}}]$ contains any absolutely irreducible factors or not. If $g(\bar{\mathbf{y}})$ does have an absolutely irreducible $\mathbb{F}_q$-factor, then from Weil's theorem we can deduce that there does exist an $\mathbb{F}_q$-rational point on $\mathtt{Y}_i$. Indeed, Weil's theorem says that any absolutely irreducible polynomial contains *a lot of* ($\Theta(q^\ell)$, provided q is large enough in comparison to the degree of the polynomial) rational points. Thus if $g(\bar{\mathbf{y}})$ has an absolutely irreducible $\mathbb{F}_q$-factor $g_1(\bar{\mathbf{y}})$ then the hypersurface $g_1(\bar{\mathbf{y}}) = 0$ has *a lot of* $\mathbb{F}_q$-rational points. Moreover, most of these points are non-singular. There is also a partial converse to Weil's theorem: if $g(\bar{\mathbf{y}}) = 0$ has no absolutely irreducible factors then any rational point on $g(\bar{\mathbf{y}}) = 0$, if it exists, is a singular point.

Thus if $g(\bar{\mathbf{y}})$ has an absolutely irreducible factor we deduce that $\mathtt{X}_i$, and hence $\mathtt{X}$, contains a $\mathbb{F}_q$-rational point and we stop. Otherwise any rational point on $\mathtt{X}_i$, if it exists, must map to a singular point on $\mathtt{Y}_i$ under ψ. Now the set of points on $\mathtt{X}_i$ that can map to a singular point on $\mathtt{Y}_i$ under ψ is a closed algebraic subset of $\mathtt{X}_i$ of dimension strictly less than ℓ. We compute the equations of this subset and then repeat the process to determine if this smaller dimensional set has a rational point or not. This process continues until the $\mathtt{X}_i$'s that we get are zero-dimensional.

6.3.3 Description of the decomposition algorithm.

The most general form of the algebraic set decomposition problem is the following -

> **Algebraic Set Decomposition Problem.** Consider a set of polynomials $f_1(\bar{\mathbf{x}}), \ldots, f_m(\bar{\mathbf{x}}) \in \mathbb{F}_q[\bar{\mathbf{x}}]$ of total degree d in n variables over the finite field $\mathbb{F}_q$. Decompose the algebraic set defined by
>
> $$f_1(\bar{\mathbf{x}}) = f_2(\bar{\mathbf{x}}) = \ldots = f_m(\bar{\mathbf{x}}) = 0$$

into $\mathbb{F}_q$-irreducible components, representing each of them by a birational hypersurface over $\mathbb{F}_q$, together with a map from the component to the hypersurface and an inverse rational map from the hypersurface to the component.

Lacking an efficient algorithm for completely factoring univariate polynomials, we cannot solve this most general form of the decomposition problem. We do solve this problem partially and as we shall see, this partial solution is good enough for deciding solvability.

```
Input: A finite field $\mathbb{F}_q$, an algebraic closed set $\mathtt{X}$ over $\mathbb{F}_q$, and polynomials
       $f_1, \ldots, f_m \in \mathbb{F}_q[x_1, \ldots, x_n]$ .
Output: A list $\langle \mathtt{X}_1, \ldots, \mathtt{X}_t \rangle$ of (possibly reducible) components of the closed set
        $\mathtt{X} \cap \left( \bigcap_{i=1}^{m} \{\bar{\mathbf{a}} \in \overline{\mathbb{F}}_q^n | f_i(\bar{\mathbf{a}}) = 0\} \right)$.
begin
    Initialize a list $L$ with the single component $\mathtt{X}$.
    for $i \leftarrow 1$ to $m$ do
        Initialize $L'$ to be the empty list.
        forall $\widehat{\mathtt{X}} := \langle \ell, \mathtt{Y}, \psi, \phi \rangle$ in the list $L$ do
            Let $\mathbb{F} \overset{\text{def}}{=} \mathbb{F}_q(y_1, \ldots, y_\ell)$, $R_{\mathtt{Y}}^{fr} := \mathbb{F}[y_{\ell+1}]/\langle g(y_1, \ldots, y_{\ell+1}) \rangle$,
                $f_i^{\phi}(y_1, \ldots, y_{\ell+1}) := \phi(f_i(x_1, \ldots, x_n)) = \frac{h_1(y_1, \ldots, y_{\ell+1})}{h_2(y_1, \ldots, y_\ell)} \in R_{\mathtt{Y}}^{fr}$
            where $h_1$ and $h_2$ are polynomials in the $y_j$'s.
            We now have two hypersurfaces $g(\bar{\mathbf{y}}) = 0$ and $h_1(\bar{\mathbf{y}}) = 0$ in the ambient
            $[y_1, \ldots, y_{\ell+1}]$-space. Compute the intersection of these two hypersurfaces
            and obtain two components $\widehat{\mathtt{Y}}_1 := \langle \ell, \mathtt{Z}_1, \psi_1, \phi_1 \rangle$ and
            $\widehat{\mathtt{Y}}_2 := \langle \ell - 1, \mathtt{Z}_2, \psi_2, \phi_2 \rangle$.
            if $\widehat{\mathtt{Y}}_1 \neq \emptyset$ then
                add the component $\widehat{\mathtt{X}}_1 := \langle \ell, \mathtt{Z}_1, \psi_1 \circ \psi, \phi \circ \phi_1 \rangle$ to $L'$.
            if $\widehat{\mathtt{Y}}_2 \neq \emptyset$ then
                add the component $\widehat{\mathtt{X}}_2 := \langle \ell - 1, \mathtt{Z}_2, \psi_2 \circ \psi, \phi \circ \phi_2 \rangle$ to $L'$.
        $L \leftarrow L'$
    Output the list $L$
end
```

Algorithm 3: **Decompose -** Compute the decomposition of an algebraic set.

```
Input: A finite field $\mathbb{F}_q$ and two $\ell$-dimensional hypersurfaces
       $\mathtt{Y}_1 : g_1(y_1, \ldots, y_{\ell+1}) = 0$ and $\mathtt{Y}_2 : g_2(y_1, \ldots, y_{\ell+1}) = 0$.
Output: The decomposition of $(\mathtt{Y}_1 \cap \mathtt{Y}_2)$ as the union of two component closed
        subsets $\widehat{\mathtt{Y}}_1 := \langle \ell, \mathtt{Z}_1, \psi_1, \phi_1 \rangle$ and $\widehat{\mathtt{Y}}_2 := \langle \ell - 1, \mathtt{Z}_2, \psi_2, \phi_2 \rangle$.
begin
    $g_1(\bar{\mathbf{y}}) \leftarrow \textsc{Rad}(g_1(\bar{\mathbf{y}}))$,   $g_2(\bar{\mathbf{y}}) \leftarrow \textsc{Rad}(g_2(\bar{\mathbf{y}}))$
    By making a suitable linear transformation $\sigma$ on the variables $y_1, \ldots y_{\ell+1}$,
    ensure that both $\sigma(g_1(\bar{\mathbf{y}}))$ and $\sigma(g_2(\bar{\mathbf{y}}))$ are monic and seperable polynomials
    with respect to $y_{\ell+1}$.
        Let  $\mathbb{F} \overset{\text{def}}{=} \mathbb{F}_q(y_1, \ldots, y_{\ell-1})$,  $R := \mathbb{F}(y_\ell)[y_{\ell+1}]/\langle \sigma(g_1(\bar{\mathbf{y}})), \sigma(g_2(\bar{\mathbf{y}})) \rangle$.
    Compute
        $h(\bar{\mathbf{y}}) := \gcd(\sigma(g_1(\bar{\mathbf{y}})), \sigma(g_2(\bar{\mathbf{y}}))), h_1(\bar{\mathbf{y}}) := \frac{\sigma(g_1(\bar{\mathbf{y}}))}{h(\bar{\mathbf{y}})}, h_2(\bar{\mathbf{y}}) := \frac{\sigma(g_2(\bar{\mathbf{y}}))}{h(\bar{\mathbf{y}})}$.
    Note that $h(\bar{\mathbf{y}}), h_1(\bar{\mathbf{y}}), h_2(\bar{\mathbf{y}}) \in \mathbb{F}_q[\bar{\mathbf{y}}]$ are all monic polynomials in $y_{\ell+1}$. The ring
    $R$ then decomposes into the direct sum of two rings:
        $R = \left(R_1 \overset{\text{def}}{=} \mathbb{F}(y_\ell)[y_{\ell+1}]/\langle h(\bar{\mathbf{y}}) \rangle\right) \oplus \left(R_2 \overset{\text{def}}{=} \mathbb{F}[y_\ell, y_{\ell+1}]/\textsc{Rad}(\langle h_1(\bar{\mathbf{y}}), h_2(\bar{\mathbf{y}}) \rangle)\right)$
    Let $\pi_1 : R \mapsto R_1$ and $\pi_2 : R \mapsto R_2$ be the projection maps. Also let $\rho_1 : R_1 \mapsto R$
    and $\rho_2 : R_2 \mapsto R$ be the natural inclusion maps.
    if $\textsc{Deg}(h(\bar{\mathbf{y}})) = 0$ then $\widehat{\mathtt{Y}}_1 \overset{\text{def}}{=} \emptyset$ else
        $\widehat{\mathtt{Y}}_1 \overset{\text{def}}{=} \langle \ell, \mathtt{Z}_1 := \{\bar{\mathbf{a}} \in \overline{\mathbb{F}}_q^{\ell+1} \mid h(\bar{\mathbf{a}}) = 0\}, \sigma^{-1} \cdot \rho_1, \pi_1 \cdot \sigma \rangle$
    Viewing $R_2$ as an algebra over $\mathbb{F}$, use the primitive element theorem to obtain a
    ring $R_{\mathtt{Z}}^{fr} := \mathbb{F}[z]/\langle \tilde{g}(z) \rangle$ such that $\phi : R_2 \mapsto R_{\mathtt{Z}}^{fr}$ is an isomorphism with inverse
    $\psi$. Here $\mathtt{Z} := \{\bar{\mathbf{a}} \in \overline{\mathbb{F}}_q^{\ell} \mid \tilde{g}(\bar{\mathbf{a}}) = 0\}$ is the algebraic closed set corresponding to
    $R_{\mathtt{Z}}^{fr}$.
    if $\mathtt{Z} = \emptyset$ then $\widehat{\mathtt{Y}}_2 = \emptyset$ else $\widehat{\mathtt{Y}}_2 \overset{\text{def}}{=} \langle \ell - 1, \mathtt{Z}, \sigma^{-1} \cdot \rho_2 \cdot \psi, \phi \cdot \pi_2 \cdot \sigma \rangle$
    return $\langle \widehat{\mathtt{Y}}_1, \widehat{\mathtt{Y}}_2 \rangle$.
end
```

Function `Intersect` - Compute the intersection of two hypersurfaces.

We now delve a little deeper and describe in more detail the process of computing the components together with their birationally equivalent hypersurfaces.

Let $\mathtt{X}^{[i]}$ be the closed set defined by the first i equations:

$$f_1(\bar{\mathbf{x}}) = f_2(\bar{\mathbf{x}}) = \ldots = f_i(\bar{\mathbf{x}}) = 0.$$

Corresponding to the closed set $\mathtt{X}^{[i]}$ we have the ring

$$R_{\mathtt{X}}^{[i]} := \mathbb{F}_q[\bar{\mathbf{x}}]/\langle f_1(\bar{\mathbf{x}}), \ldots, f_i(\bar{\mathbf{x}})\rangle.$$

Starting with $i = 1$, our algorithm successively computes the decomposition of $\mathtt{X}^{[i]}$ for $i = 2, 3, \ldots, m$ until we get the decomposition of $\mathtt{X}^{[m]} = \mathtt{X}$. Our algorithm ensures that at each stage the components that we get are all '*square-free*', i.e. each variety in the component occurs with multiplicity 1.

In order to get the decomposition of the closed set $\mathtt{X}^{[i+1]}$ from that of $\mathtt{X}^{[i]}$, we compute the intersection of each component of $\mathtt{X}^{[i]}$ with the hypersurface $\mathtt{Z}$ defined by $f_{i+1}(\bar{\mathbf{x}}) = 0$. Consider one such component $\widehat{\mathtt{X}}$ of $\mathtt{X}^{[i]}$, of dimension ℓ. Then $\widehat{\mathtt{X}} \cap \mathtt{Z}$ is the union of two components $\widehat{\mathtt{X}}_1$ and $\widehat{\mathtt{X}}_2$. $\widehat{\mathtt{X}}_1$ is the union of those ℓ-dimensional varieties in $\widehat{\mathtt{X}}$ that are a subset of $\mathtt{Z}$. Each of the remaining varieties in $\widehat{\mathtt{X}} - \widehat{\mathtt{X}}_1$ give a collection of $(\ell-1)$-dimensional varieties upon intersection with $\mathtt{Z}$, the union of which is the set $\widehat{\mathtt{X}}_2$. In this way intersecting a component of $\mathtt{X}^{[i]}$ with the hypersurface $f_{i+1}(\bar{\mathbf{x}}) = 0$ gives, in general, two components of $\mathtt{X}^{[i+1]}$. Continuing in this manner we get the decomposition of $\mathtt{X} = \mathtt{X}^{[m]}$. It remains for us to describe how to compute the intersection of a component with a hypersurface.

Computing the intersection of a component $\widehat{\mathtt{X}} := \langle \ell, \mathtt{Y}, \psi, \phi\rangle$ with a hypersurface $f_i(\bar{\mathbf{x}}) = 0$. The component $\widehat{\mathtt{X}}$ is birational to a ℓ-dimensional hypersurface $\mathtt{Y}$ with defining equation $g(y_1, \ldots, y_{\ell+1}) = 0$. We 'project' the constraint $f_i(\bar{\mathbf{x}}) = 0$ into the ambient $[y_1, \ldots, y_{\ell+1}]$-space of $\mathtt{Y}$ by using the map $\phi : R_{\widehat{\mathtt{X}}}^{fr} \mapsto R_{\mathtt{Y}}^{fr}$. Thus the problem now boils down to computing the intersection of two hypersurfaces $g_1(\bar{\mathbf{y}}) := g(\bar{\mathbf{y}})$ and $g_2(\bar{\mathbf{y}}) := \phi(f_i(\bar{\mathbf{x}}))$. After some initial preprocessing, we compute $h(\bar{\mathbf{y}}) = \gcd(g_1(\bar{\mathbf{y}}), g_2(\bar{\mathbf{y}}))$ and this captures all the varieties common to both $g_1(\bar{\mathbf{y}}) = 0$ and $g_2(\bar{\mathbf{y}}) = 0$. The hypersurface $h(\bar{\mathbf{y}}) = 0$ then gives us the representation of $\widehat{\mathtt{X}}_1$. After removing these common varieties from both $g_1(\bar{\mathbf{y}}) = 0$ and $g_2(\bar{\mathbf{y}}) = 0$, our problem boils down to computing an $(\ell - 1)$-dimensional hypersurface birational to the intersection of two 'disjoint' ℓ-dimensional hypersurfaces. We solve this problem by using the primitive element theorem as described in the next section and upon composing the relevant maps we obtain a hypersurface-representation of $\widehat{\mathtt{X}}_2$ as well.

In summary, we obtain the decomposition of the given set by introducing the constraints one by one and at each stage computing the intersection of every component with the newly introduced constraint. This completes the description of the sequential version of our algorithm.

▮ *Time complexity of the algorithm.*

The computation of the decomposition of the given algebraic set $\mathtt{X}$ can be viewed in terms of a binary tree of depth m where the nodes at depth i correspond to the components in the decomposition of the closed set $\mathtt{X}^{[i]}$. We will observe that the degree of any hypersurface is bounded by d^{c_n}. Also, the total degree of every rational function that occurs in the map from the given set $\mathtt{X}$ to the hypersurfaces that occur during the computation process is also bounded by d^{c_n}. From this it follows that the total number of $\mathbb{F}_q$-field operations that we require is $\mathrm{poly}(d^{c_n} \cdot k_m)$ where k_m is the number of components output by the decomposition algorithm. Finally, k_m is itself upper-bounded by d^{c_n} thereby implying an overall time complexity of $\mathrm{poly}(d^{c_n} \cdot m)$ field operations over $\mathbb{F}_q$. Note that both the degree and the number of components of $\mathtt{X}$ are bounded by d^{c_n}, a quantity that, remarkably, is independent of m.

▮ *Parallelizing the algorithm.*

Consider once again the binary tree corresponding to the computation of the decomposition algorithm as mentioned in the previous section.

The fundamental operations involved in the decomposition algorithm are computing the gcd of two polynomials, solving a set of linear equations and computing the characteristic polynomial of a matrix. All of these are all well-studied operations known to be efficiently parallelizable. Thus, by doing an efficient parallel implementation of these fundamental operations and a parallel traversal of the aforementioned computation tree, we get a parallel time complexity of $poly(c_n \cdot \log d \cdot m \cdot \log q)$. To make the dependence polylogarithmic in m also we need one more idea. The idea is simply to divide the given set of m equations into two sets of size $\frac{m}{2}$, compute the decomposition of the closed algebraic set induced by each set of equations recursively in parallel and then take the intersection of each pair of components to get the decomposition of the original algebraic set $\mathtt{X}$. Let $\widehat{\mathtt{X}}$ be the algebraic closed set corresponding to the equations

$$f_1(\bar{\mathbf{x}}) = f_2(\bar{\mathbf{x}}) = \ldots = f_{\frac{m}{2}}(\bar{\mathbf{x}}) = 0$$

and $\widetilde{\mathtt{X}}$ be the algebraic closed set corresponding to the rest of the equations

$$f_{\frac{m}{2}+1}(\bar{\mathbf{x}}) = f_{\frac{m}{2}+1}(\bar{\mathbf{x}}) = \ldots = f_m(\bar{\mathbf{x}}) = 0.$$

We recursively compute the decomposition of $\widehat{\mathtt{X}}$ and $\widetilde{\mathtt{X}}$ in parallel. Let $\widehat{\mathtt{X}} = \bigcup_i \widehat{\mathtt{X}}_i$ and $\widetilde{\mathtt{X}} = \bigcup_j \widetilde{\mathtt{X}}_j$ be the decomposition of $\widehat{\mathtt{X}}$ and $\widetilde{\mathtt{X}}$ respectively. Then the decomposition of $\mathtt{X}$ is given simply by $\mathtt{X} = \bigcup_{i,j}(\widehat{\mathtt{X}}_i \cap \widetilde{\mathtt{X}}_j)$. The intersection of every pair of sets $\widehat{\mathtt{X}}_i$ and $\widetilde{\mathtt{X}}_j$ is computed again in parallel and computing one such intersection again involves elementary linear algebraic operations which are also efficiently parallelized. Overall, this gives a parallel time complexity of $poly(c_n \cdot \log d \cdot \log m \cdot \log q)$.

6.3.4 The Primitive Element Theorem

We now come to the main technical section of our algorithm - computing the intersection of two hypersurfaces. In this subsection we give a very constructive version of the well known primitive element theorem (cf. Lang [Lan94]), along with explicit bounds on the sizes of the involved quantities, as required for our purposes.

Consider polynomials

$$f_1(z_1, \ldots, z_n, x) \in \mathbb{F}_q[z_1, \ldots, z_n, x, y] \quad \text{and} \quad f_2(z_1, \ldots, z_n, y) \in \mathbb{F}_q[z_1, \ldots, z_n, x, y].$$

Let $f_1(\bar{\mathbf{z}}, x)$ and $f_2(\bar{\mathbf{z}}, y)$ be squarefree polynomials of total degree d_1 and d_2 respectively over $\mathbb{F}_q$. Moreover, suppose that $f_1(\bar{\mathbf{z}}, x)$ is monic and separable with respect to the variable x while $f_2(\bar{\mathbf{z}}, y)$ is monic and separable with respect to the variable y.

Remark. If f_1 and f_2 are not monic and separable then a random linear transformation σ on the variables makes them monic and separable so that in this case we apply the appropriate linear transform on the variables and work with these new polynomials instead. See [Kal82] for a proof of the bivariate case. The proof of the general case in n variables is an easy generalization of the bivariate case. Moreover, when the number of variables is bounded such a transformation σ can be computed efficiently [Kal82].

Let $\mathbb{F}$ be the rational function field $\mathbb{F} \stackrel{\text{def}}{=} \mathbb{F}_q(z_1, \ldots, z_n)$. Let R be the ring $\mathbb{F}[x, y]/\langle f_1(x), f_2(y)\rangle$. Thus R is an algebra of dimension $d_1 \cdot d_2$ over the field $\mathbb{F}$ with basis

$$\mathcal{B}_1 \stackrel{\text{def}}{=} \{x^i y^j \mid 0 \le i < d_1, 0 \le j < d_2\}.$$

We want to express R as a ring of the form $\mathbb{F}[z]/\langle g(z)\rangle$. We will see that choosing $g(z)$ to be the minimal polynomial of some element $\alpha \in R$ of the form $\alpha = x + ty$ with $t \in \mathbb{F}_q$ works for us.

Suppose that in the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$, f_1 and f_2 factor as:

$$f_1(x) = \prod_{i=1}^{d_1}(x - \alpha_i),$$

$$f_2(y) = \prod_{j=1}^{d_2}(x - \beta_j).$$

By the squarefreeness and separability of f_1 the α_i's are all distinct. Similarly the β_j's are all distinct.

Now for some $t \in \mathbb{F}_q$, consider the element $\alpha \in R$ defined as $\alpha \overset{\text{def}}{=} (x + ty)$. Then the characteristic polynomial of α over $\mathbb{F}$ is

$$g(z) \overset{\text{def}}{=} \text{charpoly}_{\alpha/\mathbb{F}}(z) = \prod_{i=1}^{d_1}\prod_{j=1}^{d_2}(z - (\alpha_i + t\beta_j)).$$

Let $A \subset \overline{\mathbb{F}}$ be the set

$$A \overset{\text{def}}{=} \{(\alpha_{i_1} - \alpha_{i_2})/(\beta_{j_1} - \beta_{j_2}) \ \mid i_1, i_2 \in [d_1], j_1 \neq j_2 \in [d_2]\}.$$

Then for $t \notin A$, the roots of $g(z)$ are all distinct. Fix any such $t \notin A$. Then since the characteristic polynomial $g(z)$ of α is squarefree and separable, it is in fact also the minimal polynomial of α. Therefore $R = \mathbb{F}(\alpha) = \mathbb{F}[z]/\langle g(z)\rangle$. Choosing any $t \in (\mathbb{F}_q \setminus A)$ gives a suitable α. Note that $\mid A \mid < d_1^2 d_2^2$ and therefore there are at least $\mid (\mathbb{F}_q \setminus A) \mid \geq (q - d_1^2 d_2^2)$ suitable choices of t.

We now adopt a slightly different viewpoint of the above matter. The discussion above explicitly exhibits an isomorphism ψ from the ring $R_1 \overset{\text{def}}{=} \mathbb{F}[z]/\langle g(z)\rangle$ to the ring $R \overset{\text{def}}{=} \mathbb{F}[x, y]/\langle f_1(x), f_2(y)\rangle$ given by $\psi : z \mapsto (x+ty)$, where $g(\bar{\mathbf{z}}, z) \in \mathbb{F}_q[\bar{\mathbf{z}}, z]$ is the minpoly of the element $(x + ty) \in R$. Let $\phi : R \mapsto R_1$ be the inverse of ψ. Clearly then ϕ can be viewed as a map from the set of points $\mathtt{Y}$ on $g(\bar{\mathbf{z}}, z) = 0$ to the set of points $\mathtt{X}$ on $f_1(\bar{\mathbf{z}}, x) = f_2(\bar{\mathbf{z}}, y) = 0$. ψ then maps the points on $\mathtt{X}$ to points on $\mathtt{Y}$ and by the linear nature of the map, ψ is well-defined everywhere.

We now investigate the well-definedness of ϕ as a map from points in $\mathtt{Y}$ to points in $\mathtt{X}$. For $P = (\bar{\mathbf{z}}, z)$ let $\phi(P) = (\bar{\mathbf{z}}, \phi_1(P), \phi_2(P))$. Over the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$, we can obtain an explicit expression for ϕ_1 as a polynomial in z. Indeed this expression is remniscient of polynomial interpolation for the following reason. If $x = \phi_1(z) \in \overline{\mathbb{F}}[z]$ is the expression for x in terms of z then we want it to satisfy $\phi_1(\alpha_i + t\beta_j) = \alpha_i$ for all $i \in [d_1]$ and $j \in [d_2]$. Let $g_{ij}(z) \overset{\text{def}}{=} \frac{g(z)}{z-(\alpha_i+t\beta_j)} \in \overline{\mathbb{F}}[z]$. Its easy to verify that

$$\phi_1(z) := \sum_{i,j} \frac{g_{ij}(z)}{g_{ij}(\alpha_i + t\beta_j)} \alpha_i$$

works. It turns out the rhs of the above equation is actually in $\mathbb{F}[z]$ itself. From the above expression, we can deduce that $\phi_1(P)$ is well defined for all non-singular points P on $\mathtt{Y}$. Similarly, it can be shown that $\phi_2(P)$ is also well-defined for all non-singular points P on $\mathtt{Y}$.

Let us summarize the above discussion far as a theorem.

Proposition 6.3.1. *(***Primitive Element Theorem.***) Let $\mathbb{F}_q$ be a finite field. Let $f_1(z_1, \ldots, z_n, x) \in \mathbb{F}_q[z_1, \ldots, z_n, x, y]$ and $f_2(z_1, \ldots, z_n, y) \in \mathbb{F}_q[z_1, \ldots, z_n, x, y]$ be square-free polynomials of degree d_1 and d_2 respectively over $\mathbb{F}_q$. Moreover, $f_1(\bar{\mathbf{z}}, x)$ is monic and separable with respect to the variable x while $f_2(\bar{\mathbf{z}}, y)$ is monic and separable with respect to the variable y. Let $\mathbb{F}$ be the rational function field $\mathbb{F} \overset{\text{def}}{=} \mathbb{F}_q(z_1, \ldots, z_n)$. Let R be the ring $\mathbb{F}[x, y]/\langle f_1(x), f_2(y)\rangle$. Thus R is an algebra of dimension $d_1 \cdot d_2$ over the field $\mathbb{F}$. Then R is isomophic to the ring $R_1 := \mathbb{F}[z]/\langle g(z)\rangle$, where $g(z) \in \mathbb{F}_q[z_1, \ldots, z_n, z]$ is a polynomial of degree $(d_1 \cdot d_2)$ and is monic in z. The map $\psi : R_1 \mapsto R$, $\psi : z \mapsto (x + ty)$ for some $t \in \mathbb{F}_q$ is a ring isomorphism. Let $\phi : R \mapsto R_1$ be the inverse of ψ. Then ϕ maps points on the closed set of $g(\bar{\mathbf{z}}, z) = 0$ to points on the closed set of $f_1(\bar{\mathbf{z}}, x) = f_2(\bar{\mathbf{z}}, y) = 0$ in such a way that it is well-defined on all non-singular points on $g(\bar{\mathbf{z}}, z) = 0$.*

Moreover the ring R_1 together with the maps ψ and ϕ can be constructed in deterministic polynomial time (i.e. time polynomial in the size of the input and output).

6.3.5 Intersection of two hypersurfaces.

Now suppose that we are given two $(n + 1)$-dimensional hypersurfaces

$$f_1(z_1, \ldots, z_n, x, y) = 0 \quad \text{and} \quad f_2(z_1, \ldots, z_n, x, y) = 0$$

over the field $\mathbb{F}_q$. Moreover assume that f_1 and f_2 have no common varieties, i.e. the polynomials $f_1(\bar{\mathbf{z}}, x, y)$ and $f_2(\bar{\mathbf{z}}, x, y)$ are coprime. We want to compute an n-dimensional hypersurface $g(\bar{\mathbf{z}}, z) = 0$ birational to their intersection

$$f_1(\bar{\mathbf{z}}, x, y) = f_2(\bar{\mathbf{z}}, x, y) = 0.$$

Equivalently, we want to compute a ring R_1 of the form

$$R_1 = \mathbb{F}_q(\bar{\mathbf{z}})[z]/\langle g(z)\rangle$$

that is $\mathbb{F}_q(\bar{\mathbf{z}})$-isomorphic to the given ring

$$R = \mathbb{F}_q(\bar{\mathbf{z}})[x,y]/\textsc{Rad}(\langle f_1(\bar{\mathbf{z}},x,y), f_2(\bar{\mathbf{z}},x,y)\rangle).$$

We do this as follows:

1. Compute

   $$\begin{aligned} h_1(\bar{\mathbf{z}},x) &= \textsc{Rad}(\textsc{Resultant}_y(f_1(\bar{\mathbf{z}},x,y), f_2(\bar{\mathbf{z}},x,y))) \neq 0 \\ \text{and}\quad h_2(\bar{\mathbf{z}},y) &= \textsc{Rad}(\textsc{Resultant}_x(f_1(\bar{\mathbf{z}},x,y), f_2(\bar{\mathbf{z}},x,y))) \neq 0. \end{aligned}$$

   Then over the field $\mathbb{F} = \mathbb{F}_q(\bar{\mathbf{z}})$, since $h_1(x)$ and $h_2(y) \in \langle f_1(x,y), f_2(x,y)\rangle$, we have

   $$R = (\mathbb{F}[x,y]/\langle h_1(x), h_2(y)\rangle)/\langle f_1(x,y), f_2(x,y)\rangle.$$

2. Let $S \overset{\text{def}}{=} \mathbb{F}[x,y]/\langle h_1(x), h_2(y)\rangle$. Using the primitive element theorem described previously obtain a ring S' of the form $S' = \mathbb{F}[z]/\langle g_1(z)\rangle$ along with isomorphisms $\phi : S \mapsto S'$ and $\psi : S' \mapsto S$.

3. Viewing $f_1(x,y)$ and $f_2(x,y)$ as elements of S, compute

   $$f_1'(z) = \phi(f_1(x,y)) \in S', \quad f_2'(z) = \phi(f_2(x,y)) \in S'.$$

   Then $R \subseteq S$ is isomorphic to $S'/\langle f_1'(z), f_2'(z)\rangle = \mathbb{F}[z]/\langle g(z)\rangle$ where $g(z) = gcd(g_1(z), f_1'(z), f_2'(z))$. The restriction of the map ϕ to $R \subseteq S$ provides the isomorphism from R to $R_1 := \mathbb{F}[z]/\langle g(z)\rangle \subseteq S'$.

Clearly all these computations are in deterministic polynomial time. Finally, when ψ and ϕ are viewed as mappings from one algebraic closed set to another, ψ is well defined at all points whereas ϕ is well-defined at all non-singular points. We summarize this as a theorem.

Proposition 6.3.2. *Let $\mathbb{F}_q$ be a finite field. Let $f_1(z_1,\ldots,z_n,x,y) \in \mathbb{F}_q[z_1,\ldots,z_n,x,y]$ and $f_2(z_1,\ldots,z_n,x,y) \in \mathbb{F}_q[z_1,\ldots,z_n,x,y]$ be squarefree polynomials of degree d_1 and d_2 respectively over $\mathbb{F}_q$. Let $\mathbb{F}$ be the rational function field $\mathbb{F} \overset{\text{def}}{=} \mathbb{F}_q(z_1,\ldots,z_n)$. Let R be* *the ring $\mathbb{F}[x,y]/\langle f_1(x), f_2(y)\rangle$. Then R is isomophic to the ring $R_1 := \mathbb{F}[z]/\langle g(z)\rangle$, where $g(z) \in \mathbb{F}_q[z_1, \ldots, z_n, z]$. The map $\psi : R_1 \mapsto R,\quad \psi : z \mapsto (x + ty)$ for some $t \in \mathbb{F}_q$ is a ring isomorphism. Let $\phi : R \mapsto R_1$ be the inverse of ψ. Then ϕ maps points on the closed set of $g(\bar{\mathbf{z}}, z) = 0$ to points on the closed set of $f_1(\bar{\mathbf{z}}, x) = f_2(\bar{\mathbf{z}}, y) = 0$ in such a way that it is well-defined on all non-singular points on $g(\bar{\mathbf{z}}, z) = 0$.*

Moreover the ring R_1 together with the maps ψ and ϕ can be constructed in deterministic polynomial time (i.e. time polynomial in the size of the input and output).

6.3.6 Proof of Correctness

We now prove the correctness of our algorithm. The main subroutine involved in the decomposition is computing the intersection of two hypersurfaces. The important properties of this intersection algorithm and its proof of correctness has already been discussed. So we can assume that the decomposition algorithm works correctly and returns a list of components of the given algebraic set. Now consider a uniform-dimensional component $\mathtt{X}_i := \langle \ell, \mathtt{Y}_i, \psi, \phi\rangle$ in the list of components returned by the decomposition algorithm. Our algorithm then consists of two cases.

Case I: $\mathtt{Y}_i$ contains an absolutely irreducible hypersurface. We will make use of the following two results by Schmidt [Sch74].

Theorem 6.3.3. *Suppose $g(y_1, \ldots, y_{\ell+1})$ is an absolutely irreducible polynomial of total degree $d > 0$, with coefficients in the finite field $\mathbb{F}_q$. Let A be the number of $\mathbb{F}_q$-rational points on*

$$g(y_1, \ldots, y_{\ell+1}) = 0.$$

Suppose

$$q > 10^4 \ell^3 d^5 P^3(4 \lfloor logd \rfloor),$$

where $P(1) = 2, P(2) = 3, \ldots$ is the sequence of primes. In particular $P(x) \approx x \log x$, and hence the right hand side of the above inequality is $O(\ell^3 d^{5+\epsilon})$ for every $\epsilon > 0$. Then

$$A > q^\ell - (d)(d-1)q^{\ell-(1/2)}.$$

Theorem 6.3.4. *Suppose $g_1(y_1, \ldots, y_{\ell+1}), \ldots, g_m(y_1, \ldots, y_{\ell+1})$ are polynomials of degree $\leq d$ with coefficients in $\mathbb{F}_q$ and without a common factor. Then the number of $\mathbb{F}_q$-rational points on*

$$g_1(y_1, \ldots, y_{\ell+1}) = \ldots = g_m(y_1, \ldots, y_{\ell+1}) = 0$$

is $\leq 2\ell d^3 q^{\ell-1}$.

Combining these two theorems we prove the following:

Theorem 6.3.5. *Suppose that $g(y_1, \ldots, y_{\ell+1}) \in \mathbb{F}_q[y_1, \ldots, y_{\ell+1}]$ is a squarefree polynomial of total degree d having at least one absolutely irreducible $\mathbb{F}_q$-factor. If $q \geq 10^5\ell^3d^{10}$, then there exists at least one* non-singular $\mathbb{F}_q$-rational point *on the hypersurface*

$$g(y_1, \ldots, y_{\ell+1}) = 0.$$

Proof. If $d = 1$ then $g(\bar{\mathbf{y}})$ is simply a hyperplane of dimension (ℓ) and thus all the q^ℓ $\mathbb{F}_q$-rational points on $g(\bar{\mathbf{y}})$ are non-singular. So now assume $d \geq 2$.

Since $g(\bar{\mathbf{y}})$ is squarefree, therefore it must be coprime to at least one of its partial derivatives $\left(\frac{\partial g}{\partial y_i}\right)(\bar{\mathbf{y}})$. So by theorem 6.3.4 the system of equations

$$g(\bar{\mathbf{y}}) = \left(\frac{\partial g}{\partial y_1}\right)(\bar{\mathbf{y}}) = \ldots = \left(\frac{\partial g}{\partial y_{\ell+1}}\right)(\bar{\mathbf{y}}) = 0$$

has at most $2\ell d^3q^{\ell-1}$ solutions. In other words the number of $\mathbb{F}_q$-rational singular points on $g(\bar{\mathbf{y}})$ is upper bounded by $2\ell d^3q^{\ell-1}$.

Let $g_1(\bar{\mathbf{y}}) \in \mathbb{F}_q[\bar{\mathbf{y}}]$ be an absolutely irreducible $\mathbb{F}_q$-factor of $g(\bar{\mathbf{y}})$. Combining the lower bound of 6.3.3 on the number of $\mathbb{F}_q$-rational points on $g_1(\bar{\mathbf{y}}) = 0$ with this upper bound on the number of singular points on $g(\bar{\mathbf{y}}) = 0$, we get that there exists at least one non-singular $\mathbb{F}_q$-rational point on $g(\bar{\mathbf{y}}) = 0$.

□

We need to bound the number and degree of the components of various dimensions obtained as our algorithm s. We bound it as follows.

Lemma 6.3.6. *During the execution of the algorithm, the degree of any ℓ-dimensional component is at most $d^{2^{n-1-\ell}}$.*

Proof. We proceed by induction on $s \overset{\text{def}}{=} n - \ell$.

Base case $s = 1$. Any $(n-1)$ dimensional component of $\mathtt{X}$ simply corresponds to an $\mathbb{F}_q$-factor of the polynomial $f_1(\bar{\mathbf{x}})$ and thereofore its degree is bounded by d, as required.

Induction step. Now any $(\ell-1)$-dimensional component $\mathtt{X}_i$ of $\mathtt{X}$ is obtained by the intersection of a component $\widehat{\mathtt{X}}$ (obtained reviously during the computation) of dimension at most ℓ and a hypersurface $h(\bar{\mathbf{x}}) = 0$. The hypersurface $h(\bar{\mathbf{x}})$ is either one of the original input hypersurfaces $f_i(\bar{\mathbf{x}}) = 0$ or is of the form $\psi(\frac{\partial g}{\partial y_j}(\bar{\mathbf{y}}))$ for some birationally equivalent hypersurface $g(\bar{\mathbf{y}}) = 0$. In either case, the induction hypothesis implies that the degree d_h

of the hypersurface $h(\bar{\mathbf{x}}) = 0$ is bounded by $d_h \leq 2^{d^s}$. By induction hypothesis the degrees $\hat{d}$ of $\widehat{\mathtt{X}}$ is bounded by $\hat{d} \leq d^{2^s}$. By Bezout's theorem, the degree of $\mathtt{X}_i$ which is a component in their intersection is bounded by

$$\begin{aligned} d_h \cdot \hat{d} &\leq d^{2^s} \cdot d^{2^s} \\ &= d^{2^{s+1}}, \end{aligned}$$

as required.

□

Suppose that the corresponding birational hypersurface $g(\bar{\mathbf{y}}) = 0$ contains an absolutely irreducible $\mathbb{F}_q$-factor. Then the claim here is that $\mathtt{X}_i$ does indeed contain a rational point. If the dimension ℓ is zero, then the absolutely irreducible $\mathbb{F}_q$-factors of $g(y_1)$ are nothing but $\mathbb{F}_q$-points on $g(y_1) = 0$. The components output by the decomposition algorithm do not contain any singular varieties and thus no such point P is a singular point of $\mathtt{Y}_i$ and therefore $\phi(P)$ gives an $\mathbb{F}_q$-point on $\mathtt{X}_i$ as desired. If $\ell \geq 1$, then by theorem 6.3.5, $\mathtt{Y}_i$ contains a non-singular rational point P and therefore $\phi(P)$ gives a rational point on $\mathtt{X}_i$ as claimed.

Case II : $\mathtt{Y}_i$ has no absolutely irreducible $\mathbb{F}_q$-factors. We make use of the following lemma from the previous chapter.

Lemma 6.3.7. *Suppose that $h(\bar{\mathbf{y}}) \in \mathbb{F}_q[\bar{\mathbf{y}}]$ is $\mathbb{F}_q$-irreducible and it splits into absolutely irreducible factors $h_1(\bar{\mathbf{y}}), \ldots, h_t(\bar{\mathbf{y}})$ over some extension field $\mathbb{F}_{q^d}$ of $\mathbb{F}_q$. Then its absolutely irreducible factors $h_i(\bar{\mathbf{y}})$'s are all $\mathbb{F}_q$-conjugates of each other.*

Now consider a rational point P on the hypersurface $g(\bar{\mathbf{y}}) = 0$. Then P must be the zero of some $\mathbb{F}_q$-irreducible factor $h(\bar{\mathbf{y}})$ of $g(\bar{\mathbf{y}})$. That is $h(P) = 0$. Suppose $h(\bar{\mathbf{y}})$ splits completely over the extension field $\mathbb{K} \supsetneq \mathbb{F}_q$ into factors $h_1(\bar{\mathbf{y}}), \ldots, h_t(\bar{\mathbf{y}}) \in \mathbb{K}[\bar{\mathbf{y}}]$. Then P must be a rational point on some factor, say $h_1(\bar{\mathbf{y}})$, of $h(\bar{\mathbf{y}})$. Let $\sigma \in Gal_{\mathbb{K}/\mathbb{F}_q}$ be an automorphism of $\mathbb{K}$ mapping $h_1(\bar{\mathbf{y}})$ to $h_2(\bar{\mathbf{y}})$. Then since P is an $\mathbb{F}_q$-rational point, we have $\sigma(P) = P$. So P is also a zero of $h_2(\bar{\mathbf{y}})$ and hence P is a singular point on the surface $h(\bar{\mathbf{y}}) = 0$. Consequently P is also a singular point on the surface $g(\bar{\mathbf{y}}) = 0$.

Now a point P on $g(\bar{\mathbf{y}}) = 0$ is singular if and only if it is the common zero of the closed subset $\mathtt{Y}' \subsetneq \mathtt{Y}_i$ with defining equations

$$g(\bar{\mathbf{y}}) = \left(\frac{\partial g}{\partial y_1}\right)(\bar{\mathbf{y}}) = \ldots = \left(\frac{\partial g}{\partial y_{\ell+1}}\right)(\bar{\mathbf{y}}) = 0$$

By imposing the constraints $h_i(\bar{\mathbf{x}}) = \psi(\left(\frac{\partial g}{\partial y_1}\right)(\bar{\mathbf{y}}))$ on the algebraic set $\mathtt{X}_i$, our algorithm computes the preimage $\mathtt{X}' \subsetneq \mathtt{X}$ of $\mathtt{Y}'$. In this case then there is an $\mathbb{F}_q$-rational point P in $\mathtt{X}$ if and only if there is one in $\mathtt{X}'$, which our algorithm determines recursively, as required.

This completes the proof of correctness of our algorithm. We summarize it as a theorem.

Theorem 6.3.8. *Algorithm 1 is a deterministic algorithm which decides Solvability on an input consisting of a finite field $\mathbb{F}_q$ and polynomials $f_1, f_2, \cdots, f_m \in \mathbb{F}_q[x_1, x_2, \cdots, x_n]$ of total degree bounded by d in time poly($d^{c_n} \cdot m \cdot \log q$), where c_n is a constant that depends on n alone and is of size $n^{O(n)}$.*

Discussion

In this chapter we devised a deterministic algorithm for determining the existence of a rational point on a variety by using Weil estimates for the number of rational points on an absolutely irreducible curves and the deterministic factoring algorithm of the last chapter. The major open problem in this direction now is to deterministically compute a rational point if it exists and to count their number efficiently.