Chapter 9
Gröbner Bases and Buchberger's Algorithm

We next examine the problem of finding the common roots of a finite set of polynomials over a field K. To do this, we first introduce some necessary algebraic structures. *Gröbner bases* play a key role in the computational aspect of this problem.

In Chapter 10 we will see how to computationally solve arbitrary systems of polynomial equations using Gröbner bases.

9.1 Ideals and the Univariate Case

In the following we study a polynomial ring over an arbitrary field K. In Chapter 8 we defined (affine and projective) algebraic varieties for a given polynomial. We now generalize this definition to a set of polynomials. Let $S \subseteq K[x_1, \ldots, x_n]$ be an arbitrary set of polynomials. Then

$$\mathrm{V}(S) := \left\{a \in K^n : f(a_1, \ldots, a_n) = 0 \text{ for all } f \in S\right\} = \bigcap_{f \in S} \mathrm{V}(f)$$

is called the *affine variety* of S over the field K. Hence, an affine variety is an intersection of affine hyperplanes. One can immediately observe that any common root of the polynomials $f_1, \ldots, f_t \in K[x_1, \ldots, x_n]$ is also a root of $\sum_{i=1}^{t} h_i f_i$. This holds for an arbitrary choice of $h_1, \ldots, h_t \in K[x_1, \ldots, x_n]$, which motivates the following definition.

Definition 9.1 A non-empty set $I \subseteq K[x_1, \ldots, x_n]$ is called an *ideal* if for all $f, g \in I$ and all $h \in K[x_1, \ldots, x_n]$ we have $f + g \in I$ and $hf \in I$.

M. Joswig, T. Theobald, *Polyhedral and Algebraic Methods in Computational Geometry*, Universitext, DOI 10.1007/978-1-4471-4817-3_9,

For $S \subseteq K[x_1, \ldots, x_n]$ we denote by $\langle S \rangle$ the *ideal generated by* S, i.e., the smallest ideal of $K[x_1, \ldots, x_n]$ that contains S. We have

$$\langle S \rangle = \left\{ \sum_{i=1}^{t} h_i f_i : f_1, \ldots, f_t \in S,\ h_1, \ldots, h_t \in K[x_1, \ldots, x_n],\ t \in \mathbb{N} \right\}.$$

The following exercise illustrates how varieties can be defined using ideals.

Exercise 9.2 Show that $\mathrm{V}(S) = \mathrm{V}(\langle S \rangle)$.

A generating system of an ideal I is also called a *basis* of I. Here we need to stress that—unlike in the case of vector spaces—an ideal can have bases of different cardinalities: For example, every subset of an ideal I which contains a basis of I is also a basis of I. In Corollary 9.23, which is also known as the *Hilbert Basis Theorem*, we will see that every ideal $I \subseteq K[x_1, \ldots, x_n]$ is finitely generated.

Not every basis of an ideal is of equal quality. Some bases allow for the observation of more characteristics of the ideal than others. We illustrate this with an example.

Example 9.3 Let $f = x^2 y + x + 1$, $g = x^3 y + x + 1 \in \mathbb{C}[x, y]$. In order to compute the common roots of f and g, it is helpful to have a polynomial of $I = \langle f, g \rangle$ that depends on only one unknown (e.g. on x). In this case we have

$$x^2 - 1 = x \cdot f - g \in I.$$

Therefore, for every common root $(a, b)^T$ of f, g we know that $a \in \{-1, 1\}$. Substituting and solving the equations for y shows that the two points $(-1, 0)^T$ and $(1, -2)^T$ are the common roots of f and g. Figure 9.1 illustrates the real part of the curves $\mathrm{V}(f)$ and $\mathrm{V}(g)$. We have $x \cdot f - (x^2 - 1) = g$, hence $I = \langle f, x^2 - 1 \rangle$.

At this point, we briefly remark that the resultant $\mathrm{Res}_y(f, g) = -x^2(x^2 - 1)$ is also contained in I. We shall return to this connection in Chapter 10 (Proposition 10.4).

The previous example suggests the idea of solving a system of polynomial equations via step by step elimination of variables followed by backwards substitution. This corresponds to solving a linear system of equations in row echelon form. This approach motivates the following term: For an ideal $I = \langle f_1, \ldots, f_t \rangle \subseteq K[x_1, \ldots, x_n]$ and $i \in \{1, \ldots, n-1\}$ let

$$I \cap K[x_{i+1}, \ldots, x_n]$$

denote the i-th *elimination ideal*.

Exercise 9.4 Show that the i-th elimination ideal of I is indeed an ideal in $K[x_{i+1}, \ldots, x_n]$.

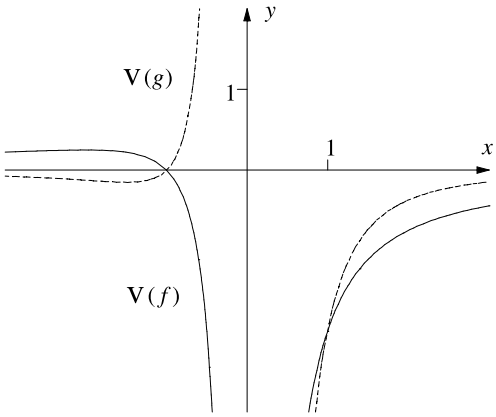

Fig. 9.1 Varieties V(f) and V(g)

We now lay the foundation for the study of elimination ideals in Chapter 10. To do this we study the question of how, given an ideal I and a polynomial f, we can determine if f is in I. This is the so-called *ideal membership problem* for which Algorithm 9.3 on p. 148 provides a solution.

We first examine the special case of the ideal membership problem with one unknown: For given polynomials $f_1, \ldots, f_t, f \in K[x]$ we ask whether $f \in \langle f_1, \ldots, f_t \rangle$. A polynomial ring $K[x]$ in one variable is a *Euclidean ring* since we can define a division algorithm. Division (via the Euclidean Algorithm 9.1) allows us to compute the greatest common divisor g of the polynomials $f_1, \ldots, f_t$, with $\langle g \rangle = \langle f_1, \ldots, f_t \rangle$. Furthermore, the Euclidean algorithm allows us to solve the ideal membership problem, since it in particular enables us to determine if the remainder of f divided by g is 0.

We assume that the reader knows the basic principles of the algorithm. However, due to the significance of these two algorithms in our further work, we will illustrate them.

For two polynomials $f, g \in K[x] \setminus \{0\}$ there exist $r, s \in K[x]$ such that

$$f = q \cdot g + r \quad \text{where } \deg r < \deg g. \tag{9.1}$$

When $\deg f \geq \deg g$, we do the following: Assume that $f = \sum_{i=0}^{n} a_i x^i$ and $g = \sum_{j=0}^{m} b_j x^j$ where $n \geq m$ and $a_n, b_m \neq 0$. Via induction over the degree we can assume that the polynomial $h := f - \frac{a_n}{b_m} \cdot x^{n-m} \cdot g$ of degree $\leq n-1$ has a decomposition $h = q' \cdot g + r$, such that $\deg r < \deg g$. This implies

$$f = h + \frac{a_n}{b_m} \cdot x^{n-m} \cdot g = \left(q' + \frac{a_n}{b_m} x^{n-m} \right) g + r.$$

Using $q := q' + \frac{a_n}{b_m} x^{n-m}$ we get the desired statement. We denote the *remainder* r as $\operatorname{rem}(f; g)$ and write $g \mid f$ if $\operatorname{rem}(f; g) = 0$.

Definition 9.5 Let K be a field. A polynomial $g \in K[x]$ is called a *greatest common divisor* (gcd) of $f_1, \ldots, f_t \in K[x] \setminus \{0\}$ if the following conditions are satisfied.

Algorithm 9.1: The Euclidean algorithm

Input: $f, g \in K[x] \setminus \{0\}$ with $\deg f \geq \deg g$
Output: $\gcd(f, g)$

```
1 r_0 ← f; r_1 ← g; i ← 1
2 while r_i ≠ 0 do
3 |   r_{i+1} ← rem(r_{i-1}; r_i)
4 |   i ← i + 1
5 return r_{i-1}
```

(a) $g \mid f_i$ for all $i \in \{1, \ldots, t\}$;
(b) if $h \mid f_1, \ldots, h \mid f_t$ then $h \mid g$ for all $h \in K[x]$.

In every unique factorization domain there exists a greatest common divisor which is unique up to multiplication by a unit (here a non-zero constant in K), see Appendix A. For uniqueness we choose the gcd with leading coefficient 1.

Analogously we can define the *least common multiple* of $f_1, \ldots, f_t$. Alternatively we can read the following computational rule for two polynomials as a definition:

$$\operatorname{lcm}(f_1, f_2) := \frac{f_1 f_2}{\gcd(f_1, f_2)}.$$

This also shows that the computation of the least common multiple can be reduced to the computation of the greatest common divisor.

A special property of the ring $K[x]$, or of any Euclidean ring, is that the gcd can be algorithmically computed.

The Euclidean Algorithm 9.1 terminates since the degrees of the polynomials r_i are strictly decreasing. We denote by q_i the polynomial such that in Step 2 we have

$$r_{i-1} = q_i \cdot r_i + r_{i+1}. \tag{9.2}$$

To prove that the algorithm is correct we show that $r := r_{i-1}$, which is returned in the last step, satisfies the two conditions from Definition 9.5. Using (9.2) we can successively deduce that r divides the remainders $r_{i-2}, r_{i-3}, \ldots, r_1 = g$ and $r_0 = f$. If h divides f as well as g, then h divides $r_2, r_3, \ldots, r_{i-1}$. This can also be deduced from (9.2).

Every remainder computed throughout the Euclidean algorithm is contained in the ideal $\langle f, g \rangle$ of the two input polynomials $f, g \in K[x]$ and hence we have $\gcd(f, g) \in \langle f, g \rangle$. Therefore,

$$\big\langle \gcd(f, g) \big\rangle = \langle f, g \rangle.$$

Example 9.6 Applying the Euclidean algorithm to the two polynomials $f = x^4 - x^3$ and $g = x^3 - x$ repeatedly yields $(q_1, r_2) = (x - 1, x^2 - x)$ and $(q_2, r_3) = (x + 1, 0)$, so that $x^2 - x$ is the gcd of f and g.

To determine a single generator of an ideal which is given by more than two polynomials it is sufficient to verify the following rule:

Exercise 9.7 For $t \geq 3$ we have $\gcd(f_1, \ldots, f_t) = \gcd(f_1, \gcd(f_2, \ldots, f_t))$.

Given a sequence of generators $f_1, \ldots, f_t$ of an ideal $I \subseteq K[x]$, the representation of I as a principal ideal $I = \langle \gcd(f_1, \ldots, f_t) \rangle$ is called a *normal form* of I. This notion is justified by the following exercise.

Exercise 9.8 For univariate polynomials $f_1, \ldots, f_t, g_1, \ldots, g_s \in K[x] \setminus \{0\}$ with $\langle f_1, \ldots, f_t \rangle$ equal to $\langle g_1, \ldots, g_s \rangle$, show that up to a constant factor the polynomials $\gcd(f_1, \ldots, f_t)$ and $\gcd(g_1, \ldots, g_s)$ coincide.

In particular, the special case $s = 1$ implies that in any representation of an ideal $I \subseteq K[x]$ as a principal ideal $I = \langle g_1 \rangle$, the polynomial g_1 is uniquely determined up to a constant factor.

The Euclidean algorithm serves to compute a normal form for a given ideal I in $K[x]$. If an ideal in $K[x]$ is given in normal form, i.e., by a single generator, then the Euclidean division solves the ideal membership problem. The goal of the following sections is to generalize these two methods to polynomial rings with an arbitrary number of unknowns.

9.2 Monomial Orders

The degree naturally defines a partial order on the polynomials in one unknown, which were studied in the previous section. The remainder polynomial, which is the result of the division of a polynomial f by g, is smaller than g with respect to this partial order. To define a proper division in the multivariate case it is necessary to first define a suitable order on the set of monomials.

A monomial $x_1^{\alpha_1} \cdots x_n^{\alpha_n}$ in $K[x_1, \ldots, x_n]$ is denoted by x^α, where $\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathbb{N}^n$ is a multi-index. In Definition 8.14 we defined the total degree of a monomial as $\operatorname{tdeg} x^\alpha := \alpha_1 + \cdots + \alpha_n$. The notation $|\alpha|$ is also used as an alternative to $\operatorname{tdeg} x^\alpha$.

Definition 9.9 A *monomial order* on $K[x_1, \ldots, x_n]$ is a relation $\prec$ on $\mathbb{N}^n$ (or equivalently a relation on the set of monomials x^α for $\alpha \in \mathbb{N}^n$), which satisfies the following properties.

(a) The relation $\prec$ is a well-ordered relation on $\mathbb{N}^n$, i.e., every non-empty subset of $\mathbb{N}^n$ has a minimal element with respect to $\prec$.
(b) $\alpha \prec \beta$ and $\gamma \in \mathbb{N}^n$ implies $\alpha + \gamma \prec \beta + \gamma$.

Every well-ordered relation is a total order. From condition (b) it follows that the zero vector (respectively the empty monomial 1) is the unique smallest element

with respect to every monomial order. The second condition requires a *compatibility with respect to multiplication*: $x^\alpha \cdot x^\gamma \prec x^\beta \cdot x^\gamma$ (when expressed in the monomial description).

Definition 9.10 (Lexicographic order) Let $\alpha, \beta \in \mathbb{N}^n$. We define $x^\alpha \prec_{\text{lex}} x^\beta$ if the leftmost non-zero coefficient in the difference $\beta - \alpha \in \mathbb{Z}^n$ is positive.

Example 9.11 We have $(4,3,1) \succ_{\text{lex}} (3,7,10)$ and $(4,3,1) \prec_{\text{lex}} (4,7,10)$. Expressed as monomials in $K[x,y,z]$, this translates to $x^4y^3z^1 \succ_{\text{lex}} x^3y^7z^{10}$ and $x^4y^3z^1 \prec_{\text{lex}} x^4y^7z^{10}$ respectively.

The relation $\prec_{\text{lex}}$ is a monomial order. It suffices to check that the relation is a well-ordered relation. If we assume that $\prec_{\text{lex}}$ is not a well-ordered relation, then we can find a strictly decreasing series

$$\alpha^{(1)} \succ_{\text{lex}} \alpha^{(2)} \succ_{\text{lex}} \alpha^{(3)} \succ_{\text{lex}} \cdots \tag{9.3}$$

of elements in $\mathbb{N}^n$. By the definition of the lexicographic order, the leftmost entries $\alpha_1^{(i)}$ define a non-increasing series in $\mathbb{N}$. Since the set of natural numbers is well-ordered, there exists an N_1 such that $(\alpha^{(i)})_1 = (\alpha^{(N_1)})_1$ for all $i \geq N_1$. Now by considering only the series elements after the index N_1, we can in the same way deduce that there exist $N_2, \ldots, N_n$ such that $(\alpha^{(i)})_j = (\alpha^{(N_j)})_j$ for all $i \geq N_j$ and $j \in \{2, \ldots, n\}$. This contradicts the series (9.3) being strictly decreasing.

A monomial order yields a unique sorted description for arbitrary polynomials. For the remaining part of this section we will fix a monomial order $\prec$ on $K[x_1, \ldots, x_n]$. For a non-zero polynomial $f = \sum_\alpha c_\alpha x^\alpha$ in $K[x_1, \ldots, x_n]$ let $\alpha^* := \max_\prec\{\alpha : c_\alpha \neq 0\}$. The *leading monomial* of f is $\text{lm}_\prec(f) := x^{\alpha^*}$ and the corresponding coefficient $\text{lc}_\prec(f) := c_{\alpha^*}$ is called the *leading coefficient*. Their product

$$\text{lt}_\prec(f) := \text{lc}_\prec(f) \cdot \text{lm}_\prec(f) = c_{\alpha^*} \cdot x^{\alpha^*}$$

is called the *leading term* of f. When the monomial order is contextually clear it is often neglected in the notation.

Example 9.12 For $f = 5x^4y^3z + 2x^3y^7z^{10}$ in $K[x,y,z]$ we have that $\text{lt}_{\prec_{\text{lex}}}(f) = x^4y^3z$, $\text{lc}_{\prec_{\text{lex}}}(f) = 5$ and $\text{lm}_{\prec_{\text{lex}}}(f) = 5x^4y^3z$ with respect to the lexicographic order.

We are now able to generalize the division algorithm to the multivariate case. Here there is a major difference in comparison to the univariate case: It is useful to describe the division of a polynomial $f \in K[x_1, \ldots x_n]$ by a set of polynomials $(f_1, \ldots, f_t)$ since ideals in $K[x_1, \ldots, x_n]$ are in general not generated by a single polynomial.

We look at the leading monomial $\text{lm}(f)$ of f and check if division by any of the leading monomials $\text{lm}(f_1), \ldots, \text{lm}(f_t)$ results in a remainder of 0. For the first

Algorithm 9.2: The multivariate division algorithm

Input: $f, f_1, \ldots, f_t \in K[x_1, \ldots, x_n]$ with $f_i \neq 0$
Output: $a_1, \ldots, a_t, r$ with $f = \sum_{i=1}^{t} a_i f_i + r$

```
a_i ← 0 for all i ∈ {1, ..., t}
p ← f
while p ≠ 0 do
    m ← lt(p)
    i ← 1
    while i ≤ t and m ≠ 0 do
        if lt(f_i) divides m then
            a_i ← a_i + m/lt(f_i) ; p ← p − m/lt(f_i) f_i
            m ← 0
        i ← i + 1
    r ← r + m ; p ← p − m
return (a_1, ..., a_t; r)
```

polynomial f_k which satisfies this condition, we subtract a suitable multiple of f_k from f,

$$f - \frac{\mathrm{lt}(f)}{\mathrm{lt}(f_k)} f_k,$$

and obtain a new polynomial which is strictly smaller than f with respect to the monomial order. We replace f by the new polynomial and repeat the process. If the leading monomial of f is not divisible by any of the leading terms $\mathrm{lt}(f_1), \ldots, \mathrm{lt}(f_t)$, we add the leading term to the remainder, subtract it from f and start again at the beginning.

The remainder r which is produced by Algorithm 9.2 is called the *remainder* of f after division by $(f_1, \ldots, f_t)$ and we denote it by $\mathrm{rem}(f; f_1, \ldots, f_t)$. In general this remainder is not independent of the order of the polynomials by which we divide.

Example 9.13 Let $f = xy^2 - y$, $f_1 = xy - 1$ and $f_2 = y^2 + 1$ be polynomials in $K[x, y]$. With respect to the lexicographic order and the ordering (f_1, f_2) of the polynomials, the division algorithm divides the leading term xy^2 by xy resulting in y. Since $f - y \cdot f_1 = 0$ the algorithm terminates and returns the decomposition

$$xy^2 - y = y \cdot (xy - 1) + 0 \cdot \left(y^2 + 1\right) + 0.$$

If we reverse the ordering of the polynomials, i.e., we divide by (f_2, f_1), the term xy^2 is divided by the leading monomial y^2, resulting in x. Since the polynomial $f - x \cdot f_2 = -y - x$ is not divisible any further by f_1 or f_2, the algorithm yields

the decomposition

$$f = x \cdot \left(y^2 + 1\right) + 0 \cdot (xy - 1) + (-x - y).$$

Our remainders are: $\operatorname{rem}(f; f_1, f_2) = 0$ and $\operatorname{rem}(f; f_2, f_1) = -x - y$.

In general, the multivariate division algorithm results in a representation of the following form.

Lemma 9.14 *For given polynomials $f, f_1, \ldots, f_t \in K[x_1, \ldots, x_n]$ the Division Algorithm 9.2 returns polynomials $a_1, \ldots, a_t$ and $r = \operatorname{rem}(f; f_1, \ldots, f_t)$, for which we have*

$$f = a_1 f_1 + \cdots + a_t f_t + r,$$

where no term of r is divisible by any of the monomials $\operatorname{lm}(f_1), \ldots, \operatorname{lm}(f_t)$. Furthermore, we have for each $i \in \{1, \ldots, t\}$ with $a_i \neq 0$ that

$$\operatorname{lm}(a_i f_i) \preceq \operatorname{lm}(f).$$

Proof It is clear by the construction of the algorithm that no term of the remainder r is divisible by any of the leading monomials $\operatorname{lm}(f_1), \ldots, \operatorname{lm}(f_t)$. The assignment $a_i \leftarrow a_i + \frac{\operatorname{lt}(p)}{\operatorname{lt}(f_i)}$ ensures that the product $a_i \operatorname{lt}(f_i)$ is a sum of terms of f. However, the terms of f are dominated by their leading term. □

Exercise 9.15 Let $\prec$ be a monomial order on $K[x_1, \ldots, x_n]$. Show that

$$\alpha \prec^{\text{tdeg}} \beta \quad :\Longleftrightarrow \quad \operatorname{tdeg}\alpha < \operatorname{tdeg}\beta \quad \text{or} \quad (\operatorname{tdeg}\alpha = \operatorname{tdeg}\beta \text{ and } \alpha \prec \beta),$$

defines a monomial order.

The construction in Exercise 9.15 can, in certain cases, yield a monomial order even if the original order does not satisfy all the axioms of a monomial order. The next exercise exhibits this phenomenon for $\prec_{\text{grevlex}}$, a monomial order that is often a very efficient one in practical computations.

Exercise 9.16 Let $\alpha, \beta \in \mathbb{N}^n$. We define $x^\alpha <_{\text{revlex}} x^\beta$ if the rightmost non-zero coefficient in the difference $\beta - \alpha \in \mathbb{Z}^n$ is negative.

(a) Show that the *reverse lexicographic order* $<_{\text{revlex}}$ is not a monomial order.
(b) Show that the *graded reverse lexicographic order* defined by

$$\begin{aligned} \alpha \prec_{\text{grevlex}} \beta \quad :\Longleftrightarrow \quad & \operatorname{tdeg}\alpha < \operatorname{tdeg}\beta \quad \text{or} \\ & (\operatorname{tdeg}\alpha = \operatorname{tdeg}\beta \text{ and } \alpha <_{\text{revlex}} \beta), \end{aligned}$$

is a monomial order.

9.3 Gröbner Bases and the Hilbert Basis Theorem

In this section we introduce the key concept for solving the ideal membership problem. We start with an example that illustrates why the multivariate case is much more complicated than the univariate case.

Example 9.17 Let $f_1 = xy + 1$, $f_2 = yz + 1$ be polynomials in $K[x, y]$. In the univariate case it would be desirable to use Euclidean division to determine if the polynomial $f = z - x$ is contained in the ideal $I = \langle f_1, f_2 \rangle$. In fact we do have

$$z - x = z \cdot (xy + 1) - x(yz + 1) \in \langle f_1, f_2 \rangle.$$

However, neither for the ordering (f_1, f_2) nor for the ordering (f_2, f_1) is the remainder zero when applying Euclidean division with respect to lexicographic order. Of course, adding $z - x$ to the ideal basis would give that division of $z - x$ by the new basis would have 0 as the remainder.

It may seem naive to enlarge the original generating system of an ideal by proper polynomials so that *every* polynomial of the ideal has a remainder of zero when divided by the basis. However, this can be algorithmically achieved. What we need for this is a criterion that determines if the generating system is large enough.

We denote the set of leading terms of an ideal I with respect to the monomial order $\prec$ by $\mathrm{lt}_\prec(I)$. The ideal $\langle \mathrm{lt}_\prec(I) \rangle$ generated by the leading terms is called the *initial ideal* of I with respect to $\prec$, and we write $\mathrm{in}_\prec(I) := \langle \mathrm{lt}_\prec(I) \rangle$.

Definition 9.18 Let I be an ideal. A finite subset $G = \{g_1, \ldots, g_t\} \subseteq I$ is called a *Gröbner basis* of I with respect to the monomial order $\prec$ if the leading terms $\mathrm{lt}_\prec(g_1), \ldots, \mathrm{lt}_\prec(g_t)$ generate the initial ideal of I, i.e.,

$$\left\langle \mathrm{lt}_\prec(g_1), \ldots, \mathrm{lt}_\prec(g_t) \right\rangle = \mathrm{in}_\prec(I).$$

Our next important intermediate goal is to show that every ideal has a Gröbner basis. We begin by proving this statement for the special case of monomial ideals. *Monomial ideals* are those ideals which have a generating system consisting only of monomials. Initial ideals are always monomial ideals.

Lemma 9.19 *Let $I = \langle x^\alpha : \alpha \in A \rangle$ where $A \subseteq \mathbb{N}^n$ is a monomial ideal. We have $x^\beta \in I$ if and only if x^β is a multiple of x^α for an $\alpha \in A$.*

Proof If x^β is a multiple of x^α for an $\alpha \in A$, then by the definition of an ideal, $x^\beta \in I$.

Conversely, if $x^\beta \in I$, then there exists a representation $x^\beta = \sum_{i=1}^t h_i x^{\alpha^{(i)}}$ with $h_i \in K[x_1, \ldots, x_n]$ and $\alpha^{(i)} \in A$ for $1 \leq i \leq t$. Every term of the polynomial on the right hand side of the equation is a multiple of a term x^α for some $\alpha \in A$. Therefore, the polynomial on the left hand side of the equation also has this property. □

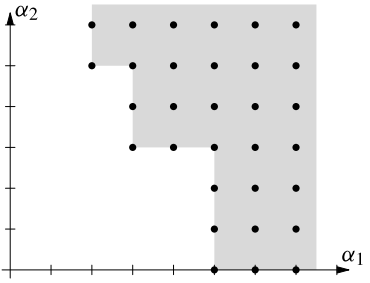

Fig. 9.2 A visualization of the Gordan–Dickson lemma for $n = 2$. Every lattice point (α_1, α_2) represents a monomial $x_1^{\alpha_1} x_2^{\alpha_2}$

The following theorem shows that monomial ideals are finitely generated.

Theorem 9.20 (Gordan–Dickson Lemma) *Every non-empty set M of monomials in $K[x_1, \ldots, x_n]$ contains a finite subset $E \subseteq M$ such that every monomial of M is a multiple of a monomial in E.*

Before beginning the proof, we illustrate the theorem for the case $n = 2$. Each point (i, j) in Fig. 9.2 represents a monomial $x^i y^j$ in $K[x, y]$. If a monomial $x^i y^j$ is contained in a monomial ideal I, then Lemma 9.19 states that every monomial $x^k y^l$ with $k \geq i$ and $l \geq j$ is contained in I as well. So the Gordan–Dickson lemma implies that the points corresponding to monomials in I can be represented as a finite union of transposed copies of the points in the positive orthant.

Proof The proof is by induction over the number of unknowns n. For $n = 1$ we have $M = \{x^\alpha : \alpha \in A\}$ for a subset $A \subseteq \mathbb{N}$. A has a smallest element β. Using Lemma 9.19 we conclude $I = \langle x^\beta \rangle$.

So let $n \geq 2$ and assume that the statement is true for $n - 1$ unknowns. Take an arbitrary monomial

$$x^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n}$$

from M.

We first show that every monomial $x^\beta \in M$ which is not a multiple of x^α belongs to at least one of the sets $M_{i,j}$, where: for $i \in \{1, \ldots, n\}$ and $j \in \{0, \ldots, \alpha_i - 1\}$, $M_{i,j}$ is the set of those monomials $x^\gamma \in M$ for which $\deg_{x_i}(x^\gamma) = j$. Since x^α does not divide the monomial x^β, we have $\beta_i < \alpha_i$ for some $i \in \{1, \ldots, n\}$. Hence, $x^\beta \in M_{i,\beta_i}$.

Let $M'_{i,j}$ be the set of monomials in $K[x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n]$ that can be obtained from monomials of $M_{i,j}$ by dropping the factor x_i^j. By the inductive hypothesis there exist finite subsets $E'_{i,j} \subseteq M'_{i,j}$ such that every monomial in $M'_{i,j}$ is a multiple of the monomial $E'_{i,j}$. We define

$$E_{i,j} := \left\{ p \cdot x_i^j : p \in E'_{i,j} \right\}.$$

Now it is clear that every monomial in M is a multiple of a monomial in the finite set

$$E := \{x^\alpha\} \cup \bigcup_{i,j} E_{i,j} \subseteq M.$$

□

Remark 9.21 This lemma will play a key role in proving the termination of several algorithms. The statement is actually purely combinatorial: Given a set $\mathcal{A}$ of subsets of $\mathbb{N}^n$ such that every $A \in \mathcal{A}$ is of the form $\alpha_A + \mathbb{N}^n$ with $\alpha_A \in \mathbb{N}^n$, the union $\bigcup_{A\in\mathcal{A}} A$ is a finite union, i.e., there exist $A_1, \ldots, A_k \in \mathcal{A}$ with $\bigcup_{A\in\mathcal{A}} A = \bigcup_{i=1}^{k} A_i$.

Using the Gordan–Dickson lemma it is now possible to prove that every non-zero ideal in $K[x_1, \ldots, x_n]$ has a Gröbner basis.

Theorem 9.22 *Let $\prec$ be a monomial order on $K[x_1, \ldots, x_n]$. Then*:

(a) *Every non-zero ideal I has a Gröbner basis.*
(b) *The elements of a Gröbner basis of I generate the ideal I.*

Proof Let $I \neq \{0\}$ be an ideal.

(a): The initial ideal $\mathrm{in}_\prec(I)$ is generated by the monomials $\mathrm{lm}_\prec(g)$, with $g \in I \setminus \{0\}$. By the Gordan–Dickson lemma 9.20 there exist finitely many $g_1, \ldots, g_t$ with

$$\langle \mathrm{lt}_\prec(g_1), \ldots, \mathrm{lt}_\prec(g_t) \rangle = \mathrm{lt}_\prec(I),$$

which ensures the existence of a Gröbner basis.

(b): The ideal J which is generated by the polynomials $g_1, \ldots, g_t$ of a Gröbner basis is clearly contained in I. To show the reverse inclusion we assume that $I \setminus J \neq \emptyset$. Let f be a polynomial in $I \setminus J$ with a leading term that is minimal with respect to $\prec$. Since $\mathrm{lm}_\prec(g_1), \ldots, \mathrm{lm}_\prec(g_t)$ generate the initial ideal $\mathrm{in}_\prec(I)$, there exist polynomials $h_1, \ldots, h_t$ with

$$\mathrm{lm}_\prec(f) = \mathrm{lm}_\prec(g_1) \cdot h_1 + \cdots + \mathrm{lm}_\prec(g_t) \cdot h_t.$$

The polynomial

$$g = f - \sum_{i=1}^{t} g_i h_i$$

is contained in I but not in J (otherwise we would have $f \in J$). We also have that the leading monomial of f does not appear in g, which means that the corresponding coefficient is zero. Hence $\mathrm{lm}_\prec(g)$ is smaller than $\mathrm{lm}_\prec(f)$ with respect to the monomial order $\prec$. This contradicts the minimality of f. We therefore have $I = J$, which proves our statement. □

As an immediate consequence of Theorem 9.22 we get the following finiteness statement.

Algorithm 9.3: A solution of the ideal membership problem

Input: $f, g_1, \ldots, g_t \in K[x_1, \ldots, x_n]$, such that $G := \{g_1, \ldots, g_t\}$ is a Gröbner basis of the ideal $I = \langle G \rangle$ with respect to the monomial order $\prec$

Output: Determine if $f \in I$

```
r ← rem_≺(f; g_1, ..., g_t)
if r = 0 then
    return "Yes"
else
    return "No"
```

Corollary 9.23 (Hilbert Basis Theorem) *Every ideal $I \subseteq K[x_1, \ldots, x_n]$ has a finite generating system.*

The important property of Gröbner bases is that they provide a solution to the ideal membership problem, as carried out in Algorithm 9.3.

Correctness of Algorithm 9.3 If $\mathrm{rem}_\prec(f; g_1, \ldots, g_t) = 0$, then f is contained in I. It remains to be shown that $\mathrm{rem}_\prec(f; g_1, \ldots, g_t) \neq 0$ implies that $f \notin I$. Assume that $\mathrm{rem}_\prec(f; g_1, \ldots, g_t) \neq 0$ and $f \in I$. Then the remainder $r = \mathrm{rem}_\prec(f; g_1, \ldots, g_t) \in I$ and therefore $\mathrm{lt}_\prec(r) \in \mathrm{in}_\prec(I)$. Since G is a Gröbner basis, it follows that $\mathrm{in}_\prec(I) = \langle \mathrm{lt}_\prec(g_1), \ldots, \mathrm{lt}_\prec(g_t) \rangle$. By Lemma 9.19, $\mathrm{lt}_\prec(r)$ is a multiple of a leading term $\mathrm{lt}_\prec(g_i)$ for an $i \in \{1, \ldots, t\}$. But by Lemma 9.14, the divisibility of $\mathrm{lt}_\prec(r)$ by $\mathrm{lt}_\prec(g_i)$ contradicts the fact that r is a remainder of the division by $g_1, \ldots, g_t$. □

For the remaining part of this section assume that $G = \{g_1, \ldots, g_t\}$ is a Gröbner basis of the ideal $I \subseteq K[x_1, \ldots, x_n]$ with respect to the monomial order $\prec$.

Exercise 9.24 Show that Euclidean division is independent of the order of polynomials in G:

$$\mathrm{rem}_\prec(f; g_1, \ldots, g_t) = \mathrm{rem}_\prec(f; g_{\sigma(1)}, \ldots, g_{\sigma(t)})$$

for all permutations σ.

We can therefore write $\mathrm{rem}_\prec(f; G)$ instead of $\mathrm{rem}_\prec(f; g, \ldots, g_t)$. The following holds for polynomials which need not form a Gröbner basis.

Exercise 9.25 Let $f_1, \ldots, f_t$ be an arbitrary finite family of polynomials in $K[x_1, \ldots, x_n]$. Show that for arbitrary $f, g \in K[x_1, \ldots, x_n]$ and $c \in K$:

(a) $\mathrm{rem}_\prec(f + g; f_1, \ldots, f_t) = \mathrm{rem}_\prec(f; f_1, \ldots, f_t) + \mathrm{rem}_\prec(g; f_1, \ldots, f_t)$;
(b) $\mathrm{rem}_\prec(cf; f_1, \ldots, f_t) = c\,\mathrm{rem}_\prec(f; f_1, \ldots, f_t)$.

This implies that a Gröbner basis defines a *normal form* for the *equivalence classes*

$$f + I = \mathrm{rem}(f; G) + I.$$

Furthermore, the normal forms of the equivalence classes for I define a K-vector space.

9.4 Buchberger's Algorithm

The proof of the existence of Gröbner bases in Theorem 9.22 was not constructive. The topic of this section is an algorithm for computing Gröbner bases that dates back to the PhD thesis of Bruno Buchberger in 1965. His method is one of the most important methods in modern computer algebra.

The following finiteness statement will later provide an argument for the termination of Buchberger's algorithm.

Proposition 9.26 (Ascending Chain Condition) *Let $I_1 \subseteq I_2 \subseteq I_3 \subseteq \cdots$ be a monotonically ascending chain of ideals in $K[x_1, \ldots, x_n]$, then there exists an $N \geq 1$ with $I_N = I_{N+1} = I_{N+2} = \cdots$.*

In other words: Every ascending chain of ideals terminates.

Proof Given an ascending chain of ideals $I_1 \subseteq I_2 \subseteq I_3 \subseteq \cdots$ we study the union $I = \bigcup_{i=1}^{\infty} I_i$. Definition 9.1 gives that I is an ideal. Hilbert's Basis Theorem 9.23 shows that I has a finite set of generators $f_1, \ldots, f_t$. Every polynomial f_i is contained in an ideal I_{j_i} for a suitable $j_i \in \mathbb{N}$. For $N = \max\{j_i : 1 \leq i \leq t\}$ we have $f_1, \ldots, f_t \in I_N$ and therefore, $I_N = I_{N+1} = \cdots = I$. □

A commutative ring is called *Noetherian* if the ascending chain condition holds. In the proof above we saw that the ascending chain condition follows from the fact that all ideals are finitely generated. The converse is also true: Hilbert's basis theorem and the ascending chain condition are equivalent.

As before we fix the monomial order $\prec$ for the following.

Definition 9.27 The *S-polynomial* of two non-zero polynomials f and g in $K[x_1, \ldots, x_n]$ is defined as

$$\mathrm{spol}_{\prec}(f, g) := \frac{\mathrm{lt}_{\prec}(g)}{m} f - \frac{\mathrm{lt}_{\prec}(f)}{m} g,$$

where m denotes the greatest common divisor of $\mathrm{lm}_{\prec}(f)$ and $\mathrm{lm}_{\prec}(g)$.

Buchberger's Gröbner basis algorithm uses the following characterization.

Theorem 9.28 (Buchberger's Criterion) *A finite set* $G = \{g_1, \ldots, g_t\} \subseteq K[x_1, \ldots, x_n]$ *is a Gröbner basis for* $\langle G \rangle$ *with respect to* $\prec$ *if and only if the remainder* $\operatorname{rem}_{\prec}(\operatorname{spol}_{\prec}(g_i, g_j); G)$ *vanishes for all* $i, j \in \{1, \ldots, t\}$.

Proof If G is a Gröbner basis, then we have $\operatorname{spol}(g_i, g_j) \in I$ and the remainder after Euclidean division by G is the zero polynomial.

For the reverse implication let $\operatorname{rem}(\operatorname{spol}(g_i, g_j); G) = 0$ for all i, j. A polynomial $f \in I$ has a representation

$$f = \sum_{i=1}^{t} h_i g_i \tag{9.4}$$

with polynomials $h_1, \ldots, h_t \in K[x_1, \ldots, x_n]$. We have to show that the leading term $\operatorname{lt}(f)$ is a multiple of $\operatorname{lt}(g_i)$ for some basis element $g_i \in G$. The representation (9.4) immediately gives that

$$\operatorname{lm}(f) \preceq \max\left\{\operatorname{lm}(h_i g_i) : 1 \leq i \leq t\right\} = x^{\alpha}$$

for an $\alpha \in \mathbb{N}^n$. Without loss of generality we can assume that $\operatorname{lm}(h_1 g_1) = x^{\alpha}$ and $\operatorname{lc}(g_i) = 1$ for all $i \in \{1, \ldots, t\}$. We distinguish between two cases.

Case 1: $\operatorname{lm}(f) = x^{\alpha}$. Here the monomial x^{α} is a multiple of $\operatorname{lm}(g_1)$ and we have nothing left to show.

Case 2: $\operatorname{lm}(f) \prec x^{\alpha}$. In this case there exists at least one other polynomial $h_i g_i$ such that $\operatorname{lt}(h_i g_i) = x^{\alpha}$, as otherwise it would be impossible to cancel the x^{α} terms through addition. Without loss of generality we can assume that $\operatorname{lm}(h_2 g_2) = x^{\alpha}$. Using the notation $\operatorname{lt}(h_1) = b_{\beta} x^{\beta}$ and $\operatorname{lt}(h_2) = c_{\gamma} x^{\gamma}$ we have

$$\begin{aligned} h_1 g_1 &= \left(b_{\beta} x^{\beta} + \cdots\right) g_1 = b_{\beta} x^{\beta} g_1 + \left(\text{terms} \prec x^{\alpha}\right) \quad \text{and} \\ h_2 g_2 &= \left(c_{\gamma} x^{\gamma} + \cdots\right) g_2 = c_{\gamma} x^{\gamma} g_2 + \left(\text{terms} \prec x^{\alpha}\right). \end{aligned}$$

By construction we have that x^{α} is a multiple of the leading monomials of g_1 and g_2 and hence also a multiple of $x^{\mu} := \operatorname{lcm}(\operatorname{lm}(g_1), \operatorname{lm}(g_2))$. This yields

$$\begin{aligned} h_1 g_1 + h_2 g_2 &= (b_{\beta} + c_{\gamma}) x^{\beta} g_1 + c_{\gamma}\left(x^{\gamma} g_2 - x^{\beta} g_1\right) + \left(\text{terms} \prec x^{\alpha}\right) \\ &= (b_{\beta} + c_{\gamma}) x^{\beta} g_1 - c_{\gamma} x^{\alpha - \mu} \operatorname{spol}(g_1, g_2) + \left(\text{terms} \prec x^{\alpha}\right). \end{aligned}$$

Our assumption implied $\operatorname{rem}(\operatorname{spol}(g_1, g_2); G) = 0$ and thus Lemma 9.14 implies that there exist polynomials $u_1, \ldots, u_t$ with

$$\operatorname{spol}(g_1, g_2) = \sum_{i=1}^{t} u_i g_i$$

Algorithm 9.4: Buchberger's algorithm

Input: finite set of polynomials $F = \{f_1, \ldots, f_t\} \subseteq K[x_1, \ldots, x_n]$
Output: Gröbner basis G for $\langle F \rangle$ with respect to $\prec$ with $F \subseteq G$

```
1 G ← F
2 repeat
3 |   G' ← G
4 |   foreach pair {p, q} ⊆ G' with p ≠ q do
5 |   |   r ← rem_≺(spol_≺(f, g); G')
6 |   |   if r ≠ 0 then
7 |   |   |   G ← G ∪ {r}
8 until G = G'
9 return (G)
```

and $\operatorname{lm}(u_i g_i) \preceq \operatorname{lm}(\operatorname{spol}(g_1, g_2)) \prec x^\mu$. In particular we have $\operatorname{lm}(x^{\alpha-\mu} u_i g_i) \prec x^\alpha$ for $1 \leq i \leq t$, which implies that there exist polynomials $h'_1, \ldots, h'_t$ with

$$f = \sum_{i=1}^{t} h'_i g_i.$$

Compared to the original representation (9.4), the number of terms $h'_i g_i$ whose leading monomial is x^α either decreases, or we have

$$\max_{\prec} \{\operatorname{lm}(h'_i g_i) : 1 \leq i \leq t\} \prec x^\alpha.$$

Therefore, after finitely many steps, the problem can be reduced to the first case. This proves the statement. □

The basic idea behind the computation of a Gröbner basis of an ideal is to successively add S-polynomials to a given generating system. By Buchberger's criterion we know that we have a Gröbner basis if all of the remainders of the S-polynomials vanish when divided by the generators. We summarize the method in Algorithm 9.4.

Theorem 9.29 *Let $f_1, \ldots, f_t \in K[x_1, \ldots, x_n]$ with $\langle f_1, \ldots, f_t \rangle \neq \{0\}$. Buchberger's algorithm computes a Gröbner basis for the ideal $I = \langle f_1, \ldots, f_t \rangle$.*

Proof Every polynomial that is added to G throughout the algorithm is contained in the ideal I. Since no polynomial is ever removed from G, we retain the property $\langle G \rangle = I$ after each step. If the algorithm terminates, Buchberger's Criterion 9.28 implies that G is a Gröbner basis.

It remains to be shown that the algorithm terminates after finitely many steps. Throughout the algorithm, when $r \neq 0$ we have that $\operatorname{lt}(r) \notin \langle \operatorname{lt}(g) : g \in G \rangle$. Hence,

adding r to the basis G makes the ideal $\langle \mathrm{lt}(g) : g \in G \rangle$ strictly larger. If the algorithm did not terminate, it would yield an infinitely ascending chain of ideals, contradicting Proposition 9.26. □

9.5 Binomial Ideals

A polynomial of the form $x^\alpha - x^{\alpha'} \in K[x_1, \ldots, x_n]$ with $\alpha, \alpha' \in \mathbb{N}^n$ is called a *binomial*, and an ideal that has a generating system consisting of binomials is called a *binomial ideal*. The previously described theories are very simple in the case of binomial ideals. This will be particularly useful in Section 10.6.

Two elementary observations illustrate the uniqueness of the situation. First, we divide two binomials. For this we fix a monomial order $\prec$. If we assume for $\alpha, \alpha', \beta, \beta' \in \mathbb{N}^n$ that $x^\alpha \succ x^{\alpha'}$, $x^\beta \succ x^{\beta'}$ and that x^β divides x^α, then we get

$$x^\alpha - x^{\alpha'} = x^{\alpha-\beta} \cdot \left(x^\beta - x^{\beta'}\right) - x^{\alpha'} + x^{\alpha-\beta+\beta'}. \tag{9.5}$$

In particular,

$$\mathrm{rem}\left(x^\alpha - x^{\alpha'}; x^\beta - x^{\beta'}\right) = x^{\alpha-\beta+\beta'} - x^{\alpha'} \tag{9.6}$$

is a binomial. From this we can deduce the following.

Lemma 9.30 *Let $b_1, \ldots, b_t$ be a family of binomials. Then*:

(a) *for every monomial x^α, $\mathrm{rem}(x^\alpha; b_1, \ldots, b_t)$ is again a monomial*; *and*
(b) *for every binomial $x^\alpha - x^{\alpha'}$, $\mathrm{rem}(x^\alpha - x^{\alpha'}; b_1, \ldots, b_t)$ is again a binomial.*

Proof For the special case $t = 1$ we explicitly showed the second statement in (9.5). The general case $t \geq 2$ follows since we can simply iterate the computation.

The first statement follows analogously. In (9.5) we can alternatively set $\alpha' = -\infty$ with the convention that $x^{-\infty} = 0$. Then $x^\alpha - x^{\alpha'} = x^\alpha$ is a monomial and $\mathrm{rem}(x^\alpha; x^\beta - x^{\beta'}) = x^{\alpha-\beta+\beta'}$. Again, a simple iteration yields the result of the division by several polynomials. □

The second observation is of similar simplicity.

Lemma 9.31 *The S-polynomial of two binomials is a binomial.*

Proof We assume $\alpha, \alpha', \beta, \beta' \in \mathbb{N}^n$ with $x^\alpha \succ x^{\alpha'}$ and $x^\beta \succ x^{\beta'}$. Furthermore, let $x^\mu = \gcd(x^\alpha, x^\beta)$. Then we have the equation

$$\begin{aligned}\mathrm{spol}\left(x^\alpha - x^{\alpha'}, x^\beta - x^{\beta'}\right) &= x^{\beta-\mu} \cdot \left(x^\alpha - x^{\alpha'}\right) - x^{\alpha-\mu} \cdot \left(x^\beta - x^{\beta'}\right) \\ &= x^{\alpha+\beta'-\mu} - x^{\alpha'+\beta-\mu}.\end{aligned}$$

□

When we examine the individual steps of Algorithm 9.4, the most important statement about binomial ideals follows directly from the above two lemmas.

Theorem 9.32 *Given a binomial generating system of a (necessarily binomial) ideal, Buchberger's algorithm computes a Gröbner basis consisting of binomials.*

9.6 Proving a Simple Geometric Fact Using Gröbner Bases

We now demonstrate how Gröbner bases can be employed to prove incidence statements and length relations in elementary geometry.

Theorem 9.33 *The three medians of a (non-degenerate) triangle* $\operatorname{conv}\{a, b, c\} \subseteq \mathbb{R}^2$ *intersect in a single point which we will call* s*. Each of the medians is divided by* s *in the relation* $2:1$.

In high school this theorem is proven directly, e.g. by setting up a system of equations that is obtained by the equations of the involved lines.

Proof Note that we can simplify our task by observing that the statement is independent of translation. That is, we can assume that the vertex a is the origin $(0, 0)$. We can choose a second point, say b, as $(1, 0)$ since the statement is independent of rotation and scaling. We denote the coordinates of the third point by $c = (x, y)$.

We use the notation from Fig. 9.3. The three midpoints of the sides have coordinates

$$p = \left(\frac{x+1}{2}, \frac{y}{2}\right), \qquad q = \left(\frac{x}{2}, \frac{y}{2}\right), \qquad r = \left(\frac{1}{2}, 0\right).$$

Let $s := (u, v)$ be the intersection of $\operatorname{aff}(a, p)$ and $\operatorname{aff}(b, q)$. The fact that s lies on $\operatorname{aff}(a, p)$ is (by comparing the slope of the lines $\operatorname{aff}(a, s)$ and $\operatorname{aff}(a, p)$) equivalent to

$$f_1 := uy - v(x+1) = 0.$$

Analogously, the relation $s \in \operatorname{aff}(b, q)$ is equivalent to

$$f_2 := (u-1)y - v(x-2) = 0.$$

s lies on $\operatorname{aff}(c, r)$ if and only if

$$g_1 := -2(u-x)y - (v-y)(1-2x) = -2uy - (v-y) + 2vx = 0.$$

The point s divides the medians in a $2:1$ relation if and only if the following three equations hold:

$$\begin{aligned}(u, v) = s - a = 2(p - s) &= (x + 1 - 2u, y - 2v),\\ (u-1, v) = s - b = 2(q - s) &= (x - 2u, y - 2v),\\ (u - x, v - y) = s - c = 2(r - s) &= (2u - 1, 2v).\end{aligned}$$

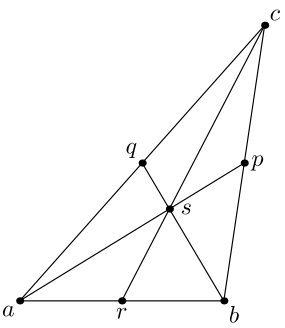

Fig. 9.3 The medians of a triangle meet in a common point s, which is in fact the *center of mass*

This reduces to

$$g_2 := 3u - x - 1 = 0,$$
$$g_3 := 3v - y = 0.$$

We have to respect the condition that our triangle $\text{conv}\{a, b, c\}$ is not degenerate, i.e., $y \neq 0$. This can be expressed by an equation if we introduce another variable z:

$$f_3 := yz - 1 = 0.$$

Now we want to show that

$$f_1 = f_2 = f_3 = 0 \quad \Longrightarrow \quad g_1 = g_2 = g_3 = 0$$

or, in other words, that $\mathrm{V}(f_1, f_2, f_3) \subseteq \mathrm{V}(g_1, g_2, g_3)$. Our proof is complete if we can show the stronger statement

$$g_1, g_2, g_3 \in \langle f_1, f_2, f_3 \rangle.$$

We compute a Gröbner basis of the ideal $I := \langle f_1, f_2, f_3 \rangle \subseteq \mathbb{R}[u, v, x, y, z]$ for, say, the graded reverse lexicographic order $\prec_{\text{grevlex}}$. Using Buchberger's Criterion 9.28 we can verify that

$$G = \{3v - y, 3u - x - 1, yz - 1\}$$

is a $\prec_{\text{grevlex}}$-Gröbner basis of I. Dividing out three candidates g_1, g_2, g_3 by G yields

$$\text{rem}(g_1; G) = \text{rem}(g_2; G) = \text{rem}(g_3; G) = 0,$$

i.e., $g_1, g_2, g_3 \in I$. □

Observe that we didn't assume x, y, u and v to be real numbers. The proof is therefore also valid over $\mathbb{C}$.

9.7 Exercises

Exercise 9.34 Show that, given two univariate polynomials $f, g \in K[x] \setminus \{0\}$, there exist polynomials $a, b \in K[x]$ such that

$$\gcd(f, g) = af + bg.$$

To do so, analyze the Euclidean Algorithm 9.1 and modify it in such a way that the polynomials a and b are computed.

The method described in Exercise 9.34 is called the *extended Euclidean algorithm*.

Exercise 9.35 Let $G = \{g_1, \ldots, g_t\}$ be a Gröbner basis of an ideal $I \subseteq K[x_1, \ldots, x_n]$ with respect to the monomial order $\prec$ and let f, g be polynomials whose difference $f - g$ lies in I. Show that $g = \operatorname{rem}_{\prec}(f; G)$ if and only if no term of g is divisible by one of the leading monomials of $\operatorname{lt}_{\prec}(g_1), \ldots, \operatorname{lt}_{\prec}(g_t)$.

For a Gröbner basis G of an ideal I, we have that every superset G' of G with $G' \subseteq I$ is a Gröbner basis of I. This leads to the question if a given Gröbner basis can have superfluous elements.

Definition 9.36 A Gröbner basis G of an ideal I is called *reduced* if for all $g \in G$:

(a) The leading coefficient is normalized: $\operatorname{lc}_{\prec}(g) = 1$.
(b) No monomial of g lies in $\operatorname{in}_{\prec}(G \setminus \{p\})$.

Exercise 9.37 Show that every non-zero ideal has a unique reduced Gröbner basis for the monomial order $\prec$.

9.8 Remarks

The structure of our presentation is based on the beautiful and comprehensive introduction to the theory of Gröbner bases by Cox, Little and O'Shea [28]. Another text worth reading is the monograph of Adams and Loustaunau [1]. The example of the geometric proof was taken from zur Gathen and Gerhard [97].

Gröbner bases were introduced in the 1960s by Hironaka [60, 61] (who called them "standard bases") and independently by Buchberger in his dissertation [17] in 1965. The term "Gröbner basis" was established by Buchberger in honor of his PhD advisor Wolfgang Gröbner. The exact origin of the "S" in the term "S-polynomial" is not clear. It is sometimes interpreted as "subtraction" or "syzygy".

The statement of the Gordan–Dickson Lemma 9.20 was (re)discovered several times. Its first explicit appearances are usually credited to the German mathematician Paul Gordan [50] and to the American mathematician Leonard Eugene Dickson [35].

If the coefficients of two polynomials f and g are rational numbers, then the computation of the greatest common divisor via the Euclidean algorithm is performed in polynomial time. As stated in Appendix C, polynomial time performance refers to the total length of the input coded as a series of bits. In contrast to this, the ideal membership problem, as well as the problem of computing a Gröbner basis, are intrinsically difficult problems. Mayr and Meyer [76] showed that, with respect to complexity theory, every problem that can be solved with an exponentially large memory can be reduced to an ideal membership problem. Since an exponentially large memory is sufficient, we have that the ideal membership problem is EXPSPACE-*complete*. EXPSPACE-complete problems are significantly more difficult than NP-complete problems: All known algorithms for EXPSPACE-complete problems have at least double-exponential worst-case run-time.

From a practical viewpoint, Buchberger's algorithm can be more efficient in several ways, e.g. by avoiding the computation of superfluous S-polynomials (besides the aforementioned books, see also *Using Algebraic Geometry* by Cox, Little and O'Shea [29] as well as the book by Becker and Weispfenning [11]).

Algorithmic concepts which occur in the solution of problems in the field of *real algebraic geometry* include a variety of methods which are not mentioned in this book. For an overview we refer to the monograph by Basu, Pollack and Roy [10]. Additionally, over the real numbers the question of how to deal with systems of polynomial *inequalities* arises. This leads to *semi-algebraic geometry*. For this, Collins developed an important approach called the *cylindric algebraic decomposition* [25] (for quantifier elimination over real-closed fields). This method is implemented in `QEPCAD` [65].