

Solving Equations via Algebras

David A. Cox
Department of Mathematics and Computer Science
Amherst College
Amherst, MA 01002 USA

Abstract. These notes, written for the CIMPA Graduate School on *Systems of Polynomial Equations (Argentina, 2003)*, will discuss the rich mathematics involved in solving polynomial equations via eigenvalues and eigenvectors.

Contents

1 Solving Equations	3
1.1 Finite-Dimensional Algebras and Gröbner Bases	3
1.2 Eigenvalues of Multiplication Maps	5
1.3 Eigenvectors of Multiplication Maps	7
1.4 Single-Variable Representation	11
1.5 Generalized Eigenspaces and Multiplicities	12
2 Ideals Defined By Linear Conditions	13
2.1 Duals and Dualizing Modules	13
2.2 Differential Conditions Defining Ideals	14
2.3 Two Algorithms	18
2.4 Ideals of Points and Basis Conversion	22
3 Resultants	23
3.1 Solving Equations	23
3.2 The U-Resultant	27
4 Factoring	31
4.1 Factoring over Number Fields	31
4.2 Finite Fields and Primitive Elements	34
4.3 Primary Decomposition	36
5 Galois Theory	42
5.1 Splitting Algebras	42
5.2 The Galois Group	44
5.3 Primary Decomposition	50

Introduction

These notes will consider the quotient rings

$$A = F[x_1, \dots, x_n]/\langle f_1, \dots, f_s \rangle$$

where F is a field and f_1, \dots, f_s are polynomials with coefficients in F in the variables x_1, \dots, x_n . Of course, besides being a ring, A is also a vector space over F in a compatible way. We express this by saying that A is an *algebra* over F .

We will be most interested in the case when A has finite-dimension as a vector space over F . This is what we mean by the word “algebras” in the title “Solving Equations via Algebras” of these notes.

What’s Covered. Our basic claim is that some wonderful mathematics applies to this situation, including the following:

- Solving equations.
- Resultants.
- Factoring over number fields and finite fields.
- Primary decomposition.
- Galois theory.

The main reason for this richness is that an element a of a finite commutative algebra A gives a *multiplication map*

$$m_a : A \longrightarrow A$$

defined by $m_a(b) = ab$ for $b \in A$. This is a linear map from a finite-dimensional vector space to itself, which means that many tools of linear algebra can be brought to bear to study m_a . Furthermore, since A is commutative, the linear maps m_a all commute as we vary $a \in A$.

What’s Omitted. These notes do not discuss everything of interest connected with finite commutative algebras. The two main topics not covered are:

- Gorenstein duality.
- Real solutions.

A careful treatment of duality can be found in [EM1] and [EM2], and a discussion of real solutions of polynomial equations appears in [CLO2].

Some Notation. Let A be a finite commutative algebra as above. If $f \in F[x_1, \dots, x_n]$, then we will use the following notation:

- $[f] \in A$ is the coset of f in the quotient ring A shown above.
- m_f is the multiplication map $m_{[f]}$. Thus $m_f([g]) = [fg]$ for all $[g] \in A$.
- M_f is the matrix of m_f relative to a chosen basis of A over F .

Other notation will be introduced as needed.

1 Solving Equations

1.1 Finite-Dimensional Algebras and Gröbner Bases

Consider a system of polynomial equations

$$\begin{aligned} f_1(x_1, \dots, x_n) &= 0 \\ f_2(x_1, \dots, x_n) &= 0 \\ &\vdots \\ f_s(x_1, \dots, x_n) &= 0 \end{aligned} \tag{1.1.1}$$

in n variables x_1, \dots, x_n with coefficients in a field F . In this section, we will address the following questions:

- When does (1.1.1) have only finitely many solutions over the algebraic closure \overline{F} of F ?
- When (1.1.1) has only finitely many solutions over \overline{F} , how do we find them?

As we will see, the algebra

$$A = F[x_1, \dots, x_n]/\langle f_1, \dots, f_s \rangle \tag{1.1.2}$$

has a crucial role to play in both of these questions.

Before going any further, let's give an example taken from [MS] which we will use throughout this section and the next.

Example 1.1.1 Consider the equations

$$\begin{aligned} f_1 &= x^2 + 2y^2 - 2y = 0 \\ f_2 &= xy^2 - xy = 0 \\ f_3 &= y^3 - 2y^2 + y = 0 \end{aligned} \tag{1.1.3}$$

over the complex numbers \mathbb{C} . If we write the the third equation as

$$f_3 = y(y-1)^2 = 0$$

and the first equation as

$$f_1 = x^2 + 2y(y-1) = 0,$$

then it follows easily that the only solutions are the points

$$(0, 0) \quad \text{and} \quad (0, 1).$$

However, this ignores multiplicities, which as we will see are perfectly captured by the algebra $A = \mathbb{C}[x, y]/\langle f_1, f_2, f_3 \rangle$. \square

Our first major result is a necessary and sufficient condition for the algebra A corresponding to the equations (1.1.1) to be finite-dimensional over F .

Theorem 1.1.2 *The algebra A defined in (1.1.2) is finite-dimensional over F if and only if the equations (1.1.1) have only finitely many solutions over the algebraic closure \overline{F} .*

Proof. We will sketch the main ideas since this result is so important. A complete proof can be found in Chapter 5, §3 of [CLO1].

First suppose that A is finite-dimensional over F . Then, for each i , the set $\{[1], [x_i], [x_i^2], \dots\} \subset A$ must be linearly dependent, so that there is a nonzero polynomial $p_i(x_i)$ such that $[p_i(x_i)] = [0]$ in A . This means that

$$p_i(x_i) \in \langle f_1, \dots, f_s \rangle,$$

which easily implies that p_i vanishes at all common solutions of (1.1.1). It follows that for each i , the solutions have only finitely many distinct i th coordinates. Hence the number of solutions is finite.

Going the other way, suppose that there are only finitely many solutions over \overline{F} . Then in particular there are only finitely many i th coordinates, so that we can find a nonzero polynomial $q_i(x_i)$ which vanishes on all solutions of (1.1.1) over \overline{F} . In this situation, *Hilbert's Nullstellensatz* (see Chapter 4, §1 of [CLO1] for a proof) asserts that

$$p_i(x_i) = q_i^N(x_i) \in \langle f_1, \dots, f_s \rangle$$

for some sufficiently large integer N .

Now consider the lexicographic order $>_{\text{lex}}$ on monomials $x^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n}$. Recall that $x^\alpha > x^\beta$ if $a_1 > b_1$, or $a_1 = b_1$ and $a_2 > b_2$, or \dots (in other words, the left-most nonzero entry of $\alpha - \beta \in \mathbb{Z}^n$ is positive). This allows us to define the *leading term* of any nonzero polynomial in $F[x_1, \dots, x_n]$.

The theory of Gröbner bases (explained in Chapters 2 and 5 of [CLO1]) implies that $\langle f_1, \dots, f_s \rangle$ has a *Gröbner basis* g_1, \dots, g_t with the following properties:

- g_1, \dots, g_t are a basis of $\langle f_1, \dots, f_s \rangle$.
- The leading term of every nonzero element of $\langle f_1, \dots, f_s \rangle$ is divisible by the leading term of one of the g_j .
- The set of *remainder monomials*

$$\mathcal{B} = \{x^\alpha \mid x^\alpha \text{ is not divisible by the leading term of any } g_j\}$$

gives the cosets $[x^\alpha]$, $x^\alpha \in \mathcal{B}$, which form a basis of the quotient algebra A over F .

Since the leading term of $p_i(x_i)$ is a power of x_i , the second bullet implies that the leading term of some g_j is a power of x_i . It follows that in any $x^\alpha \in \mathcal{B}$, x_i must appear to strictly less than this power. Since this is true for all i , it follows that \mathcal{B} is finite, so that A is finite-dimensional by the third bullet. \square

Let's apply this to our example.

Example 1.1.3 For the equations $f_1 = f_2 = f_3 = 0$ of Example 1.1.1, one can show that f_1, f_2, f_3 form a Gröbner basis for lexicographic order with $x > y$. Thus the leading terms of the polynomials in the Gröbner basis are

$$x^2, xy^2, y^3,$$

so that the remainder monomials (= monomials not divisible by any of these leading terms) are

$$\mathcal{B} = \{1, y, y^2, x, xy\}.$$

Hence A has dimension 5 over \mathbb{C} in this case. \square

1.2 Eigenvalues of Multiplication Maps

For the remainder of this section, we will assume that

$$A = F[x_1, \dots, x_n]/\langle f_1, \dots, f_s \rangle$$

is finite-dimensional over F . For simplicity of exposition, we will also assume that

$$F = \overline{F}.$$

Thus F will always be algebraically closed.

As in the introduction, a polynomial $f \in F[x_1, \dots, x_n]$ gives a multiplication map

$$m_f : A \longrightarrow A.$$

Our main result is the following observation first noticed by Lazard in 1981 (see [Lazard]).

Theorem 1.2.1 *Assume that (1.1.1) has a finite positive number of solutions. The eigenvalues of m_f are the values of f at the solutions of (1.1.1) over F .*

Proof. We will sketch the proof and refer to Theorem 4.5 of Chapter 2 of [CLO2] for the details.

First suppose $\lambda \in F$ is not a value of f at a solution of (1.1.1). Then the equations

$$f - \lambda = f_1 = \dots = f_s = 0$$

have no solutions over $F = \overline{F}$, so that by the Nullstellensatz, we can write

$$1 = h \cdot (f - \lambda) + \sum_{i=1}^s h_i f_i$$

for some polynomials $h, h_1, \dots, h_s \in F[x_1, \dots, x_n]$. Since the multiplication map m_1 is the identity 1_A and each m_{f_i} is the zero map, it follows that

$$m_f - \lambda 1_A = m_{f-\lambda} : A \longrightarrow A$$

is an isomorphism with inverse given by m_h . Thus λ is not an eigenvalue of m_f .

Going the other way, let $p \in F^n$ be a solution of (1.1.1). As in the proof of Theorem 1.1.2, the remainder monomials $\mathcal{B} = \{x^{\alpha(1)}, \dots, x^{\alpha(m)}\}$ give the basis $[x^{\alpha(1)}], \dots, [x^{\alpha(m)}]$ of A . The matrix of m_f relative to this basis is denoted M_f . For $j = 1, \dots, m$, let $p^{\alpha(j)}$ be the element of F obtained by evaluating $x^{\alpha(j)}$ at p . Then we claim that

$$M_f^T (p^{\alpha(1)}, \dots, p^{\alpha(m)})^T = f(p) (p^{\alpha(1)}, \dots, p^{\alpha(m)})^T, \quad (1.2.1)$$

where T denotes transpose. Since 1 is an element of \mathcal{B} , the vector $(p^{\alpha(1)}, \dots, p^{\alpha(m)})^T$ is nonzero. Thus (1.2.1) implies that $f(p)$ is an eigenvalue of M_f^T and hence also of M_f and m_f .

To prove (1.2.1), suppose that $M_f = (m_{ij})$. This means that

$$[x^{\alpha(j)} f] = \sum_{i=1}^m m_{ij} [x^{\alpha(i)}]$$

for $j = 1, \dots, m$. Then $x^{\alpha(j)} f \equiv \sum_{i=1}^m m_{ij} x^{\alpha(i)} \pmod{\langle f_1, \dots, f_s \rangle}$. Since f_1, \dots, f_s all vanish at p , evaluating this congruence at p implies that

$$p^{\alpha(j)} f(p) = \sum_{i=1}^m m_{ij} p^{\alpha(i)}$$

for $j = 1, \dots, m$. This easily implies (1.2.1). \square

Example 1.2.2 For the polynomials of Examples 1.1.1 and 1.1.3, the set $\mathcal{B} = \{1, y, y^2, x, xy\}$ gives a basis of A . Then one easily sees that the matrix of m_x is

$$M_x = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 2 \\ 0 & 0 & 0 & -2 & -2 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \end{pmatrix}.$$

The first and second columns are especially easy to see since here m_x maps basis elements to basis elements. For the third column, one uses $f_2 = xy^2 - xy$ to show that

$$m_x([y^2]) = [xy^2] = [xy].$$

The fourth and fifth columns are obtained similarly. Using Maple or Mathematica, one finds that the characteristic polynomial of M_x is $\text{CharPoly}_{M_x}(u) = u^5$. By Theorem 1.2.1, it follows that all solutions of the equations (1.1.3) have x -coordinate equal to 0.

In a similar way, one finds that m_y has matrix

$$M_y = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 & 0 \\ 0 & 1 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{pmatrix}.$$

with characteristic polynomial $\text{CharPoly}_{M_y}(u) = u^2(u-1)^3$. Thus the y -coordinate of a solution of (1.1.3) is 0 or 1. For later purposes we also note that

$$M_x \text{ has minimal polynomial } u^3$$

$$M_y \text{ has minimal polynomial } u(u-1)^2.$$

We will see that later that since y takes distinct values 0 and 1 at the solution, the characteristic polynomial $u^2(u-1)^3$ of M_y tells us the *multiplicities* of the solutions of (1.1.3). \square

In general, the matrix of $m_f : A \rightarrow A$ is easy to compute once we have a Gröbner basis G of $\langle f_1, \dots, f_s \rangle$. This is true because of the following:

- As we saw in the proof of Theorems 1.1.2 and 1.2.1, G determines the remainder monomials \mathcal{B} which give a basis of A .
- Given $g \in F[x_1, \dots, x_n]$, the division algorithm from Chapter 2, §3 of [CLO1] constructs a *normal form*

$$\text{NF}(g) \in \text{Span}(\mathcal{B})$$

with the property that $g \equiv \text{NF}(g) \pmod{\langle f_1, \dots, f_s \rangle}$.

This gives an easy algorithm for computing M_f with respect to the basis of A given by \mathcal{B} : For each $x^\alpha \in \mathcal{B}$, simply compute $\text{NF}(x^\alpha f)$ using the division algorithm. This is what we did in Example 1.2.2.

In particular, we can compute the matrix M_{x_i} for each i , and then Theorem 1.2.1 implies that the x_i -coordinates of the solutions are given by the eigenvalues of M_{x_i} . But how do we put these coordinates together to figure out the actual solutions? This was trivial to do in Example 1.2.2. In the general case, one could simply try all possible combinations of the coordinates to find the solutions. But this is very inefficient.

1.3 Eigenvectors of Multiplication Maps

A better method for solving equations, first described in [AS], is to use the *eigenvectors* of M_f^T given by (1.2.1), namely

$$M_f^T(p^{\alpha(1)}, \dots, p^{\alpha(m)})^T = f(p)(p^{\alpha(1)}, \dots, p^{\alpha(m)})^T.$$

In this equation, p is a solution of (1.1.1), $\mathcal{B} = \{x^{\alpha(1)}, \dots, x^{\alpha(m)}\}$, and $p^{\alpha(j)}$ is the element of F obtained by evaluating $x^{\alpha(j)}$ at p . As we noted in the proof of Theorem 2.1, (1.2.1) implies that $(p^{\alpha(1)}, \dots, p^{\alpha(m)})^T$ is an eigenvector of M_f^T for the eigenvalue $f(p)$.

This allows us to use eigenvalues to find solutions as follows. Suppose that all eigenspaces of M_f^T have dimension 1 (we say that M_f^T is *non-derogatory* in this case). Then suppose that λ is an eigenvalue of M_f^T with eigenvector

$$\vec{v} = (u_1, \dots, u_m)^T.$$

By assumption, we know that \vec{v} is unique up to a scalar. At this point, we know that $\lambda = f(p)$ for some solution p , but we don't know what p is.

To determine p , observe that $(p^{\alpha(1)}, \dots, p^{\alpha(m)})^T$ is also an eigenvector of M_f^T for λ . Since we may assume that $x^{\alpha(1)} = 1$, the first coordinate of this eigenvector is 1. Since λ has a 1-dimensional eigenspace, our computed eigenvector \vec{v} is a scalar multiple of $(p^{\alpha(1)}, \dots, p^{\alpha(m)})^T$. Hence, if we rescale \vec{v} so that its first coordinate is 1, then

$$\vec{v} = (1, u_2, \dots, u_m)^T = (1, p^{\alpha(2)}, \dots, p^{\alpha(m)})^T. \quad (1.3.1)$$

The key point is that the monomials $x^{\alpha(j)} \in \mathcal{B}$ include some (and often all) of the variables x_1, \dots, x_n . This means that we can read off the corresponding coordinates of p from \vec{v} . Here is an example of how this works.

Example 1.3.1 Consider the matrices M_x^T and M_y^T from Example 1.2.2. Neither is non-derogatory since their eigenspaces all have dimension 2. However, if we set $f = 2x + 3y$, then

$$M_f^T = 2M_x^T + 3M_y^T$$

is non-derogatory such that

$$\text{the eigenvalue } 0 \text{ has eigenbasis } \vec{v} = (1, 0, 0, 0, 0)^T$$

$$\text{the eigenvalue } 3 \text{ has eigenbasis } \vec{v} = (1, 1, 1, 0, 0)^T.$$

Since $\mathcal{B} = \{1, y, y^2, x, xy\}$ has the variables x and y in the fourth and second positions respectively, it follows from (1.3.1) that the x - and y -coordinates of the solutions are the fourth and second entries of the eigenvectors. This gives the solutions

$$(0, 0) \quad \text{and} \quad (0, 1)$$

found in Example 1.1.1. □

We should also mention that being non-derogatory is equivalent to saying that the minimal polynomial equals the characteristic polynomial. This will have a nice consequence in Section 1.4 below.

In order for this method to work in general, we need to answer the following questions:

- What happens when some variables are missing from \mathcal{B} ?
- Can we find $f \in F[x_1, \dots, x_n]$ such that M_f^T is non-derogatory? What happens if we can't?

The remainder of Section 1.3 will be devoted to answering these questions.

Missing Variables. Once we fix a monomial ordering, the ideal $\langle f_1, \dots, f_s \rangle$ has a Gröbner basis G . We assume that G is *reduced*, which means the following:

- The leading coefficient of every $g \in G$ is 1.
- For any $g \in G$, its non-leading terms are not divisible by the leading terms of the remaining polynomials in G .

As we've already explained, G then determines the remainder monomials

$$\mathcal{B} = \{x^\alpha \mid x^\alpha \text{ is not divisible by the leading term of any } g \in G\}.$$

We will assume that $G \neq \{1\}$, which implies that $1 \in \mathcal{B}$ and that (1.1.1) has solutions in F (the latter is true by the Consistency Algorithm described in Chapter 4, §1 of [CLO1] since $F = \overline{F}$).

We need to understand which variables lie in \mathcal{B} . We say that x_i is *known* if $x_i \in \mathcal{B}$ and *missing* otherwise (this is not standard terminology). As explained above, if M_f^T is non-derogatory, then the eigenvectors determine the known coordinates of all solutions. It remains to find the missing coordinates. We will analyze this using the arguments of [MS].

A variable x_i is missing if it is divisible by the leading term of some element of G . Since $G \neq \{1\}$ and is reduced, it follows that there is some $g_i \in G$ such that

$$g_i = x_i + \text{terms strictly smaller according to the term order.}$$

Furthermore, since this is true for every missing variable and G is reduced, it follows that other terms in the above formula for g_i involve only known variables (if a missing variable appeared in some term, it would be a missing variable $x_j \neq x_i$, so that the term would be divisible by the leading term of $g_j = x_j + \dots \in G$). Thus

$$g_i = x_i + \text{terms involving only known variables.}$$

Now let p be a solution of (1.1.1). Then $g_i(p) = 0$, so that the above analysis implies that

$$0 = p_i + \text{terms involving only known coordinates.}$$

Hence the $g_i \in G$ tell us how to find the missing coordinates in terms of the known ones.

Derogatory Polynomials. Our first observation is that there are systems of equations such that M_f^T is derogatory for *all* polynomials $f \in F[x_1, \dots, x_n]$. Here is a simple example.

Example 1.3.2 Consider the equations

$$x^2 = y^2 = 0.$$

The only solution is $p = (0, 0)$ and $\mathcal{B} = \{1, x, y, xy\}$. Given $f = a + bx + cy + dx^2 + exy + \dots$ in $F[x, y]$, we have $f(p) = a$ and

$$M_f^T = \begin{pmatrix} a & b & c & e \\ 0 & a & 0 & c \\ 0 & 0 & a & b \\ 0 & 0 & 0 & a \end{pmatrix}.$$

Since $M_f^T - aI_4$ has rank ≤ 2 , it follows that the eigenspace of M_f^T for the eigenvalue $f(p) = a$ has dimension at least 2. Thus M_f^T is always derogatory. \square

To describe what happens in general, we need to discuss the local structure of solutions. A basic result from commutative algebra states that the ideal $\langle f_1, \dots, f_s \rangle$ has a *primary decomposition*. Since we are over algebraically closed field and the equations (1.1.1) have only finitely many solutions, the primary decomposition can be written

$$\langle f_1, \dots, f_s \rangle = \bigcap_p I_p \quad (1.3.2)$$

where the intersection is over all solutions p of (1.1.1) and each I_p is defined by

$$I_p = \{f \in F[x_1, \dots, x_n] \mid gf \in \langle f_1, \dots, f_s \rangle \text{ for some } g \in F[x_1, \dots, x_n] \text{ with } g(p) \neq 0\}. \quad (1.3.3)$$

One can show that I_p is a primary ideal, which in this case means that $\sqrt{I_p}$ is the maximal ideal $\langle x_1 - p_1, \dots, x_n - p_n \rangle$, $p = (p_1, \dots, p_n)$. We will explain how to compute primary decompositions in Section 4.

Given the above primary decomposition of $\langle f_1, \dots, f_s \rangle$, we set

$$A_p = F[x_1, \dots, x_n]/I_p.$$

Then (1.3.2) and the Chinese Remainder Theorem give an algebra isomorphism

$$A = F[x_1, \dots, x_n]/\langle f_1, \dots, f_s \rangle \simeq \prod_p F[x_1, \dots, x_n]/I_p = \prod_p A_p. \quad (1.3.4)$$

We call A_p the *local ring* of the solution p . The idea is that A_p reflects the local structure of the solutions. For example, the *multiplicity* of p as a solution of (1.1.1) is defined to be

$$\text{mult}(p) = \dim_F A_p.$$

We now define a special kind of solution.

Definition 1.3.3 *A solution p of (1.1.1) is curvilinear if $A_p \simeq F[x]/\langle x^k \rangle$ for some integer $k \geq 1$.*

Over the complex numbers, p is curvilinear if and only if we can find local analytic coordinates u_1, \dots, u_n at p and an integer $k \geq 1$ such that the equations are equivalent to

$$u_1 = u_2 = \dots = u_{n-1} = u_n^k = 0.$$

Alternatively, let \mathfrak{m}_p be the maximal ideal of A_p . The integer

$$e_p = \dim_F \mathfrak{m}_p / \mathfrak{m}_p^2 = \# \text{ minimal generators of } \mathfrak{m}_p \quad (1.3.5)$$

is called the *embedding dimension* of A_p . Then p is curvilinear if and only if A_p has embedding dimension $e_p \leq 1$.

Here is the characterization of those systems of equations for which M_f^T is non-derogatory for some $f \in F[x_1, \dots, x_n]$.

Theorem 1.3.4 *There exists $f \in F[x_1, \dots, x_n]$ such that M_f^T (or M_f) is non-derogatory if and only if every solution of (1.1.1) is curvilinear. Furthermore, if this happens, then M_f^T and M_f are non-derogatory when f is a generic linear combination of x_1, \dots, x_n .*

Proof. First observe that since there are only finitely many solutions and F is infinite (being algebraically closed), a generic choice of a_1, \dots, a_n guarantees that $f = a_1x_1 + \dots + a_nx_n$ takes distinct values at the solutions p .

Next observe that m_f is compatible with the algebra isomorphism (1.3.4). If we also assume that f takes distinct values at the solutions, then it follows that M_f is non-derogatory if and only if

$$m_f : A_p \rightarrow A_p$$

is non-derogatory for every p .

To prove the theorem, first suppose that $m_f : A_p \rightarrow A_p$ is non-derogatory and let

$$u = [f - f(p)] \in A_p$$

be the element of A_p determined by $f - f(p)$. Then the kernel of $m_{f-f(p)}$ has dimension 1, which implies the following:

- u lies in the maximal ideal \mathfrak{m}_p since elements in $A_p \setminus \mathfrak{m}_p$ are invertible.
- The image of $m_{f-f(p)}$ has codimension 1.

Since the image is $\langle u \rangle \subset \mathfrak{m}_p$ and \mathfrak{m}_p also has codimension 1 in A_p , it follows that

$$\langle u \rangle = \mathfrak{m}_p$$

This proves that p is curvilinear.

Conversely, if every p is curvilinear, then it is easy to see that m_f is non-derogatory when f is a generic linear combinations of the variables. We leave the details as an exercise to the reader. \square

Since not all systems of equations have curvilinear solutions, it follows that the above method for finding solutions needs to be modified. There are two ways to proceed:

- First, one can compute the *radical*

$$\sqrt{\langle f_1, \dots, f_s \rangle} = \{f \in F[x_1, \dots, x_n] \mid f^k \in \langle f_1, \dots, f_s \rangle \text{ for some } k \geq 1\}.$$

The radical gives a system of equations with the same solutions as (1.1.1), except that all solutions now have multiplicity 1 and hence are curvilinear. Thus Theorem 1.3.4 applies to the radical system. Furthermore, Proposition 2.7 of Chapter 2 of [CLO2] states that

$$\sqrt{\langle f_1, \dots, f_s \rangle} = \langle f_1, \dots, f_s, (p_1(x_1))_{\text{red}}, \dots, (p_n(x_n))_{\text{red}} \rangle,$$

where $p_i(x_i)$ is the minimal polynomial of M_{x_i} written as a polynomial in x_i and $(p_i(x_i))_{\text{red}}$ is the squarefree polynomial with the same roots as $p_i(x_i)$.

- Second, one can intersect eigenspaces of the $M_{x_i}^T$. Let p_1 be an eigenvalue of $M_{x_1}^T$, so that p_1 is the first coordinate of a solution of (1.1.1). Then, since $M_{x_1}^T$ and $M_{x_2}^T$ commute, $M_{x_2}^T$ induces a linear map

$$M_{x_2}^T : E_A(p_1, M_{x_1}^T) \rightarrow E_A(p_1, M_{x_1}^T).$$

where $E_A(p_1, M_{x_1}^T)$ is the eigenspace of $M_{x_1}^T$ for the eigenvalue p_1 . The eigenvalues of this map give the second coordinates of all solutions which have p_1 as their first coordinate. Continuing in this way, one gets all solutions. This method is analyzed carefully in [MT].

1.4 Single-Variable Representation

One nice property of the non-derogatory case is that when m_f is non-derogatory, we can represent the algebra A using one variable. Here is the precise result.

Proposition 1.4.1 *Assume that $f \in F[x_1, \dots, x_n]$ and that m_f is non-derogatory. Then there is an algebra isomorphism*

$$F[u]/\langle \text{CharPoly}_{m_f}(u) \rangle \simeq A.$$

Proof. Consider the map $F[u] \rightarrow A$ defined by $P(u) \mapsto [P(f)]$. Then $P(u)$ is in the kernel if and only if $[P(f)] = [0]$, which implies that for any $g \in F[x_1, \dots, x_n]$, we have

$$P(m_f)([g]) = m_{P(f)}([g]) = [P(f) \cdot g] = [P(f)][g] = [0].$$

It follows that $P(u)$ must be divisible by the minimal polynomial of m_f . Furthermore, applying the above equation with $g = 1$ shows that the minimal polynomial is in the kernel. Thus we get an injective algebra homomorphism

$$F[u]/\langle \text{MinPoly}_{m_f}(u) \rangle \longrightarrow A.$$

But $\text{MinPoly}_{m_f}(u) = \text{CharPoly}_{m_f}(u)$ since m_f is non-derogatory, and we also know that

$$\dim_F F[u]/\langle \text{CharPoly}_{m_f}(u) \rangle = \deg \text{CharPoly}_{m_f}(u) = \dim_F A.$$

It follows that the above injection is the desired isomorphism. \square

Notice that this proof applies over any field F . Thus, when F is infinite and all of the solutions are curvilinear (e.g., all have multiplicity 1), then Proposition 1.4.1 applies when f is a generic linear combination of the variables.

We can use the single variable representation to give an alternate method for finding solutions. The idea is that the isomorphism

$$F[u]/\langle \text{CharPoly}_{m_f}(u) \rangle \simeq A$$

enables us to express the coset $[x_i] \in A$ as a polynomial in $[f]$, say

$$[x_i] = P_i([f]). \tag{1.4.1}$$

Furthermore, it is an easy exercise to show that P_i can be explicitly computed using a Gröbner basis for $\langle f_1, \dots, f_s \rangle$. Now we get all solutions as follows.

Proposition 1.4.2 *Assume that m_f be non-derogatory and let P_1, \dots, P_n be constructed as above. Then for any root λ of $\text{CharPoly}_{m_f}(u)$, the n -tuple*

$$(P_1(\lambda), \dots, P_n(\lambda))$$

is a solution of (1.1.1), and all solutions of (1.1.1) arise this way.

Proof. A solution $p = (p_1, \dots, p_n)$ of (1.1.1) corresponds to an algebra homomorphism $A \rightarrow F$ which takes $[x_i]$ to p_i . Similarly, algebra homomorphisms $F[u]/\langle \text{CharPoly}_{m_f}(u) \rangle \rightarrow F$ are given by evaluation at roots of $\text{CharPoly}_{m_f}(u)$. Thus we are done by Proposition 1.4.1 and (1.4.1). \square

The single-variable representation will have some unexpected consequences in later sections.

1.5 Generalized Eigenspaces and Multiplicities

The final observation of Section 1 relates multiplicities and generalized eigenspaces. Given an eigenvalue λ of a linear map $T : V \rightarrow V$, recall that its *generalized eigenspace* is

$$G_V(\lambda, T) = \{v \in V \mid (T - \lambda I)^N(v) = 0 \text{ for some } N \geq 1\}.$$

It is well-known that the dimension of $G(\lambda, T)$ is the multiplicity of λ as a root of the characteristic polynomial of T .

Proposition 1.5.1 *Let $f \in F[x_1, \dots, x_n]$ take distinct values at the solutions of (1.1.1). Then, if p is one of the solutions, then the generalized eigenspace $G_A(f(p), m_f)$ is naturally isomorphic to the local ring A_p . Furthermore, the characteristic polynomial of $m_f : A \rightarrow A$ is*

$$\text{CharPoly}_{m_f}(u) = \prod_p (u - f(p))^{\text{mult}(p)}.$$

Proof. First observe that m_f is compatible with the isomorphism

$$A \simeq \prod_p A_p = \prod_p F[x_1, \dots, x_n]/I_p$$

where $\langle f_1, \dots, f_s \rangle = \bigcap_p I_p$ is the primary decomposition and the product and intersection are over all solutions of (1.1.1).

Now fix one solution p and consider the behavior of $m_{f-f(p)}$. This is invertible on A_q for $q \neq p$ since $f(q) \neq f(p)$ and nilpotent on A_p since $f(p)$ is the only eigenvalue of m_f on A_p . It follows immediately that we can identify A_p with the generalized eigenspace $G_A(f(p), m_f)$.

The second assertion of the proposition follows from $\text{mult}(p) = \dim_F A_p = \dim_F G_A(f(p), m_f)$ and, as already observed, the latter is the multiplicity of $f(p)$ as a root of the characteristic polynomial of m_f . \square

Since F is infinite, it follows that if $f = a_1 x_1 + \dots + a_n x_n$ for randomly chosen $a_1, \dots, a_n \in F$, then f will take distinct values at the solutions of (1.1.1) with a very high probability. In this situation, the factorization of the minimal polynomial of m_f gives the number of solutions and their multiplicities. Thus, given a system of equations (1.1.1), we have a probabilistic algorithm for finding both the number of solutions and their respective multiplicities.

Numerical Issues. A serious numerical issue is that it is sometimes hard to distinguish between a single solution of multiplicity $k > 1$ and a cluster of k very close solutions of multiplicity 1. Several people, including Hans Stetter, are trying to come up with numerically stable methods for understanding such clusters. For example:

- While the individual points in a cluster are not stable, their center of gravity is.
- When the cluster consists of two points, the slope of the line connecting them is numerically stable.

More details can be found in [HS]. We should also note that from a sophisticated point of view, this can be considered as studying the numerical stability points in the punctal Hilbert scheme of fixed finite length.

Other Notions of Multiplicity. The multiplicity $\text{mult}(p)$ defined in this section is sometimes called the *geometric multiplicity*. There is also a more subtle version of multiplicity called the *algebraic multiplicity* $e(p)$. A discussion of this notion of multiplicity can be found in [Cox].

2 Ideals Defined By Linear Conditions

2.1 Duals and Dualizing Modules

As in the previous section, we will assume that

$$A = F[x_1, \dots, x_n] / \langle f_1, \dots, f_s \rangle$$

is a finite dimensional algebra over an algebraically closed field F . The dual space

$$A^* = \text{Hom}_F(A, F),$$

becomes an A module via

$$(a\ell)(b) = \ell(ab)$$

for $a, b \in A$ and $\ell \in A^*$.

If $\{\ell_1, \dots, \ell_m\}$ is a basis of A^* , then composing with the quotient map $F[x_1, \dots, x_n] \rightarrow A$ gives linear maps

$$L_1, \dots, L_m : F[x_1, \dots, x_n] \longrightarrow F$$

with the property that

$$\langle f_1, \dots, f_s \rangle = \{f \in F[x_1, \dots, x_n] \mid L_i(f) = 0, i = 1, \dots, m\}$$

Thus the ideal $\langle f_1, \dots, f_s \rangle$ is defined by the linear conditions given by the L_i . In this section, we will explore some interesting ways in which this can be done.

But first, we need to explain how A^* relates to commutative algebra. Recall from Section 1 that we have the product decomposition

$$A \simeq \prod_p A_p$$

induced by the primary decomposition $\langle f_1, \dots, f_s \rangle = \bigcap_p I_p$, where the product and intersection are over all solutions in F^n of the equations

$$f_1 = f_2 = \dots = f_s = 0.$$

The product induces a natural isomorphism

$$A^* \simeq \prod_p A_p^*. \tag{2.1.1}$$

Since A_p is a 0-dimensional local ring, the dual space A_p^* is the *dualizing module* of A_p . The theory of dualizing modules is explained in Chapter 21 of [Eisenbud].

One feature of (2.1.1) is the following. For each solution p , let $\{\ell_{p,i}\}_{i=1}^{\text{mult}(p)}$ be a basis of A_p^* . As usual, every $\ell_{p,i}$ gives a linear map

$$L_{p,i} : F[x_1, \dots, x_n] \longrightarrow F.$$

Then:

- If we fix p , then

$$I_p = \{f \in F[x_1, \dots, x_n] \mid L_{p,i}(f) = 0 \text{ for } i = 1, \dots, \text{mult}(p)\}$$

is the primary ideal such that $A_p = F[x_1, \dots, x_n]/I_p$.

- If we vary over all p and i , then the $L_{p,i}$ define the ideal

$$\langle f_1, \dots, f_s \rangle = \bigcap_p I_p.$$

Thus this way of thinking of the linear conditions gives not only the ideal but also its primary decomposition.

Finally, we say that the local ring A_p is *Gorenstein* if there is an A_p -module isomorphism

$$A_p^* \simeq A_p$$

This is equivalent to the existence of a nondegenerate bilinear form

$$\langle \cdot, \cdot \rangle : A_p \times A_p \rightarrow F$$

with the property that

$$\langle ab, c \rangle = \langle a, bc \rangle$$

for $a, b, c \in A_p$. See Chapter 21 of [Eisenbud] for more on duality in this situation.

2.2 Differential Conditions Defining Ideals

So far, we've seen that each primary ideal I_p can be described using $\text{mult}(p)$ linear conditions. We will now explain how to represent these linear conditions using constant coefficient differential operators evaluated at p . We will assume that F has characteristic 0.

Let's begin with some examples.

Example 2.2.1 The equation

$$x^2(x-1)^3 = 0$$

has the solutions 0 of multiplicity 2 and 1 of multiplicity 3. In terms of derivatives, we have

$$\langle x^2(x-1)^3 \rangle = \{f \in F[x] \mid f(0) = f'(0) = 0, f(1) = f'(1) = f''(1) = 0\}.$$

Notice that the multiplicities correspond to the number of conditions defining the ideal at 0 and 1 respectively. \square

Example 2.2.2 Consider the three sets of equations

$$\begin{aligned} \text{(a)} : x^2 = xy = y^2 = 0 \\ \text{(b)} : x^2 = y^2 = 0 \\ \text{(c)} : x = y^3 = 0. \end{aligned}$$

One easily sees that $(0, 0)$ is the only solution with multiplicity 3 in case (a), 4 in case (b), and 3 in case (c). In terms of partial derivatives, the corresponding ideals are given by

$$\begin{aligned} \text{(a)} : \langle x^2, xy, y^2 \rangle &= \{f \in F[x, y] \mid f(0, 0) = f_x(0, 0) = f_y(0, 0) = 0\} \\ \text{(b)} : \langle x^2, y^2 \rangle &= \{f \in F[x, y] \mid f(0, 0) = f_x(0, 0) = f_y(0, 0) = f_{xy}(0, 0) = 0\} \\ \text{(c)} : \langle x, y^3 \rangle &= \{f \in F[x, y] \mid f(0, 0) = f_y(0, 0) = f_{yy}(0, 0) = 0\}. \end{aligned}$$

In each case, the multiplicity is the number of conditions defining the ideal. \square

We will now generalize this description and use it to obtain interesting information about the local rings. For instance, in the above example, we will see that the descriptions of the ideals in terms of partial derivatives imply the following:

- In cases (b) and (c), the ring is Gorenstein but not in case (a).
- In case (c) the ring is curvilinear but not in cases (a) and (b).

We begin by setting up some notation. Consider the polynomial ring $F[\partial_1, \dots, \partial_n]$. Then an exponent vector $\alpha = (a_1, \dots, a_n)$ gives the monomial ∂^α , which we regard as the partial derivative

$$\partial^\alpha = \frac{\partial^{a_1 + \dots + a_n}}{\partial x_1^{a_1} \dots \partial x_n^{a_n}}.$$

Thus elements of $F[\partial_1, \dots, \partial_n]$ become constant coefficient differential operators on $F[x_1, \dots, x_n]$.

In examples, we sometimes write ∂^α as ∂_{x^α} . Thus

$$\partial_{xy^2} = \partial^{(1,2)} = \frac{\partial^3}{\partial x \partial y^2}$$

when operating on $F[x, y]$. Also note that Example 2.2.2 involves the operators

$$\begin{aligned} \text{(a)} : 1, \partial_x, \partial_y \\ \text{(b)} : 1, \partial_x, \partial_y, \partial_{xy} \\ \text{(c)} : 1, \partial_y, \partial_{y^2}. \end{aligned} \tag{2.2.1}$$

applied to polynomials in $F[x, y]$ and evaluated at $(0, 0)$. Here, 1 is the identity operator on $F[x, y]$.

We next define the *deflation* or *shift* of $D = \sum_\alpha c_\alpha \partial^\alpha \in F[\partial_1, \dots, \partial_n]$ by an exponent vector β to be the operator

$$\sigma_\beta D = \sum_\alpha c_\alpha \binom{\alpha}{\beta} \partial^{\alpha-\beta},$$

where $\binom{\alpha}{\beta} = \binom{a_1}{b_1} \dots \binom{a_n}{b_n}$ and $\partial^{\alpha-\beta} = 0$ whenever $\alpha - \beta$ has a negative coordinate. The reason for the binomial coefficients in the formula for $\sigma_\beta D$ is that they give the Leibniz formula

$$D(fg) = \sum_\beta \partial^\beta(f) \sigma_\beta D(g).$$

for $f, g \in F[x_1, \dots, x_n]$. Here are some simple examples of deflations.

Example 2.2.3 Observe that

$$\begin{aligned} \partial_{xy} \text{ has nonzero deflations } \partial_{xy}, \partial_x, \partial_y, 1 \\ \partial_{y^2} \text{ has nonzero deflations } \partial_{y^2}, 2\partial_y, 1. \end{aligned}$$

These correspond to cases (b) and (c) of Example 2.2.2. On the other hand, the operators of case (a) are not deflations of a single operator. As we will see, this is why case (a) is not Gorenstein. \square

Definition 2.2.4 A subspace $L \subset F[\partial_1, \dots, \partial_n]$ is **closed** if it is finite-dimensional over F and closed under deflation, i.e., $\sigma_\beta(L) \subset L$ for all β .

The reader can easily check that the differential operators in cases (a), (b) and (c) of (2.2.1) span closed subspaces. Here is the main result of Section 2.2 (see [MMM1] for a proof).

Theorem 2.2.5 For a solution p of (1.1.1), there is a unique closed subspace $L_p \subset F[\partial_1, \dots, \partial_n]$ of dimension $\text{mult}(p)$ such that

$$\langle f_1, \dots, f_s \rangle = \{f \in F[x_1, \dots, x_n] \mid D(f)(p) = 0 \text{ for all solutions } p \text{ and all } D \in L_p\},$$

where $D(f)(p)$ means the evaluation of the polynomial $D(f)$ at the point p . Furthermore, the primary component of $\langle f_1, \dots, f_s \rangle$ corresponding to a solution p is

$$I_p = \{f \in F[x_1, \dots, x_n] \mid D(f)(p) = 0 \text{ for all } D \in L_p\},$$

and conversely,

$$L_p = \{D \in F[\partial_1, \dots, \partial_n] \mid D(f)(p) = 0 \text{ for all } f \in I_p\}.$$

It should not be surprising that Examples 2.2.1 and 2.2.2 are examples of this theorem. Here is a more substantial example.

Example 2.2.6 Consider the equations

$$\begin{aligned} f_1 &= x^2 + 2y^2 - 2y = 0 \\ f_2 &= xy^2 - xy = 0 \\ f_3 &= y^3 - 2y^2 + y = 0 \end{aligned}$$

from Example 1.1.1. There, we saw that the only solutions were $(0, 0)$ and $(0, 1)$. In [MS], it is shown that

$$\begin{aligned} L_{(0,0)} &= \text{Span}(1, \partial_x) \\ L_{(0,1)} &= \text{Span}(1, \partial_x, \partial_{x^2} - \partial_y). \end{aligned} \tag{2.2.2}$$

Thus

$$\langle f_1, f_2, f_3 \rangle = \{f \in F[x, y] \mid f(0, 0) = f_x(0, 0) = f(0, 1) = f_x(0, 1) = 0, f_{xx}(0, 1) = f_y(0, 1)\},$$

and looking at the conditions for $(0, 0)$ and $(0, 1)$ separately gives the primary decomposition of $\langle f_1, f_2, f_3 \rangle$. In Section 2.3 we will describe how these operators were found. \square

Gorenstein and Curvilinear Points. We conclude Section 2.2 by explaining how special properties of the local ring A_p can be determined from the representation given in Theorem 2.2.5.

Theorem 2.2.7 A_p is Gorenstein if and only if there is $D \in L_p$ whose deflations span L_p .

Proof. Let L_p^{ev} denote the linear forms $F[x_1, \dots, x_n] \rightarrow F$ obtained by composing elements of L_p with evaluation at p . Each such map vanishes on I_p and thus gives an element of A_p^* . Hence

$$L_p \simeq L_p^{\text{ev}} \simeq A_p^*.$$

Furthermore, if $D \in L_p$ maps to $\tilde{D} \in A_p^*$, then the Leibniz formula makes it easy to see that the deflation $\sigma_\beta D$ maps to $(x-p)^\beta \tilde{D}$, where

$$(x-p)^\beta = (x_1 - p_1)^{b_1} \cdots (x_n - p_n)^{b_n}$$

for $p = (p_1, \dots, p_n)$ and $\beta = (b_1, \dots, b_n)$. It follows that these deflations span L_p if and only if A_p^* is generated by a single element as an A_p -module. In the latter case, we have a surjective A_p -module homomorphism $A_p \rightarrow A_p^*$ which must be an isomorphism since A_p and A_p^* have the same dimension over F . Then we are done by the definition of Gorenstein given in Section 2.1. \square

Before stating our next result, we need the following definition from [MS].

Definition 2.2.8 *The order of $D = \sum_\alpha c_\alpha \partial^\alpha$ is the degree of D as a polynomial in $F[\partial_1, \dots, \partial_n]$. A basis $D_1, \dots, D_{\text{mult}(p)}$ of L_p is **consistently ordered** if for every $r \geq 1$, there is $j \geq 1$ such that*

$$\text{Span}(D \in L_p \mid D \text{ has order } \leq r) = \text{Span}(D_1, \dots, D_j).$$

Note that every consistently ordered basis has $D_1 = 1$. Also observe that the bases listed in (2.2.1) and (2.2.2) are consistently ordered.

We can now characterize when A_p is curvilinear.

Theorem 2.2.9 *The embedding dimension e_p of A_p is the number of operators of order 1 in a consistently ordered basis of L_p . In particular, A_p is curvilinear if and only if any such basis has a unique operator of order 1.*

Proof. Let \mathfrak{m}_p be the maximal ideal of A_p . Recall from equation (1.3.5) that

$$\begin{aligned} e_p &= \dim_F \mathfrak{m}_p / \mathfrak{m}_p^2 \\ &= \# \text{ minimal generators of } \mathfrak{m}_p. \end{aligned}$$

Also let $L_p^r = \text{Span}(D \in L_p \mid D \text{ has order } \leq r)$. Then $L_p^0 \subset L_p^1 \subset \cdots$ and, for $r \geq 0$, we have

$$\dim_F L_p^r / L_p^{r-1} = \# \text{ operators of order } r \text{ in a consistently ordered basis of } L_p. \quad (2.2.3)$$

We claim that there is a natural isomorphism

$$L_p^1 / L_p^0 \simeq \text{Hom}_F(\mathfrak{m}_p / \mathfrak{m}_p^2, F). \quad (2.2.4)$$

Assuming this for the moment, the first assertion of the theorem follows immediately from (2.2.3) for $r = 1$ and the above formula for e_p . Then the final assertion follows since by definition A_p is curvilinear if and only if it has embedding dimension $e_p = 1$.

To prove (2.2.4), let $\mathfrak{M}_p = \langle x_1 - p_1, \dots, x_n - p_n \rangle \subset F[x_1, \dots, x_n]$ be the maximal ideal of p . Then any operator $D = \sum_{i=1}^n a_i \partial_i$ induces the linear map

$$\mathfrak{M}_p \longrightarrow F$$

which sends $f \in \mathfrak{M}_p$ to $D(f)(p)$. By the product rule, this vanishes if $f \in \mathfrak{M}_p^2$, so that we get an element of the dual space

$$\mathrm{Hom}_F(\mathfrak{M}_p/\mathfrak{M}_p^2, F).$$

Furthermore, it is easy to see that every element of the dual space arises in this way.

The isomorphism $F[x_1, \dots, x_n]/I_p \simeq A_p$ induces exact sequences

$$0 \longrightarrow I_p \longrightarrow \mathfrak{M}_p \longrightarrow \mathfrak{m}_p \longrightarrow 0$$

and

$$0 \longrightarrow \mathrm{Hom}_F(\mathfrak{m}_p/\mathfrak{m}_p^2, F) \longrightarrow \mathrm{Hom}_F(\mathfrak{M}_p/\mathfrak{M}_p^2, F) \longrightarrow \mathrm{Hom}_F(I_p/I_p \cap \mathfrak{M}_p^2, F) \longrightarrow 0.$$

It follows that $D = \sum_{i=1}^n a_i \partial_i$ gives an element of $\mathrm{Hom}_F(\mathfrak{m}_p/\mathfrak{m}_p^2, F)$ if and only if D vanishes on I_p , which is equivalent to $D \in L_p$. Since these operators represent L_p^1/L_p^0 , the theorem follows. \square

We get the following corollary when we combine this result with Theorem 1.3.4.

Corollary 2.2.10 *The matrix M_f^T is non-derogatory when f is a generic linear combination of x_1, \dots, x_n if and only if for every solution p , a consistently ordered basis of L_p has a unique operator of order 1.*

Since the bases in Example 2.2.6 are consistently ordered, (2.2.2) shows that M_f^T is non-derogatory when f is a generic linear combination of x, y . Of course, we computed a specific instance of this in Example 1.3.1, but now we know the systematic reason for our success.

Note also that if we apply Theorems 2.2.7 and 2.2.9 to Example 2.2.2, then we see that the ring is Gorenstein in cases (b) and (c) (but not (a)) and curvilinear in case (c) (but not (a) and (b)). This proves the claims made in the two bullets on page 13.

2.3 Two Algorithms

We've seen that the ideal $\langle f_1, \dots, f_s \rangle$ can be described using Gröbner bases and using conditions on partial derivatives. As we will now explain, going from one description to the other is a simple matter of linear algebra.

Gröbner Bases to Partial Derivatives. If we have a Gröbner basis for $\langle f_1, \dots, f_s \rangle$, then we obtain the required closed subspaces L_p in a three-step process. The first step is to compute the primary decomposition

$$\langle f_1, \dots, f_s \rangle = \bigcap_p I_p.$$

In particular, this means knowing a Gröbner basis for each I_p . We will explain how to compute such a primary decomposition in Section 4.

Given this, we fix a primary ideal I_p . We next recall a useful fact which relates I_p to the maximal ideal $\mathfrak{M}_p = \langle x_1 - p_1, \dots, x_n - p_n \rangle$ of p in $F[x_1, \dots, x_n]$.

Lemma 2.3.1 *If $m = \mathrm{mult}(p)$, then $\mathfrak{M}_p^m \subset I_p$.*

Proof. It suffices to prove that $\mathfrak{m}_p^m = \{0\}$, where \mathfrak{m}_p is the maximal ideal of A_p . By Nakayama's Lemma, we know that that $\mathfrak{m}_p^k \neq \mathfrak{m}_p^{k+1}$ whenever $\mathfrak{m}_p^k \neq \{0\}$. Using

$$A_p \supset \mathfrak{m}_p \supset \mathfrak{m}_p^2 \supset \cdots \supset \mathfrak{m}_p^k \supset \{0\},$$

it follows that $\dim_F A_p \geq k + 1$ whenever $\mathfrak{m}_p^k \neq \{0\}$. The lemma now follows immediately. \square

This lemma will enable us to describe I_p in terms of differential operators of order at most m . However, this description works best when $p = 0$. So the second step is to translate so that $p = 0$. Hence for the rest of our discussion, we will assume that $p = 0$. Thus Lemma 2.3.1 tells us that

$$\mathfrak{M}_0^m \subset I_0, \quad m = \text{mult}(0).$$

The third step is to write down the differential operators in I_0 as follows. Let \mathcal{B}_0 be the set of remainder monomials for the Gröbner basis of I_0 and set

$$\text{Mon}_m = \{x^\alpha \mid x^\alpha \notin \mathcal{B}_0, \deg(x^\alpha) < m\}.$$

For each $x^\alpha \in \text{Mon}_m$, let

$$x^\alpha \equiv \sum_{x^\beta \in \mathcal{B}_0} c_{\alpha\beta} x^\beta \pmod{I_0} \quad (2.3.1)$$

In other words, $\sum_{x^\beta \in \mathcal{B}_0} c_{\alpha\beta} x^\beta$ is the remainder of x^α on division by the Gröbner basis of I_0 . Then, for each $x^\beta \in \mathcal{B}_0$, set

$$D_\beta = \partial^\beta + \sum_{x^\alpha \in \text{Mon}_m} c_{\alpha\beta} \frac{\beta!}{\alpha!} \partial^\alpha.$$

where $\alpha! = a_1! \cdots a_n!$ for $\alpha = (a_1, \dots, a_n)$ and similarly for $\beta!$.

Proposition 2.3.2 *$f \in F[x_1, \dots, x_n]$ lies in I_p if and only if $D_\beta(f)(0) = 0$ for all $x^\beta \in \mathcal{B}_0$.*

Proof. Let $f = \sum_{\alpha} a_\alpha x^\alpha$. Since $\mathfrak{M}_0^m \subset I_0$, we can assume that $f = \sum_{\deg(\alpha) < m} a_\alpha x^\alpha$. Using (2.3.1), it is straightforward to show that

$$f \in I_p \iff a_\beta + \sum_{x^\alpha \in \text{Mon}_m} c_{\alpha\beta} a_\alpha = 0 \text{ for all } x^\beta \in \mathcal{B}_0.$$

However, since

$$\partial^\gamma(x^\delta)(0) = \begin{cases} \gamma! & \text{if } \gamma = \delta \\ 0 & \text{otherwise,} \end{cases} \quad (2.3.2)$$

one easily sees that for $x^\beta \in \mathcal{B}_0$,

$$D_\beta(f)(0) = (\partial^\beta + \sum_{x^\alpha \in \text{Mon}_m} c_{\alpha\beta} \frac{\beta!}{\alpha!} \partial^\alpha) \left(\sum_{\deg(\gamma) < m} a_\gamma x^\gamma \right) (0) = \beta! (a_\beta + \sum_{x^\alpha \in \text{Mon}_m} c_{\alpha\beta} a_\alpha).$$

The proposition now follows immediately since F has characteristic 0. \square

Here is an example of this result.

Example 2.3.3 For the polynomials of Example 2.2.6, we will show in Section 4 that the primary decomposition is

$$\langle x^2 + 2y^2 - 2y, xy^2 - xy, y^3 - 2y^2 + y \rangle = \langle x^2, y \rangle \cap \langle x^2 + 2(y-1), x(y-1), (y-1)^2 \rangle = I_{(0,0)} \cap I_{(0,1)}.$$

Let's focus on $I_{(0,1)}$. If we translate this to the origin, we get the ideal

$$I_0 = \langle x^2 + 2y, xy, y^2 \rangle.$$

The generators are a Gröbner basis for lex order with $x > y$, the remainder monomials are $\mathcal{B}_0 = \{1, x, y\}$, and the multiplicity is $m = 3$. Thus

$$\text{Mon}_3 = \{x^2, xy, y^2\}.$$

The coefficients $c_{\alpha\beta}$ are given by

$$\begin{aligned} x^2 &\equiv 0 \cdot 1 + 0 \cdot x + (-2) \cdot y \pmod{I_0} \\ xy &\equiv 0 \cdot 1 + 0 \cdot x + 0 \cdot y \pmod{I_0} \\ y^2 &\equiv 0 \cdot 1 + 0 \cdot x + 0 \cdot y \pmod{I_0}. \end{aligned}$$

so that

$$D_1 = 1, \quad D_x = \partial_x, \quad D_y = \partial_y + (-2)\frac{1!}{2!}\partial_{x^2} = \partial_y - \partial_{x^2}.$$

Up to a sign, this gives the basis of $L_{(0,1)}$ that appeared in Example 2.2.6. The treatment for $L_{(0,0)}$ is even easier and is omitted. \square

We also want to remark on an alternate way to view the construction $L_0 = \text{Span}(D_\beta \mid x^\beta \in \mathcal{B}_0)$. Here, we are in the situation where $p = 0$, so that by Theorem 2.2.5, we have

$$I_0 = \{f \in F[x_1, \dots, x_n] \mid D_\beta(f)(0) = 0 \text{ for all } x^\beta \in \mathcal{B}_0\}$$

and

$$L_0 = \{D \in F[\partial_1, \dots, \partial_n] \mid D_\beta(f)(0) = 0 \text{ for all } f \in I_0\}.$$

Now we will do something audacious: switch x_i with ∂_i . This means I_0 becomes an ideal

$$\hat{I}_0 \subset F[\partial_1, \dots, \partial_n]$$

and L_0 becomes a subspace

$$\hat{L}_0 \subset F[x_1, \dots, x_n].$$

The key observation is that the pairing (2.3.2) is unchanged under $x_i \leftrightarrow \partial_i$. Thus:

\hat{L}_0 gives the polynomial solutions of the infinitely many differential operators in \hat{I}_0 !

This is the point of view taken in Chapter 10 of [Sturmfels]. Here is an example.

Example 2.3.4 In Example 2.3.3, we showed that

$$I_0 = \langle x^2 + 2y, xy, y^2 \rangle \implies L_0 = \text{Span}(1, \partial_x, \partial_y - \partial_{x^2}).$$

This means that under the switch $x \leftrightarrow \partial_x, y \leftrightarrow \partial_y$, the subspace

$$\hat{L}_0 = \text{Span}(1, x, y - x^2) \subset F[x, y]$$

is the space of all polynomial solutions of the infinitely many operators in the ideal

$$\hat{I}_0 = \langle \partial_x^2 + 2\partial_y, \partial_x \partial_y, \partial_y^2 \rangle \subset F[\partial_x, \partial_y].$$

Other examples can be found in [Sturmfels]. □

We should also note that the description of I_p given by differential conditions can require a lot of space. Examples plus more efficient methods can be found in Section 3.3 of [MMM2].

Partial Derivatives to Gröbner Bases. Now suppose that conversely, we are given the data of Theorem 2.2.5. This means that for each solution p we have a closed subspace L_p of dimension $\text{mult}(p)$ such that

$$\langle f_1, \dots, f_s \rangle = \{f \in F[x_1, \dots, x_n] \mid D(f)(p) = 0 \text{ for all } p \text{ and } D \in L_p\}.$$

If we pick a basis $D_{p,i}$ of each L_p , then the linear forms $f \mapsto D_{p,i}(f)(p)$ give a linear map

$$L : F[x_1, \dots, x_n] \longrightarrow F^m \tag{2.3.3}$$

where $m = \dim_F A$. This map is surjective and its kernel is $\langle f_1, \dots, f_s \rangle$. Given an order $>$ (with some restrictions to be noted below), our goal is to find a Gröbner basis of $\langle f_1, \dots, f_s \rangle$ with respect to $>$ using the linear map (2.3.3).

The idea is to simultaneously build up the Gröbner basis G and the set of remainder monomials \mathcal{B} . So we begin with both lists being empty. We then feed in monomials, beginning with 1. The main loop of the algorithm is described as follows.

Main Loop: Given a monomial x^α , compute $L(x^\alpha)$ together with $L(x^\beta)$ for all $x^\beta \in \mathcal{B}$.

- If $L(x^\alpha)$ is *linearly dependent* on the $L(x^\beta)$, then compute a linear relation

$$L(x^\alpha) = \sum_{x^\beta \in \mathcal{B}} a_\beta L(x^\beta), \quad a_\beta \in F$$

(hence $x^\alpha - \sum_{x^\beta \in \mathcal{B}} a_\beta x^\beta \in \langle f_1, \dots, f_s \rangle$) and add $x^\alpha - \sum_{x^\beta \in \mathcal{B}} a_\beta x^\beta$ to G .

- If $L(x^\alpha)$ is *linearly independent* from the $L(x^\beta)$, then add x^α to \mathcal{B} .

Once this loop is done for x^α , we feed in the next monomial, which is the minimal element (with respect to $>$) of the set

$$N(x^\alpha, G) = \{x^\gamma \mid x^\gamma > x^\alpha, x^\gamma \text{ not divisible by any leading term of } g \in G\}. \tag{2.3.4}$$

Hence we need to find the minimal element of $N(x^\alpha, G)$. As explained in [BW], this is easy to do whenever $>$ is a lex or total degree order. The algorithm terminates when (2.3.4) becomes empty.

In [MMM1], it is shown that this algorithm always terminates and that when this happens, G is the desired Gröbner basis and \mathcal{B} is the corresponding set of remainder monomials. Here is an example.

Example 2.3.5 As in Example 2.3.3, consider $p = (0, 0)$ and $L_0 = \text{Span}(1, \partial_x, \partial_y - \partial_{x^2})$. It follows that I_0 is the kernel of the map

$$L : F[x, y] \longrightarrow F^3$$

defined by $L(f) = (f(0, 0), f_x(0, 0), f_y(0, 0) - f_{xx}(0, 0))$. If we use lex order with $x > y$, then the above algorithm starts with $B = G = \emptyset$ and proceeds as follows:

x^α	$L(x^\alpha)$	\mathcal{B}	G	$\min(N(x^\alpha, G))$
1^*	$(1, 0, 0)$	$\{1\}$	\emptyset	y
y^*	$(0, 0, 1)$	$\{1, y\}$	\emptyset	y^2
y^2	$(0, 0, 0)$	$\{1, y\}$	$\{y^2\}$	x
x^*	$(0, 1, 0)$	$\{1, y, x\}$	$\{y^2\}$	xy
xy	$(0, 0, 0)$	$\{1, y, x\}$	$\{y^2, xy\}$	x^2
x^2	$(0, 0, -2)$	$\{1, y, x\}$	$\{y^2, xy, x^2 + 2y\}$	none!

(2.3.5)

In this table, an asterisk denotes those monomials which become remainder monomials. The other monomials are leading terms of the Gröbner basis. It is worth checking the steps of the algorithm to make sure you see how it is working. \square

A proof of correctness, together with a complexity analysis, can be found in [MMM1].

2.4 Ideals of Points and Basis Conversion

We conclude Section 2 by observing that the algorithm illustrated in Example 2.3.5 applies to many situations besides partial derivatives. The key point is that if

$$L : F[x_1, \dots, x_n] \longrightarrow F^m$$

is *any* surjective linear map whose kernel is an ideal I , then the algorithm described in the discussion following (2.3.3) gives a Gröbner basis for I . Here are two situations where this is useful.

Ideals of Points. Suppose we have a finite list of points $p_1, \dots, p_m \in F^n$. Then we want to compute a Gröbner basis of the ideal

$$I = \{f \in F[x_1, \dots, x_n] \mid f(p_1) = \dots = f(p_m) = 0\}$$

consisting of all polynomials which vanish at p_1, \dots, p_m . This is now easy, for the points give the linear map $L : F[x_1, \dots, x_n] \rightarrow F^m$ defined by

$$L(f) = (f(p_1), \dots, f(p_m))$$

whose kernel is the ideal I . Furthermore, it is easy to see that L is surjective (see the proof of Theorem 2.10 of Chapter 2 of [CLO2]). Thus we can find a Gröbner basis of I using the above algorithm.

Example 2.4.1 Consider the points $(0, 0), (1, 0), (0, 1) \in F^2$. This gives $L : F[x, y] \rightarrow F^3$ defined by

$$L(f) = (f(0, 0), f(1, 0), f(0, 1)).$$

If you apply the algorithm for lex order with $x > y$ as in Example 2.3.5, you will obtain a table remarkably similar to (2.3.5), except that the Gröbner basis will be $\{y^2 - y, xy, x^2 - x\}$. We recommend this exercise to the reader. \square

A more complete treatment of this problem appears in [MMM1]. The harder problem of computing the homogeneous ideal of a finite set of points in projective space is discussed in [ABKR].

Basis Conversion. Suppose that we have a Gröbner basis G' for $\langle f_1, \dots, f_s \rangle$ with respect to one order $>'$ and want to find a Gröbner basis G with respect to a second order $>$.

We can do this as follows. Let \mathcal{B}' be the set of remainder monomials with respect to G' . Then taking the remainder on division by G' gives a linear map

$$L : F[x_1, \dots, x_n] \longrightarrow \text{Span}(\mathcal{B}') \simeq F^m.$$

The kernel is $\langle f_1, \dots, f_s \rangle$ and the map is surjective since $L(x^\beta) = x^\beta$ for $x^\beta \in \mathcal{B}'$. Then we can apply the above method to find the desired Gröbner basis G . This is the basis conversion algorithm of [FGLM].

Example 2.4.2 By Example 2.3.5, we know that $\{y^2, xy, x^2 + xy\}$ is a Gröbner basis with respect to lex order with $x > y$. If you apply the above method to convert this to lex order with $y > x$, you will obtain a table similar to (2.3.5) which gives the new Gröbner basis $\{x^3, y + \frac{1}{2}x^2\}$. \square

Besides ideals of points and basis conversion, this algorithm has other interesting applications. See [MMM1] for details.

3 Resultants

3.1 Solving Equations

The method for solving equations discussed in Section 1 assumed that we had a Gröbner basis available. In this section, we will see that when our equations have more structure, we can sometimes compute the required matrices directly.

We will work in $F[x_1, \dots, x_n]$, F algebraically closed, but we will now assume that we have n equations in n unknowns, i.e.,

$$f_1(x_1, \dots, x_n) = \dots = f_n(x_1, \dots, x_n) = 0. \quad (3.1.1)$$

Bézout's Theorem tells us that if f_i has degree d_i and f_1, \dots, f_n are generic, then (3.1.1) has precisely $\mu = d_1 \cdots d_n$ solutions of multiplicity 1.

We first describe a solution method due to Auzinger and Stetter [AS]. The idea is to construct a $\mu \times \mu$ matrix whose eigenvectors will determine the solutions. For this purpose, let

$$d = d_1 + \dots + d_n - n + 1 \quad (3.1.2)$$

and divide the monomials of degree $\leq d$ into $n + 1$ disjoint sets as follows:

$$\begin{aligned} S_n &= \{x^\gamma : \deg(x^\gamma) \leq d, x_n^{d_n} \text{ divides } x^\gamma\} \\ S_{n-1} &= \{x^\gamma : \deg(x^\gamma) \leq d, x_n^{d_n} \text{ doesn't divide } x^\gamma \text{ but } x_{n-1}^{d_{n-1}} \text{ does}\} \\ &\vdots \\ S_0 &= \{x^\gamma : \deg(x^\gamma) \leq d, x_n^{d_n}, \dots, x_1^{d_1} \text{ don't divide } x^\gamma\}. \end{aligned}$$

In particular, note that

$$S_0 = \{x_1^{a_1} \cdots x_n^{a_n} \mid 0 \leq a_i \leq d_i - 1 \text{ for } i = 1, \dots, n\}. \quad (3.1.3)$$

Since S_0 plays a special role in what follows, we will use x^α to denote elements of S_0 and x^β to denote elements of $S_1 \cup \dots \cup S_n$. Then observe that

$$\begin{aligned} \text{if } x^\alpha \in S_0, \text{ then } x^\alpha \text{ has degree } \leq d-1, \\ \text{if } x^\beta \in S_i, i > 0, \text{ then } x^\beta/x_i^{d_i} \text{ has degree } \leq d-d_i, \end{aligned}$$

where the first assertion uses $d-1 = d_1 + \dots + d_n - n = \sum_{i=1}^n (d_i - 1)$.

Now let $f_0 = a_1 x_1 + \dots + a_n x_n$, $a_i \in F$, and consider the equations:

$$\begin{aligned} x^\alpha f_0 &= 0 \quad \text{for all } x^\alpha \in S_0 \\ (x^\beta/x_1^{d_1}) f_1 &= 0 \quad \text{for all } x^\beta \in S_1 \\ &\vdots \\ (x^\beta/x_n^{d_n}) f_n &= 0 \quad \text{for all } x^\beta \in S_n. \end{aligned}$$

Since the $x^\alpha f_0$ and $x^\beta/x_i^{d_i} f_i$ have degree $\leq d$, we can write these polynomials as linear combinations of the x^α and x^β . We will order these monomials so that the elements $x^\alpha \in S_0$ come first, followed by the elements $x^\beta \in S_1 \cup \dots \cup S_n$. This gives a square matrix M_0 such that

$$M_0 \begin{pmatrix} x^{\alpha_1} \\ x^{\alpha_2} \\ \vdots \\ x^{\beta_1} \\ x^{\beta_2} \\ \vdots \end{pmatrix} = \begin{pmatrix} x^{\alpha_1} f_0 \\ x^{\alpha_2} f_0 \\ \vdots \\ x^{\beta_1}/x_1^{d_1} f_1 \\ x^{\beta_2}/x_1^{d_1} f_1 \\ \vdots \end{pmatrix}, \quad (3.1.4)$$

where, in the column on the left, the first two elements of S_0 and the first two elements of S_1 are listed explicitly. The situation is similar for the column on the right.

We next partition M_0 so that the rows and columns of M_0 corresponding to elements of S_0 lie in the upper left hand corner. This gives

$$M_0 = \begin{pmatrix} M_{00} & M_{01} \\ M_{10} & M_{11} \end{pmatrix}, \quad (3.1.5)$$

where M_{00} is a $\mu \times \mu$ matrix for $\mu = d_1 \cdots d_n$, and M_{11} is also a square matrix. One can show that M_{11} is invertible for a generic choice of f_1, \dots, f_n . Hence we can define the $\mu \times \mu$ matrix

$$\widetilde{M}_{f_0} = M_{00} - M_{01} M_{11}^{-1} M_{10}. \quad (3.1.6)$$

Also, given a point $p \in F^n$, let \mathbf{p}^α be the column vector $(p^{\alpha_1}, p^{\alpha_2}, \dots)^T$ obtained by evaluating all monomials in S_0 at p (where T means transpose).

Theorem 3.1.1 *Let f_1, \dots, f_n be generic polynomials of total degree d_1, \dots, d_n and construct \widetilde{M}_{f_0} as in (3.1.6) with $f_0 = a_1 x_1 + \dots + a_n x_n$. Then \mathbf{p}^α is an eigenvector of \widetilde{M}_{f_0} with eigenvalue $f_0(p)$ whenever p is a solution of (3.1.1). Furthermore, the \mathbf{p}^α are linearly independent as p ranges over all solutions of (3.1.1).*

Proof. Let \mathbf{p}^β be the column vector $(p^{\beta_1}, p^{\beta_2}, \dots)^T$ given by evaluating all monomials in $S_1 \cup \dots \cup S_n$ at p . Then evaluating (3.1.4) at a solution p of (3.1.1) gives

$$M_0 \begin{pmatrix} \mathbf{p}^\alpha \\ \mathbf{p}^\beta \end{pmatrix} = \begin{pmatrix} f_0(p) \mathbf{p}^\alpha \\ \mathbf{0} \end{pmatrix},$$

which in terms of (3.1.5) becomes

$$\begin{pmatrix} M_{00} & M_{01} \\ M_{10} & M_{11} \end{pmatrix} \begin{pmatrix} \mathbf{p}^\alpha \\ \mathbf{p}^\beta \end{pmatrix} = \begin{pmatrix} f_0(p) \mathbf{p}^\alpha \\ \mathbf{0} \end{pmatrix}. \quad (3.1.7)$$

It is straightforward to show that (3.1.7) implies the equation

$$\widetilde{M}_{f_0} \mathbf{p}^\alpha = f_0(p) \mathbf{p}^\alpha, \quad (3.1.8)$$

so that for each solution p , $f_0(p)$ is an eigenvalue of \widetilde{M}_{f_0} with \mathbf{p}^α as eigenvector. For a generic choice of a_1, \dots, a_n , $f_0 = a_1 x_1 + \dots + a_n x_n$ takes distinct values at the solutions, i.e., the eigenvalues $f_0(p)$ are distinct. This shows that the corresponding eigenvectors \mathbf{p}^α are linearly independent. \square

We can now solve (3.1.1) by the method of Section 1.3. We know that there are $\mu = d_1 \cdots d_n$ solutions p . Furthermore, the values $f_0(p)$ are distinct for a generic choice of $f_0 = a_1 x_1 + \dots + a_n x_n$. Then Theorem 3.1.1 implies that the $\mu \times \mu$ matrix \widetilde{M}_{f_0} has μ eigenvectors \mathbf{p}^α . Hence all of the eigenspaces must have dimension 1, i.e., \widetilde{M}_{f_0} is non-derogatory.

Also notice that $1 \in S_0$ by (3.1.3). It follows that we can assume that every \mathbf{p}^α is of the form

$$\mathbf{p}^\alpha = (1, p^{\alpha(2)}, \dots, p^{\alpha(\mu)})^T.$$

Thus, once we compute an eigenvector \vec{v} of \widetilde{M}_{f_0} for the eigenvalue $f_0(p)$, we know how to rescale \vec{v} so that $\vec{v} = \mathbf{p}^\alpha$.

As in Section 1.3, the idea is to read off the solution p from the entries of the eigenvector \mathbf{p}^α . If f_i has degree $d_i > 1$, then $x_i \in S_0$, so that p_i appears as a coordinate of \mathbf{p}^α . Hence we can recover all coordinates of p_i except for those corresponding to equations with $d_i = 1$. These were called the “missing variables” in Section 1.3. In this situation, the missing variables correspond to linear equations. Since we can find the coordinates of the solution p for all of the other variables, we simply substitute these known values into the linear equations corresponding to the missing variables. Hence we find all coordinates of the solution by linear algebra. Details of this procedure are described in Exercise 5 of Section 3.6 of [CLO2].

All of this is very nice but seems to ignore the quotient algebra

$$A = F[x_1, \dots, x_n] / \langle f_1, \dots, f_n \rangle.$$

In fact, what we did above has a deep relation to A as follows.

Theorem 3.1.2 *If f_1, \dots, f_n are generic polynomials of total degree d_1, \dots, d_n , then the cosets of the monomials*

$$S_0 = \{x_1^{a_1} \cdots x_n^{a_n} \mid 0 \leq a_i \leq d_i - 1 \text{ for } i = 1, \dots, n\}$$

form a basis of the quotient algebra A . Furthermore, if $f_0 = a_1 x_1 + \dots + a_n x_n$ and \widetilde{M}_{f_0} is the matrix constructed in (3.1.6) using f_0, f_1, \dots, f_n , then

$$\widetilde{M}_{f_0} = M_{f_0}^T,$$

where M_{f_0} is the matrix of the multiplication map $m_{f_0} : A \rightarrow A$ relative to the basis given by S_0 .

Proof. Recall from Bézout's Theorem that when f_1, \dots, f_n are generic, the equations (3.1.1) have $\mu = d_1 \cdots d_n$ solutions of multiplicity 1 in F^n . It follows that A has dimension μ over F . Since this is also the cardinality of S_0 , the first part of the theorem will follow once we show that the cosets of the monomials in S_0 are linearly independent.

Write the elements of S_0 as $x^{\alpha(1)}, \dots, x^{\alpha(\mu)}$ and suppose we have a linear relation among the cosets $[x^{\alpha(j)}]$, say

$$c_1[x^{\alpha(1)}] + \cdots + c_\mu[x^{\alpha(\mu)}] = 0.$$

Evaluating this equation at a solution p makes sense and implies that

$$c_1 p^{\alpha(1)} + \cdots + c_\mu p^{\alpha(\mu)} = 0. \quad (3.1.9)$$

In the generic case, our equations have $\mu = d_1 \cdots d_n$ solutions, so that (3.1.9) gives μ equations in μ unknowns c_1, \dots, c_μ . But the coefficients of the rows give the transposes of the vectors \mathbf{p}^α , which are linearly independent by Theorem 1.1. It follows that $c_1 = \cdots = c_\mu = 0$. This proves that the cosets $[x^{\alpha(1)}], \dots, [x^{\alpha(\mu)}]$ are linearly independent. Thus S_0 gives a basis of A as claimed.

For the second assertion of the theorem, observe that equation (2.1) shows that

$$M_{f_0}^T \mathbf{p}^\alpha = f_0(p) \mathbf{p}^\alpha$$

for each solution p . Comparing this to (3.1.8), we get

$$M_{f_0}^T \mathbf{p}^\alpha = \widetilde{M} \mathbf{p}^\alpha$$

for all solutions p . Since f_1, \dots, f_n are generic, we have μ solutions p , and the corresponding eigenvectors \mathbf{p}^α are linearly independent by Theorem 3.1.1. This implies $M_{f_0}^T = \widetilde{M}_{f_0}$. \square

It is satisfying to see how the method described in this section relates to what we did in Section 1. However, there is a *lot* more going on here. Here are a couple of items of interest.

Multiplication Matrices. By setting $f_0 = x_i$ in Theorem 3.1.2, we can construct the matrix of multiplication by x_i as $M_{x_i} = \widetilde{M}_{x_i}^T$. However, it is possible to compute all of these maps simultaneously by using $f_0 = u_1 x_1 + \cdots + u_n x_n$, where u_1, \dots, u_n are variables. In the decomposition (3.1.5), the matrices M_{10} and M_{11} don't involve the coefficients of f_0 . Thus, we can still form the matrix \widetilde{M}_{f_0} from (3.1.6), and it is easy to see that

$$\widetilde{M}_{f_0}^T = u_1 M_{x_1} + \cdots + u_n M_{x_n}.$$

Thus one computation gives all of the M_{x_i} .

Solving via Multivariate Factorization. As above, suppose that $f_0 = u_1 x_1 + \cdots + u_n x_n$, where u_1, \dots, u_n are variables. In this case, $\det(\widetilde{M}_{f_0})$ becomes a polynomial in $F[u_1, \dots, u_n]$. The results of this section imply that for f_1, \dots, f_n generic, the eigenvalues of \widetilde{M}_{f_0} are $f_0(p)$ as p ranges over all solutions of (3.1.1). Since all of the eigenspaces have dimension 1, we obtain

$$\det(\widetilde{M}_{f_0}) = \prod_p (u_1 p_1 + \cdots + u_n p_n). \quad (3.1.10)$$

It follows that if we can factor $\det(\widetilde{M}_{f_0})$ into irreducibles in $F[u_1, \dots, u_n]$, then we get all solutions of (3.1.1). We will see in Section 3.2 that (3.1.10) is closely related to resultants.

Ideal Membership. Given $f \in F[x_1, \dots, x_n]$, how do we tell if $f \in \langle f_1, \dots, f_n \rangle$? This is the *Ideal Membership Problem*. How do we do this without a Gröbner basis? One method (probably not very efficient) uses the above matrices M_{x_i} as follows:

$$f \in \langle f_1, \dots, f_n \rangle \iff f(M_{x_1}, \dots, M_{x_n}) \text{ is the zero matrix.}$$

To prove this criterion, note that $f(M_{x_1}, \dots, M_{x_n}) = M_f$ since the M_{x_i} commute. Using $m_f([1]) = [f] \in A$, it follows easily that M_f is the zero matrix if and only if f is in the ideal.

Sparse Polynomials. It is also possible to develop a sparse version of the solution method described in this section. The idea is that one fixes in advance the terms which appear in each f_i and then considers what happens when f_i is generic relative to these terms. One gets results similar to Theorems 3.1.1 and 3.1.2, and there are also nice relations to polyhedral geometry. This material is discussed in Chapter 7 of [CLO2].

Duality. The assumption that f_1, \dots, f_n have only finitely many solutions in F^n implies that these polynomials form a *regular sequence*. This allows us to apply the duality theory of complete intersections. There are also interesting relations with multidimensional residues. This material is discussed in [EM1] and [EM2].

Finally, Section 3.2 will discuss relations with the theory of multivariate resultants.

3.2 The U-Resultant

The classical multivariable resultant $\text{Res}_{d_0, \dots, d_n}$ in an irreducible polynomial in the coefficients of $n + 1$ homogeneous polynomials

$$F_0, \dots, F_n \in F[x_0, \dots, x_n]$$

of degrees d_0, \dots, d_n with the property that number

$$\text{Res}_{d_0, \dots, d_n}(F_0, \dots, F_n) = 0$$

if and only if the F_i have a common solution in the projective space $\mathbb{P}^n(F)$ (as usual, $F = \overline{F}$).

This resultant can also have an affine version as follows. If we dehomogenize F_i by setting $x_0 = 1$, then we get polynomials $f_i \in F[x_1, \dots, x_n]$ of degree at most d_i . Since F_i and f_i have the same coefficients, we can write the resultant as

$$\text{Res}_{d_0, \dots, d_n}(f_0, \dots, f_n).$$

Then the vanishing of this resultant means that either the equations $f_0 = \dots = f_n$ have a solution in F^n or they have a solution “at infinity” (i.e., a projective solution with $x_0 = 0$).

In the situation of Section 3.1, we have n polynomials f_1, \dots, f_n of degrees d_1, \dots, d_n in x_1, \dots, x_n . To compute a resultant, we need one more equation. Not surprisingly, we will use

$$f_0 = a_1 x_1 + \dots + a_n x_n.$$

We will usually assume $a_i \in F$ though (as illustrated at the end of Section 3.1) it is sometimes useful to replace a_i with a variable u_i .

In order to compute the resultant $\text{Res}_{1,d_1,\dots,d_n}(f_0, f_1, \dots, f_n)$, we need to study the behavior of the system $f_1 = \dots = f_n = 0$ at ∞ . Write

$$f_i = \sum_{j=0}^{d_i} f_{i,j}$$

where $f_{i,j}$ is homogeneous of degree j in x_1, \dots, x_n . Then f_i homogenizes to

$$F_i = \sum_{j=0}^{d_i} f_{i,j} x_0^{d_i-j}$$

of degree d_i in x_0, x_1, \dots, x_n . Then (3.1.1) has a *solution at ∞* when the homogenized system

$$F_1 = \dots = F_n = 0$$

has a nontrivial solution with $x_0 = 0$.

The following result relates solutions at ∞ to the algebra $A = F[x_1, \dots, x_n]/\langle f_1, \dots, f_n \rangle$.

Lemma 3.2.1 *The following are equivalent:*

$$\begin{aligned} f_1 = \dots = f_n = 0 \text{ has no solutions at } \infty &\iff \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \neq 0 \\ &\iff A \text{ has dimension } d_1 \cdots d_n \text{ over } F. \end{aligned}$$

Proof. Note that F_i reduces to f_{i,d_i} when $x_0 = 0$. Thus the f_i have a solution at ∞ if and only if the system of homogeneous equations

$$f_{1,d_1} = \dots = f_{n,d_n} = 0$$

has a nontrivial solution. This gives the first equivalence. The second uses Bézout's Theorem and some facts from algebraic geometry. See Section 3 of Chapter 3 of [CLO2] for the details. \square

When there are no solutions at ∞ , it follows that we get our algebra A of dimension $d_1 \cdots d_n$ over F . But unlike Section 3.1, the solutions may have multiplicities > 1 . In this case, we can relate resultants and multiplication maps as follows.

Theorem 3.2.2 *If $f_0 = u_1 x_1 + \dots + u_n x_n$ and (3.1.1) has no solutions at ∞ , then*

$$\begin{aligned} \text{Res}_{1,d_1,\dots,d_n}(f_0, f_1, \dots, f_n) &= \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \det(m_{f_0}) \\ &= \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \prod_p (u_1 p_1 + \dots + u_n p_n)^{\text{mult}(p)}. \end{aligned}$$

and

$$\begin{aligned} \text{Res}_{1,d_1,\dots,d_n}(u - f_0, f_1, \dots, f_n) &= \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \text{CharPoly}_{m_{f_0}}(u) \\ &= \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \prod_p (u - (u_1 p_1 + \dots + u_n p_n))^{\text{mult}(p)}. \end{aligned}$$

Proof. In each case, the first equality uses Theorem 3.4 of Chapter 3 of [CLO2] and the second equality uses Proposition 1.5.1 of these notes. \square

While this is nice (and will have some unexpected consequences when we discuss Galois theory in Section 5), the relation between resultants and what we did in Section 3.1 goes much deeper. We can explore this as follows.

Computing Resultants. First recall that the method given in Section 3.1 for computing multiplication maps used the equality

$$\widetilde{M}_{f_0} = M_{f_0}^T$$

from Theorem 3.1.2. This in turn requires that all solutions have multiplicity 1 and that $\det(M_{11}) \neq 0$ (since $\det(M_{11})^{-1}$ is used in the construction of \widetilde{M}_{f_0}). This relates to resultants as follows.

A standard method for computing $\text{Res}_{1,d_1,\dots,d_n}(f_0, f_1, \dots, f_n)$ involves the quotient of two determinants. In our situation, the relevant formula is

$$\det(M_0) = \text{Res}_{1,d_1,\dots,d_n}(f_0, f_1, \dots, f_n) \det(M'_0), \quad (3.2.1)$$

where M_0 is *precisely* the matrix appearing in (3.1.4) and M'_0 is the submatrix described in Section 4 of Chapter 3 of [CLO2]. It follows that

$$\text{Res}_{1,d_1,\dots,d_n}(f_0, f_1, \dots, f_n) = \frac{\det(M_0)}{\det(M'_0)}$$

whenever $\det(M'_0) \neq 0$. The subtle point is that $\det(M_0)$ and $\det(M'_0)$ can both vanish even though $\text{Res}_{1,d_1,\dots,d_n}(f_0, f_1, \dots, f_n)$ is nonzero. So to calculate the resultant using M_0 , we definitely need $\det(M'_0) \neq 0$. Yet for \widetilde{M}_{f_0} , we need $\det(M_{11}) \neq 0$. Here is the nice relation between these determinants.

Proposition 3.2.3 *We have*

$$\det(M_{11}) = \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \det(M'_0).$$

Furthermore, if $\det(M_{11}) \neq 0$, then \widetilde{M}_{f_0} and m_{f_0} have the same determinant and same characteristic polynomial.

Proof. First observe that by (3.1.4) and the definition of \widetilde{M}_{f_0} , we have

$$\det(M_0) = \det \begin{pmatrix} I & M_{01}M_{11}^{-1} \\ 0 & I \end{pmatrix} \det \begin{pmatrix} M_{00} & M_{01} \\ M_{10} & M_{11} \end{pmatrix} = \det \begin{pmatrix} \widetilde{M}_{f_0} & 0 \\ M_{10} & M_{11} \end{pmatrix} = \det(\widetilde{M}_{f_0}) \det(M_{11}).$$

whenever $\det(M_{11}) \neq 0$. Using this together with (3.2.1) and Theorems 3.2.2 and 3.1.2, we obtain

$$\begin{aligned} \det(\widetilde{M}_{f_0}) \det(M_{11}) &= \det(M_0) \\ &= \text{Res}_{1,d_1,\dots,d_n}(f_0, f_1, \dots, f_n) \det(M'_0) \\ &= \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \det(m_{f_0}) \det(M'_0) \\ &= \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \det(\widetilde{M}_{f_0}) \det(M'_0) \end{aligned}$$

when f_1, \dots, f_n are sufficiently generic. Cancelling $\det(\widetilde{M}_{f_0})$ (which is nonzero generically) shows that the equality $\det(M_{11}) = \text{Res}_{d_1,\dots,d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \det(M'_0)$ holds generically. Since each side is a polynomial in the coefficients of the f_i , this equality must hold unconditionally.

For the final assertion of the proposition, observe that

$$\det(M_{11}) \neq 0 \implies \text{Res}_{d_1, \dots, d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \neq 0 \quad (3.2.2)$$

by what we just proved. It follows that both \widetilde{M}_{f_0} and m_{f_0} are defined when $\det(M_{11}) \neq 0$. Unfortunately, we no longer know that S_0 provides a basis of A , and even if it does, the proof that $\widetilde{M}_{f_0}^T = M_{f_0}$ given in Theorem 3.1.2 breaks down since the eigenvectors \mathbf{p}^α no longer span A when there are solutions of multiplicity > 1 . But we can still relate $\det(\widetilde{M}_{f_0})$ and $\det(m_{f_0})$ as follows. If $\det(M_{11}) \neq 0$, then (3.2.1), (3.2.2), Theorem 3.2.2 and the first assertion of the proposition imply that

$$\det(M_0) = \det(m_{f_0}) \det(M_{11}).$$

However, we showed above that $\det(M_0) = \det(\widetilde{M}_{f_0}) \det(M_{11})$ when $\det(M_{11}) \neq 0$. The desired equality now follows immediately, and we get the statement concerning characteristic polynomials by replacing f_0 with $u - f_0$. \square

Proposition 3.2.3 shows that the assumption $\det(M_{11}) \neq 0$ needed to define \widetilde{M}_{f_0} in Section 3.1 guarantees three things:

- $\text{Res}_{d_1, \dots, d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \neq 0$, so that (3.1.1) has no solutions at ∞ .
- $\det(M'_0) \neq 0$, so that M_0 and M'_0 can be used to compute $\text{Res}_{1, d_1, \dots, d_n}(f_0, f_1, \dots, f_n)$.
- Even if there are multiplicities, \widetilde{M}_{f_0} and m_{f_0} have the same characteristic polynomial.

So the link between Section 3.1 and resultants is very strong.

For the experts, note that (3.2.1) and Proposition 3.2.3 imply that if $\det(M'_0) \neq 0$, then

$$\begin{aligned} \text{Res}_{1, d_1, \dots, d_n}(f_0, f_1, \dots, f_n) &= \frac{\det M_0}{\det M'_0} \\ \text{Res}_{d_1, \dots, d_n}(f_{1,d_1}, \dots, f_{n,d_n}) &= \frac{\det M_{11}}{\det M'_0}, \end{aligned}$$

where M'_0 is a submatrix of M_{11} , which in turn is a submatrix of M_0 . So M_0 allows us to compute not one but two resultants. Has this been noticed before?

Genericity. In Section 3.1, we required that f_1, \dots, f_n be “generic”, which upon careful rereading means first, that the system (3.1.1) has $d_1 \cdots d_n$ solutions of multiplicity 1, and second, that $\det(M_{11}) \neq 0$. In terms of resultants, this means the following:

- $\text{Res}_{d_1, \dots, d_n}(f_{1,d_1}, \dots, f_{n,d_n}) \neq 0$.
- $\text{Res}_{d-1, d_1, \dots, d_n}(\det(\frac{\partial f_i}{\partial x_j}), f_1, \dots, f_n) \neq 0$, where d is defined in (3.1.2).
- $\det(M'_0) \neq 0$.

The first item guarantees that A has the correct dimension by Lemma 3.2.1 and the second guarantees that the Jacobian is nonvanishing at all solutions, so that every solution has multiplicity 1 by the implicit function theorem. Finally, the first and third conditions are equivalent to $\det(M_{11}) \neq 0$ by Proposition 3.2.3.

One historical remark is that while the formula (3.2.1) is due to Macaulay in 1902, many ideas of Section 3.2 are present in the work of Kronecker in 1882. For example, Kronecker defines

$$\text{Res}_{1, d_1, \dots, d_n}(u - f_0, f_1, \dots, f_n)$$

and shows that as a polynomial in u , its roots are $f_0(p)$ for p a solution of (3.1.1). He also noted that the discriminant condition of the second bullet is needed to get solutions of multiplicity 1 and that when this is true, the ideal $\langle f_1, \dots, f_n \rangle$ is radical (see pp. 276 and 330 of [Kronecker, Vol. II]).

4 Factoring

4.1 Factoring over Number Fields

Near the end of Section 3.1, we gave the formula (3.1.10)

$$\det(\widetilde{M}_{f_0}) = \prod_p (u_1 p_1 + \dots + u_n p_n)$$

where $f_0 = u_1 x_1 + \dots + u_n x_n$. The point was that we could compute the left-hand side, so that if we knew how to factor multivariable polynomials over an algebraically closed field, *then* we could find all of the solutions. Of course, such factoring is very difficult. But as we will now see, it is sometimes possible to turn the tables and use the finite algebras and their multiplication maps to do factoring over number fields. We begin with a lovely result of Dedekind.

Dedekind Reciprocity. Suppose that $f(x), g(x) \in \mathbb{Q}[x]$ are irreducible with roots $\alpha, \beta \in \mathbb{C}$ such that $f(\alpha) = g(\beta) = 0$. Then factor $f(x)$ into irreducibles over $\mathbb{Q}(\beta)$, say

$$f(x) = f_1(x) \cdots f_r(x), \quad f_i(x) \in \mathbb{Q}(\beta)[x]. \quad (4.1.1)$$

The $f_i(x)$ are distinct (i.e., none is a constant multiple of any of the others) since f is separable. Then the *Dedekind Reciprocity Theorem* describes the factorization of $g(x)$ over $\mathbb{Q}(\alpha)$ as follows.

Theorem 4.1.1 *Given the above factorization of $f(x)$ into irreducibles over $\mathbb{Q}(\beta)$, the factorization of $g(x)$ into irreducibles over $\mathbb{Q}(\alpha)$ can be written as*

$$g(x) = g_1(x) \cdots g_r(x), \quad g_i(x) \in \mathbb{Q}(\alpha)[x]$$

where

$$\frac{\deg(f_1)}{\deg(g_1)} = \frac{\deg(f_2)}{\deg(g_2)} = \dots = \frac{\deg(f_r)}{\deg(g_r)} = \frac{\deg(f)}{\deg(g)}.$$

Proof. Consider the \mathbb{Q} -algebra

$$A = \mathbb{Q}[x, y] / \langle f(x), g(y) \rangle.$$

Since $y \mapsto \beta$ induces $\mathbb{Q}[y] / \langle g(y) \rangle \simeq \mathbb{Q}(\beta)$ and the $f_i(x)$ are distinct irreducibles, we get

$$A \simeq \mathbb{Q}(\beta)[x] / \langle f(x) \rangle \simeq \prod_{i=1}^r K_i, \quad (4.1.2)$$

where $K_i = \mathbb{Q}(\beta)[x] / \langle f_i(x) \rangle$ is a field. Since $[K_i : \mathbb{Q}(\beta)] = \deg(f_i)$, the degree of K_i over \mathbb{Q} is

$$[K_i : \mathbb{Q}] = [K_i : \mathbb{Q}(\beta)][\mathbb{Q}(\beta) : \mathbb{Q}] = \deg(f_i) \deg(g). \quad (4.1.3)$$

In the same way, the factorization of $g(y)$ into r' irreducibles $g_i(y)$ over $\mathbb{Q}(\alpha) \simeq \mathbb{Q}[x] / \langle f \rangle$ gives an isomorphism between A and a product of s fields $A \simeq \prod_{i=1}^{r'} K'_i$ such that

$$[K_i : \mathbb{Q}] = \deg(g_i) \deg(f). \quad (4.1.4)$$

However, the decomposition of A into a product of fields is unique up to isomorphism. Hence we must have $r = r'$ and (after a suitable permutation) $K_i \simeq K'_i$. It follows that (4.1.3) and (4.1.4) must equal for all i , and the result follows. \square

According to [Edwards], Dedekind discovered this result in 1855, though his version wasn't published until 1982. Kronecker found this theorem independently and stated it in his university lectures. Theorem 4.1.1 was first published by Kneser in 1887.

A Factorization Algorithm. The algebra A used in the proof of Theorem 4.1.1 can also be used to construct the factorization of $f(x)$ over $\mathbb{Q}(\beta)$. The idea is to compute

$$\Phi(u) = \text{CharPoly}_{m_{f_0}}(u), \quad f_0 = x + ty,$$

for a carefully chosen $t \in \mathbb{Q}$. This polynomial in $\mathbb{Q}[u]$ can be computed using the methods of Section 1 and factored using known algorithms for factoring over \mathbb{Q} . Kronecker observed that these factors determine the factorization of $f(x)$ over $\mathbb{Q}(\beta)$.

To see how this works, we first need to see how the factorization of $f(x)$ influences $\Phi(u)$. Under the isomorphism $\mathbb{Q}(\beta) \simeq \mathbb{Q}[y]/\langle g(y) \rangle$, a polynomial $h(x, y) \in \mathbb{Q}[x, y]$ gives $h(x, \beta) \in \mathbb{Q}(\beta)[x]$, and all polynomials in $\mathbb{Q}(\beta)[x]$ can be represented this way. In particular, the factorization (4.1.1) of $f(x)$ over $\mathbb{Q}(\beta)$ is written

$$f(x) = f_1(x, \beta) \cdots f_r(x, \beta) \in \mathbb{Q}(\beta)[x], \quad f_i(x, y) \in \mathbb{Q}[x, y], \quad (4.1.5)$$

and the decomposition (4.1.2) can be written

$$A \simeq \prod_{i=1}^r K_i = \prod_{i=1}^r \mathbb{Q}[x, y]/\langle g(y), f_i(x, y) \rangle. \quad (4.1.6)$$

This decomposition induces a factorization

$$\Phi(u) = \prod_{i=1}^r \Phi_i(u) \quad (4.1.7)$$

over \mathbb{Q} , where $\Phi_i(u)$ is the characteristic polynomial of m_{f_0} on K_i .

Theorem 4.1.2 *If $f_0 = x + ty$ takes distinct values at the solutions of $f(x) = g(y) = 0$, then for all $i = 1, \dots, r$, $\Phi_i(u)$ is irreducible over \mathbb{Q} and the irreducible factor $f_i(x, \beta)$ from (4.1.5) is*

$$f_i(x, \beta) = \text{GCD}(\Phi_i(x + t\beta), f(x)),$$

where the GCD is computed in $\mathbb{Q}(\beta)[x]$.

Proof. We will prove the irreducibility of $\Phi_i(u)$ using the methods introduced in Section 1. If $n = \deg(f(x))$ and $m = \deg(g(y))$, then Bézout's Theorem implies that the equations $f(x) = g(y) = 0$ have at most nm solutions in x and y counted with multiplicity. But in fact there are exactly nm solutions since $f(x)$ and $g(y)$ are separable. It follows that all of the multiplicities are 1.

By assumption, $f_0 = x + ty$ takes distinct values at all solutions of $f(x) = g(y) = 0$. Since they have multiplicity 1, it follows that m_{f_0} is non-derogatory. Then the single-variable representation given by Proposition 1.4.1 implies the map sending u to $[x + ty] \in A$ induces an isomorphism

$$\mathbb{Q}[u]/\langle \Phi(u) \rangle \simeq A$$

since $\Phi(u)$ is the characteristic polynomial of multiplication by $f_0 = x + ty$ on A . Notice also that

$$\text{Disc}(\Phi(u)) \neq 0$$

since the eigenvalues all have multiplicity 1. This implies that $\Phi(u)$ is a product of distinct irreducibles. Such a factorization gives a decomposition of $\mathbb{Q}[u]/\langle \Phi(u) \rangle \simeq A$ into a product of fields. This decomposition must coincide with (4.1.6) since each K_i is a field. It follows that (4.1.7) is the irreducible factorization of $\Phi(u)$.

We next observe that $f_i(x, \beta)$ divides $\Phi_i(x + t\beta)$ in $\mathbb{Q}(\beta)[x]$. By the Cayley-Hamilton Theorem, $\Phi_i(m_{f_0})$ is the zero linear map on $K_i = \mathbb{Q}[x, y]/\langle g(y), f_i(x, y) \rangle$ since $\Phi_i(u)$ is the characteristic polynomial of m_{f_0} on K_i . Applying this to $[1] \in K_i$, we get

$$[0] = \Phi_i(m_{f_0})([1]) = [\Phi_i(f_0)]$$

in $K_i = \mathbb{Q}[x, y]/\langle g(y), f_i(x, y) \rangle$. This implies

$$\Phi_i(x + ty) \in \langle g(y), f_i(x, y) \rangle \subset \mathbb{Q}[x, y]. \quad (4.1.8)$$

Then the substitution $y \mapsto \beta$ gives

$$\Phi_i(x + t\beta) \in \langle f_i(x, \beta) \rangle \subset \mathbb{Q}(\beta)[x].$$

Since $f_i(x, \beta)$ divides $f(x)$ by (4.1.5), we conclude that $f_i(x, \beta)$ divides the GCD in the statement of the theorem.

To see that $f_i(x, \beta) = \text{GCD}(\Phi_i(x + t\beta), f(x))$, note that $f(x) = \prod_{i=1}^r f_i(x, \beta)$ divides $\Phi(x + t\beta) = \prod_{i=1}^r \Phi_i(x + t\beta)$. Thus

$$f(x) = \text{GCD}(\Phi(x + t\beta), f(x)).$$

However, we proved above that $\Phi(u)$ has distinct roots, so that the same is true for $\Phi(x + t\beta)$. It follows that in the factorization $\Phi(x + t\beta) = \prod_{i=1}^r \Phi_i(x + t\beta)$, the factors $\Phi_i(x + t\beta)$ are mutually relatively prime. Hence the above GCD calculation may be written as

$$f(x) = \prod_{i=1}^r \text{GCD}(\Phi_i(x + t\beta), f(x)).$$

In other words,

$$\prod_{i=1}^r f_i(x, \beta) = \prod_{i=1}^r \text{GCD}(\Phi_i(x + t\beta), f(x)).$$

Since $f_i(x, \beta)$ divides $\text{GCD}(\Phi_i(x + t\beta), f(x))$ for each i , we get the desired equality. \square

This theorem leads to the following algorithm for factoring $f(x)$ over $\mathbb{Q}(\beta)$:

- Pick a random $t \in \mathbb{Q}$ and compute $\Phi(u) = \text{CharPoly}_{m_{f_0}}(u)$ for $f_0 = x + ty$. Also compute $\text{Disc}(\Phi(u))$.
- If $\text{Disc}(\Phi(u)) \neq 0$, then factor $\Phi(u) = \prod_{i=1}^r \Phi_i(u)$ into irreducibles in $\mathbb{Q}[u]$ and for each i compute $\text{GCD}(\Phi_i(x + t\beta), f(x))$ in $\mathbb{Q}(\beta)[x]$. This gives the desired factorization.
- If $\text{Disc}(\Phi(u)) = 0$, then pick a new $t \in \mathbb{Q}$ and return to the first bullet.

Since $\text{Disc}(\Phi(u)) \neq 0$ if and only if $x + ty$ takes distinct values at the solutions of $f(x) = g(y) = 0$, Theorem 4.1.2 implies that the second bullet correctly computes the required factorization when the discriminant is nonzero. Notice that it only uses the Euclidean algorithm in $\mathbb{Q}(\beta)[x]$, which can be done constructively using the representation $\mathbb{Q}(\beta) \simeq \mathbb{Q}[y]/\langle g(y) \rangle$.

As for the third bullet, it is an easy exercise to show that number of $t \in \mathbb{Q}$ which satisfy the equation $\text{Disc}(\Phi(u)) = 0$ is bounded above by $\frac{1}{2}nm(nm - 1)$ ($n = \deg(f(x))$, $m = \deg(g(y))$). Thus the third bullet can occur at most $\frac{1}{2}nm(nm - 1)$ times. It follows that the above algorithm is deterministic.

An alternate approach would be to follow what Kronecker does on pages 258–259 of [Kronecker, Vol. II] and regard t as a variable in $f_0 = x + ty$. Then $\Phi(u)$ becomes a polynomial $\Phi(u, t) \in \mathbb{Q}[x, t]$. If one can factor $\Phi(u, t)$ into irreducibles $\mathbb{Q}[u, t]$, say $\Phi(u, t) = \prod_{i=1}^r \Phi_i(u, t)$, then it is straightforward to recover $f_i(x, \beta)$ from $\Phi_i(x + t\beta, t)$. A rigorously constructive version of this is described in [Edwards]. We should also mention that for Kronecker, the crucial observation (4.1.8) was a consequence of the properties of resultants (see page 330 of [Kronecker Vol. II]).

4.2 Finite Fields and Primitive Elements

Here, we will give two further applications of finite commutative algebras.

Factoring over Finite Fields. We begin with a brief description for factoring a polynomial $f(x) \in \mathbb{F}_q[x]$, where \mathbb{F}_q is a finite field with $q = p^\ell$ elements. We will use the algebra

$$A = \mathbb{F}_q[x]/\langle f(x) \rangle$$

and the Frobenius map

$$F : A \rightarrow A, \quad F(a) = a^q.$$

This map is linear over \mathbb{F}_q and can be used to detect whether or not $f(x)$ is irreducible as follows.

Proposition 4.2.1 *If $f(x)$ has no multiple roots (i.e., $\text{GCD}(f(x), f'(x)) = 1$), then the dimension of the eigenspace $E_A(F, 1)$ is the number of irreducible factors of $f(x)$.*

Proof. Since $f(x)$ has no multiple roots, a factorization $f(x) = f_1(x) \cdots f_r(x)$ into irreducible polynomials in $\mathbb{F}_q[x]$ gives an algebra isomorphism

$$A \simeq \prod_{i=1}^r K_i = \prod_{i=1}^r \mathbb{F}_q[x]/\langle f_i(x) \rangle$$

which is compatible with the Frobenius map F . If $a \in K_i$, then since K_i is a field, we have

$$F(a) = a \iff a^q = a \iff a \in \mathbb{F}_q.$$

It follows that on K_i , the eigenvalue 1 has a 1-dimensional eigenspace $E_{K_i}(F, 1)$. Since the eigenspace $E_A(F, 1)$ is the direct sum of the $E_{K_i}(F, 1)$, the result follows. \square

Here is a simple example of this result.

Example 4.2.2 Let $f(x) = x^5 + x^4 + 1 \in \mathbb{F}_2[x]$. One easily sees that $f(x)$ is separable. Then $A = \mathbb{F}_2[x]/\langle f(x) \rangle$ is a vector space over \mathbb{F}_2 of dimension 5 with basis $[1], [x], [x^2], [x^3], [x^4]$, which for simplicity we write as $1, x, x^2, x^3, x^4$.

Note that $F : A \rightarrow A$ is the squaring map since $q = 2$. To compute the matrix of F , we apply F to each basis element and represent the result in terms of the basis:

$$\begin{aligned} 1 &\mapsto 1 \\ x &\mapsto x^2 \\ x^2 &\mapsto x^4 \\ x^3 &\mapsto x^6 = 1 + x + x^4 \\ x^4 &\mapsto x^8 = 1 + x + x^2 + x^3 + x^4. \end{aligned}$$

Here, $x^6 = 1 + x + x^4$ means that $1 + x + x^4$ is the remainder of x^6 on division by $f(x) = x^5 + x^4 + 1$, and similarly for the last line. Hence the matrix of $F - 1_A$ is

$$\begin{pmatrix} 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{pmatrix} - \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 \end{pmatrix}$$

(remember that we are in characteristic 2). This matrix has rank 3 since the first column is zero and the sum of the last three columns is zero. Hence we have two linearly independent eigenvectors for the eigenvalue 1. By Proposition 4.2.1, $f(x)$ is not irreducible over \mathbb{F}_2 . \square

Besides giving the number of irreducible factors of $f(x)$, one can also use the eigenspace $E_A(F, 1)$ to construct the irreducible factorization of $f(x)$. The rough idea is that if $[h(x)] \in E_A(F, 1)$ is a nonzero element of $E_A(F, 1)$, then $\text{GCD}(h(x), f(x))$ is a factor of f and (if $h(x)$ chosen correctly) is actually one of the irreducible factors of $f(x)$. This is Berlekamp's Algorithm, which is described in Section 4.1 of [LN].

Theorem of the Primitive Element. The single-variable representation used in the proof of Theorem 4.1.2 may remind the reader of the Theorem of the Primitive Element. As we will now show, this is no accident.

Theorem 4.2.3 *Let $F \subset L = F(\alpha_1, \dots, \alpha_n)$ be an extension such that F is infinite and each α_i is separable over F . Then there are $t_1, \dots, t_n \in F$ such that*

$$L = F(\alpha), \quad \alpha = t_1\alpha_1 + \dots + t_n\alpha_n.$$

Proof. Let f_i be the minimal polynomial of α_i over F and let

$$A = F[x_1, \dots, x_n] / \langle f_1(x_1), \dots, f_n(x_n) \rangle.$$

Note that we use a separate variable x_i for each polynomial f_i . Then arguing as in the proof of Theorem 4.1.2, one easily sees that by Bézout's Theorem, all solutions of

$$f_1(x_1) = f_2(x_2) = \dots = f_n(x_n) = 0 \tag{4.2.1}$$

have multiplicity 1. Since F is infinite, we can pick $t_1, \dots, t_n \in F$ such that $f_0 = t_1x_1 + \dots + t_nx_n$ takes distinct values at all solutions of (4.2.1). It follows that $m_{f_0} : A \rightarrow A$ is non-derogatory, so that by Proposition 1.4.1, the map $u \mapsto [t_1x_1 + \dots + t_nx_n] \in A$ induces a surjection

$$F[u] \longrightarrow A.$$

But $x_i \mapsto \alpha_i$ induces well-defined map $A \rightarrow L$ which is a surjection since $L = F(\alpha_1, \dots, \alpha_n)$. Then the theorem follows since the composed map $F[u] \rightarrow A \rightarrow L$ is surjective and maps u to $\alpha = t_1\alpha_1 + \dots + t_n\alpha_n$. \square

Here is an example to illustrate the role of separability.

Example 4.2.4 Let $F = \mathbb{F}_p(t, u)$, where t and u are variables, and let $F \subset L$ be the field obtained by adjoining the the p th roots of t and u . It is easy to see that $F \subset L$ is purely inseparable of degree p^2 and that $L \neq F(\alpha)$ for all $\alpha \in L$ (the latter follows from $\alpha \in L \Rightarrow \alpha^p \in F$).

Hence the single-variable representation of Proposition 1.4.1 must fail. To see how this works, first observe that

$$L \simeq F[x, y]/\langle x^p - t, y^p - u \rangle$$

is the algebra we used in the proof of Theorem 4.2.3. In the algebraic closure \overline{F} of F , the only solution of

$$x^p - t = y^p - u = 0$$

is given by $x = \sqrt[p]{t}$ and $y = \sqrt[p]{u}$. The local ring at this point is

$$\overline{F}[x, y]/\langle x^p - t, y^p - u \rangle = \overline{F}[x, y]/\langle (x - \sqrt[p]{t})^p, (y - \sqrt[p]{u})^p \rangle \simeq \overline{F}[x, y]/\langle x^p, y^p \rangle,$$

which clearly has embedding dimension 2 and hence is not curvilinear. It follows that m_f is derogatory for *all* $f \in F[x, y]$. Since the single-variable representation requires that m_f be non-derogatory, we can see why the Theorem of the Primitive Element fails in this case. \square

4.3 Primary Decomposition

The final task of Section 4 is to extend the factorizations introduced in Section 4.1 to the realm of ideals. Suppose that

$$f_1(x_1, \dots, x_n) = \dots = f_s(x_1, \dots, x_n) = 0 \tag{4.3.1}$$

is a system of equations with coefficients in a field F and only finitely many solutions over the algebraic closure \overline{F} . (Thus we are back in the situation where the number of equations need not equal the number of variables.) We say that $\langle f_1, \dots, f_s \rangle$ is *zero-dimensional* since a finite set of points has dimension 0. Our goal is to give an algorithm for computing the primary decomposition of a zero-dimensional ideal.

Theoretical Results. An ideal $I \subset F[x_1, \dots, x_n]$ is *primary* if $fg \in I$ always implies that either $f \in I$ or $g^N \in I$ for some $N \geq 1$. It is easy to see that the radical \sqrt{I} of a primary ideal is prime. by Chapter 4, §7 of [CLO1], every ideal $I \subset F[x_1, \dots, x_n]$ has a *primary decomposition*

$$I = I_1 \cap \dots \cap I_r \tag{4.3.2}$$

into an intersection of primary ideals. We say that (4.3.2) is *minimal* when r is as small as possible.

In the zero-dimensional case, the primary components I_i of $\langle f_1, \dots, f_s \rangle$ can be obtained from the given ideal by adding one more carefully chosen polynomial u_i . Here is the precise result.

Lemma 4.3.1 *A zero-dimensional ideal $\langle f_1, \dots, f_s \rangle$ has a minimal primary decomposition*

$$\langle f_1, \dots, f_s \rangle = I_1 \cap \dots \cap I_r$$

such that $\sqrt{I_1}, \dots, \sqrt{I_r}$ are distinct maximal ideals. Furthermore, for each i ,

$$I_i \not\subset \bigcup_{j \neq i} \sqrt{I_j},$$

and any $u_i \in I_i \setminus \bigcup_{j \neq i} \sqrt{I_j}$ has the property that

$$I_i = \langle f_1, \dots, f_s, u_i \rangle.$$

Proof. Let $\langle f_1, \dots, f_s \rangle = I_1 \cap \dots \cap I_r$ be a minimal primary decomposition. Note that I_i and hence $\sqrt{I_i}$ are zero-dimensional since $\langle f_1, \dots, f_s \rangle$ is. But we also know that $\sqrt{I_i}$ is prime. An easy argument shows that the only zero-dimensional prime ideals are maximal. Also, if $\sqrt{I_i} = \sqrt{I_j}$ for some $i \neq j$, then one can easily show that $I_i \cap I_j$ is primary, which contradicts the minimality of our representation. Hence the $\sqrt{I_i}$ are distinct.

If $I_i \subset \bigcup_{j \neq i} \sqrt{I_j}$, then $I_i \subset \sqrt{I_j}$ for some $j \neq i$ by the Prime Avoidance Theorem (Theorem 3.61 of [Sharp]). This implies $\sqrt{I_i} \subset \sqrt{I_j}$ and hence $\sqrt{I_i} = \sqrt{I_j}$ since the radicals are maximal. This contradiction proves that $I_i \not\subset \bigcup_{j \neq i} \sqrt{I_j}$.

Now let $u_i \in I_i \setminus \bigcup_{j \neq i} \sqrt{I_j}$. Then we certainly have $\langle f_1, \dots, f_s, u_i \rangle \subset I_i$. For the opposite inclusion, take $j \neq i$ and note that $u_i \notin \sqrt{I_j}$ implies that $1 + u_i r_j \in \sqrt{I_j}$ for some r_j since $\sqrt{I_j}$ is maximal. Thus $(1 + u_i r_j)^{N_j} \in I_j$ for some $N_j \geq 1$. Expanding the product

$$\prod_{j \neq i} (1 + u_i r_j)^{N_j} \in \prod_{j \neq i} I_j \subset \bigcap_{j \neq i} I_j,$$

we see that $1 + u_i r \in \bigcap_{j \neq i} I_j$ for some r . Now take $a \in I_i$. Then

$$a(1 + u_i r) \in I_i \cap \bigcap_{j \neq i} I_j = \langle f_1, \dots, f_s \rangle.$$

Hence $a = a(1 + u_i r) + u_i(-ar) \in \langle f_1, \dots, f_s \rangle + \langle u_i \rangle = \langle f_1, \dots, f_s, u_i \rangle$, as desired. \square

In the zero-dimensional case, one can also prove that the ideals I_i in the primary decomposition are unique. For general ideals, uniqueness need not hold (see Exercise 6 of Chapter 4, §7 of [CLO1] for an example) due to the phenomenon of *embedded components*.

The most commonly used algorithm for computing the primary decomposition of a zero-dimensional ideal is described in [GTZ] and uses Gröbner bases plus a change of coordinates to find the u_i of Lemma 4.3.1. However, the recent paper [Monico] shows how to find the u_i using the quotient algebra

$$A = F[x_1, \dots, x_n] / \langle f_1, \dots, f_s \rangle.$$

We will describe Monico's method, beginning the following special case.

The Rational Case. The solutions of (4.3.1) are *rational over F* if all solutions in \overline{F}^n actually lie in F^n . In this situation, it is easy to see that the primary decomposition is

$$\langle f_1, \dots, f_s \rangle = \bigcap_p I_p,$$

where the intersection is over all solutions p of (4.3.1). Furthermore, as we noted in (1.3.3), the primary component I_p is

$$I_p = \{f \in F[x_1, \dots, x_n] \mid gf \in \langle f_1, \dots, f_s \rangle \text{ for some } g \in F[x_1, \dots, x_n] \text{ with } g(p) \neq 0\},$$

and $\sqrt{I_p}$ is the maximal ideal $\langle x_1 - p_1, \dots, x_n - p_n \rangle$ when $p = (p_1, \dots, p_n)$. Unfortunately, this elegant description of I_p is not useful for computational purposes. But we can use the methods of Section 1 to find the polynomials u_i of Lemma 4.3.1 as follows.

Proposition 4.3.2 *Suppose that $\langle f_1, \dots, f_s \rangle$ is zero-dimensional and all solutions of (4.3.1) are rational over F . If $f \in F[x_1, \dots, x_n]$ takes distinct values at the solutions of (4.3.1), then for each solution p , the corresponding primary component is*

$$I_p = \langle f_1, \dots, f_s, (f - f(p))^{\text{mult}(p)} \rangle.$$

Proof. Let $u_p = (f - f(p))^{\text{mult}(p)}$. By Lemma 4.3.1, it suffices to show that $u_p \in I_p$ and $u_p \notin \sqrt{I_q}$ for all solutions $q \neq p$. Since $\sqrt{I_q}$ is the maximal ideal of q , the latter condition is equivalent to the non-vanishing of u_p at q , which follows since f takes distinct values at the solutions.

To prove that $u_p \in I_p$, let $v_p = \prod_{q \neq p} (f - f(q))^{\text{mult}(q)}$. By Proposition 1.5.1,

$$u_p v_p = \text{CharPoly}_{m_f}(f) \tag{4.3.3}$$

since f takes distinct values at the solutions. However, the Cayley-Hamilton Theorem tells us that $\text{CharPoly}_{m_f}(m_f)$ is the zero operator on A . Applied to $[1] \in A$, we obtain

$$[0] = \text{CharPoly}_{m_f}(m_f)[1] = \text{CharPoly}_{m_f}([f]) = [\text{CharPoly}_{m_f}(f)].$$

Combined with (4.3.3), this implies

$$u_p v_p = \text{CharPoly}_{m_f}(f) \in \langle f_1, \dots, f_s \rangle \subset I_p.$$

Since I_p is primary, either u_p or some power of v_p lies in I_p . But

$$v_p(p) = \prod_{q \neq p} (f(p) - f(q))^{\text{mult}(q)} \neq 0$$

since f takes distinct values at the solutions. Hence no power of v_p lies in I_p , so that $u_p \in I_p$. \square

Here is an example of this proposition.

Example 4.3.3 Consider the ideal $\langle x^2 + 2y^2 - 2y, xy^2 - xy, y^3 - 2y^2 + y \rangle \subset F[x, y]$. We saw in Example 1.1.1 that the corresponding equations have solutions $(0, 0)$ and $(0, 1)$, which are rational over F . Since y takes distinct values at the solutions, we can use $f = y$ in Proposition 4.3.2 to compute the primary decomposition.

By Example 1.2.2, the characteristic polynomial of m_y is $u^2(u-1)^3$. It follows that the primary components are

$$I_{(0,0)} = \langle x^2 + 2y^2 - 2y, xy^2 - xy, y^3 - 2y^2 + y, y^2 \rangle = \langle x^2, y \rangle$$

$$I_{(0,1)} = \langle x^2 + 2y^2 - 2y, xy^2 - xy, y^3 - 2y^2 + y, (y-1)^3 \rangle = \langle x^2 + 2(y-1), x(y-1), (y-1)^2 \rangle,$$

where we leave the final equality of each line as an exercise for the reader (for $I_{(0,1)}$, the congruences

$$y(y-1)^2 \equiv (y-1)^2 \pmod{(y-1)^3} \quad \text{and} \quad y(y-1) \equiv y-1 \pmod{(y-1)^2}$$

will be useful). Putting these together, we obtain the primary decomposition

$$\langle x^2 + 2y^2 - 2y, xy^2 - xy, y^3 - 2y^2 + y \rangle = \langle x^2, y \rangle \cap \langle x^2 + 2(y-1), x(y-1), (y-1)^2 \rangle = I_{(0,0)} \cap I_{(0,1)}$$

given in Example 2.3.3. \square

We note that in Proposition 4.3.2, one can replace the characteristic polynomial with the minimal polynomial. Here is the precise result.

Proposition 4.3.4 *Suppose that $\langle f_1, \dots, f_s \rangle$ is zero-dimensional and all solutions of (4.3.1) are rational over F . If $f \in F[x_1, \dots, x_n]$ takes distinct values at the solutions of (4.3.1), then for each solution p , the corresponding primary component is*

$$I_p = \langle f_1, \dots, f_s, (f - f(p))^{n(p)} \rangle,$$

where $\text{MinPoly}_{m_f}(u) = \prod_p (u - f(p))^{n(p)}$.

We leave the proof as an exercise. Here is an example.

Example 4.3.5 For the ideal of Example 4.3.3, recall from Example 1.2.2 that the minimal polynomial of y is $u(u-1)^2$. Thus

$$I_{(0,0)} = \langle x^2 + 2y^2 - 2y, xy^2 - xy, y^3 - 2y^2 + y, y \rangle = \langle x^2, y \rangle$$

$$I_{(0,1)} = \langle x^2 + 2y^2 - 2y, xy^2 - xy, y^3 - 2y^2 + y, (y-1)^2 \rangle = \langle x^2 + 2(y-1), x(y-1), (y-1)^2 \rangle.$$

This gives the same primary decomposition as Example 4.3.3, though the initial description of the primary components is simpler because the minimal polynomial has smaller exponents than the characteristic polynomial. \square

The General Case. Now suppose that F is a field and that the equations (4.3.1) have solutions whose coordinates may lie in a strictly larger field. This means that in the primary decomposition over F , the number of primary components no longer equals the number of solutions. Here is an example taken from [Monico].

Example 4.3.6 The equations $x^2 - 2 = y^2 - 2 = 0$ have four solutions $(\pm\sqrt{2}, \pm\sqrt{2})$, none of which is rational over \mathbb{Q} . We will see below that the primary decomposition of $\langle x^2 - 2, y^2 - 2 \rangle \subset \mathbb{Q}[x, y]$ is

$$\langle x^2 - 2, y^2 - 2 \rangle = I_1 \cap I_2 = \langle x^2 - 2, x - y \rangle \cap \langle x^2 - 2, x + y \rangle.$$

Note that I_1 corresponds to $(\sqrt{2}, \sqrt{2})$ while I_2 corresponds to $(\sqrt{2}, -\sqrt{2})$. \square

Here is a description of the primary decomposition of an arbitrary zero-dimensional ideal.

Proposition 4.3.7 *Suppose that $\langle f_1, \dots, f_s \rangle$ is zero-dimensional and $f \in F[x_1, \dots, x_n]$ takes distinct values at the solutions of (4.3.1). If the irreducible factorization of $\text{CharPoly}_{m_f}(u)$ is*

$$\text{CharPoly}_{m_f}(u) = \prod_{i=1}^r p_i(u)^{m_i},$$

where $p_1(u), \dots, p_r(u)$ are distinct monic irreducible polynomials, then the primary decomposition of $\langle f_1, \dots, f_s \rangle$ is given by

$$\langle f_1, \dots, f_s \rangle = I_1 \cap \dots \cap I_r,$$

where

$$I_i = \langle f_1, \dots, f_s, p_i(f) \rangle.$$

Proof. We will use Galois theory to prove the proposition in the special case when F is perfect (see [Monico] for the general case). This means that either F has characteristic zero, or F has characteristic $p > 0$ and every element of F is a p th power. Every finite extension of a perfect field is separable.

If I_i is a primary component of $\langle f_1, \dots, f_s \rangle$, then its radical $\sqrt{I_i}$ is prime in $F[x_1, \dots, x_n]$. Then the following are true:

- The variety $\mathbf{V}(I_i) = \mathbf{V}(\sqrt{I_i}) \subset \overline{F}^n$ is irreducible over F .
- The Galois group $\text{Gal}(\overline{F}/F)$ acts on $\mathbf{V}(I_i)$.

These bullets imply that the action of $\text{Gal}(\overline{F}/F)$ acts on $\mathbf{V}(I_i)$ is transitive. Hence all $p \in \mathbf{V}(I_i)$ have the same multiplicity, denoted m_i . It is also easy to show that $\mathbf{V}(I_i) \cap \mathbf{V}(I_j) = \emptyset$ for $i \neq j$.

By Proposition 1.5.1, we see that

$$\text{CharPoly}_{m_f}(u) = \prod_{i=1}^r \prod_{p \in \mathbf{V}(I_i)} (u - f(p))^{m_i}.$$

Since f has coefficients in F , we see that $\sigma(f(p)) = f(q)$ whenever $\sigma \in \text{Gal}(\overline{F}/F)$ takes p to q . But we also know that the $f(p)$ are all distinct and F is perfect. Thus standard arguments from Galois theory imply that $p_i(u) = \prod_{p \in \mathbf{V}(I_i)} u - f(p)$ is irreducible over F . It follows that the above factorization coincides with the one in the statement of the proposition.

From here, the rest of the proof is similar to what we did in the proof of Proposition 4.3.2. The key point as always is that f takes distinct values at the solutions. We leave the details as an exercise for the reader. \square

The above proof shows that when F is perfect, the m_i 's compute the multiplicities of the corresponding points. However, this can fail when F is not perfect. We should also mention that one can weaken the hypothesis that f takes distinct values at the solutions: an analysis of the proof in [Monico] reveals that it is sufficient to assume that $f(p) \neq f(q)$ whenever p and q are solutions of (4.3.1) lying in different orbits of the $\text{Gal}(\overline{F}/F)$ -action. When this happens, however, the exponent m_i may fail to equal the multiplicity.

Here is an example of Proposition 4.3.7.

Example 4.3.8 For the ideal $\langle x^2 - 2, y^2 - 2 \rangle \subset \mathbb{Q}[x, y]$ of Example 4.3.6, one easily sees that $f = x + 2y$ takes distinct values at the solutions and has characteristic polynomial

$$\text{CharPoly}_{m_f}(u) = (u^2 - 18)(u^2 - 2),$$

where $u^2 - 18$ and $u^2 - 2$ are irreducible over \mathbb{Q} . By Proposition 4.3.7, we get the primary decomposition $\langle x^2 - 2, y^2 - 2 \rangle = I_1 \cap I_2$, where

$$\begin{aligned} I_1 &= \langle x^2 - 2, y^2 - 2, (x + 2y)^2 - 18 \rangle = \langle x^2 - 2, x - y \rangle \\ I_2 &= \langle x^2 - 2, y^2 - 2, (x + 2y)^2 - 2 \rangle = \langle x^2 - 2, x + y \rangle. \end{aligned}$$

This is the primary decomposition given in Example 4.3.6.

We could instead have used $f = x + y$, which has characteristic polynomial $u^2(u^2 - 8)$. The function f does not take distinct values on the roots but does separate orbits of the Galois action. As noted above, the conclusion of Proposition 4.3.7 still holds for such an f . The reader should check that $u^2(u^2 - 8)$ leads to the above primary decomposition. \square

We next relate primary decomposition to the factorizations discussed in Section 4.1.

Example 4.3.9 As in Section 4.1, suppose that $f(x), g(x) \in \mathbb{Q}[x]$ are irreducible and $\alpha, \beta \in \mathbb{C}$ satisfy $f(\alpha) = g(\beta) = 0$. Also suppose that we have the irreducible factorization

$$f(x) = f_1(x, \beta) \cdots f_r(x, \beta) \text{ over } \mathbb{Q}(\beta).$$

We can relate this to primary decomposition as follows. Pick $t \in \mathbb{Q}$ such that $f = x + ty$ takes distinct values at the solutions of $f(x) = g(y) = 0$. In the proof of Theorem 4.1.2, we showed that all solutions have multiplicity 1. Hence the proof of Proposition 4.3.7 shows that

$$\text{CharPoly}_{m_f}(u) = \prod_{i=1}^r \Phi_i(u),$$

where $\Phi_i(u) \in \mathbb{Q}[u]$ is irreducible. Then the primary decomposition of $\langle f(x), g(y) \rangle \subset \mathbb{Q}[x, y]$ is

$$\langle f(x), g(y) \rangle = \bigcap_{i=1}^r \langle f(x), g(y), \Phi_i(x + ty) \rangle.$$

However, Theorem 4.1.2 asserts that

$$f_i(x, \beta) = \text{GCD}(\Phi_i(x + t\beta), f(x)).$$

Since $\mathbb{Q}(\beta) \simeq \mathbb{Q}[y]/\langle g(y) \rangle$, it is an easy exercise to show that

$$\langle f(x), g(y), \Phi_i(x + ty) \rangle = \langle g(y), f_i(x, y) \rangle.$$

Hence the above primary decomposition can be written

$$\langle f(x), g(y) \rangle = \bigcap_{i=1}^r \langle g(y), f_i(x, y) \rangle.$$

This shows the close relation between primary decomposition and factorization. □

There is also a version of Proposition 4.3.7 which uses minimal polynomials instead of characteristic polynomials.

Proposition 4.3.10 *Suppose that $\langle f_1, \dots, f_s \rangle$ is zero-dimensional and $f \in F[x_1, \dots, x_n]$ takes distinct values at the solutions of (4.3.1). If the irreducible factorization of $\text{MinPoly}_{m_f}(u)$ is*

$$\text{MinPoly}_{m_f}(u) = \prod_{i=1}^r p_i(u)^{n_i},$$

where $p_1(u), \dots, p_r(u)$ are distinct monic irreducible polynomials, then the primary decomposition of $\langle f_1, \dots, f_s \rangle$ is given by

$$\langle f_1, \dots, f_s \rangle = I_1 \cap \dots \cap I_r,$$

where

$$I_i = \langle f_1, \dots, f_s, p_i(f)^{n_i} \rangle.$$

Proof. See [ABRW] or [YNT]. □

Algorithmic Aspects. From the point of view of algorithmic primary decomposition, one weakness of Proposition 4.3.7 is that f needs to take distinct values at the solutions. How do we do this without knowing the solutions? One way would be to make a random choice of $f = a_1x_1 + \dots + a_nx_n$. But this gives only a probabilistic algorithm. Another weakness of this method is that computing the characteristic polynomial of a large matrix can be time-consuming. The timings reported in [Monico] indicate that as the number of solutions increases, methods based on [GTZ] outperform the algorithm using Proposition 4.3.7.

Other approaches to primary decomposition are given in [EHV] and [MMM2].

5 Galois Theory

5.1 Splitting Algebras

Solving equations has been one of our main topics of discussion. Since Galois theory is also concerned with the solutions of equations, it makes sense that there should be some link. As we will see, turning a polynomial equation $f(x) = 0$ of degree n into n equations in n unknowns is a very useful thing to do.

Let F be an infinite field. We will assume that $f(x) \in F[x]$ is a monic polynomial of degree n with distinct roots. We will write $f(x)$ as

$$f(x) = x^n - c_1x^{n-1} + \cdots + (-1)^n c_n, \quad c_i \in F.$$

The elementary symmetric polynomials $\sigma_1, \dots, \sigma_n \in F[x_1, \dots, x_n]$ are defined by the identity

$$(x - x_1) \cdots (x - x_n) = x^n - \sigma_1 x^{n-1} + \cdots + (-1)^n \sigma_n. \quad (5.1.1)$$

Consider the system of n equations in x_1, \dots, x_n given by

$$\sigma_1(x_1, \dots, x_n) - c_1 = \sigma_2(x_1, \dots, x_n) - c_2 = \cdots = \sigma_n(x_1, \dots, x_n) - c_n = 0. \quad (5.1.2)$$

The associated algebra is

$$A = F[x_1, \dots, x_n] / \langle \sigma_1 - c_1, \dots, \sigma_n - c_n \rangle.$$

This is called the *splitting algebra* of f over F . The system (5.1.2) and the algebra A were first written down by Kronecker in 1882 and 1887 respectively (see page 282 of [Kronecker, Vol. II] for the equations and page 213 of [Kronecker, Vol. III] for the algebra). A very nice modern treatment of the splitting algebra appears in the recent preprint [EL].

The Universal Property. We first explain why the splitting algebra deserves its name. The natural map $F[x_1, \dots, x_n] \rightarrow A$ takes σ_i to c_i , so that by (5.1.1), the cosets $[x_i] \in A$ become roots of $f(x)$. It follows that $f(x)$ splits completely in $A[x]$. But more is true, for the factorization of $f(x)$ in $A[x]$ controls *all possible ways* in which $f(x)$ splits. Here is the precise statement.

Proposition 5.1.1 *Suppose that R is an F -algebra such that $f(x)$ splits completely in $R[x]$ via*

$$f(x) = (x - \alpha_1) \cdots (x - \alpha_n), \quad \alpha_1, \dots, \alpha_n \in R.$$

Then there is an F -algebra homomorphism $\varphi : A \rightarrow R$ such that this splitting is the image under φ of the splitting of $f(x)$ in $A[x]$.

Proof. Consider the F -algebra homomorphism $\Phi : F[x_1, \dots, x_n] \rightarrow R$ determined by $x_i \mapsto \alpha_i$. This maps (5.1.1) to the splitting in the statement of the proposition, so that Φ maps σ_i to c_i . Hence $\Phi(\sigma_i - c_i) = 0$ for all i , which implies that Φ induces an F -algebra homomorphism $\varphi : A \rightarrow R$. It follows easily that φ has the desired property. \square

The splitting of $f(x)$ in $A[x]$ is thus the “universal splitting” in the sense that any other splitting is a homomorphic image of this one.

The Structure of A . Our next task is to understand the structure of the algebra A and explain how it relates to the splitting field of f over F . Let \overline{F} be the algebraic closure of F and fix a splitting

$$f(x) = (x - \alpha_1) \cdots (x - \alpha_n) \in \overline{F}[x].$$

Using this, we can describe the solutions of (5.1.2) as follows. If $(\beta_1, \dots, \beta_n) \in \overline{F}^n$ is a solution, then the substitutions $x_i \mapsto \beta_i$ take (5.1.1) to

$$f(x) = (x - \beta_1) \cdots (x - \beta_n) \in \overline{F}[x].$$

Thus the β_i 's are some permutation of the α_i . Since $f(x)$ has distinct roots by hypothesis, there is a unique $\sigma \in S_n$ such that $\beta_i = \alpha_{\sigma(i)}$ for all i . It follows easily that (5.1.2) has precisely $n!$ solutions given by

$$(\alpha_{\sigma(1)}, \dots, \alpha_{\sigma(n)}), \quad \sigma \in S_n.$$

We can determine the multiplicities of these solutions as follows. Since $\sigma_i - c_i$ has degree i as a polynomial in x_1, \dots, x_n , Bézout's theorem tells us that (5.1.2) has at most $1 \cdot 2 \cdot 3 \cdots n = n!$ solutions, counting multiplicity. Since we have $n!$ solutions, the multiplicities must all be 1.

Since F is infinite, we can find $t_1, \dots, t_n \in F$ such that $f_0 = t_1 x_1 + \cdots + t_n x_n$ takes distinct values at the solutions of (5.1.2). Thus, as σ varies over the elements of S_n ,

$$f_0(\alpha_{\sigma(1)}, \dots, \alpha_{\sigma(n)}) = t_1 \alpha_{\sigma(1)} + \cdots + t_n \alpha_{\sigma(n)} \quad (5.1.3)$$

gives $n!$ distinct elements of \overline{F} . Since all solutions of (5.1.2) have multiplicity 1, we conclude that:

- The characteristic polynomial of m_{f_0} on A is

$$\text{CharPoly}_{m_{f_0}}(u) = \prod_{\sigma \in S_n} (u - (t_1 \alpha_{\sigma(1)} + \cdots + t_n \alpha_{\sigma(n)})).$$

Our choice of t_1, \dots, t_n implies that this polynomial has distinct roots.

- The linear map m_{f_0} is non-derogatory, so that by Proposition 1.4.1, the map sending u to $[t_1 x_1 + \cdots + t_n x_n] \in A$ induces an F -algebra isomorphism

$$F[u]/\langle \text{CharPoly}_{m_{f_0}}(u) \rangle \simeq A.$$

- A has dimension $n!$ over F .

Now factor $\text{CharPoly}_{m_{f_0}}(u)$ into a product of monic irreducible polynomials in $F[u]$, say

$$\text{CharPoly}_{m_{f_0}}(u) = \prod_{i=1}^r G_i(u).$$

Since $\text{CharPoly}_{m_{f_0}}(u)$ has distinct roots, the $G_i(u)$ are distinct. When combined with the third bullet, we get F -algebra isomorphisms

$$A \simeq F[u]/\langle \text{CharPoly}_{m_{f_0}}(u) \rangle \simeq \prod_{i=1}^r F[u]/\langle G_i(u) \rangle = \prod_{i=1}^r K_i. \quad (5.1.4)$$

Each K_i is a field, and since the projection map $A \rightarrow K_i$ is surjective, we see that K_i is a splitting field of f over F . Thus the factorization of the characteristic polynomial of m_{f_0} shows that A is isomorphic to a product of fields, each of which is a copy of the splitting field of f over F .

5.2 The Galois Group

We now use the above description of A to compute the Galois group of f and prove some of its basic properties. We will also describe an algorithm for computing the Galois group.

The Galois Group and the Symmetric Group. An important observation is that the action of S_n on $F[x_1, \dots, x_n]$ given by permuting variables descends to an action of S_n on A . Since the decomposition $A \simeq \prod_{i=1}^r K_i$ is unique up to isomorphism, it follows that for $1 \leq i \leq r$ and $\sigma \in S_n$, we have $\sigma(K_i) = K_j$ for some j . Then we get the following result.

Theorem 5.2.1 *There is a natural isomorphism*

$$\text{Gal}(K_i/F) \simeq \{\sigma \in S_n \mid \sigma(K_i) = K_i\}.$$

Furthermore,

$$|\text{Gal}(K_i/F)| = [K_i : F]$$

and, when the characteristic of F doesn't divide the order of the Galois group,

$$F = \{\alpha \in K_i \mid \gamma(\alpha) = \alpha \text{ for all } \gamma \in \text{Gal}(K_i/F)\}.$$

Proof. Let $G_i = \{\sigma \in S_n \mid \sigma(K_i) = K_i\}$. Since every σ induces an automorphism of K_i , we get an injective group homomorphism $G_i \rightarrow \text{Gal}(K_i/F)$. To show that this map is surjective, take $\gamma \in \text{Gal}(K_i/F)$. Under the projection $A \rightarrow K_i$, the cosets $[x_i]$ map to roots of $f(x)$ lying in K_i . Then γ must permute these according to some $\sigma \in S_n$. Since the roots generate K_i over F and σ permutes the roots, we have $\sigma(K_i) = K_i$. It follows that $\sigma \in G_i$ maps to γ . This gives the desired isomorphism.

For the second assertion, we first show that S_n permutes the K_i transitively. Under the isomorphism

$$F[u]/\langle \text{CharPoly}_{m_{f_0}}(u) \rangle \simeq A, \tag{5.2.1}$$

S_n permutes the factors of $\text{CharPoly}_{m_{f_0}}(u) = \prod_{i=1}^r G_i(u)$. Over \overline{F} , the factorization

$$\text{CharPoly}_{m_{f_0}}(u) = \prod_{\sigma \in S_n} (u - (t_1 \alpha_{\sigma(1)} + \dots + t_n \alpha_{\sigma(n)})) \tag{5.2.2}$$

shows that S_n must permute the $G_i(u)$ transitively. By (5.1.4), we conclude that S_n permutes the K_i transitively. Since G_i is the isotropy subgroup of K_i under this action, we see that

$$|\text{Gal}(K_i/F)| = |G_i| = \frac{n!}{r}.$$

However, the transitivity also shows that the K_1, \dots, K_r are mutually isomorphic. Thus

$$n! = \dim_F(A) = [K_1 : F] + \dots + [K_r : F] = r [K_i : F].$$

Combining this with the previous equation gives the desired equality $|\text{Gal}(K_i/F)| = [K_i : F]$.

Finally, suppose that $\alpha \in K_i$ is in the fixed field of $\text{Gal}(K_i/F)$. We may assume that $\alpha \neq 0$. Let $p \in F[x_1, \dots, x_n]$ map to $\alpha_i \in K_i$ and to $0 \in K_j$ for $j \neq i$. Then $P = \sum_{\sigma \in S_n} \sigma \cdot p$ is symmetric and hence is a polynomial in the σ_i . In A , this means that $[P] \in F$, which means that P projects to an element of F in each of K_1, \dots, K_r .

Observe that $\sigma(\alpha) \in K_i$ implies $\sigma(\alpha) \in K_i \cap \sigma(K_i)$. Thus the intersection is nonzero, which by (5.1.4) implies that $\sigma \in G_i$. But then $\sigma(\alpha) = \alpha$ since α is in the fixed field. It follows that the projection of P onto K_i is $|G_i|\alpha$. Thus $|G_i|\alpha \in F$, and then $\alpha \in F$ follows by hypothesis. \square

While the above theorem is not as general as possible, it shows that basic properties of the Galois group follow from the splitting algebra. See [EL] for a more general version of Theorem 5.2.1.

Computing the Galois Group. Our final topic is to explain an algorithm for computing $G_i \simeq \text{Gal}(K_i/F)$. We will use the characteristic polynomial $\text{CharPoly}_{m_{f_0}}(u)$, where as usual $f_0 = t_1x_1 + \cdots + t_nx_n$ is chosen so that the $n!$ numbers are distinct. Of course, since we don't know the roots of f , in practice we simply make a random choice of $t_1, \dots, t_n \in F$.

One way to compute $\text{CharPoly}_{m_{f_0}}(u)$ would be to use the product formula (5.2.2) together with the classical theory of symmetric polynomials. But we can also use Gröbner bases as follows. Given variables u_1, \dots, u_s , set $h_0(u_1, \dots, u_s) = 1$ and, for $k > 0$, let $h_k(u_1, \dots, u_s)$ be the sum of all monomials of degree k in these variables. Thus

$$h_k(u_1, \dots, u_s) = \sum_{\deg(u^\alpha) = k} u^\alpha$$

when $k > 0$. The proof of Proposition 5 of Chapter 7, Section 1 of [CLO1] implies the following.

Proposition 5.2.2 *For the lexicographic order defined in Section 1.1, the polynomials*

$$g_k = h_k(x_k, \dots, x_n) + \sum_{i=1}^k (-1)^i c_i h_{k-i}(x_k, \dots, x_n), \quad k = 1, \dots, n$$

form a Gröbner basis for the ideal $\langle \sigma_1 - c_1, \dots, \sigma_n - c_n \rangle \subset F[x_1, \dots, x_n]$.

For this Gröbner basis, we leave it as an exercise for the reader to show that the leading term of g_k is x_k^k and the set of remainder monomials is

$$\mathcal{B} = \{x_1^{a_1} \cdots x_n^{a_n} \mid 0 \leq a_k \leq k - 1 \text{ for } k = 1, \dots, n\}. \quad (5.2.3)$$

It is easy to see that $|\mathcal{B}| = n!$, which gives another proof that $\dim_F(A) = n!$.

Using this Gröbner basis, one can compute the matrix of m_{f_0} with respect to \mathcal{B} and hence its characteristic polynomial $\text{CharPoly}_{m_{f_0}}(u)$. Then computing the discriminant of $\text{CharPoly}_{m_{f_0}}(u)$ reveals whether or not we made a good choice of t_1, \dots, t_n . If not, try again. Furthermore, by Exercise 6 of Chapter 1, Section 1 of [CLO1], one can describe a procedure whereby a good choice of t_1, \dots, t_n can be found after finitely many attempts.

Once we have $\text{CharPoly}_{m_{f_0}}(u)$ with distinct roots, the next step is to factor it over F . Let $G(u)$ be an irreducible factor. Then

$$K = F[u]/\langle G(u) \rangle$$

is a splitting field of f over F and the degree of $G(u)$ is the order of the Galois group.

Example 5.2.3 Consider $f(x) = x^3 + x^2 - 2x - 1 \in \mathbb{Q}[x]$. Let's use the above procedure to construct the splitting field of $f(x)$. We will write the equations (5.1.2) as

$$x_1 + x_2 + x_3 + 1 = x_1x_2 + x_1x_3 + x_2x_3 + 2 = x_1x_2x_3 + 1 = 0.$$

For lexicographic order with $x_1 > x_2 > x_3$, Proposition 5.2.2 implies that the associated Gröbner basis consists of the polynomials

$$x_1 + x_2 + x_3 + 1, \quad x_2^2 + x_2x_3 + x_3^2 + x_2 + x_3 - 2, \quad x_3^3 + x_3^2 - 2x_3 - 1, \quad (5.2.4)$$

with remainder monomials $1, x_2, x_3, x_2x_3, x_3^2, x_2x_3^2$. The “random” choice $t_1 = 0, t_2 = 1, t_3 = 2$ gives $f_0 = x_2 + 2x_3$. Using the above Gröbner basis, one computes that the matrix of m_{f_0} is

$$M_{f_0} = \begin{pmatrix} 0 & 2 & 0 & -1 & 2 & 0 \\ 1 & -1 & 0 & 0 & 0 & 1 \\ 2 & -1 & 0 & 0 & 4 & -1 \\ 0 & 1 & 1 & -1 & 0 & 2 \\ 0 & -1 & 2 & 0 & -2 & 0 \\ 0 & 0 & 0 & 1 & 1 & -2 \end{pmatrix}.$$

Then the corresponding characteristic polynomial and its irreducible factors over \mathbb{Q} are

$$\begin{aligned} \text{CharPoly}_{m_{f_0}}(u) &= u^6 + 6u^5 + u^4 - 36u^3 - 20u^2 + 48u - 13 \\ &= (u^3 + 3u^2 - 4u - 13)(u^3 + 3u^2 - 4u + 1). \end{aligned}$$

This polynomial has discriminant $18078415936 \neq 0$, so that it has distinct roots (this also follows from the above factorization—do you see why?). By (5.1.4), it follows that

$$K_1 = \mathbb{Q}[u]/\langle u^3 + 3u^2 - 4u - 13 \rangle \quad \text{and} \quad K_2 = \mathbb{Q}[u]/\langle u^3 + 3u^2 - 4u + 1 \rangle \quad (5.2.5)$$

are splitting fields of $f(x)$ over \mathbb{Q} . Hence $\text{Gal}(K_1/\mathbb{Q}) \simeq \text{Gal}(K_2/\mathbb{Q})$ has order 3. \square

In order for splitting fields such as those constructed in Example 5.2.3 to be useful, we need to find an explicit representation of the roots. This is done as follows. Let

$$\alpha = [t_1x_1 + \cdots + t_nx_n] \in A. \quad (5.2.6)$$

The isomorphism (5.2.1) shows that $1, \alpha, \dots, \alpha^{n-1}$ is a basis of A over F . But we also have the basis given by \mathcal{B} of (5.2.3). The division algorithm using the Gröbner basis of Proposition 5.2.2 enables us to express each α^i in terms of \mathcal{B} . Thus we can compute the transition matrix from one basis to the other, and using the inverse of this matrix, we can express the cosets coming from \mathcal{B} as linear combinations of powers of α . Since $x_2, \dots, x_n \in \mathcal{B}$ and $[x_1] = [c_1 - x_2 - \cdots - x_n]$, it follows that we can compute polynomials $P_j(u) \in \mathbb{Q}[u]$ such that

$$[x_j] = P_j(\alpha), \quad j = 1, \dots, n. \quad (5.2.7)$$

Using the isomorphism (5.2.1) followed by the projection onto $K = F[u]/\langle G(u) \rangle$, it follows that

$$[P_j(u)] \in K = F[u]/\langle G(u) \rangle, \quad j = 1, \dots, n$$

represent the roots of $f(x)$ in the splitting field. Hence we have complete computational control of the roots. Here is an example.

Example 5.2.4 Consider $f(x) = x^3 + x^2 - 2x - 1 \in \mathbb{Q}[x]$ from Example 5.2.3. Since $t_1 = 0, t_2 = 1, t_3 = 2$, we have $\alpha = [x_2 + 2x_3] \in A$. Using the Gröbner basis (5.2.4), one easily computes the

matrix

$$\begin{pmatrix} 1 & 0 & 2 & 1 & -3 & -61 \\ 0 & 1 & -1 & 3 & 4 & -31 \\ 0 & 2 & -1 & 17 & -35 & 194 \\ 0 & 0 & 3 & -5 & 37 & -116 \\ 0 & 0 & 3 & -7 & 45 & -164 \\ 0 & 0 & 0 & 6 & -24 & 130 \end{pmatrix}$$

whose columns express the remainders of $1, \alpha, \alpha^2, \alpha^3, \alpha^4, \alpha^5$ on division by the Gröbner basis as linear combinations of the remainder monomials $1, x_2, x_3, x_2x_3, x_3^2, x_2x_3^2$. Using the inverse of this matrix, one computes that

$$\begin{aligned} [x_1] &= \frac{1}{14}(-45 + 51\alpha + 17\alpha^2 - 12\alpha^3 - 3\alpha^4) \\ [x_2] &= \frac{1}{14}(62 - 116\alpha - 34\alpha^2 + 24\alpha^3 + 6\alpha^4) \\ [x_3] &= \frac{1}{14}(-31 + 65\alpha + 17\alpha^2 - 12\alpha^3 - 3\alpha^4) \end{aligned}$$

in the algebra A . Mapping this to the field $K_1 = \mathbb{Q}[u]/\langle u^3 + 3u^2 - 4u - 13 \rangle$ of (5.2.5) and reducing modulo $u^3 + 3u^2 - 4u - 13$, we see that $f(x) = x^3 + x^2 - 2x - 1$ has roots

$$\begin{aligned} a_1 &= [-6 + u^2] \\ a_2 &= [10 - u - 2u^2] \\ a_3 &= [-5 + u + u^2] \end{aligned} \tag{5.2.8}$$

in $K_1 = \mathbb{Q}[u]/\langle u^3 + 3u^2 - 4u - 13 \rangle$. Similarly, one computes that the roots are

$$\begin{aligned} b_1 &= [-3 + 3u + u^2] \\ b_2 &= [4 - 7u - 2u^2] \\ b_3 &= [-2 + 4u + u^2] \end{aligned} \tag{5.2.9}$$

in $K_2 = \mathbb{Q}[u]/\langle u^3 + 3u^2 - 4u + 1 \rangle$. □

Our final task is to compute the Galois group. Consider the set of permutations defined by

$$G = \{\sigma \in S_n \mid G(u) \text{ divides } G(t_1P_{\sigma(1)}(u) + \cdots + t_nP_{\sigma(n)}(u)) \text{ in } \mathbb{Q}[u]\}. \tag{5.2.10}$$

where P_1, \dots, P_n are as in (5.2.7). There is an obvious (though not very efficient) algorithm for computing G . Then G determines $\text{Gal}(K/F)$ as follows.

Proposition 5.2.5 *The set G defined in (5.2.10) is a subgroup of S_n and*

$$G \simeq \text{Gal}(K/F).$$

Proof. The field $K = F[u]/\langle G(u) \rangle$ appears in the decomposition (5.1.4). Under the projection map

$$A \simeq F[u]/\langle \text{CharPoly}_{m_f}(u) \rangle \longrightarrow K = F[u]/\langle G(u) \rangle, \tag{5.2.11}$$

suppose that $[x_i]$ maps to $a_i \in K$. Then $\alpha = [t_1x_1 + \cdots + t_nx_n] \in A$ maps to $a = t_1\alpha_1 + \cdots + t_n\alpha_n \in K$. This is a primitive element of K over F and $G(u)$ is its minimal polynomial over F . It follows easily that

$$\sigma \in S_n \text{ comes from } \text{Gal}(K/F) \iff G(t_1\alpha_{\sigma(1)} + \cdots + t_n\alpha_{\sigma(n)}) = 0 \text{ in } K.$$

But now recall from (5.2.7) that $[x_i] = P_i(\alpha)$ in A . Applying (5.2.11), we see that $a_i = P_i(\alpha)$ in K . If we substitute this into the above equivalence, then we obtain

$$\sigma \in S_n \text{ comes from } \text{Gal}(K/F) \iff G(t_1 P_{\sigma(1)}(a) + \cdots + t_n P_{\sigma(n)}(a)) = 0 \text{ in } K.$$

Since $a \in K$ is really $[u] \in F[u]/\langle G(u) \rangle$, we see that

$$G(t_1 P_{\sigma(1)}(a) + \cdots + t_n P_{\sigma(n)}(a)) = 0 \iff G(u) \text{ divides } G(t_1 P_{\sigma(1)}(u) + \cdots + t_n P_{\sigma(n)}(u)).$$

The proposition now follows immediately. \square

Here is an example of how Proposition (5.2.5) can be used to compute a Galois group.

Example 5.2.6 Continuing Example 5.2.4, we have the splitting field $K_1 = \mathbb{Q}[u]/\langle G(u) \rangle$, $G(u) = u^3 + 3u^2 - 4u - 13$, with roots given by (5.2.8). This means that

$$P_1(u) = -6 + u^2, \quad P_2(u) = 10 - u - 2u^2, \quad P_3(u) = -5 + u + u^2.$$

Also recall that $t_1 = 0, t_2 = 1, t_3 = 2$, so that

$$G(t_1 P_{\sigma(1)}(u) + t_2 P_{\sigma(2)}(u) + t_3 P_{\sigma(3)}(u)) = G(P_{\sigma(2)}(u) + 2P_{\sigma(3)}(u)).$$

Then one easily computes the following table:

$\sigma \in S_3$	$G(P_{\sigma(2)}(u) + 2P_{\sigma(3)}(u)) \bmod G(u)$
e	0
(123)	0
(132)	0
(12)	-14
(23)	-14
(13)	-14

It follows that $\text{Gal}(K/F) \simeq \{e, (123), (132)\} = A_3 \subset S_3$. \square

Of course, we already knew that the group $\text{Gal}(K/F)$ of this example was cyclic of order 3 since $|\text{Gal}(K/F)| = 3$ by Example 5.2.3. However, it is still nice to see how the calculation works in this case.

In practice, this algorithm described above is hopelessly inefficient except for polynomials of very small degree, and for even these, much better methods are available (see, for example, Section 6.3 of [Cohen]).

History. The methods described here date back to Galois and Kronecker. For example, in 1830 Galois chose t_1, \dots, t_n such that the $n!$ values (5.1.3) are distinct and showed that

$$V = t_1 \alpha_1 + \cdots + t_n \alpha_n$$

is a primitive element of the splitting field. The minimal polynomial $G(u)$ of V over F is the same polynomial that we used in Proposition 5.2.5, and Galois uses both $G(u)$ and the representation of the roots given by

$$\alpha_i = P_i(V)$$

(which is what (5.2.7) looks like in this situation). In all of this, Galois simply assumed the existence of the roots.

In 1887, Kronecker gave the first rigorous construction of the splitting field. His method was to prove the existence of t_1, \dots, t_n as above and then factor $\text{CharPoly}_{m_{f_0}}(u)$ into irreducibles. As we have seen, Kronecker knew how to do this constructively. Then, if $G(u)$ is one of the factors, then the splitting field is exactly the field $K = F[u]/\langle G(u) \rangle$ used above.

When I first read Kronecker, I remember thinking “what about the other irreducible factors of the characteristic polynomial?” The Galois theory given here answers this question nicely, for we see that the other factors give other models of the splitting field which when taken together give the splitting algebra

$$A \simeq K_1 \times \cdots \times K_r.$$

In Section 5.3, we will use primary decomposition to help understand why multiple copies of the splitting field are necessary.

Resultants. We next observe that the characteristic polynomial $\text{CharPoly}_{m_{f_0}}(u)$ can be interpreted as a resultant. Namely, we claim that

$$\text{Res}_{1,1,2,\dots,n}(u - f_0, \sigma_1 - c_1, \dots, \sigma_n - c_n) = -\text{CharPoly}_{m_{f_0}}(u). \quad (5.2.12)$$

To prove this, recall from Theorem 3.2.2 of Section 3 that this resultant equals the characteristic polynomial multiplied by

$$\text{Res}_{1,2,\dots,n}((\sigma_1 - c_1)_1, \dots, (\sigma_n - c_n)_n),$$

where $(\sigma_i - c_i)_i$ consists of the terms of $\sigma_i - c_i$ of degree i . This is obviously just σ_i , so that this multiplier reduces to

$$\text{Res}_{1,2,\dots,n}(\sigma_1, \dots, \sigma_n).$$

This resultant equals -1 by Exercise 11 of Section 3 of Chapter 3 of [CLO2]. Hence we obtain (5.2.12) as claimed.

Action of the Symmetric Group. Finally, we will describe the action of S_n on the product decomposition $A = K_1 \times \cdots \times K_r$. If we let $\alpha_{ij} \in K_j$ be the projection of $[x_i] \in A$ onto the j th factor, then $\alpha_{1j}, \dots, \alpha_{nj}$ are the roots of $f(x)$ in K_j . So we have r isomorphic copies of the splitting field together with an ordered list of the roots in each field. Let

$$\mathbf{e}_j = (0, \dots, 0, 1, 0, \dots, 0) \in K_1 \times \cdots \times K_r,$$

where the 1 is in the j th position. Then by abuse of notation we can write $\alpha_{ij} \mathbf{e}_j \in A$. Now take $\sigma \in S_n$ and suppose that $\sigma(\mathbf{e}_j) = \mathbf{e}_\ell$ (this is a precise way of saying that $\sigma(K_j) = K_\ell$). Then one can show without difficulty that $\sigma([x_i]) = [x_{\sigma(i)}]$ implies that

$$\sigma(\alpha_{ij} \mathbf{e}_j) = \alpha_{\sigma(i)\ell} \mathbf{e}_\ell.$$

In the special case when σ comes from an element of $\text{Gal}(K_j/F)$, we have $\sigma(\mathbf{e}_j) = \mathbf{e}_j$. Then the above formula gives the action of the Galois group on the roots. But now we know what happens when we apply permutations which don't come from the Galois group!

5.3 Primary Decomposition

For our final topic, we will use primary decomposition to describe how algebraic relations among the roots influence the Galois group and the splitting field. We will work over \mathbb{Q} for simplicity.

Given $f = x^n - c_1x^{n-1} + \cdots + (-1)^nc_n \in \mathbb{Q}[x]$ as in Section 5.1, the splitting algebra is

$$A = \mathbb{Q}[x_1, \dots, x_n] / \langle \sigma_1 - c_1, \dots, \sigma_n - c_n \rangle,$$

where as usual σ_i is the i th elementary symmetric polynomial. We've seen that A is a product

$$A = \prod_{i=1}^r K_i,$$

where each K_i is a splitting field of f over \mathbb{Q} . But we also have the primary decomposition

$$\langle \sigma_1 - c_1, \dots, \sigma_n - c_n \rangle = \bigcap_{i=1}^r I_i,$$

where each I_i is a maximal ideal in $\mathbb{Q}[x_1, \dots, x_n]$. These decompositions are related via

$$K_i = \mathbb{Q}[x_1, \dots, x_n] / I_i, \quad i = 1, \dots, r.$$

Each I_i is larger than $\langle \sigma_1 - c_1, \dots, \sigma_n - c_n \rangle$. The idea is that the polynomials we add to get from I_i to $\langle \sigma_1 - c_1, \dots, \sigma_n - c_n \rangle$ reflect the algebraic relations between the roots which hold in the splitting field K_i . Having more relations among the roots means that I_i is larger and hence K_i and the Galois group are smaller.

For instance, if the Galois group of f is all of S_n , then $\langle \sigma_1 - c_1, \dots, \sigma_n - c_n \rangle$ is a maximal ideal and the splitting algebra is the splitting field. This means that *all* relations among the roots are consequences of the fact that the coefficients of f are (up to sign) the elementary symmetric polynomials of the roots.

Let's see what happens when the Galois group is smaller than S_n .

Example 5.3.1 Let $f = x^3 - c_1x^2 + c_2x - c_3 \in \mathbb{Q}[x]$ be an irreducible cubic. The splitting algebra of f is $A = \mathbb{Q}[x_1, x_2, x_3] / \langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3 \rangle$. It is well-known that

$$\text{the Galois group of } f \text{ is isomorphic to } \begin{cases} S_3 & \text{if } \Delta(f) \notin \mathbb{Q}^2 \\ \mathbb{Z}/3\mathbb{Z} & \text{if } \Delta(f) \in \mathbb{Q}^2, \end{cases}$$

where $\Delta(f) \in \mathbb{Q}$ is the discriminant of f . By the above analysis, it follows that A is the splitting field of f when $\Delta(f) \notin \mathbb{Q}^2$.

Now suppose that $\Delta(f) = a^2$ for some $a \in \mathbb{Q}$. In this case, the splitting algebra is a product of two copies of the splitting field, i.e., $A = K_1 \times K_2$. Let

$$\sqrt{\Delta} = (x_1 - x_2)(x_1 - x_3)(x_2 - x_3) \in \mathbb{Q}[x_1, x_2, x_3].$$

In the splitting algebra A , we have $[\sqrt{\Delta}]^2 = [\Delta(f)]$, so that

$$[\sqrt{\Delta}]^2 = [a]^2.$$

Since A is not an integral domain, this does not imply $[\sqrt{\Delta}] = \pm[a]$. In fact, $[\sqrt{\Delta}] \in A$ cannot have a numerical value since $[\sqrt{\Delta}]$ is not invariant under S_3 . Yet once we map to a field, the value

must be $\pm a$. But which sign do we choose? The answer is *both*, which explains why we need two fields in the splitting algebra.

We leave it as an exercise for the reader to show that we have the primary decomposition

$$\langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3 \rangle = I_1 \cap I_2,$$

where

$$\begin{aligned} I_1 &= \langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3, \sqrt{\Delta} - a \rangle \\ I_2 &= \langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3, \sqrt{\Delta} + a \rangle. \end{aligned}$$

It is also easy to see that this is compatible with the action of S_3 . For example, $(12) \in S_3$ maps I_1 to I_2 since $(12) \cdot \sqrt{\Delta} = -\sqrt{\Delta}$. It follows that (12) maps K_1 to K_2 in the decomposition $A = K_1 \times K_2$. This is consistent with the description of the S_n action given at the end of Section 5.2. \square

Example 5.3.1 is analogous to what happens in quantum mechanics when an observation forces a mixed state (such as a superposition of pure states with different energy levels) to become a pure state (with a fixed energy level). In Example 5.3.1, the idea is that $[\sqrt{\Delta}]^2 = [D(f)]^2 = [a]^2$ means that $[\sqrt{\Delta}]$ is somehow a “mixed state” which becomes a “pure state” (i.e., $\pm a \in \mathbb{Q}$) when “observed” (i.e., when mapped to a field).

The quartic is slightly more complicated since there are five possibilities for the Galois group of an irreducible quartic. Hence we will only discuss the following special case.

Example 5.3.2 Let $f = x^4 - c_1x^3 + c_2x^2 - c_3x + c_4 \in \mathbb{Q}[x]$ be an irreducible quartic with splitting algebra $A = \mathbb{Q}[x_1, x_2, x_3, x_4]/\langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3, \sigma_4 - c_4 \rangle$. One of the tools used in solving the quartic is the *Ferrari resolvent*

$$x^3 - c_2x^2 + (c_1c_3 - 4c_4)x - c_3^2 - c_1^2c_4 + 4c_2c_4. \quad (5.3.1)$$

Euler showed that if $\beta_1, \beta_2, \beta_3$ are the roots of (5.3.1), then the roots of f are

$$\frac{1}{4} \left(c_1 \pm \sqrt{\beta_1 + c_1^2 - 4c_2} \pm \sqrt{\beta_2 + c_1^2 - 4c_2} \pm \sqrt{\beta_3 + c_1^2 - 4c_2} \right),$$

provided the signs are chosen so that the product of the square roots is $c_1^3 - 4c_1c_2 + 8c_3$. Also, as shown by Lagrange, the roots of the resolvent (5.3.1) are

$$\alpha_1\alpha_2 + \alpha_3\alpha_4, \quad \alpha_1\alpha_3 + \alpha_2\alpha_4, \quad \alpha_1\alpha_4 + \alpha_2\alpha_3. \quad (5.3.2)$$

As is well-known, the Galois group G of f over \mathbb{Q} is isomorphic to one of the following 5 groups:

$$S_4, \quad A_4, \quad D_8, \quad \mathbb{Z}/4\mathbb{Z}, \quad \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z},$$

where D_8 is the dihedral group of order 8. Three of these cases are easy to distinguish:

$$G \simeq \begin{cases} S_4 & \text{if } \Delta(f) \notin \mathbb{Q}^2 \text{ and (5.3.1) is irreducible over } \mathbb{Q} \\ A_4 & \text{if } \Delta(f) \in \mathbb{Q}^2 \text{ and (5.3.1) is irreducible over } \mathbb{Q} \\ \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} & \text{if } \Delta(f) \in \mathbb{Q}^2 \text{ and (5.3.1) is reducible over } \mathbb{Q}. \end{cases}$$

The remaining case is when $\Delta(f) \notin \mathbb{Q}^2$ and (5.3.1) has a root in \mathbb{Q} . Here, the Galois group is D_8 or $\mathbb{Z}/4\mathbb{Z}$. We state without proof the following nice fact:

$$G \simeq D_8 \iff \Delta(f) \notin \mathbb{Q}^2, \text{ (5.3.1) has a root } b \in \mathbb{Q}, \text{ and we have the primary decomposition}$$

$$\langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3, \sigma_4 - c_4 \rangle = I_1 \cap I_2 \cap I_3,$$

where

$$I_1 = \langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3, \sigma_4 - c_4, x_1x_2 + x_3x_4 - b \rangle$$

$$I_2 = \langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3, \sigma_4 - c_4, x_1x_3 + x_2x_4 - b \rangle$$

$$I_3 = \langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3, \sigma_4 - c_4, x_1x_4 + x_2x_3 - b \rangle.$$

The reason for three ideals is that b is one of the three combinations of roots given in (5.3.2). To get a field out of the ideal $\langle \sigma_1 - c_1, \sigma_2 - c_2, \sigma_3 - c_3, \sigma_4 - c_4 \rangle$, we must commit to which combination gives b . This gives I_1, I_2, I_3 as above. \square

We also observe that primary decomposition gives another algorithm for computing the Galois group. Namely, suppose that $f = x^n - c_1x^{n-1} + \dots + (-1)^n c_n \in \mathbb{Q}[x]$ and that we have the primary decomposition $\langle \sigma_1 - c_1, \dots, \sigma_n - c_n \rangle = \bigcap_{i=1}^r I_i$ in $\mathbb{Q}[x_1, \dots, x_n]$. One easily sees that S_n permutes the I_i and that

$$\text{Gal}(K_i/\mathbb{Q}) \simeq \{\sigma \in S_n \mid \sigma(I_i) = I_i\}.$$

Using a Gröbner basis of I_i , we can easily determine whether $\sigma(I_i)$ equals I_i for any given $\sigma \in S_n$. Hence, by going through the elements of S_n one-by-one, we get a (horribly inefficient) algorithm for computing the Galois group.

Finally, we should note that many of the ideas in this section are well-known to researchers in computational Galois theory. See, for example, [AV] and [PZ].

Acknowledgements

I would like to thank Hal Schenck for inviting me to lecture on the preliminary version of these notes at Texas A&M and Alicia Dickenstein for inviting me to use this expanded version for my lectures at the CIMPA Summer School in Argentina.

I am grateful to Michael Möller for piquing my interest in eigenvector methods and to Alicia Dickenstein and Hal Schenck for useful comments on earlier drafts of the notes.

References

- [ABKR] J. Abbott, A. Bigatti, M. Kreuzer and L. Robbiano, *Computing ideals of points*, J. Symbolic Comput. **30** (2000), 341–356.
- [ABRW] M.-E. Alonso, E. Becker, M.-F. Roy and T. Wörmann, *Zeros, multiplicities, and idempotents for zero-dimensional systems*, in *Algorithms in algebraic geometry and applications (Santander, 1994)*, Progr. Math. **143**, Birkhäuser, Basel, 1996, 1–15.
- [AV] P. Aubry and A. Valibouze, *Using Galois ideals for computing relative resolvents*, J. Symbolic Comput. **30** (2000), 635–651.
- [AS] W. Auzinger and H. J. Stetter, *An elimination algorithm for the computation of all zeros of a system of multivariate polynomial equations*, in *Numerical Mathematics, Singapore 1988* (R.

- Agarwal, Y. Chow and S. Wilson, editors), Intern. Series of Numerical Math. **86**, Birkhäuser, Basel, 1988, 11–30.
- [Cohen] H. Cohen, *A Course in Computational Algebraic Number Theory*, Springer-Verlag, New York Berlin Heidelberg, 1993.
- [Cox] D. Cox, *What is the multiplicity of a basepoint?*, lecture notes, available at the URL <http://www.amherst.edu/~dacox>.
- [CLO1] D. Cox, J. Little and D. O’Shea, *Ideals, Varieties and Algorithms*, Second Edition, Springer-Verlag, New York Berlin Heidelberg, 1997.
- [CLO2] D. Cox, J. Little and D. O’Shea, *Using Algebraic Geometry*, Springer-Verlag, New York Berlin Heidelberg, 1998.
- [Edwards] H. M. Edwards, *Essays in Constructive Mathematics*, in preparation.
- [Eisenbud] D. Eisenbud, *Commutative Algebra with a View Toward Algebraic Geometry*, Springer-Verlag, New York Berlin Heidelberg, 1995.
- [EHV] D. Eisenbud, C. Huneke and W. Vasconcelos, *Direct methods for primary decomposition*, Invent. Math. **110** (1992), 207–235.
- [EL] T. Ekedahl and D. Laksov, *Splitting algebras, symmetric functions, and Galois theory*, preprint, 2002, math.AC/0211125.
- [EM1] M. Elkadi and B. Mourrain, *Approche effective des résidus algébriques*, Rapport de recherche **2884**, INRIA, Sophia Antipolis, 1996.
- [EM2] M. Elkadi and B. Mourrain, *Some applications of Bezoutians in effective algebraic geometry*, Rapport de recherche **3572**, INRIA, Sophia Antipolis, 1998.
- [GTZ] P. Gianni, B. Trager and G. Zacharias, *Gröbner bases and primary decomposition of polynomial ideals*, J. Symbolic Comput. **6** (1988), 149–167. Reprinted in *Computational Aspects of Commutative Algebra* (L. Robbiano, editor), Academic Press, San Diego, CA, 1989, 15–33.
- [FGLM] J. C. Faugère, P. Gianni, D. Lazard and T. Mora, *Efficient computation of zero-dimensional Gröbner bases by change of ordering*, J. Symbolic Comput. **16** (1993), 329–344.
- [HS] V. Hribernic and H. J. Stetter, *Detection and validation of clusters of polynomial zeros*, J. Symbolic Comput. **24** (1997), 667–681.
- [Kronecker] L. Kronecker, *Leopold Kronecker’s Werke, Volumes II and III* (K. Hensel, editor), B. G. Teubner, Leipzig, 1897 (Volume II), 1899 and 1931 (Volume III). Reprint by Chelsea, New York, 1968.
- [Lazard] D. Lazard, *Résolution des systèmes d’équations algébriques*, Theor. Comp. Sci. **15** (1981), 77–110.
- [LN] R. Lidl and H. Niederreiter, *Finite Fields*, Encyclopedia of Mathematics and its Applications **20**, Addison-Wesley, Reading, MA, 1983.
- [MMM1] M. Marinari, H. M. Möller and T. Mora, *Gröbner bases of ideals defined by functionals with an application to ideals of projective points*, Appl. Algebra Eng. Commun. Comput. **4** (1993), 103–145.

- [MMM2] M. Marinari, H. M. Möller and T. Mora, *On multiplicities in polynomial system solving*, Trans. Am. Math. Soc. **348** (1996), 3283–3321.
- [MS] H. M. Möller and H. J. Stetter, *Multivariate polynomial systems with multiple zeros solved by matrix eigenproblems*, Numer. Math. **70** (1995), 311–329.
- [MT] H. M. Möller and R. Tenberg, *Multivariate polynomial system solving using intersections of eigenspaces*, J. Symbolic Comput. **30** (2001), 1–19.
- [Monico] C. Monico, *Computing the primary decomposition of zero-dimensional ideals*, J. Symbolic Comput. **34** (2002), 451–459.
- [PZ] M. Pohst and H. Zassenhaus, *Algorithmic Algebraic Number Theory*, Cambridge Univ. Press, Cambridge, 1989.
- [Sharp] R. Y. Sharp, *Steps in Commutative Algebra*, Cambridge Univ. Press, Cambridge, 1990.
- [Sturmfels] B. Sturmfels, *Solving Systems of Polynomial Equations*, AMS, Providence, 2003.
- [YNT] K. Yokoyama, M. Noro and T. Takeshima, *Solutions of systems of algebraic equations and linear maps on residue class rings*, J. Symbolic Comput. **14** (1992), 399–417.