Computational Commutative Algebra 1

Bearbeitet von
Martin Kreuzer, Lorenzo Robbiano

Weitere Fachgebiete > EDV, Informatik > Datenbanken, Informationssicherheit, Geschäftssoftware > Zeichen- und Zahlendarstellungen

Zu Inhaltsverzeichnis

schnell und portofrei erhältlich bei

beck-shop.de
DIE FACHBUCHHANDLUNG

2. Gröbner Bases

A threat is stronger than its execution.
(Aaron Nimzowitch)

Towards the end of Chapter 1 we encountered Macaulay's Basis Theorem. It says that, given a polynomial ring $P = K[x_1, \ldots, x_n]$ over a field K and a P-submodule M of P^r, one can attack the problem of computing a K-basis of the quotient module P^r/M if one knows $\mathrm{LT}_\sigma(M)$ for some term ordering σ. But we saw that the leading terms of a set of generators of M do not necessarily generate $\mathrm{LT}_\sigma(M)$.

Thus the opening sections of this chapter are variations on the theme that *not all systems of generators of a module are equal. Some are more special than others.* In Section 2.1 we find that the leading terms of a system of non-zero generators $\{g_1, \ldots, g_s\}$ of M generate $\mathrm{LT}_\sigma(M)$ if and only if it is special in the following sense: for every $m \in M \setminus \{0\}$ there exists a representation $m = \sum_{i=1}^s f_i g_i$ with $f_1, \ldots, f_s \in P$ such that $\mathrm{LT}_\sigma(m) \geq_\sigma \mathrm{LT}_\sigma(f_i g_i)$ for all $i = 1, \ldots, s$ such that $f_i \neq 0$.

Then we change our strategy and attack systems of generators from another side. Given a term ordering σ on $\mathbb{T}^n \langle e_1, \ldots, e_r \rangle$, every element $g \in P^r \setminus \{0\}$ can be split as $g = \mathrm{LM}_\sigma(g) - g'$. By looking at this equation modulo $\langle g \rangle$, we can view g as a *rewrite rule*, namely the rule which substitutes $\mathrm{LM}_\sigma(g)$ with the element g' which represents the same residue class. If we have a bunch of non-zero vectors $\{g_1, \ldots, g_s\}$, we get a bunch of rewrite rules. What kind of game can we play with those rules?

Suppose a vector $m \in P^r$ contains a term in its support which is a multiple of $\mathrm{LT}_\sigma(g_i)$ for some $i \in \{1, \ldots, s\}$. Then we can use the rule associated to g_i and rewrite m. The element obtained in this way is congruent to m modulo M. The procedure of moving from one representative of this residue class to another resembles the division algorithm. However, at each point we may have several moves available, and a different order of those moves could lead to a different result. A generating set $\{g_1, \ldots, g_s\}$ of M is special if, no matter which order you choose, you always arrive at the same result. In Section 2.2, we treat rewrite rules and prove the surprising fact that this new kind of specialty is equivalent to the ones described before.

However, the most fundamental motive for looking at special systems of generators is still missing. The notion of a syzygy of a tuple $(g_1, \ldots, g_s)$ is

one of the decisive ideas for successful applications of Computational Commutative Algebra. Using the theory of gradings developed in Section 1.7, we show that every syzygy of $(\mathrm{LT}_\sigma(g_1),\dots,\mathrm{LT}_\sigma(g_s))$ can be lifted to a syzygy of $(g_1,\dots,g_s)$ if and only if $\{g_1,\dots,g_s\}$ has the special properties discussed earlier.

After threatening to do it for a long time, we finally combine all those ideas and introduce Gröbner bases. A Gröbner basis of a submodule M of P^r is a set of generators which is special in one (and therefore all) of the above ways. In Section 2.4 we launch an investigation into their properties and uses by showing that their existence can be viewed as a consequence of Dickson's Lemma. Most applications of Gröbner bases will be treated in Chapter 3 and Volume 2, but some rewards for our careful preparations can be reaped immediately, for instance a proof of Hilbert's Basis Theorem, the notion of normal forms, the submodule membership test, and a new version of Macaulay's Basis Theorem.

Next we put a great emphasis on the derived notion of a reduced Gröbner basis. It has the astonishing property that, given a submodule M of P^r and a term ordering σ, it is a unique system of generators of M satisfying certain natural conditions. We believe that this is one of the most ubiquitous theoretical tools in Computational Commutative Algebra. Just to give the flavour of its importance, we show how one can use it to deduce a seemingly unrelated result about the existence and uniqueness of the field of definition of submodules of P^r.

After all this theory, it is time to explain how one can actually step into action and compute a Gröbner basis of M from a given finite set of generators. The power of our study of syzygies enables us to capture the spirit of Buchberger's Algorithm in Section 2.5. Not only shall we prove and improve its basic procedure, but we shall also finally achieve our goal of effectively computing in residue class modules via Macaulay's Basis Theorem and normal forms.

As sometimes happens in real life, including science, the discovery of a tool which enables us to solve one problem opens the door to many other discoveries. Gröbner bases are certainly one of those tools, but before delving into the realm of their applications, we close the chapter with another one, namely Hilbert's Nullstellensatz. This theorem is one of the milestones in the process of translating algebra into geometry and geometry into algebra and forms the background for many applications in algebraic geometry. Section 2.6 is entirely devoted to its proof, which also uses some pieces of Gröbner basis theory. It highlights the importance of switching from one ground field to a field extension, so that the geometric notion of an affine variety gets its proper perspective.

Once more the chapter closes with an *opening theme.* Besides being a metaphor of life, this end of one struggle already lays the groundwork for successful applications in subsequent chapters.

2.1 Special Generation

> *All animals are equal.*
> *But some animals are more equal than others.*
> (George Orwell)

Let f be a non-zero polynomial and g a non-zero polynomial in the principal ideal generated by f, i.e. let $g = hf$ for a suitable polynomial h. If σ is a monoid ordering on $\mathbb{T}^n$, then $\mathrm{LT}_\sigma(g) = \mathrm{LT}_\sigma(hf) = \mathrm{LT}_\sigma(h)\,\mathrm{LT}_\sigma(f)$. In other words, the leading term of every element in the principal ideal generated by f is in the ideal generated by $\mathrm{LT}_\sigma(f)$.

On the other hand, let us go back for a moment to Example 1.5.5. We saw that for $f = y(x^2 - 1) - x(xy - 1) = x - y$ and $\sigma = \mathtt{DegLex}$ the leading monomials of the two summands cancel out, so that x, the leading term of the result, is smaller than the leading terms of the summands. This shows that some generators have a special behaviour with respect to the leading terms of the elements they generate. More precisely, we see that $x = \mathrm{LT}_\sigma(f) \notin (\mathrm{LT}_\sigma(x^2 - 1), \mathrm{LT}_\sigma(xy - 1)) = (x^2, xy)$. However, if we add in this example the elements guaranteed by Proposition 1.5.6.b, we get another set of generators of the ideal $(x^2 - 1, xy - 1)$ whose leading terms generate the leading term ideal.

This is the prototypical case of the phenomenon that *not all systems of generators of an ideal or module are equal* alluded to in the introduction of this chapter. Some systems of generators have special properties which we want to describe in this and the following sections. Later it will become clear that all of those properties are incarnations of the same concept, namely the concept of Gröbner bases.

As usual, we let K be a field, $n \geq 1$, $P = K[x_1, \ldots, x_n]$ a polynomial ring, $r \geq 1$, and σ a module term ordering on $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$.

Proposition 2.1.1. (Special Generation of Submodules)
Let $M \subseteq P^r$ be a P-submodule, and let $g_1, \ldots, g_s \in P^r \setminus \{0\}$. Then the following conditions are equivalent.

A_1) *For every element $m \in M \setminus \{0\}$, there are $f_1, \ldots, f_s \in P$ such that $m = \sum_{i=1}^s f_i g_i$ and $\mathrm{LT}_\sigma(m) \geq_\sigma \mathrm{LT}_\sigma(f_i g_i)$ for all $i = 1, \ldots, s$ such that $f_i g_i \neq 0$.*

A_2) *For every element $m \in M \setminus \{0\}$, there are $f_1, \ldots, f_s \in P$ such that $m = \sum_{i=1}^s f_i g_i$ and $\mathrm{LT}_\sigma(m) = \max_\sigma\{\mathrm{LT}_\sigma(f_i g_i) \mid i \in \{1, \ldots, s\}, f_i g_i \neq 0\}$.*

Proof. Since Condition A_2) obviously implies A_1), it suffices to prove the reverse direction. The inequality "$\geq_\sigma$" in A_2) follows immediately from A_1). The inequality "$\leq_\sigma$" in A_2) follows from Proposition 1.5.3.a.   □

If $M \subseteq P^r$ is a P-submodule and $g_1, \ldots, g_s \in M \setminus \{0\}$, then Conditions A_1) and A_2) say that $\{g_1, \ldots, g_s\}$ is a special system of generators of M. Using the example mentioned above, we see that it is not true that

Conditions A_1) and A_2) hold for every system of generators of M, because $\mathrm{LT}_\sigma(f) <_\sigma \max_\sigma\{x^2, xy\} \leq_\sigma \max_\sigma\{\mathrm{LT}_\sigma(f_1(x^2-1)), \mathrm{LT}_\sigma(f_2(xy-1))\}$, independent of which elements $f_1, f_2 \in P \setminus \{0\}$ we choose.

It is also interesting to observe that if τ is a term ordering on $\mathbb{T}^n$ and σ is a module term ordering on $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$ which is compatible with τ, then we can expand $\mathrm{LT}_\sigma(f_i g_i) = \mathrm{LT}_\tau(f_i)\,\mathrm{LT}_\sigma(g_i)$ in the above statements.

The intuitive meaning of Conditions A_1) and A_2) is that every element $m \in M \setminus \{0\}$ should have a representation $m = \sum_{i=1}^s f_i g_i$ such that the highest term which occurs in the computation of the right-hand side does not cancel. Consequently, the leading term of m is a multiple of one of the terms $\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)$. Now we examine this last property more closely.

Proposition 2.1.2. (Generation of Leading Term Modules)
Let $M \subseteq P^r$ be a P-submodule and $g_1, \ldots, g_s \in M \setminus \{0\}$. Then the following conditions are equivalent.

B_1) *The set $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$ generates the $\mathbb{T}^n$-monomodule $\mathrm{LT}_\sigma\{M\}$.*
B_2) *The set $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$ generates the P-submodule $\mathrm{LT}_\sigma(M)$ of P^r.*

Proof. Since B_1) implies B_2) by definition, it suffices to show the reverse direction. Let $m \in M \setminus \{0\}$, and let $\mathrm{LT}_\sigma(m) = f_1\,\mathrm{LT}_\sigma(g_1) + \cdots + f_s\,\mathrm{LT}_\sigma(g_s)$ for some polynomials $f_1, \ldots, f_s \in P$. By Proposition 1.5.3.a, the term $\mathrm{LT}_\sigma(m)$ is in the support of one of the vectors $f_1\,\mathrm{LT}_\sigma(g_1), \ldots, f_s\,\mathrm{LT}_\sigma(g_s)$. Thus there is an index $i \in \{1, \ldots, s\}$ and a term $t \in \mathrm{Supp}(f_i)$ such that $\mathrm{LT}_\sigma(m) = t \cdot \mathrm{LT}_\sigma(g_i)$. $\qquad\square$

Finally, we show the first important link between the two properties of special systems of generators which we have described so far.

Proposition 2.1.3. *Let $M \subseteq P^r$ be a P-submodule, and let $g_1, \ldots, g_s$ be non-zero elements of M. Then Conditions A_1), A_2) of Proposition 2.1.1 and Conditions B_1), B_2) of Proposition 2.1.2 are equivalent.*

Proof. Condition A_2) implies B_1) by Proposition 1.5.3.d. Thus we show B_1) $\Rightarrow A_1$). Suppose there exists an element $m \in M \setminus \{0\}$ which cannot be represented in the desired way. By Theorem 1.4.19, there exists such an element m with minimal leading term with respect to σ. By B_1), we have $\mathrm{LT}_\sigma(m) = t \cdot \mathrm{LT}_\sigma(g_i)$ for some $i \in \{1, \ldots, s\}$ and some $t \in \mathbb{T}^n$. Clearly, we have $m - \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_i)}tg_i \neq 0$, since $m = \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_i)}tg_i$ would be a representation satisfying A_1). Therefore we find $\mathrm{LT}_\sigma(m - \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_i)}tg_i) <_\sigma \mathrm{LT}_\sigma(m)$, and the element $m - \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_i)}tg_i \in M \setminus \{0\}$ can be represented as required in A_1). But then also m can be represented as required in A_1), in contradiction with our assumption. $\qquad\square$

Exercise 1. Give an example of a term ordering σ, a module $M \subseteq P^r$, and a set of elements $\{g_1, \ldots, g_s\} \subseteq P^r \setminus M$ which satisfies Conditions A_1) and A_2).

Exercise 2. Let K be a field, let $P = K[x_1, \ldots, x_n]$, let σ be a term ordering on $\mathbb{T}^n$, let $g_1, g_2 \in P$ be two K-linearly independent linear polynomials, and let $i_1, i_2 \in \{1, \ldots, n\}$ be such that $x_{i_1} = \mathrm{LT}_\sigma(g_1)$ and $x_{i_2} = \mathrm{LT}_\sigma(g_2)$. Prove that the following conditions are equivalent.

a) Conditions A_1) and A_2) hold for g_1, g_2.
b) $x_{i_1} \neq x_{i_2}$

Exercise 3. Prove that for $r = 1$ and $\sigma = \mathtt{RevLex}$, Conditions B_1) and B_2) are strictly weaker than A_1) and A_2).

Exercise 4. Let $r = 1$ and $\sigma = \mathtt{DegLex}$. Show that the polynomials $g_1 = x_1 x_2 - x_2$ and $g_2 = x_1^2 - x_2$ do not have properties B_1) and B_2). Find $\mathrm{LT}_{\mathtt{DegLex}}((g_1, g_2))$ and a third polynomial $g_3 \in (g_1, g_2)$ such that $\{g_1, g_2, g_3\}$ satisfies B_1) and B_2).

Exercise 5. Let σ be a term ordering on $\mathbb{T}^2$, and let $g_1 = x^3 - 1$ and $g_2 = y^3 - y$. Prove that $\{g_1, g_2\}$ satisfies B_1) and B_2). Represent $f = x^3 y + xy^3 - x^3 - xy - y + 1$ as a combination of g_1 and g_2 according to Condition A_1).

Tutorial 17: Minimal Polynomials of Algebraic Numbers

In this tutorial we let K be a field and $L = K[x]/(f)$ a finite extension field of K, where $f \in K[x]$ is an irreducible polynomial of degree d. We represent an element $\ell \in L$ as the residue class of a polynomial $g \in K[x]$ and ask the following question.

How can one compute the minimal polynomial of ℓ over K?

Below we shall develop two elementary approaches to this question. In Section 3.6, we shall see a more general method for determining the minimal polynomial of an element in an arbitrary finitely generated K-algebra.

a) Let $\bar{x}$ be the residue class of x in L. Show that $\{1, \bar{x}, \ldots, \bar{x}^{d-1}\}$ is a K-basis of L and conclude that the minimal polynomial of ℓ over K has degree $\leq d$.
b) For $i = 0, \ldots, d$, let $a_i \in K$ be the coefficient of x^i in the minimal polynomial of ℓ over K and $h_i \in K[x]$ the remainder of the division of g^i by f. Prove $a_0 + a_1 h_1 + \cdots + a_d h_d = 0$ and show that this yields a system of d linear equations for $a_0, \ldots, a_d$. Explain how we can use its solution space to answer our question.
c) Implement the method developed in b) in a CoCoA function $\mathtt{LinAlgMP}(\ldots)$ which takes f and g and computes the minimal polynomial of ℓ over K. *Hint:* You may use the CoCoA function $\mathtt{Syz}(\ldots)$ to find the solution space of a system of linear equations.

d) Apply your function `LinAlgMP(...)` to compute the minimal polynomials over $\mathbb{Q}$ of the following algebraic numbers.

   1) $\frac{1}{2} + \frac{i}{2}\sqrt{3}$
   2) $\sqrt[4]{2} + \sqrt{2} + 2$
   3) $(\bar{x}^3 + \bar{x} - 1)/\bar{x}$, where $f = x^5 - x - 2$. (*Hint:* Notice that $\frac{1}{\bar{x}} = \frac{1}{2}\bar{x}^4 - \frac{1}{2}$.)

e) Now we consider the ideal $I = (x_2 - g(x_1), f(x_1)) \subseteq K[x_1, x_2]$. Prove that a polynomial $h \in K[x_2]$ satisfies $h(\ell) = 0$ if and only if $h \in I$. Conclude that the minimal polynomial of ℓ over K is an element of minimal degree in the principal ideal $I \cap K[x_2]$. (*Hint:* Show that $K[x_1, x_2]/I \cong L$.)

f) Prove that $\mathrm{LT}_{\mathtt{Lex}}(I)$ contains a power of x_2. Conclude that, in order to find the minimal polynomial of ℓ over K, it suffices to compute a system of generators of I which satisfies Conditions B_1) and B_2) with respect to $\mathtt{Lex}$.

g) Write a CoCoA function `LexMP(...)` which takes f and g and computes the minimal polynomial of ℓ over K using the method developed in f). (*Hint:* You may assume that the base ring is $\mathbb{Q}[\mathtt{x}[1], \mathtt{x}[2]], \mathtt{Lex}$ and apply the CoCoA function `LT(I)`.) Use your function `LexMP(...)` to check your results in d).

h) Compute the minimal polynomial of $(\bar{x}^3 + \bar{x} - 1)/\bar{x}^5$ over $\mathbb{Q}$ in the case $f = x^7 - x - 1$ using both `LinAlgMP(...)` and `LexMP(...)`. Write down the two polynomials whose leading terms generate $\mathrm{LT}_{\mathtt{Lex}}(I) = (x_1, x_2^7)$. Which of the two methods is in general more efficient? Why?

i) Develop different methods for computing the representation of ℓ^{-1} in the K-basis $\{1, \bar{x}, \ldots, \bar{x}^{d-1}\}$ of L using the following ideas.

   1) Linear Algebra
   2) The Extended Euclidean Algorithm

   Prove the correctness of your methods. Then write two CoCoA functions `LinAlgInv(...)` and `ExtEucInv(...)`, and compare the results in the cases of d).

2.2 Rewrite Rules

All roads lead to Rome.
(Roman Proverb)
All roads do not lead to Rome.
(Slovenian Proverb)

Let us go back to the Division Algorithm discussed in Section 1.6 and try to understand its working more deeply. What is its essence? If we look at Theorem 1.6.4, we see that the event which triggers steps 2) and 3) is the detection of a term in the support of m which is a multiple of one of the leading terms $\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)$. Once such a term is found, the basic operation is to replace it by smaller terms.

A closer look at what happens is provided by the following example. Let $f = x^2y$, $g_1 = x^2 - x + 1$, $g_2 = xy - x - y + 3$, and let $\sigma = \mathtt{DegLex}$. Since x^2y is a multiple of $\mathrm{LT}_\sigma(g_1)$, the first step of the Division Algorithm applied to f and (g_1, g_2) yields $f = y \cdot g_1 + 0 \cdot g_2 + (f - yg_1)$, and we find $f - yg_1 = y(x - 1)$. In this first step we have replaced f by $f - yg_1$. The core of this operation is to take g_1, write it as $x^2 - (x - 1)$, and replace x^2 by $x - 1$. Thus we use g_1 as a rule for replacing its *head*, namely x^2, by its *tail*, namely $x - 1$. Clearly, if a polynomial g_1 is written as $a - b$, we have $a = b \bmod (g_1)$, but here we emphasize the fact that $a = b \bmod (g_1)$ can be viewed as a rule for replacing a by b. In other words, we orient the equality by destroying its symmetry in order to use a polynomial as a *rewrite rule*.

Now we continue with the Division Algorithm. First we observe that $\mathrm{LT}_\sigma(xy - y) = xy$ is a multiple of $\mathrm{LT}_\sigma(g_2)$. So the second step yields $f = y \cdot g_1 + 1 \cdot g_2 + (f - yg_1 - g_2)$ and $f - yg_1 - g_2 = x - 3$. Again we stress the point that the core of this operation is to use g_2 as a rewrite rule in the sense that its leading term xy is replaced by its tail $x + y - 3$. Here the Division Algorithm stops.

Suppose instead that we perform the Division Algorithm with respect to f and (g_2, g_1). Then we get $f = (x+1) \cdot g_2 + 1 \cdot g_1 + (f - (x+1)g_2 - g_1)$, and we see that $f - (x+1)g_2 - g_1 = -x + y - 4$. The algorithm stops and returns an output which is not the same as before.

Summarizing, we can say that the core of the Division Algorithm is to use the elements $g_1, \ldots, g_s$ as rewrite rules. To use g_i as a rewrite rule means to replace the leading term of g_i by the remaining part of it, with the obvious adjustment if g_i is not monic. Of course we should be allowed to use the rewrite rules repeatedly. But in the Division Algorithm the rewrite rules have a well defined hierarchy, i.e. the application of the first rewrite rule is preferred to the second one, and so on. If we have the possibility of using several rewrite rules at a certain point, the Division Algorithm forces us to use the first one in the hierarchy.

What happens if we destroy this hierarchy? Then we are allowed to use at each step any applicable rewrite rule, but the drawback is immediately clear. A look at the previous example convinces us that different possible paths

may lead to different results. So the natural question is whether there are sets of rewrite rules such that all possible paths can be continued until they reach the same result. "Confluence" is the name of this game and the essence of this section, a modern version of the motto *"all roads lead to Rome"*.

And there is a final surprising result. We will discover that for a set of polynomials or vectors of polynomials, being special in the sense of confluence is equivalent to being special in the sense of Conditions A) and B) described in Section 2.1. Thus rewrite rules provide a different aspect of the same phenomenon. Although it is beyond the scope of this book, it turns out that this view is most suitable for generalizations in a number of directions, e.g. to the non-commutative case.

Now it is time to study these ideas in a more technical manner. Let K be a field, $n \neq 1$, $P = K[x_1, \ldots, x_n]$ a polynomial ring, $r \geq 1$, and σ a module term ordering on $\mathbb{T}^n \langle e_1, \ldots, e_r \rangle$.

Definition 2.2.1. Let $g_1, \ldots, g_s \in P^r \setminus \{0\}$ and $G = \{g_1, \ldots, g_s\}$.

a) Let $m_1, m_2 \in P^r$, and suppose there exist a constant $c \in K$, a term $t \in \mathbb{T}^n$, and an index $i \in \{1, \ldots, s\}$ such that $m_2 = m_1 - c\,t g_i$ and $t \cdot \mathrm{LT}_\sigma(g_i) \notin \mathrm{Supp}(m_2)$. Then we say that m_1 **reduces to** m_2 **in one step** using the **rewrite rule** defined by g_i (or simply that m_1 reduces to m_2 in one step using g_i), and we write $m_1 \xrightarrow{g_i} m_2$. The passage from m_1 to m_2 is also called a **reduction step**.

b) The transitive closure of the relations $\xrightarrow{g_1}, \ldots, \xrightarrow{g_s}$ is called the **rewrite relation** defined by G and is denoted by $\xrightarrow{G}$. In other words, for $m_1, m_2 \in P^r$, we let $m_1 \xrightarrow{G} m_2$ if and only if there exist indices $i_1, \ldots, i_t \in \{1, \ldots, s\}$ and elements $m_0', \ldots, m_t' \in P^r$ such that

$$m_1 = m_0' \xrightarrow{g_{i_1}} m_1' \xrightarrow{g_{i_2}} \cdots \xrightarrow{g_{i_t}} m_t' = m_2$$

c) An element $m_1 \in P^r$ with the property that there is no $i \in \{1, \ldots, s\}$ and no $m_2 \in P^r \setminus \{m_1\}$ such that $m_1 \xrightarrow{g_i} m_2$ is called **irreducible** with respect to $\xrightarrow{G}$.

d) The equivalence relation defined by $\xrightarrow{G}$ will be denoted by $\xleftarrow{G}$.

In part a) of this definition, we can choose $c = 0$ and $t \in \mathbb{T}^n$ such that $t \cdot \mathrm{LT}_\sigma(g_i) \notin \mathrm{Supp}(m_1)$. This is called a **trivial reduction**. By using it we see that $m_1 \xrightarrow{g_i} m_1$. In the example mentioned in the introduction, we have for instance $f \xrightarrow{g_1} xy - y$ and $xy - y \xrightarrow{g_2} x - 3$. Thus $f \xrightarrow{G} x - 3$ and $x - 3 \xleftarrow{G} f$ hold, while $x - 3 \xrightarrow{G} f$ is not true, because the leading term of f is larger than x.

Proposition 2.2.2. (Properties of Rewrite Relations)
Let $g_1, \ldots, g_s \in P^r \setminus \{0\}$, and let $G = \{g_1, \ldots, g_s\}$.

a) *If $m_1, m_2 \in P^r$ satisfy $m_1 \xrightarrow{G} m_2$ and $m_2 \xrightarrow{G} m_1$, then $m_1 = m_2$.*

b) *If $m_1, m_2 \in P^r$ satisfy $m_1 \xrightarrow{G} m_2$, and if $t \in \mathbb{T}^n$, then we have $tm_1 \xrightarrow{G} tm_2$.*

c) *Every chain $m_1 \xrightarrow{G} m_2 \xrightarrow{G} \cdots$ such that $m_1, m_2, \ldots \in P^r$ becomes eventually stationary.*

d) *If $m_1, m_2 \in P^r$ satisfy $m_1 \xrightarrow{g_i} m_2$ for $i \in \{1, \ldots, s\}$, and if $m_3 \in P^r$, then there exists an element $m_4 \in P^r$ such that $m_1 + m_3 \xrightarrow{G} m_4$ and $m_2 + m_3 \xrightarrow{G} m_4$.*

e) *If $m_1, m_2, m_3, m_4 \in P^r$ satisfy $m_1 \xleftrightarrow{G} m_2$ and $m_3 \xleftrightarrow{G} m_4$, then we have $m_1 + m_3 \xleftrightarrow{G} m_2 + m_4$.*

f) *If $m_1, m_2 \in P^r$ satisfy $m_1 \xleftrightarrow{G} m_2$, and if $f \in P$, then we have $fm_1 \xleftrightarrow{G} fm_2$.*

g) *For $m \in P^r$, we have $m \xleftrightarrow{G} 0$ if and only if $m \in \langle g_1, \ldots, g_s \rangle$.*

h) *For $m_1, m_2 \in P^r$, we have $m_1 \xleftrightarrow{G} m_2$ if and only if $m_1 - m_2 \in \langle g_1, \ldots, g_s \rangle$.*

Proof. To show claim a), we consider a chain of reduction steps which represents $m_1 \xrightarrow{G} m_2 \xrightarrow{G} m_1$, i.e. a chain $m_1 = m_0' \xrightarrow{g_{i_1}} \cdots \xrightarrow{g_{i_t}} m_t' = m_1$ such that $i_1, \ldots, i_t \in \{1, \ldots, s\}$ and $m_j' = m_2$ for some $j \in \{1, \ldots, t-1\}$. The effect of a reduction step is that a term is replaced by other terms, all of which are smaller with respect to σ. So let te_k with $t \in \mathbb{T}^n$ and $k \in \{1, \ldots, s\}$ be the largest term with respect to σ which is reduced in this chain. This term is not contained in the support of the result anymore, unless each reduction step is trivial, i.e. unless $m_1 = m_2$.

Claim b) holds, since it holds at each reduction step. Thus we prove c) now. Suppose there exist $i_1, i_2, \ldots \in \{1, \ldots, s\}$ and $m_1, m_2, \ldots \in P^r$ such that we have a chain of reduction steps $m_1 \xrightarrow{g_{i_1}} m_2 \xrightarrow{g_{i_2}} \cdots$ which does not become stationary. The first claim is that each m_i must have a term in its support which reduces eventually. Indeed we observe that if this does not happen, it means that starting from m_i the sequence of reductions is actually a sequence of equalities. Therefore there exists a term t_i in $\mathrm{Supp}(m_i)$ which is the largest term with respect to σ which is reduced later in the chain. Then we have $t_1 \geq_\sigma t_2 \geq_\sigma \cdots$, and since every term t_i is reduced eventually, this chain does not become stationary either, in contradiction with Theorem 1.4.19.

For the proof of d), we let $c \in K$, $t \in \mathbb{T}^n$, and $i \in \{1, \ldots, s\}$ be such that $m_2 = m_1 - c\,tg_i$ and $t\,\mathrm{LT}_\sigma(g_i) \notin \mathrm{Supp}(m_2)$. Clearly we may assume $c \neq 0$. We let c' be the coefficient of $t\,\mathrm{LT}_\sigma(g_i)$ in m_3 and distinguish two cases. When $c' = -c$, we have $m_1 + m_3 = m_2 + m_3 + c\,tg_i = m_2 + m_3 - c'tg_i$. Since the coefficient of $t\,\mathrm{LT}_\sigma(g_i)$ in $m_2 + m_3 - c'tg_i$ vanishes, we get

$m_2 + m_3 \xrightarrow{g_i} m_1 + m_3$, and we can choose $m_4 = m_1 + m_3$. When $c' \neq -c$ we define m_4 by

$$m_4 = m_1 + m_3 - (c + c')t g_i = m_2 + m_3 - c' t g_i$$

and obtain the claim, because the coefficient of $t \, \mathrm{LT}_\sigma(g_i)$ vanishes in m_4.

Next, claim e) follows from d), and f) follows from b) and e) by representing f as a sum of monomials. Since h) is an immediate consequence of e) and g), it remains to show g). If $m \xleftrightarrow{G} 0$, we collect the terms used in the various reduction steps and get a representation $m = f_1 g_1 + \cdots + f_s g_s$ with $f_1, \ldots, f_s \in P$. Conversely, given an element $m \in P^r$ with such a representation, it suffices by e) to prove $f_i g_i \xleftrightarrow{G} 0$ for $i = 1, \ldots, s$. This follows from $g_i \xleftrightarrow{G} 0$ and f). $\qquad\square$

Unfortunately, it is not clear how we could use part g) of the above proposition to check whether a given element $m \in P^r$ is contained in the submodule $\langle g_1, \ldots, g_s \rangle$, because we do not know the direction of the reduction steps used in $m \xleftrightarrow{G} 0$. In other words, if we use only reduction steps $m = m_0 \xrightarrow{g_{i_1}} m_1 \xrightarrow{g_{i_2}} \cdots$, we might get stuck at some point with an irreducible element with respect to $\xrightarrow{G}$. The next example shows that this can really happen.

Example 2.2.3. Let $n = 3$, $r = 1$, $G = \{g_1, g_2\}$ with $g_1 = x_1^2 - x_2$ and $g_2 = x_1 x_2 - x_3$, and let σ be the term ordering `DegRevLex`. Then the polynomial $f = x_1^2 x_2 - x_1 x_3$ is contained in the ideal (g_1, g_2), since $f = x_1 g_2$. But if we use the reduction step $f \xrightarrow{g_1} x_2^2 - x_1 x_3$, we arrive at an irreducible element with respect to $\xrightarrow{G}$.

It is also important to notice that if σ is not a term ordering, then claim c) of Proposition 2.2.2 may fail to hold, as the following example shows.

Example 2.2.4. Let $n = 2$, let $r = 1$, let $G = \{g\}$ with $g = x - xy$, and let $\sigma = $ `RevLex`. Then the chain $x \xrightarrow{g} xy \xrightarrow{g} xy^2 \xrightarrow{g} \cdots$ does not become stationary.

After seeing the main properties of rewrite relations, we want to investigate the property of confluence which, as we said before, is crucial for later applications.

Proposition 2.2.5. Let $g_1, \ldots, g_s \in P^r \setminus \{0\}$, let $G = \{g_1, \ldots, g_s\}$, and let $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$. Then the following conditions are equivalent.

C_1) For an element $m \in P^r$, we have $m \xrightarrow{G} 0$ if and only if $m \in M$.

C_2) If $m \in M$ is irreducible with respect to $\xrightarrow{G}$, then we have $m = 0$.

C_3) For every element $m_1 \in P^r$, there is a unique element $m_2 \in P^r$ such that $m_1 \xrightarrow{G} m_2$ and m_2 is irreducible with respect to $\xrightarrow{G}$.

C_4) If $m_1, m_2, m_3 \in P^r$ satisfy $m_1 \xrightarrow{G} m_2$ and $m_1 \xrightarrow{G} m_3$, then there exists an element $m_4 \in P^r$ such that $m_2 \xrightarrow{G} m_4$ and $m_3 \xrightarrow{G} m_4$. (A relation $\xrightarrow{G}$ with this property is called **confluent**.)
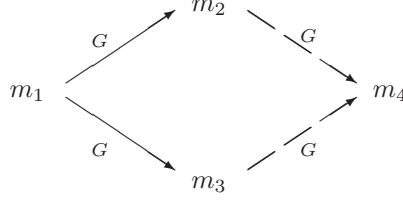
Proof. For the proof of C_1) $\Rightarrow C_2$), we note that if $m \in M$, then C_1) implies $m \xrightarrow{G} 0$. Thus if m is irreducible with respect to $\xrightarrow{G}$, we get $m = 0$. Next we show that C_2) implies C_3). By Proposition 2.2.2.c, there is an element $m_2 \in P^r$ which is irreducible with respect to $\xrightarrow{G}$ and which satisfies $m_1 \xrightarrow{G} m_2$. Suppose $m_2' \in P^r$ is another element with those properties. Then we have $m_2 - m_2' \in M$, since $m_1 \xrightarrow{G} m_2$ and $m_1 \xrightarrow{G} m_2'$. Furthermore, the element $m_2 - m_2'$ is irreducible with respect to $\xrightarrow{G}$, since no term in $\mathrm{Supp}(m_2) \cup \mathrm{Supp}(m_2')$ is a multiple of one of the terms $\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)$. By C_2), we conclude $m_2 = m_2'$.

Now we prove C_3) $\Rightarrow C_4$). By Proposition 2.2.2.c, there are elements $m_2', m_3' \in P^r$ which are irreducible with respect to $\xrightarrow{G}$ and which satisfy $m_2 \xrightarrow{G} m_2'$ as well as $m_3 \xrightarrow{G} m_3'$. From $m_1 \xrightarrow{G} m_2'$, $m_1 \xrightarrow{G} m_3'$, and C_3), we conclude $m_2' = m_3'$. Then the claim follows for $m_4 = m_2' = m_3'$.

Finally, to show C_4) $\Rightarrow C_1$), it suffices, by Proposition 2.2.2.g, to prove $m \xrightarrow{G} 0$ for $m \in M$, where we already know $m \xleftarrow{G} 0$. Let $m_1, \ldots, m_t \in P^r$ be such that $m_1 = m$, $m_t = 0$, and for all $i = 1, \ldots, t-1$ we either have $m_i \xrightarrow{G} m_{i+1}$ or $m_{i+1} \xrightarrow{G} m_i$. Let $\ell \in \{1, \ldots, t-2\}$ be the largest index such that $m_{\ell+1} \xrightarrow{G} m_\ell$. Then we have $m_{\ell+1} \xrightarrow{G} 0$ and $m_{\ell+1} \xrightarrow{G} m_\ell$, and C_4) yields $m_\ell \xrightarrow{G} 0$. If we replace the sequence $m = m_1, \ldots, m_t = 0$ by the shorter sequence $m = m_1, \ldots, m_\ell, 0$, we see that the claim follows by induction. $\square$

The remainder of this section deals with connections between confluent rewrite relations and the previous section. First we prove a useful technical result.

Lemma 2.2.6. *Let $g_1, \ldots, g_s \in P^r \setminus \{0\}$, let $G = \{g_1, \ldots, g_s\}$, and let $M = \langle g_1, \ldots, g_s \rangle$. Assume that an element $m \in M \setminus \{0\}$ satisfies $m \xrightarrow{G} 0$.*

a) *There exist an index $\alpha \in \{1, \ldots, s\}$ and a term $t \in \mathbb{T}^n$ such that $\mathrm{LT}_\sigma(m) = t \cdot \mathrm{LT}_\sigma(g_\alpha)$.*

b) *By collecting all reduction steps in $m \xrightarrow{G} 0$, we get $f'_1, \ldots, f'_s \in P$ such that $m - \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_\alpha)} t\, g_\alpha = \sum_{i=1}^s f'_i g_i$ and such that $\mathrm{LT}_\sigma(m) >_\sigma \mathrm{LT}_\sigma(f'_i g_i)$ for $i = 1, \ldots, s$ with $f'_i g_i \neq 0$.*

c) *If we put $f_i = f'_i$ for $i \in \{1, \ldots, s\} \setminus \{\alpha\}$ and $f_\alpha = f'_\alpha + \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_\alpha)} t$, then we obtain an element $m = \sum_{i=1}^s f_i g_i$ whose leading term satisfies $\mathrm{LT}_\sigma(m) = \max_\sigma \{\mathrm{LT}_\sigma(f_i g_i) \mid i \in \{1, \ldots, s\},\ f_i g_i \neq 0\}$.*

Proof. Claim a) follows immediately from the fact that $\mathrm{LT}_\sigma(m)$ has to be eliminated at one of the reduction steps.

Now we prove b). Let $m_1, \ldots, m_t \in P^r$ be such that $m_1 = m$, $m_t = 0$, and for all $i = 1, \ldots, t - 1$ we have $m_i \xrightarrow{G} m_{i+1}$ using one reduction step. By a), there exists a reduction step where the leading term of m is reduced. This step is unique, since it substitutes $\mathrm{LT}_\sigma(m)$ with smaller terms. So, let $\ell \in \{1, \ldots, t - 1\}$ be such that $m_{\ell+1} = m_\ell - \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_\alpha)} t\, g_\alpha$. Then

$$m - \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_\alpha)} t\, g_\alpha = m - (m_\ell - m_{\ell+1}) = \sum_{i=1}^{\ell-1} (m_i - m_{i+1}) + \sum_{i=\ell+1}^{t-1} (m_i - m_{i+1})$$

is of the form $\sum_{i=1}^s f'_i g_i$. Here the polynomials f'_i are obtained by collecting the elements of type ct appearing in the two sums, where each difference $m_i - m_{i+1}$ is of the form $m_i - m_{i+1} = ct\, g_\beta$ for some $c \in K$, $t \in \mathbb{T}^n$, and $\beta \in \{1, \ldots, s\}$. To conclude the proof it suffices to observe that when we write $m_i - m_{i+1} = ct\, g_\beta$, we get $t\, \mathrm{LT}_\sigma(g_\beta) \leq_\sigma \mathrm{LT}_\sigma(m_i)$ by the definition of a reduction step.

Finally, we see that c) is an immediate consequence of b).   □

Let us examine the claims of this lemma in a concrete case.

Example 2.2.7. Let $g_1 = x^2 - xy$, $g_2 = xy - x - z$, and $g_3 = xy + xz$ be polynomials in $\mathbb{Q}[x, y, z]$, let $G = \{g_1, g_2, g_3\}$, and let $\sigma = \mathtt{DegLex}$. Suppose we want to reduce the polynomial x^3 with respect to the rewrite relation $\xrightarrow{G}$.

One possibility is to apply the following chain of reduction steps.

$$x^3 \xrightarrow{g_1} x^2 y \xrightarrow{g_2} x^2 + xz \xrightarrow{g_1} xy + xz \xrightarrow{g_3} 0$$

As predicted by part a) of the lemma, we find $x^3 = \mathrm{LT}_\sigma(x^3) = x\, \mathrm{LT}_\sigma(g_1)$. Furthermore, by collecting the reduction steps, we get $x^3 - xg_1 = xg_2 + g_1 + g_3$, where $x^3 = \mathrm{LT}_\sigma(x^3)$ is strictly bigger than $x^2 y = \mathrm{LT}_\sigma(xg_2)$, $x^2 = \mathrm{LT}_\sigma(g_1)$, and $xy = \mathrm{LT}_\sigma(g_3)$ with respect to σ.

Finally, to check part c) of the lemma, we also bring xg_1 to the other side and write

$$x^3 = (x + 1)g_1 + xg_2 + g_3$$

Here we have $x^3 = \max_\sigma \{\mathrm{LT}_\sigma((x + 1)g_1), \mathrm{LT}_\sigma(xg_2), \mathrm{LT}_\sigma(g_3)\}$, as claimed.

Unfortunately, the lemma requires that the element reduces to zero. In our case, we could have followed a different sequence of reduction steps, for instance

$$x^3 \xrightarrow{g_1} x^2 y \xrightarrow{g_1} xy^2 \xrightarrow{g_2} xy + yz \xrightarrow{g_2} yz + x + z$$

Here we end up with an element which cannot be reduced further and which is non-zero. By looking at this sequence of instruction steps, we cannot decide whether x^3 satisfies the hypothesis of the lemma.

Both in the introduction to this section and in the previous example we have seen that the property of being confluent is not shared by all rewrite relations. Is there a better way of understanding it? The following proposition gives a somehow unexpected answer.

Proposition 2.2.8. *Let* $g_1, \ldots, g_s \in P^r \setminus \{0\}$, *let* $G = \{g_1, \ldots, g_s\}$, *and let* $M = \langle g_1, \ldots, g_s \rangle$. *Then Conditions* $A_1)$, $A_2)$ *of Proposition 2.1.1 are equivalent with Conditions* $C_1)$, $C_2)$, $C_3)$, *and* $C_4)$ *of Proposition 2.2.5.*

Proof. To prove $A_2) \Rightarrow C_2)$ by contradiction, we suppose that there is an element $m \in M \setminus \{0\}$ which is irreducible with respect to $\xrightarrow{G}$. By Condition $A_2)$, the element m has a representation $m = \sum_{i=1}^{s} f_i g_i$ such that $f_1, \ldots, f_s \in P$ and $\mathrm{LT}_\sigma(m) = \max_\sigma \{\mathrm{LT}_\sigma(f_i g_i) \mid i \in \{1, \ldots, s\}, \ f_i g_i \neq 0\}$. Let $t\,\mathrm{LT}_\sigma(g_i)$ be the term which achieves this maximum. Then the element $m' = m - \frac{\mathrm{LC}_\sigma(m)}{\mathrm{LC}_\sigma(g_i)} t g_i$ satisfies $m \xrightarrow{G} m'$ and $m' \neq m$, a contradiction.

Conversely, $C_1) \Rightarrow A_2)$ follows directly from Lemma 2.2.6. $\qquad \square$

Exercise 1. Let σ be a module term ordering, let $g \in P^r \setminus \{0\}$, and let $G = \{g\}$. Show that the rewrite relation $\xrightarrow{G}$ is confluent.

Exercise 2. Let σ be a module term ordering, and let G be a finite set of terms in P^r. Show that Conditions $C)$ of Proposition 2.2.5 hold for the rewrite relation $\xrightarrow{G}$.

Exercise 3. Give an example of a rewrite relation $\xrightarrow{G}$ which is not confluent.

Exercise 4. Let σ be a monoid ordering on $\mathbb{T}^n$, let $t_1, t_2 \in \mathbb{T}^n$ be terms with $t_1 >_\sigma t_2$, and let $g = t_1 - t_2$. Consider the rewrite relation defined by $G = \{g\}$. (Observe that here we do not assume that σ is a term ordering.) Prove that the following conditions are equivalent.

a) $t_1 \nmid t_2$

b) Every chain $f_1 \xrightarrow{G} f_2 \xrightarrow{G} \cdots$ such that $f_1, f_2, \ldots \in P$ becomes eventually stationary.

Tutorial 18: Algebraic Numbers

In this tutorial, we want to use CoCoA to give some hints about how one can effectively compute in the field $\overline{\mathbb{Q}}$ of algebraic numbers, i.e. the algebraic closure of $\mathbb{Q}$. We shall compute only up to conjugates, i.e. we shall represent an algebraic number by its minimal polynomial over $\mathbb{Q}$. To distinguish between conjugate algebraic numbers, we would also have to provide reasonably good approximations in $\mathbb{Q}[i]$. Furthermore, we shall be content to find *some* polynomial which has a certain algebraic number as one of its zeros. After factoring this polynomial using the CoCoA function `Factor(...)` one could then try to use methods of numerical analysis to find the factor which is the minimal polynomial of the desired algebraic number.

Let $a_1, a_2 \in \overline{\mathbb{Q}}$ be two algebraic numbers represented by irreducible polynomials $g_1, g_2 \in \mathbb{Q}[x]$ of degrees d_1, d_2, respectively.

a) Use Macaulay's Basis Theorem 1.5.7 to show that the residue classes of $\{x_1^i x_2^j \mid 0 \le i < d_1, \ 0 \le j < d_2\}$ form a $\mathbb{Q}$-basis of the $\mathbb{Q}$-algebra $\mathbb{Q}[x_1, x_2]/(g_1(x_1), g_2(x_2))$.

b) Show that one can find a polynomial having $a_1 + a_2$ as one of its zeros in the following way.

   1) Represent the residue classes of the powers $1, x_1 + x_2, (x_1 + x_2)^2, \ldots$ in the basis given in a). Use the rewrite relation $\xrightarrow{G}$ corresponding to $G = \{g_1(x_1), g_2(x_2)\}$ to find such representations.

   2) Continue with step 1) until there is a linear relation between the representations of $1, x_1 + x_2, \ldots, (x_1 + x_2)^d$ for some $d \ge 0$. Then there is a polynomial of degree d which vanishes at $a_1 + a_2$.

c) Write a CoCoA program `AlgSum(...)` which takes the pair (g_1, g_2) and computes a polynomial which vanishes at $a_1 + a_2$ using the algorithm developed in b).

d) Repeat parts b) and c) for the product $a_1 a_2$. In particular, write a CoCoA program `AlgMult(...)` which finds a polynomial which vanishes at $a_1 a_2$.

e) Given an algebraic number $a \in \overline{\mathbb{Q}}$ represented by an irreducible polynomial $g \in \mathbb{Q}[x]$, what is the minimal polynomial of $-a$? Write a CoCoA program `AlgNeg(...)` which takes g and computes the minimal polynomial of $-a$.

f) Given a non-zero algebraic number $a \in \overline{\mathbb{Q}}$ represented by an irreducible polynomial $g \in \mathbb{Q}[x]$, what is the minimal polynomial of $\frac{1}{a}$? Write a CoCoA program `AlgInv(...)` which takes g and computes the minimal polynomial of $\frac{1}{a}$.

g) Apply your CoCoA programs `AlgSum(...)`, `AlgMult(...)`, `AlgNeg(...)`, and `AlgInv(...)` in the following cases. (You'll have to find g_1, g_2 first!)

   1) $a_1 = \sqrt{2}$, $a_2 = \sqrt{3}$
   2) $a_1 = \sqrt[3]{3}$, $a_2 = \frac{1}{2} + \frac{i}{2}\sqrt{3}$
   3) $a_1 = \sqrt{2} + \sqrt{3}$, $a_2 = -i$

2.3 Syzygies

> *Not in the beauty of the words*
> *lies the persuasion of an explanation,*
> *but in their combination ($\sigma\upsilon\zeta\upsilon\gamma\acute{\iota}\alpha$, syzygía).*
> (Dionysius Halicarnassensis)

In the previous two sections we saw a number of conditions satisfied by certain *special* systems of generators of an ideal or module, but not by all of them. Although Proposition 1.5.6 says that such special systems of generators exist always, we do not yet know how to replace a given system of generators with another one having those additional properties.

In this section we change our point of view once more and look at these phenomena from the perspective of syzygies. Despite the exotic name, a syzygy is a very simple object to define. Namely, given a ring R and a tuple of elements $(g_1, \ldots, g_s)$ of an R-module, every tuple $(f_1, \ldots, f_s)$ of elements of R such that $f_1 g_1 + \cdots + f_s g_s = 0$ is called a *syzygy* of $(g_1, \ldots, g_s)$. The introduction of syzygies will eventually achieve several goals. First of all, we see in this section that the failure of Conditions A), B), C) can be better understood in terms of syzygies. Even more important is the fact that in subsequent sections we shall use syzygies to find an algorithmic way to replace a given set of generators, which does not satisfy the conditions, with another one, which does.

What we have said so far suggests the importance of syzygies, and in fact they turn out to be one of the most fundamental algebraic objects. Consequently, the computation of a system of generators for the module of syzygies of a given tuple is one of the central problems in Computational Commutative Algebra. It is also the key to many applications studied in Chapter 3.

But for the moment, let us get down to earth and start digging for the hidden treasures in the land of syzygies. To find the set of all syzygies of a given tuple $\mathcal{G} = (g_1, \ldots, g_s)$ of non-zero polynomial vectors $g_1, \ldots, g_s \in P^r$, where $P = K[x_1, \ldots, x_n]$ is a polynomial ring over a field K, we use the same strategy which brought us rich rewards before: reduce questions about polynomials or vectors of polynomials to questions about their leading terms. Thus we start out by connecting the defining exact sequence of the module of syzygies $\mathrm{Syz}(\mathcal{G})$ of $\mathcal{G}$ and the defining exact sequence of $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$, the syzygy module of $\mathrm{LM}_\sigma(\mathcal{G}) = (\mathrm{LM}_\sigma(g_1), \ldots, \mathrm{LM}_\sigma(g_s))$, via a *fundamental diagram*.

Then we compute an explicit system of generators for $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$, and finally we try to *lift* those syzygies to syzygies of $\mathcal{G}$. This means that we try to find syzygies of $\mathcal{G}$ whose highest homogeneous components (in some sense) are the syzygies generating $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$. When we try to lift the treasures of syzygies in this way, we encounter another unexpected gem: A system of generators of a module has the property that the syzygies of their leading

terms can be lifted if and only if the set of generators satisfies Conditions
$A)$, $B)$, and $C)$!

Definition 2.3.1. Let R be a ring, M an R-module, and $\mathcal{G} = (g_1, \ldots, g_s)$
a tuple of elements of M.

a) A **syzygy** of $\mathcal{G}$ is a tuple $(f_1, \ldots, f_s) \in R^s$ such that $f_1 g_1 + \cdots + f_s g_s = 0$.

b) The set of all syzygies of $\mathcal{G}$ forms an R-module which we call the
(first) syzygy module of $\mathcal{G}$ and which we denote by $\mathrm{Syz}_R(\mathcal{G})$ or by
$\mathrm{Syz}_R(g_1, \ldots, g_s)$. If no confusion can arise, we shall also write $\mathrm{Syz}(\mathcal{G})$ or
$\mathrm{Syz}(g_1, \ldots, g_s)$.

As in the previous sections, we let K be a field, $n \geq 1$, $P = K[x_1, \ldots, x_n]$
a polynomial ring, $r \geq 1$, and σ a module term ordering on $\mathbb{T}^n \langle e_1, \ldots, e_r \rangle$.
Furthermore, we let $g_1, \ldots, g_s \in P^r \setminus \{0\}$, we let $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$, and
we denote the s-tuple $(g_1, \ldots, g_s)$ by $\mathcal{G}$. Then we consider the P-module
P^s with canonical basis $\{\varepsilon_1, \ldots, \varepsilon_s\}$ and the homomorphism $\lambda : P^s \longrightarrow M$
given by $\varepsilon_j \mapsto g_j$ for $j = 1, \ldots, s$. In this situation we can also describe the
syzygy module of $\mathcal{G}$ by $\mathrm{Syz}_P(\mathcal{G}) = \ker(\lambda)$.

The nature of many facts explained in this section is not elementary, so
the inexperienced reader might have some difficulties. For instance, it is clear
that even if we start with an ideal, given by a set of polynomial generators, the
set of their syzygies is a module. So the theory is described in the framework
of modules. Moreover, we shall need to introduce a fine grading on the module
of syzygies in order to detect the correct "highest homogeneous component"
when we follow the above approach.

Since we do not want any reader running away from this book at this
point, we decided to use a didactic tool: *a running example*. This is an
example which we will revisit several times during the section, and which
we will use to make all definitions and constructions as lucid as possible. Let
us start our running example by introducing its basic objects.

Example 2.3.2. Let $n = 3$, let $r = 1$, and let us equip $P = \mathbb{Q}[x, y, z]$
with the degree-lexicographic term ordering σ. Then we consider the ideal
$M = \langle g_1, g_2 \rangle$ generated by $g_1 = x^2 - y^2 - x$ and $g_2 = xy^2 - z^3$, and the pair
$\mathcal{G} = (g_1, g_2)$. Of course the reason why we call this ideal M (and not I) is
to have a better way of comparing the example with the general theory.

The syzygy module of $\mathcal{G}$ is the submodule $\mathrm{Syz}(\mathcal{G}) = \{(f_1, f_2) \in P^2 \mid
f_1 g_1 + f_2 g_2 = 0\} = \{(f_1, f_2) \in P^2 \mid f_1(x^2 - y^2 - x) + f_2(xy^2 - z^3) = 0\}$ of P^2.
Some syzygies of $\mathcal{G}$ are obviously given by $(g_2, -g_1)$ and its multiples, but
are there others?

When we combine the exact sequence $0 \longrightarrow M \longrightarrow P^r \longrightarrow P^r/M \longrightarrow 0$
with the description of $\mathrm{Syz}(\mathcal{G})$ as the kernel of λ, we obtain a long exact
sequence

$$0 \longrightarrow \mathrm{Syz}(\mathcal{G}) \longrightarrow P^s \xrightarrow{\lambda} P^r \longrightarrow P^r/M \longrightarrow 0$$

Now let $N \subseteq P^r$ be the P-submodule of P^r generated by the vectors $\{\mathrm{LM}_\sigma(g_1), \ldots, \mathrm{LM}_\sigma(g_s)\}$, let $\mathrm{LM}_\sigma(\mathcal{G})$ be the tuple $(\mathrm{LM}_\sigma(g_1), \ldots, \mathrm{LM}_\sigma(g_s))$, and let $\Lambda : P^s \longrightarrow N$ denote the homomorphism given by $\varepsilon_j \mapsto \mathrm{LM}_\sigma(g_j)$ for $j = 1, \ldots, s$. Then $\mathrm{Ker}(\Lambda)$ is the syzygy module of $\mathrm{LM}_\sigma(\mathcal{G})$. Consequently, it will be denoted by $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$. We obtain another long exact sequence

$$0 \longrightarrow \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) \longrightarrow P^s \xrightarrow{\Lambda} P^r \longrightarrow P^r/N \longrightarrow 0$$

Recall from Example 1.7.5 that P^r carries a natural structure of a $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-graded module over the $\mathbb{T}^n$-graded ring P. More precisely, we have $(P^r)_{te_i} = K \cdot te_i$ and $P_t = K \cdot t$ for $t \in \mathbb{T}^n$ and $i = 1, \ldots, r$. If we look at the definition of Λ, we see that $\Lambda(\sum_{j=1}^s f_j \varepsilon_j) = \sum_{j=1}^s f_j \mathrm{LM}_\sigma(g_j)$. This fact suggests that we should try to equip the P-module P^s with a $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-grading which is somehow compatible with Λ. By using this approach we find, in the next proposition, that the second sequence carries more structure than the first one.

Proposition 2.3.3. *In the above situation we define*

$$(P^s)_{te_i} = \{\sum_{j=1}^s c_j t_j \varepsilon_j \in P^s \mid c_j = 0 \quad \text{or} \quad t_j \mathrm{LT}_\sigma(g_j) = te_i \ \text{ for } j = 1, \ldots, s\}$$

for all $te_i \in \mathbb{T}^n\langle e_1, \ldots, e_r \rangle$.

a) *We have $P^s = \oplus_{te_i \in \mathbb{T}^n\langle e_1, \ldots, e_r \rangle}(P^s)_{te_i}$. In this way, P^s becomes a $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-graded module over the $\mathbb{T}^n$-graded ring P.*

b) *The map Λ is a homomorphism of $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-graded P-modules. In fact, the sequence $0 \longrightarrow \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) \longrightarrow P^s \xrightarrow{\Lambda} P^r \longrightarrow P^r/N \longrightarrow 0$ consists of homomorphisms of $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-graded modules.*

Proof. In order to show a), we first observe that $(P^s)_{te_i}$ is a group for every $t \in \mathbb{T}^n$ and every $i \in \{1, \ldots, r\}$. Then we verify $P^s = \oplus_{te_i \in \mathbb{T}^n\langle e_1, \ldots, e_r \rangle}(P^s)_{te_i}$. Every element $\sum_{j=1}^s f_j \varepsilon_j \in P^s$ is a sum of elements of the form $c\,t'\varepsilon_j$ with $c \in K \setminus \{0\}$ and $t' \in \mathbb{T}^n$. By definition, we have $c\,t'\varepsilon_j \in (P^s)_{t' \mathrm{LT}_\sigma(g_j)}$, so that it remains to show that the sum is a direct sum.

To this end we notice that, for each $j \in \{1, \ldots, s\}$, there exists at most one term t' in the support of f_j such that $t' \mathrm{LT}_\sigma(g_j) = te_i$. Therefore every term in the support of $\sum_{j=1}^s f_j \varepsilon_j$ is contained precisely in one summand $(P^s)_{te_i}$. Finally, we observe that $t \cdot (P^s)_{t'e_i} \subseteq (P^s)_{tt'e_i}$ shows that our definition actually yields a $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-graded module over the $\mathbb{T}^n$-graded ring P.

Now we prove b). For every $te_i \in \mathbb{T}^n\langle e_1, \ldots, e_r \rangle$ and every element $\sum_{j=1}^s c_j t_j \varepsilon_j \in (P^s)_{te_i}$ we have $\Lambda(\sum_{j=1}^s c_j t_j \varepsilon_j) = \sum_{j=1}^s c_j t_j \mathrm{LM}_\sigma(g_j) = (\sum_{j=1}^s c_j \mathrm{LC}_\sigma(g_j))te_i \in (P^r)_{te_i}$. Therefore the map Λ is a homomorphism of $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-graded modules, and $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) = \mathrm{ker}(\Lambda)$ inherits the structure of a $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-graded module. Since N is a monomial submodule of P^r, it is a $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$-graded submodule by Proposition 1.7.10,

and the canonical homomorphism $P^r \longrightarrow P^r/N$ is a homomorphism of $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$-graded P-modules by Remark 1.7.9. Thus the whole sequence consists of homomorphisms of $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$-graded modules. $\qquad\square$

Example 2.3.2 (continued) In our example we have $\mathrm{LM}_\sigma(\mathcal{G}) = (x^2, xy^2)$. Then for instance $(P^2)_{x^2y^2} = \{(c_1t_1, c_2t_2) \in P^2 \mid c_1t_1x^2, c_2t_2xy^2 \in \mathbb{Q}\cdot x^2y^2\}$. Examples of elements which belong to $(P^2)_{x^2y^2}$ are $(y^2, 0)$, $(-y^2, x)$, and $(\frac{1}{2}y^2, -4x)$.

The intrinsic meaning of the new concepts which we are now going to introduce will be discussed more thoroughly in Volume 2. For the time being, they are only defined with the purpose of better dealing with the $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$-gradings described above.

Definition 2.3.4. Let m be a non-zero element of a $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$-graded module, and let $m = \sum_{\mu \in \mathbb{T}^n\langle e_1,\ldots,e_r\rangle} m_\mu$ be the decomposition of m into its homogeneous components. The term $\max_\sigma\{\mu \in \mathbb{T}^n\langle e_1, \ldots, e_r\rangle \mid m_\mu \neq 0\}$ is called the σ-**degree** of m, and the homogeneous component of m of this degree is called the σ-**leading form** of m.

In the case of the $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$-grading on P^s defined in Proposition 2.3.3, we denote the σ-degree of an element $m \in P^s \setminus \{0\}$ by $\deg_{\sigma,\mathcal{G}}(m)$, and its σ-leading form by $\mathrm{LF}_{\sigma,\mathcal{G}}(m)$. In the next proposition we show how to determine $\deg_{\sigma,\mathcal{G}}(m)$ and $\mathrm{LF}_{\sigma,\mathcal{G}}(m)$ for a non-zero element $m \in P^s$.

Proposition 2.3.5. *Let the module P^s be equipped with the $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$-grading defined above, let $f_1, \ldots, f_s \in P$, and let $m = \sum_{j=1}^s f_j\varepsilon_j \in P^s \setminus \{0\}$.*

a) *We have $\deg_{\sigma,\mathcal{G}}(m) = \max_\sigma\{\mathrm{LT}_\sigma(f_jg_j) \mid j \in \{1, \ldots, s\},\ f_jg_j \neq 0\}$.*

b) *We have $\mathrm{LF}_{\sigma,\mathcal{G}}(m) = \sum_{j=1}^s \bar{f}_j\varepsilon_j$, where*

$$\bar{f}_j = \begin{cases} 0 & \text{if } f_j = 0 \text{ or } \mathrm{LT}_\sigma(f_jg_j) <_\sigma \deg_{\sigma,\mathcal{G}}(m) \\ c_jt_j & \text{if } \mathrm{LT}_\sigma(f_jg_j) = \deg_{\sigma,\mathcal{G}}(m) \text{ and } c_j \in K,\ t_j \in \mathrm{Supp}(f_j) \\ & \text{are such that } \mathrm{LM}_\sigma(f_jg_j) = c_jt_j\,\mathrm{LM}_\sigma(g_j) \end{cases}$$

Proof. Claim a) follows from Proposition 1.5.3 and Definition 2.3.4. To show b), we use that $\deg_{\sigma,\mathcal{G}}(m) = \max_\sigma\{t\,\mathrm{LT}_\sigma(g_j) \mid 1 \leq j \leq s,\ t \in \mathrm{Supp}(f_j)\}$ by a), and this maximum is achieved precisely for the terms described in the formula. $\qquad\square$

Sometimes we are dealing with the case $r = 1$, or we can pick a monoid ordering τ on $\mathbb{T}^n$ such that σ is compatible with τ. In this case, we have $\bar{f}_j = c_jt_j = \mathrm{LM}_\tau(f_j)$ in part b) of this proposition.

Example 2.3.2 (continued) Let us compute both the σ-degree and the σ-leading form of some elements of P^s in our running example. For instance, if we consider the pair $(\frac{1}{2}y^2z, -4xz)$, we have $\deg_{\sigma,\mathcal{G}}(\frac{1}{2}y^2z, -4xz) = x^2y^2z$ and $\mathrm{LF}_{\sigma,\mathcal{G}}(\frac{1}{2}y^2z, -4xz) = (\frac{1}{2}y^2z, -4xz)$. Alternatively, if we start with the pair $(y^2z - x, -4x^2 - y - 3) \in P^2$, we get $\deg_{\sigma,\mathcal{G}}(y^2z - x, -4x^2 - y - 3) = x^3y^2$ and $\mathrm{LF}_{\sigma,\mathcal{G}}(y^2z - x, -4x^2 - y - 3) = (0, -4x^2)$.

Our next goal is to connect the two long exact sequences constructed above. We define a map $\mathrm{LM} : P^r \longrightarrow P^r$, which sends 0 to 0 and m to $\mathrm{LM}_\sigma(m)$ if $m \neq 0$. Analogously we define a map $\mathrm{LF} : P^s \longrightarrow P^s$ which sends 0 to 0 and m to $\mathrm{LF}_{\sigma,\mathcal{G}}(m)$ if $m \neq 0$. In this way we get the following **fundamental diagram**.

$$
\begin{array}{ccccccccc}
0 & \longrightarrow & \mathrm{Syz}(\mathcal{G}) & \longrightarrow & P^s & \overset{\lambda}{\longrightarrow} & P^r & \longrightarrow & P^r/M & \longrightarrow & 0 \\
 & & & & \big\downarrow {\scriptstyle \mathrm{LF}} & & \big\downarrow {\scriptstyle \mathrm{LM}} & & & & \\
0 & \longrightarrow & \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) & \longrightarrow & P^s & \overset{\Lambda}{\longrightarrow} & P^r & \longrightarrow & P^r/N & \longrightarrow & 0
\end{array}
$$

This diagram suggests natural questions, for instance whether the vertical maps are homomorphisms (clearly they aren't), and whether the diagram commutes (it doesn't). A more precise answer to the second question is provided by our next proposition.

Proposition 2.3.6. *In the situation described above, let $m \in P^s \setminus \mathrm{Syz}(\mathcal{G})$.*

a) We have $\mathrm{LT}_\sigma(\lambda(m)) \leq_\sigma \deg_{\sigma,\mathcal{G}}(m)$.
b) We have $\mathrm{LF}(m) \in \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ if and only if $\mathrm{LT}_\sigma(\lambda(m)) <_\sigma \deg_{\sigma,\mathcal{G}}(m)$.
c) We have $\Lambda(\mathrm{LF}(m)) = \mathrm{LM}(\lambda(m))$ if and only if $\mathrm{LT}_\sigma(\lambda(m)) = \deg_{\sigma,\mathcal{G}}(m)$.

Now, let $m \in \mathrm{Syz}(\mathcal{G})$ instead.

d) We have $\mathrm{LF}(m) \in \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$. Therefore the map LF induces a map

$$
\mathrm{LF}\,|_{\mathrm{Syz}(\mathcal{G})} : \mathrm{Syz}(\mathcal{G}) \longrightarrow \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))
$$

which we denote by LF again.

Proof. Claim a) follows from the rules for computing with leading terms (see Proposition 1.5.3) and from Proposition 2.3.5.a. Namely, for the element $m = \sum_{j=1}^{s} f_j \varepsilon_j \in P^s \setminus \{0\}$ we calculate

$$
\mathrm{LT}_\sigma(\lambda(m)) = \mathrm{LT}_\sigma\Big(\sum_{j=1}^{s} f_j g_j\Big) \leq_\sigma \max_\sigma\{\mathrm{LT}_\sigma(f_j g_j) \mid j \in \{1,\ldots,s\},\ f_j g_j \neq 0\}
$$
$$
= \deg_{\sigma,\mathcal{G}}(m)
$$

To prove b), we write $m = \sum_{j=1}^{s} f_j \varepsilon_j \in P^s \setminus \{0\}$ and $\mathrm{LF}_{\sigma,\mathcal{G}}(m) = \sum_{j=1}^{s} \bar{f}_j \varepsilon_j$ as in Proposition 2.3.5. Then $\Lambda(\mathrm{LF}(m)) = \sum_{j=1}^{s} \bar{f}_j \, \mathrm{LM}_\sigma(g_j) = 0$ is equivalent to the vanishing of the coefficient of $\deg_{\sigma,\mathcal{G}}(m)$ in $\sum_{j=1}^{s} f_j g_j$, i.e. it is equivalent to $\mathrm{LT}_\sigma(\lambda(m)) <_\sigma \deg_{\sigma,\mathcal{G}}(m)$.

To prove c), we note that $\mathrm{LT}_\sigma(\lambda(m)) \neq \deg_{\sigma,\mathcal{G}}(m)$ implies by a) and b) that we have $\Lambda(\mathrm{LF}(m)) = 0$. Since $\lambda(m) \neq 0$, we then get $\mathrm{LM}(\lambda(m)) = \mathrm{LM}_\sigma(\lambda(m)) \neq 0 = \Lambda(\mathrm{LF}(m))$. Conversely, if $\mathrm{LT}_\sigma(\lambda(m)) = \deg_{\sigma,\mathcal{G}}(m)$, then $\mathrm{LM}(\lambda(m)) = \mathrm{LM}_\sigma(\sum_{j=1}^{s} f_j g_j) = \sum_{\{j \mid \bar{f}_j \neq 0\}} \mathrm{LM}_\sigma(f_j g_j) = \sum_{j=1}^{s} \bar{f}_j \, \mathrm{LM}_\sigma(g_j) = \Lambda(\mathrm{LF}(m))$.

Finally we show claim d). Let $m = \sum_{j=1}^{s} f_j \varepsilon_j \in \mathrm{Syz}(\mathcal{G}) \setminus \{0\}$. Starting with $\lambda(m) = 0$, we get that the coefficient of $\deg_{\sigma,\mathcal{G}}(m)$ in $\sum_{j=1}^{s} f_j g_j$ vanishes, and hence $\sum_{\{j \mid \bar{f}_j \neq 0\}} \mathrm{LM}_\sigma(f_j g_j) = \sum_{j=1}^{s} \bar{f}_j \mathrm{LM}_\sigma(g_j) = \Lambda(\mathrm{LF}(m)) = 0$.
$\square$

Let us check the claims of this proposition in our running example.

Example 2.3.2 (continued) Recall that $M = \langle g_1, g_2 \rangle$ is the ideal generated by $g_1 = x^2 - y^2 - x$ and $g_2 = xy^2 - z^3$, and that $\sigma = \mathtt{DegLex}$.

a) The element $m = (y^2, -x)$ of P^2 satisfies $\lambda(m) = y^2 g_1 - x g_2 = -y^4 - xy^2 + xz^3$, and thus $\deg_{\sigma,\mathcal{G}}(m) = x^2 y^2$ is not a scalar multiple of $\mathrm{LM}(\lambda(m)) = \mathrm{LM}_\sigma(\lambda(m)) = xz^3$. Going the other way in the fundamental diagram, we calculate $\mathrm{LF}(m) = (y^2, -x)$ and $\Lambda(\mathrm{LF}(m)) = y^2 \mathrm{LM}_\sigma(g_1) - x \mathrm{LM}_\sigma(g_2) = 0$. In particular $\mathrm{LF}(m) \in \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$. Here we have a case where $\mathrm{LT}_\sigma(\lambda(m)) <_\sigma \deg_{\sigma,\mathcal{G}}(m)$ and where $\mathrm{LM}(\lambda(m)) \neq \Lambda(\mathrm{LF}(m))$.

b) The element $m = (x, y)$ of P^2 satisfies $\lambda(m) = xg_1 + yg_2 = x^3 - xy^2 - x^2 + xy^3 - yz^3$, and thus $\deg_{\sigma,\mathcal{G}}(m) = xy^3$ as well as $\mathrm{LM}(\lambda(m)) = xy^3$. On the other hand, we calculate $\mathrm{LF}(m) = (0, y)$ and $\Lambda(\mathrm{LF}(m)) = y \mathrm{LM}_\sigma(g_2) = xy^3$. Here we have a case where $\mathrm{LT}_\sigma(\lambda(m)) = \deg_{\sigma,\mathcal{G}}(m)$ and $\mathrm{LM}(\lambda(m)) = \Lambda(\mathrm{LF}(m))$.

In this example the element $m = (y^2, -x)$ satisfies $m \notin \mathrm{Syz}(\mathcal{G})$, whereas $\mathrm{LF}(m) \in \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$. The fact that $\mathrm{LF}(m)$ is a syzygy of $\mathrm{LM}_\sigma(\mathcal{G})$ may be considered as a sort of first step in the construction of a syzygy of $\mathcal{G}$. Thus a possible approach to our problem of computing a system of generators for $\mathrm{Syz}(\mathcal{G})$ could be to find elements which generate $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ and to "lift" them to elements of $\mathrm{Syz}(\mathcal{G})$ in some way. The remainder of this section is devoted to studying the feasibility of such an approach. As a first step we see how to obtain an explicit finite set of generators of $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$.

Theorem 2.3.7. (Syzygies of Elements of Monomial Modules)
For $j = 1, \ldots, s$, we write $\mathrm{LM}_\sigma(g_j)$ in the form $\mathrm{LM}_\sigma(g_j) = c_j t_j e_{\gamma_j}$ with $c_j \in K$, $t_j \in \mathbb{T}^n$, and $\gamma_j \in \{1, \ldots, r\}$. For all $i, j \in \{1, \ldots, s\}$, we define $t_{ij} = \frac{\mathrm{lcm}(t_i, t_j)}{t_i}$.

*a) For all $i, j \in \{1, \ldots, s\}$ such that $i < j$ and $\gamma_i = \gamma_j$, the element $\sigma_{ij} = \frac{1}{c_i} t_{ij} \varepsilon_i - \frac{1}{c_j} t_{ji} \varepsilon_j \in P^s$ is a syzygy of $\mathrm{LM}_\sigma(\mathcal{G})$, called a **fundamental syzygy**, and is homogeneous of σ-degree $\deg_{\sigma,\mathcal{G}}(\sigma_{ij}) = \mathrm{lcm}(t_i, t_j) e_{\gamma_i}$.*

b) We have

$$\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) = \langle \sigma_{ij} \mid 1 \leq i < j \leq s, \gamma_i = \gamma_j \rangle$$

In particular, $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ is a finitely generated $\mathbb{T}^n \langle e_1, \ldots, e_r \rangle$-graded submodule of P^s.

Proof. To prove a), we note that $\Lambda(\sigma_{ij}) = 0$ and that

$$\deg_{\sigma,\mathcal{G}}(t_{ij} \varepsilon_i) = \frac{\mathrm{lcm}(t_i, t_j)}{t_i} \mathrm{LT}_\sigma(g_i) = \mathrm{lcm}(t_i, t_j) e_{\gamma_i}$$

$$= \mathrm{lcm}(t_i, t_j)e_{\gamma_j} = \tfrac{\mathrm{lcm}(t_i,t_j)}{t_j} \mathrm{LT}_\sigma(g_j) = \deg_{\sigma,\mathcal{G}}(t_{ji}\varepsilon_j)$$

Now we prove b). In view of a), it is clear that $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) \neq 0$ if and only if there exist $i, j \in \{1, \ldots, s\}$ such that $i < j$ and $\gamma_i = \gamma_j$. Since Λ is a homomorphism of $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$-graded P-modules, its kernel is a $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$-graded submodule of P^s and has a homogeneous system of generators. Let us consider one of those homogeneous generators and write it as $m = \sum_{j=1}^s a_j \bar{t}_j \varepsilon_j \in P^s \setminus \{0\}$ with $a_j \in K$ and $\bar{t}_j \in \mathbb{T}^n$. There are an index $\mu \in \{1, \ldots, s\}$ and a term $t \in \mathbb{T}^n$ such that $\bar{t}_j \mathrm{LT}_\sigma(g_j) = te_\mu$ whenever $a_j \neq 0$, which is another way of saying that m is homogeneous and $\deg_{\sigma,\mathcal{G}}(m) = te_\mu$. Next, let $\mathrm{size}(m)$ denote the cardinality of the set $\{i \in \{1, \ldots, s\} \mid a_i \neq 0\}$. Since $\Lambda(m) = 0$, we have $\sum_{j=1}^s a_j c_j = 0$, and since $m \neq 0$, it follows that $\mathrm{size}(m) \geq 2$. Hence there are at least two indices α, β such that $a_\alpha \neq 0$ and $a_\beta \neq 0$. From $t = \bar{t}_\alpha t_\alpha = \bar{t}_\beta t_\beta$ we see that t is a multiple of $\mathrm{lcm}(t_\alpha, t_\beta)$, hence

$$\bar{t}_\alpha = \tfrac{t}{t_\alpha} = \tfrac{t}{\mathrm{lcm}(t_\alpha,t_\beta)}t_{\alpha\beta} \quad \text{and} \quad \bar{t}_\beta = \tfrac{t}{t_\beta} = \tfrac{t}{\mathrm{lcm}(t_\alpha,t_\beta)}t_{\beta\alpha}$$

We deduce that the syzygy $\tfrac{t}{\mathrm{lcm}(t_\alpha,t_\beta)}\sigma_{\alpha\beta}$ has the same σ-degree as m. Moreover we see that if $m' = m - a_\alpha c_\alpha \tfrac{t}{\mathrm{lcm}(t_\alpha,t_\beta)}\sigma_{\alpha\beta}$, then $\mathrm{size}(m') < \mathrm{size}(m)$. An obvious inductive argument concludes the proof. $\qquad\square$

As an immediate consequence of the above theorem, it follows that there are no non-zero syzygies if $\gamma_i \neq \gamma_j$ for all $1 \leq i < j \leq s$. This observation is amplified in Exercise 7. Clearly, the proof of the theorem can be used as an algorithm for computing the representation of an element of $\mathrm{Syz}(\mathrm{LM}(\mathcal{G}))$ in terms of the generators σ_{ij}.

Example 2.3.8. Let $n = 3$, let $r = 1$, and let us equip $P = \mathbb{Q}[x,y,z]$ with the term ordering $\sigma = \mathtt{DegRevLex}$. We consider the vector $\mathcal{G} = (g_1, g_2, g_3)$, where $g_1 = 4x^2y - x$, $g_2 = 3xy^3$, and $g_3 = yz - x - 1$. Then we have $\mathrm{LM}_\sigma(\mathcal{G}) = (4x^2y, 3xy^3, yz)$.

The module element $m = (y^2z, -2xz, 2x^2y^2) = y^2z\varepsilon_1 - 2xz\varepsilon_2 + 2x^2y^2\varepsilon_3$ is contained in $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$, since $y^2z \cdot 4x^2y - 2xz \cdot 3xy^3 + 2x^2y^2 \cdot yz = 0$. Moreover, the element m is homogeneous of σ-degree $\deg_{\sigma,\mathcal{G}}(m) = x^2y^3z$, and we have $\mathrm{size}(m) = 3$.

According to Theorem 2.3.7, we should be able to express m as a combination of σ_{12}, σ_{13}, and σ_{23}. Using the notation of the proof of the theorem, we see that a_1, a_2, a_3 are different from zero. So, let $\alpha = 1$ and $\beta = 2$. We get $\mathrm{lcm}(t_1, t_2) = x^2y^3$, and therefore $\tfrac{x^2y^3z}{\mathrm{lcm}(t_1,t_2)} = \tfrac{x^2y^3z}{x^2y^3} = z$. Thus we form the element $m' = m - a_1c_1 \tfrac{x^2y^3z}{\mathrm{lcm}(t_1,t_2)}\sigma_{12} = m - 4z\sigma_{12}$. Now we compute $\sigma_{12} = \tfrac{1}{4}y^2\varepsilon_1 - \tfrac{1}{3}x\varepsilon_2 = (\tfrac{1}{4}y^2, -\tfrac{1}{3}x, 0)$ and get $m' = (y^2z, -2xz, 2x^2y^2) - 4z(\tfrac{1}{4}y^2, -\tfrac{1}{3}x, 0) = (0, -\tfrac{2}{3}xz, 2x^2y^2)$. Finally, we determine $\sigma_{23} = \tfrac{1}{3}z\varepsilon_2 - xy^2\varepsilon_3 = (0, \tfrac{1}{3}z, -xy^2)$. It is clear that $(0, -\tfrac{2}{3}xz, 2x^2y^2) = -2x\sigma_{23}$. In conclusion, we find the desired representation $m = 4z\sigma_{12} - 2x\sigma_{23}$.

The next steps in our program are to give a meaning to the process of "lifting" a syzygy of $\mathrm{LM}_\sigma(\mathcal{G})$ to a syzygy of $\mathcal{G}$, and then to study whether such liftings can always be found.

Definition 2.3.9. An element $m \in P^s$ is called a **lifting** of an element $\overline{m} \in P^s$ if we have $\mathrm{LF}(m) = \overline{m}$.

Proposition 2.3.10. *The following conditions are equivalent.*

D_1) *Every homogeneous element of* $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ *has a lifting in* $\mathrm{Syz}(\mathcal{G})$.

D_2) *There exists a homogeneous system of generators of* $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ *consisting entirely of elements which have a lifting in* $\mathrm{Syz}(\mathcal{G})$.

D_3) *There exists a finite homogeneous system of generators of* $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ *consisting entirely of elements which have a lifting in* $\mathrm{Syz}(\mathcal{G})$.

Proof. Since $D_1) \Rightarrow D_3)$ as an immediate consequence of Theorem 2.3.7, and since $D_3) \Rightarrow D_2)$ holds trivially, it suffices to prove that $D_1)$ follows from $D_2)$. Let I be a set, let $\{\overline{m}_i\}_{i \in I}$ be a homogeneous system of generators of $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ indexed over I, and let $m_i \in \mathrm{Syz}(\mathcal{G})$ be a lifting of $\overline{m}_i$ for every $i \in I$. Given a homogeneous element $\overline{m} \in \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) \setminus \{0\}$, there exists a natural number h such that we have $\overline{m} = \sum_{j=1}^h c_j t_j \overline{m}_{i_j}$ with $c_j \in K \setminus \{0\}$, with $t_j \in \mathbb{T}^n$, and with $i_j \in I$ for $j = 1, \ldots, h$. Clearly, we may assume $\deg_{\sigma,\mathcal{G}}(t_j \overline{m}_{i_j}) = \deg_{\sigma,\mathcal{G}}(\overline{m})$ for $j = 1, \ldots, h$. From the fact that $\mathrm{LF}(t_j m_{i_j}) = t_j \overline{m}_{i_j}$ we conclude $\deg_{\sigma,\mathcal{G}}(t_j m_{i_j}) = \deg_{\sigma,\mathcal{G}}(\overline{m})$. This, in turn, implies $\mathrm{LF}(\sum_{j=1}^h c_j t_j m_{i_j}) = \sum_{j=1}^h c_j t_j \overline{m}_{i_j} = \overline{m}$, which concludes the proof. $\square$

If we want to find all elements of $\mathrm{Syz}(\mathcal{G})$ using this process of lifting, we need to ascertain that there exists a system of generators of $\mathrm{Syz}(\mathcal{G})$ consisting of liftings. This is achieved by the following proposition whose proof demonstrates once more the power of term orderings.

Proposition 2.3.11. *Let* $\{\overline{m}_1, \ldots, \overline{m}_t\}$ *be a homogeneous system of generators of the module* $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$, *and let* $m_1, \ldots, m_t \in \mathrm{Syz}(\mathcal{G})$ *be elements such that* $\mathrm{LF}(m_i) = \overline{m}_i$ *for* $i = 1, \ldots, t$. *Then* $\{m_1, \ldots, m_t\}$ *is a system of generators of* $\mathrm{Syz}(\mathcal{G})$.

Proof. For contradiction we assume that the subset S of $\mathrm{Syz}(\mathcal{G})$ of syzygies which are not generated by $\{m_1, \ldots, m_t\}$ is not empty. By the fundamental property of term orderings (see Theorem 1.4.19), there exists $m \in S$ with minimal $\deg_{\sigma,\mathcal{G}}$. Then there exists a natural number h such that we have $\mathrm{LF}(m) = \sum_{j=1}^h c_j t_j \overline{m}_{i_j}$ with $c_j \in K \setminus \{0\}$, with $t_j \in \mathbb{T}^n$, and with $i_j \in \{1, \ldots, t\}$ for $j = 1, \ldots, h$. The element $m' = m - \sum_{i=1}^h c_i t_i m_{i_j}$ satisfies either $m' = 0$ or $\deg_{\sigma,\mathcal{G}}(m') <_\sigma \deg_{\sigma,\mathcal{G}}(m)$. In both cases we get a contradiction, and the proof is complete. $\square$

The final proposition in this section is the gem we promised in the introduction.

Proposition 2.3.12. *Let $g_1, \ldots, g_s \in P^r \setminus \{0\}$ and $M = \langle g_1, \ldots, g_s \rangle$. Then Conditions A_1), A_2) of Proposition 2.1.1 and Conditions D_1), D_2), D_3) of Proposition 2.3.10 are equivalent.*

Proof. First we show that Condition A_2) implies D_1). Let $m = \sum_{j=1}^{s} f_j \varepsilon_j$ be a non-zero homogeneous element of $\operatorname{Syz}(\operatorname{LM}_\sigma(\mathcal{G}))$. We may suppose that $\lambda(m) \neq 0$, since in case $\lambda(m) = 0$ we have $m \in \operatorname{Syz}(\mathcal{G})$ and $\operatorname{LF}(m) = m$, i.e. the element m is a lifting of itself. By Condition A_2), the element $\lambda(m)$ has a representation $\lambda(m) = \sum_{i=1}^{s} h_i g_i$ with polynomials $h_1, \ldots, h_s \in P$ such that $\operatorname{LT}_\sigma(\lambda(m)) = \max_\sigma \{\operatorname{LT}_\sigma(h_i g_i) \mid i \in \{1, \ldots, s\}, \ h_i g_i \neq 0\}$. Now we consider the element $h = \sum_{j=1}^{s} h_j \varepsilon_j \in P^s$. We have $m - h \in \operatorname{Syz}(\mathcal{G})$ and $\operatorname{LT}_\sigma(\lambda(m)) = \operatorname{LT}_\sigma(\lambda(h)) = \deg_{\sigma, \mathcal{G}}(h)$. On the other hand, since $\operatorname{LF}(m) = m$ and $\Lambda(\operatorname{LF}(m)) = 0$, Proposition 2.3.6.b yields $\operatorname{LT}_\sigma(\lambda(m)) <_\sigma \deg_{\sigma, \mathcal{G}}(m)$. Altogether, we get $\deg_{\sigma, \mathcal{G}}(m) >_\sigma \deg_{\sigma, \mathcal{G}}(h)$ and $\operatorname{LF}(m - h) = \operatorname{LF}(m) = m$. Thus the element $m - h$ is a lifting of m.

Now let us show the reverse implication. We assume for contradiction that there exists an element $v \in M \setminus \{0\}$ which cannot be represented as requested by Condition A_2). We observe that if $v = \sum_{i=1}^{s} f_i g_i$ for some polynomials $f_1, \ldots, f_s \in P$ and if $m = \sum_{j=1}^{s} f_j \varepsilon_j$, then we have $v = \lambda(m)$. In other words, the element m is a preimage of v under λ. By the fundamental property of term orderings (see Theorem 1.4.19), we know that among all preimages of v under λ, there exists one preimage m with minimal $\deg_{\sigma, \mathcal{G}}(m)$. We cannot have $\deg_{\sigma, \mathcal{G}}(m) = \operatorname{LT}_\sigma(v)$, because otherwise the representation $v = \sum_{i=1}^{s} f_i g_i$ is already of the form required by Condition A_2). Therefore Proposition 2.3.6.a shows that we must have $\operatorname{LT}_\sigma(v) <_\sigma \deg_{\sigma, \mathcal{G}}(m)$. Next, Proposition 2.3.6.b yields $\operatorname{LF}(m) \in \operatorname{Syz}(\operatorname{LM}_\sigma(\mathcal{G}))$. Thus Condition D_1) gives us an element $m' = \sum_{j=1}^{s} f'_j \varepsilon_j \in \operatorname{Syz}(\mathcal{G})$ such that $\operatorname{LF}(m') = \operatorname{LF}(m)$. In particular, this means that $\deg_{\sigma, \mathcal{G}}(m - m') <_\sigma \deg_{\sigma, \mathcal{G}}(m)$ and $\lambda(m - m') = \lambda(m) = v$, which contradicts the minimality of the σ-degree of m. $\square$

Exercise 1. Find a term ordering σ and elements $g_1, \ldots, g_s \in P^r \setminus \{0\}$ which generate a submodule $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$ such that Conditions D_1), D_2), and D_3) are not satisfied.

Exercise 2. Find a 2×3-matrix over $P = K[x, y, z]$ whose associated ideal of 2×2-minors is generated by $\{x^2 - y, xy - z, y^2 - xz\}$. By adding suitable rows to this matrix, show how one can produce non-trivial syzygies of the triple $\mathcal{G} = (x^2 - y, xy - z, y^2 - xz)$.

Exercise 3. In the case $n = 2$, $P = \mathbb{Q}[x, y]$, $r = 2$, compute a system of generators of the syzygy module of the tuple $\mathcal{G} = ((xy + y, x), (x - y, y), (x, x + y), (-x, y))$ by hand.

Exercise 4. Let $P = K[x, y, z]$ be a polynomial ring over a field K, let $r = 1$, and let $\mathcal{G} = (x, y, z)$. Compute the syzygy module of a set of generators of $\operatorname{Syz}_P(\mathcal{G})$.

Exercise 5. Give a direct proof for the fact that Condition D_1) of Proposition 2.3.10 implies Condition B_2) of Proposition 2.1.2.
Hint: If $m \in M \setminus \{0\}$ has a leading term outside N, pick a preimage of m under λ of smallest σ-degree and look at the fundamental diagram.

Exercise 6. Let $g_1, \ldots, g_s \in P^r \setminus \{0\}$, let $M = \langle g_1, \ldots, g_s \rangle$, and let $\mathcal{G}$ be the tuple $(g_1, \ldots, g_s)$.
  a) Prove that $\mathrm{Syz}(\mathcal{G}) = 0$ if and only if M is a free P-module with basis $\{g_1, \ldots, g_s\}$.
  b) Let $s = 3$, let $n = 3$, let $r = 2$, let $g_1 = (x^2, x - y)$, let $g_2 = (0, y)$, and let $g_3 = (xy, z)$. Then show that $\mathrm{Syz}(\mathcal{G}) \neq 0$.

Exercise 7. Let $g_1, \ldots, g_s \in P^r \setminus \{0\}$, let $M = \langle g_1, \ldots, g_s \rangle$, let $\mathcal{G}$ be the s-tuple $(g_1, \ldots, g_s)$, let σ be a module term ordering on $\mathbb{T}^n \langle e_1, \ldots, e_r \rangle$, and let $\mathrm{LT}_\sigma(g_i) = t_i e_{\gamma_i}$ with $t_i \in \mathbb{T}^n$ and $\gamma_i \in \{1, \ldots, r\}$ for $i = 1, \ldots, s$.
  a) Prove that M is a free P-module if $\gamma_i \neq \gamma_j$ for all $i \neq j$.
  b) Deduce that the submodule of P^3 generated by the set of vectors $\{(x, y - z, x), (z, y^2 - x, x), (z^2 - y + 1, y^2 - x, x - 3)\}$ is free.

Tutorial 19: Syzygies of Elements of Monomial Modules

Let K be a field, $n \geq 1$, $P = K[x_1, \ldots, x_n]$, $r \geq 1$, and $M \subseteq P^r$ a monomial submodule generated by $\{t_1 e_{\gamma_1}, \ldots, t_s e_{\gamma_s}\}$, where $t_1, \ldots, t_s \in \mathbb{T}^n$ and $\gamma_1, \ldots, \gamma_s \in \{1, \ldots, r\}$.

 a) Use Theorem 2.3.7 to give an explicit system of generators of the syzygy module of $(t_1 e_{\gamma_1}, \ldots, t_s e_{\gamma_s})$. Write a CoCoA function `MonomialSyz(...)` which takes a system of generators of a monomial module M as above and computes its first syzygy module.
 b) Show by example that the system of generators of the syzygy module given in a) is in general not minimal, even if $\{t_1 e_{\gamma_1}, \ldots, t_s e_{\gamma_s}\}$ is minimal.
 c) Apply your function `MonomialSyz(...)` to compute the syzygy modules of the following tuples.
     1) $(x^{34} y^7, x^{23} y^{19}) \in \mathbb{Q}[x, y]^2$
     2) $(x, y, z) \in \mathbb{Q}[x, y, z]^3$
     3) $(xy, yz, xz) \in \mathbb{Q}[x, y, z]^3$
     4) $(xe_1, ye_1, ye_2, ze_2, xe_3, ze_3) \in (\mathbb{Q}[x, y, z]^3)^6$
 d) Show that if $r = 1$, $1 \leq i < j < k \leq s$, and t_k divides $\mathrm{lcm}(t_i, t_j)$, then the fundamental syzygy σ_{ij} (as defined in Theorem 2.3.7) is in the module generated by σ_{ik} and σ_{jk}.
 e) Write an improved version `MonomialIdealSyz(...)` of your program from a) which works for systems of generators of monomial ideals and takes the optimization of part d) into account.
 f) Apply the function `MonomialIdealSyz(...)` in the appropriate cases of c). Each time, try to determine whether the computed system of generators of the syzygy module is minimal.

Tutorial 20: Lifting of Syzygies

In this tutorial we shall try to program the lifting of syzygies discussed in the last part of the current section. As usual, let K be a field, let $n \geq 1$, let $P = K[x_1, \ldots, x_n]$ be a polynomial ring, let $r \geq 1$, let σ be a module term ordering on $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$, let $\mathcal{G} = (g_1, \ldots, g_s) \in (P^r)^s$ be a tuple of non-zero vectors, and let $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$. We assume that Conditions D_1), D_2), and D_3) are satisfied. For $i = 1, \ldots, s$, we write $\mathrm{LT}_\sigma(g_i) = t_i e_{\gamma_i}$ with $t_i \in \mathbb{T}^n$ and $1 \leq \gamma_i \leq r$, and, for $i, j \in \{1, \ldots, s\}$ such that $\gamma_i = \gamma_j$, we let $t_{ij} = \mathrm{lcm}(t_i, t_j)/t_i = t_j/\gcd(t_i, t_j)$.

a) Show that, for $1 \leq i < j \leq s$ such that $\gamma_i = \gamma_j$, there are representations

$$\mathrm{LC}_\sigma(g_i)^{-1} t_{ij} g_i - \mathrm{LC}_\sigma(g_j)^{-1} t_{ji} g_j = \sum_{k=1}^{s} f_{ijk} g_k$$

where $f_{ij1}, \ldots, f_{ijs} \in P$, and where $\mathrm{LT}_\sigma(f_{ijk} g_k) <_\sigma \mathrm{LT}_\sigma(t_{ij} g_i)$ for all $k \in \{1, \ldots, s\}$ such that $f_{ijk} \neq 0$.

b) Let $\{\sigma_{ij} \mid 1 \leq i < j \leq s, \ \gamma_i = \gamma_j\}$ be the system of generators of the kernel of the map $\Lambda : P^s \longrightarrow P^r$, $e_i \longmapsto \mathrm{LM}_\sigma(g_i)$ introduced in Theorem 2.3.7. Prove that the elements $s_{ij} = \sigma_{ij} - \sum_{k=1}^{s} f_{ijk} \varepsilon_k$ are liftings of σ_{ij} for all i, j as above.

c) Conclude that the set $\{s_{ij} \mid 1 \leq i < j \leq s, \ \gamma_i = \gamma_j\}$ is a system of generators of the syzygy module $\mathrm{Syz}(\mathcal{G})$.

d) Using the program $\mathtt{Division}(\ldots)$ from Tutorial 14 as a subfunction, write a CoCoA program $\mathtt{StdRepr}(\ldots)$ which takes the tuple $\mathcal{G}$ and indices i, j as above and computes a list of polynomials $[f_{ij1}, \ldots, f_{ijs}]$ corresponding to the representation in a).

e) Using the program $\mathtt{MonomialSyz}(\ldots)$ from Tutorial 19 and $\mathtt{StdRepr}(\ldots)$ as subfunctions, write a CoCoA program $\mathtt{LiftSyz}(\ldots)$ which takes the tuple $\mathcal{G}$ and computes the list of all syzygies s_{ij} as in b).

f) Using the module term ordering $\mathtt{DegRevLexPos}$, compute the lists of all syzygies σ_{ij} and all s_{ij} in the following cases.

   1) $\mathcal{G} = (x_1^2 - x_2, x_2^2 - x_3, x_3^2 - x_1) \in \mathbb{Q}[x_1, x_2, x_3]^3$
   2) $\mathcal{G} = (x_1 e_1, x_2 e_1, x_3 e_2, x_1 e_3) \in (\mathbb{Q}[x_1, x_2, x_3]^3)^4$
   3) $\mathcal{G} = (x_1 x_4 - x_2 x_3, x_1 x_3^2 - x_2^2 x_4, x_1^2 x_3 - x_2^3, x_2 x_4^2 - x_3^3) \in \mathbb{Q}[x_1, x_2, x_3, x_4]^4$

2.4 Gröbner Bases of Ideals and Modules

> *The motifs of a combination, in themselves simple,*
> *are often interwoven with each other. [...]*
> *The idea which links the motifs is artistic,*
> *it creates something that had never before been there.*
> (Emanuel Lasker)

In the previous three sections we saw many conditions arising from a number of different motifs, and all of them turned out to be equivalent. Whenever such a phenomenon shows up, it is clear that something very important is going on: there must be some fundamental idea behind the scene which needs to be brought to center stage. In our case it is the notion of a *Gröbner basis*. It is one of those rare notions in the history of modern mathematics which was able to deviate the main stream of events. It became a fundamental tool, both for its theoretical and practical consequences.

The section opens by linking the different motifs studied before through the idea of a Gröbner basis. The natural search for the existence of such objects leads to a fairly easy positive answer (see Proposition 2.4.3). Part of this existence result is Hilbert's Basis Theorem 2.4.6 for finitely generated modules over finitely generated K-algebras. Of course it is not necessary to develop the theory of Gröbner bases to achieve that result, but we decided to include it here as an application in order to highlight the theoretical power of Gröbner bases.

Then we become more ambitious and try to solve the problem of computing in residue class modules. Using a Gröbner basis, we define the *normal form* of an element with respect to a submodule and show that it is independent of the Gröbner basis chosen. It agrees with the normal remainder given by the Division Algorithm 1.6.4. Thus it is a unique representative of the residue class of the given element which can be computed by performing the Division Algorithm with respect to any Gröbner basis of the submodule. Consequently, we get a *submodule membership test*, also called *ideal membership test* when $r = 1$, and a new formulation of Macaulay's Basis Theorem.

But what is really striking is another form of uniqueness. In our opinion, it is one of the most important theoretical results of this theory. Given a Gröbner basis of a submodule M of P^r, we can modify its elements in such a way that we get another Gröbner basis with the extra properties of being monic, minimal, and interreduced. Surprisingly, this *reduced Gröbner basis* of M depends only on the module and the chosen term ordering. As we shall see, the possibility of representing a submodule by a unique system of generators has numerous theoretical and practical applications. To give a first support to this claim, we devote the last part of this section to the proof of the existence and uniqueness of the *field of definition* of a given submodule M, i.e. a minimal subfield of K which contains the coefficients of some system of generators of M.

Now we start the main part of this section by recalling that, as usual, we let K be a field, $n \geq 1$, $P = K[x_1, \ldots, x_n]$ a polynomial ring, $r \geq 1$, and σ a module term ordering on $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$. In the following theorem we collect all the conditions studied in the previous sections.

Theorem 2.4.1. (Characterization of Gröbner Bases)
For a set of elements $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$ which generates a submodule $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$, let $\xrightarrow{G}$ be the rewrite rule defined by G, let $\mathcal{G}$ be the tuple $(g_1, \ldots, g_s)$, let λ be the map $\lambda : P^s \longrightarrow P^r$ defined by $\varepsilon_i \mapsto g_i$, and let $\Lambda : P^s \longrightarrow P^r$ be the map defined by $\varepsilon_i \mapsto \mathrm{LM}_\sigma(g_i)$. Then the following conditions are equivalent.

A_1) *For every element $m \in M \setminus \{0\}$, there are $f_1, \ldots, f_s \in P$ such that $m = \sum_{i=1}^s f_i g_i$ and $\mathrm{LT}_\sigma(m) \geq_\sigma \mathrm{LT}_\sigma(f_i g_i)$ for all $i = 1, \ldots, s$ such that $f_i g_i \neq 0$, i.e. such that $\mathrm{LT}_\sigma(m) \geq_\sigma \deg_{\sigma, \mathcal{G}}(\sum_{i=1}^s f_i \varepsilon_i)$.*

A_2) *For every element $m \in M \setminus \{0\}$, there are $f_1, \ldots, f_s \in P$ such that $m = \sum_{i=1}^s f_i g_i$ and $\mathrm{LT}_\sigma(m) = \max_\sigma\{\mathrm{LT}_\sigma(f_i g_i) \mid i \in \{1, \ldots, s\}, f_i g_i \neq 0\}$, i.e. such that $\mathrm{LT}_\sigma(m) = \deg_{\sigma, \mathcal{G}}(\sum_{i=1}^s f_i \varepsilon_i)$.*

B_1) *The set $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$ generates the $\mathbb{T}^n$-monomodule $\mathrm{LT}_\sigma\{M\}$.*

B_2) *The set $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$ generates the P-submodule $\mathrm{LT}_\sigma(M)$ of P^r.*

C_1) *For an element $m \in P^r$, we have $m \xrightarrow{G} 0$ if and only if $m \in M$.*

C_2) *If $m \in M$ is irreducible with respect to $\xrightarrow{G}$, then we have $m = 0$.*

C_3) *For every element $m_1 \in P^r$, there is a unique element $m_2 \in P^r$ such that $m_1 \xrightarrow{G} m_2$ and m_2 is irreducible with respect to $\xrightarrow{G}$.*

C_4) *If $m_1, m_2, m_3 \in P^r$ satisfy $m_1 \xrightarrow{G} m_2$ and $m_1 \xrightarrow{G} m_3$, then there exists an element $m_4 \in P^r$ such that $m_2 \xrightarrow{G} m_4$ and $m_3 \xrightarrow{G} m_4$.*

D_1) *Every homogeneous element of $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ has a lifting in $\mathrm{Syz}(\mathcal{G})$.*

D_2) *There exists a homogeneous system of generators of $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ consisting entirely of elements which have a lifting in $\mathrm{Syz}(\mathcal{G})$.*

D_3) *There exists a finite homogeneous system of generators of $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$ consisting entirely of elements which have a lifting in $\mathrm{Syz}(\mathcal{G})$.*

Proof.   This follows from Propositions 2.1.3, 2.2.8, and 2.3.12.   □

Definition 2.4.2. Let $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$ be a set of elements which generates a submodule $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$. If the conditions of Theorem 2.4.1 are satisfied, then G is called a **Gröbner basis** of M with respect to σ or a **σ-Gröbner basis** of M. In the case $M = \langle 0 \rangle$, we shall say that $G = \emptyset$ is a σ-Gröbner basis of M.

2.4.A   Existence of Gröbner Bases

Our first task is to show the existence of Gröbner bases. If we recall Proposition 1.5.6.b, it is clear that there are elements $g_1, \ldots, g_s \in M$ satisfying Condition B_2). But do they generate M? Our next proposition answers this question affirmatively.

Proposition 2.4.3. (Existence of a σ-Gröbner Basis)
Let M be a non-zero P-submodule of P^r.

a) *Given $g_1, \ldots, g_s \in M \setminus \{0\}$ such that $\mathrm{LT}_\sigma(M) = \langle \mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s) \rangle$, we have $M = \langle g_1, \ldots, g_s \rangle$, and the set $G = \{g_1, \ldots, g_s\}$ is a σ-Gröbner basis of M.*

b) *The module M has a σ-Gröbner basis $G = \{g_1, \ldots, g_s\} \subseteq M \setminus \{0\}$.*

Proof. First we show claim a) by contradiction. Suppose $\langle g_1, \ldots, g_s \rangle \subset M$. By Theorem 1.4.19, there exists an element $m \in M \setminus \langle g_1, \ldots, g_s \rangle$ whose leading term $\mathrm{LT}_\sigma(m)$ is minimal with respect to σ among all elements of that set. Since we have $\mathrm{LT}_\sigma(m) \in \mathrm{LT}_\sigma(M) = \langle \mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s) \rangle$, there are $c \in K \setminus \{0\}$, $t \in \mathbb{T}^n$, and $i \in \{1, \ldots, s\}$ such that $\mathrm{LM}_\sigma(m) = c\, t\, \mathrm{LM}_\sigma(g_i)$. Thus we get $\mathrm{LT}_\sigma(m - c\, t\, g_i) <_\sigma \mathrm{LT}_\sigma(m)$, and hence $m - c\, t\, g_i \in \langle g_1, \ldots, g_s \rangle$, contradicting $m \notin \langle g_1, \ldots, g_s \rangle$.

Claim b) follows from a) using Proposition 1.5.6.b. $\square$

The existence of Gröbner bases implies one of the most important properties of polynomial rings over fields. In Section 1.3 we described the property of being Noetherian in the case of monoideals. Using a similar formulation, we extend it to ideals and modules.

Definition 2.4.4. A ring (resp. module) is called **Noetherian** if every ascending chain of ideals (resp. submodules) becomes eventually stationary.

The following characterizations of Noetherian modules are in complete analogy with the case of Noetherian monoids and can be shown exactly as Proposition 1.3.4.

Proposition 2.4.5. *Let R be a ring and M an R-module. The following conditions are equivalent.*

a) *Every submodule of M is finitely generated.*

b) *Every ascending chain $N_1 \subseteq N_2 \subseteq \cdots$ of submodules of M is eventually stationary.*

c) *Every non-empty set of submodules of M has a maximal element (with respect to inclusion).*

As a consequence of Proposition 2.4.3, we obtain a version of Hilbert's Basis Theorem for finitely generated modules over finitely generated K-algebras.

Theorem 2.4.6. (Hilbert's Basis Theorem)
Every finitely generated module over a finitely generated K-algebra is Noetherian. In particular, $P = K[x_1, \dots, x_n]$ is a Noetherian ring.

Proof. If we represent the K-algebra in the form P/I with a polynomial ring $P = K[x_1, \dots, x_n]$ and an ideal $I \subseteq P$, we can view the module M as a finitely generated P-module via the canonical map $P \longrightarrow P/I$. Obviously it suffices to show that every P-submodule of M is finitely generated. Since M is finitely generated, we can represent M in the form $M = P^r/U$ with $r \geq 1$ and a submodule $U \subseteq P^r$. Since every submodule of M is of the form N/U with a submodule $N \subseteq P^r$, it suffices to show that every P-submodule of P^r is finitely generated, and this is an immediate consequence of Proposition 2.4.3. $\square$

2.4.B   Normal Forms

Our next application of Gröbner bases is to show how they help us to perform effective calculations in a residue class module P^r/M. Several attempts to solve this question have failed so far, because we were not able to find a unique representative in P^r for a residue class in P^r/M. Using a Gröbner basis, we now find that all those attempts lead to the same unique answer.

Let $G = \{g_1, \dots, g_s\} \subseteq P^r \setminus \{0\}$ be a σ-Gröbner basis of $M = \langle g_1, \dots, g_s \rangle \subseteq P^r$, and let $m \in P^r$. By Condition C_3), there exists a unique element $m_G \in P^r$ such that $m \xrightarrow{G} m_G$ and such that m_G is irreducible with respect to $\xrightarrow{G}$. A priori this element seems to depend on the Gröbner basis chosen, but indeed it does not, as the following proposition shows.

Proposition 2.4.7. *In the above situation, m_G is the unique element of P^r with the properties that $m - m_G \in M$ and $\mathrm{Supp}(m_G) \cap \mathrm{LT}_\sigma\{M\} = \emptyset$. In particular, it does not depend on the particular σ-Gröbner basis chosen.*

Proof. We know that $m - m_G \in M$ and that the support of m_G does not intersect $\mathrm{LT}_\sigma\{M\}$. Uniqueness follows from the observation that, for two such elements m_G and m_H, the support of $m_G - m_H \in M$ does not intersect $\mathrm{LT}_\sigma\{M\}$, and this is, by Condition C_2), only possible if $m_G - m_H = 0$. $\square$

Definition 2.4.8. Let $M \subseteq P^r$ be a non-zero module, and let $m \in P^r$. The element $m_G \in P^r$ described above is called the **normal form** of m with respect to σ. It is denoted by $\mathrm{NF}_{\sigma,M}(m)$, or simply by $\mathrm{NF}_\sigma(m)$ if it is clear which submodule is considered.

Below we collect some properties of normal forms. In particular, we see that the Division Algorithm with respect to a Gröbner basis provides an effective method for computing normal forms.

Corollary 2.4.9. *In the above situation, let $\mathcal{G} = (g_1, \ldots, g_s)$.*

a) *If $m \in P^r$, then $\mathrm{NR}_{\sigma,\mathcal{G}}(m)$ agrees with $\mathrm{NF}_\sigma(m)$. In particular, the normal remainder does not depend on the order of the elements $g_1, \ldots, g_s$.*

b) *For $m_1, m_2 \in P^r$, we have $\mathrm{NF}_\sigma(m_1 - m_2) = \mathrm{NF}_\sigma(m_1) - \mathrm{NF}_\sigma(m_2)$.*

c) *For $m \in P^r$, we have $\mathrm{NF}_\sigma(\mathrm{NF}_\sigma(m)) = \mathrm{NF}_\sigma(m)$.*

Proof. Claim a) follows from $m - \mathrm{NR}_{\sigma,\mathcal{G}}(m) \in M$ and from the fact that the support of $\mathrm{NR}_{\sigma,\mathcal{G}}(m)$ does not meet $\mathrm{LT}_\sigma\{M\}$. Next we show b). We have
$$m_1 - m_2 - (\mathrm{NF}_\sigma(m_1) - \mathrm{NF}_\sigma(m_2)) = (m_1 - \mathrm{NF}_\sigma(m_1)) - (m_2 - \mathrm{NF}_\sigma(m_2)) \in M$$
and $\mathrm{NF}_\sigma(m_1) - \mathrm{NF}_\sigma(m_2)$ is irreducible with respect to $\xrightarrow{G}$. The uniqueness of such an element yields the conclusion. Claim c) follows similarly, because $\mathrm{NF}_\sigma(m) - \mathrm{NF}_\sigma(m) = 0 \in M$ and $\mathrm{NF}_\sigma(m)$ is irreducible with respect to $\xrightarrow{G}$. $\qquad\square$

For the purposes of actual computations, one of the most useful applications of normal forms is the possibility to check whether an element is contained in a submodule or whether one submodule is contained in another.

Proposition 2.4.10. (Submodule Membership Test)
Let $\{g_1, \ldots, g_s\} \subseteq P^r$ generate a P-submodule $M = \langle g_1, \ldots, g_s \rangle$ of P^r, and let $\{h_1, \ldots, h_t\} \subseteq P^r$ generate a P-submodule $N = \langle h_1, \ldots, h_t \rangle \subseteq P^r$.

a) *For $m_1, m_2 \in P^r$, we have $m_1 - m_2 \in M$ if and only if $\mathrm{NF}_{\sigma,M}(m_1) = \mathrm{NF}_{\sigma,M}(m_2)$. In particular, an element $m \in P^r$ satisfies $m \in M$ if and only if $\mathrm{NF}_{\sigma,M}(m) = 0$.*

b) *We have $N \subseteq M$ if and only if $\mathrm{NF}_{\sigma,M}(h_i) = 0$ for $i = 1, \ldots, t$.*

c) *The condition $M = N$ is equivalent to $\mathrm{NF}_{\sigma,N}(g_i) = \mathrm{NF}_{\sigma,M}(h_j) = 0$ for $i = 1, \ldots, s$ and $j = 1, \ldots, t$.*

d) *If $N \subseteq M$ and $\mathrm{LT}_\sigma\{M\} \subseteq \mathrm{LT}_\sigma\{N\}$, then $M = N$.*

Proof. To show the first claim, let $m_1, m_2 \in P^r$ such that $m_1 - m_2 \in M$. Then $0 = \mathrm{NF}_{\sigma,M}(m_1 - m_2) = \mathrm{NF}_{\sigma,M}(m_1) - \mathrm{NF}_{\sigma,M}(m_2)$ by Corollary 2.4.9.b. Conversely, let $\mathrm{NF}_{\sigma,M}(m_1) = \mathrm{NF}_{\sigma,M}(m_2)$. In this case, the claim follows from $m_1 - m_2 = (m_1 - \mathrm{NF}_{\sigma,M}(m_1)) - (m_2 - \mathrm{NF}_{\sigma,M}(m_2)) \in M$.

Clearly, claim b) is a consequence of a), and claim c) follows from b). Thus it remains to prove claim d). Since we have $N \subseteq M$, it is clear that $\mathrm{LT}_\sigma\{N\} \subseteq \mathrm{LT}_\sigma\{M\}$. Thus the hypothesis that we have the other inclusion $\mathrm{LT}_\sigma\{M\} \subseteq \mathrm{LT}_\sigma\{N\}$ implies equality $\mathrm{LT}_\sigma\{N\} = \mathrm{LT}_\sigma\{M\}$. Now take an element $m \in M$. We have $\mathrm{Supp}(\mathrm{NF}_{\sigma,N}(m)) \cap \mathrm{LT}_\sigma\{N\} = \emptyset$, and therefore $\mathrm{Supp}(\mathrm{NF}_{\sigma,N}(m)) \cap \mathrm{LT}_\sigma\{M\} = \emptyset$. The uniqueness in Proposition 2.4.7 shows that $\mathrm{NF}_{\sigma,N}(m) = 0$, i.e. we get $m \in N$. $\qquad\square$

As an important application of the notion of Gröbner basis, we get a new version of Macaulay's Basis Theorem 1.5.7.

Corollary 2.4.11. (New Version of Macaulay's Basis Theorem)
Let $M \subseteq P^r$ be a P-submodule, let $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$ be a σ-Gröbner basis of M, and let B be the set of all terms in $\mathbb{T}^n\langle e_1, \ldots, e_r\rangle$ which are not a multiple of any term in the set $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$. Then the residue classes of the elements of B form a K-basis of P^r/M.

Proof. The fact that G is a σ-Gröbner basis of M implies that $\mathrm{LT}_\sigma\{M\}$ is generated by $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$ by Condition B_1) of Theorem 2.4.1. So the statement follows immediately from Theorem 1.5.7. $\square$

2.4.C   Reduced Gröbner Bases

In the last part of this section we address the question of uniqueness of Gröbner bases and provide an application of it. Given a module term ordering σ, a submodule $M \subseteq P^r$ has many σ-Gröbner bases. For instance, we can add arbitrary elements of M to a σ-Gröbner basis and it remains a σ-Gröbner basis of M. However, there is a unique one which satisfies the following additional conditions.

Definition 2.4.12. Let $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$ and $M = \langle g_1, \ldots, g_s\rangle$. We say that G is a **reduced σ-Gröbner basis** of M if the following conditions are satisfied.
  a) For $i = 1, \ldots, s$, we have $\mathrm{LC}_\sigma(g_i) = 1$.
  b) The set $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$ is a minimal system of generators of $\mathrm{LT}_\sigma(M)$.
  c) For $i = 1, \ldots, s$, we have $\mathrm{Supp}(g_i - \mathrm{LT}_\sigma(g_i)) \cap \mathrm{LT}_\sigma\{M\} = \emptyset$.

Theorem 2.4.13. (Existence and Uniqueness of Reduced Gröbner Bases)
For every P-submodule $M \subseteq P^r$, there exists a unique reduced σ-Gröbner basis.

Proof. We start by proving existence. Let $G = \{g_1, \ldots, g_s\}$ be any σ-Gröbner basis of M. If we replace g_i by $\mathrm{LC}_\sigma(g_i)^{-1}g_i$ for $i = 1, \ldots, s$, we obtain a Gröbner basis with property a). By Condition B_2) of Theorem 2.4.1, the monomial module $\mathrm{LT}_\sigma(M)$ is generated by $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$. Then we use Proposition 1.3.11.b to get from this set the unique minimal system of generators of $\mathrm{LT}_\sigma(M)$. After possibly renumbering the vectors we may assume that this minimal system of generators is $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_t)\}$, where $t \leq s$. And using again Condition B_2) and Proposition 2.4.3.a, we see that the set $G' = \{g_1, \ldots, g_t\}$ is a σ-Gröbner basis of M which satisfies conditions a) and b) of the definition.

Now we write $g_i = \mathrm{LT}_\sigma(g_i) + h_i$, and if we let $g_i' = \mathrm{LT}_\sigma(g_i) + \mathrm{NF}_\sigma(h_i)$ for $i = 1, \ldots, t$, we can form the set $G'' = \{g_1', \ldots, g_t'\}$. We claim that G'' is a reduced σ-Gröbner basis of M. Since $g_i' = g_i - (h_i - \mathrm{NF}_\sigma(h_i))$, we use

Proposition 2.4.7 and get $g_i' \in M$ for $i = 1, \ldots, t$. By Condition B_2), the set G'' is a σ-Gröbner basis of M. Since it clearly satisfies conditions a) and b) of the definition, it remains to prove that it also satisfies condition c). Indeed, for every $i \in \{1, \ldots, t\}$, no term in $\mathrm{Supp}(\mathrm{NF}_\sigma(h_i))$ lies in $\mathrm{LT}_\sigma\{M\}$, because $\mathrm{NF}_\sigma(h_i)$ is irreducible with respect to $\xrightarrow{G'}$.

Finally, to show uniqueness, we assume that $G = \{g_1, \ldots, g_s\}$ and $H = \{h_1, \ldots, h_t\}$ are two reduced σ-Gröbner bases of M. From the fact that the minimal monomial system of generators of a monomial module is unique (see Proposition 1.3.11.b), we conclude $s = t$ and that we can renumber the elements of H such that $\mathrm{LT}_\sigma(g_i) = \mathrm{LT}_\sigma(h_i)$ for $i = 1, \ldots, s$. Moreover, for $i = 1, \ldots, s$, we have $g_i - h_i \in M$, and $g_i - h_i$ is, by condition c) of the definition, irreducible with respect to $\xrightarrow{G}$. Thus property C_2) of Theorem 2.4.1 proves $g_i = h_i$ for $i = 1, \ldots, s$. $\square$

As an application of the existence and uniqueness of reduced σ-Gröbner bases we can show the existence and uniqueness of a field of definition for submodules of P^r.

Definition 2.4.14. Let K be a field, $P = K[x_1, \ldots, x_n]$ a polynomial ring, and $M \subseteq P^r$ a P-submodule.

a) Let $k \subseteq K$ be a subfield. We say that M is **defined over** k if there exist elements in $k[x_1, \ldots, x_n]^r$ which generate M as a P-module.

b) A subfield $k \subseteq K$ is called a **field of definition** of M if M is defined over k and there exists no proper subfield $k' \subset k$ such that M is defined over k'.

It is clear that if a field of definition of a P-submodule $M \subseteq P^r$ exists, it has to contain the prime field of K. Let us look at a concrete example.

Example 2.4.15. Let $I \subseteq \mathbb{C}[x_1, x_2, x_3]$ be the ideal generated by the set $\{x_1^2 - \sqrt{5}x_1x_2 + 3x_1x_3 + 2\sqrt{5}x_3^2, \ x_1x_2 - \sqrt{2}x_3^2, \ 2x_1x_2 + \sqrt{3}x_3^2\}$. Obviously, the ideal I is defined over $\mathbb{Q}[\sqrt{2}, \sqrt{3}, \sqrt{5}]$.

But it is also easy to check that $I = (x_1^2 + 3x_1x_3, \ x_1x_2, \ x_3^2)$. Therefore, the ideal I is defined over the prime field $\mathbb{Q}$ of $\mathbb{C}$, and the unique field of definition of I is $\mathbb{Q}$.

The following lemma captures one important aspect of the proof of the existence and uniqueness of the field of definition.

Lemma 2.4.16. *Let $K' \subseteq K$ be a field extension, let $P' = K'[x_1, \ldots, x_n]$, let $M' \subseteq (P')^r$ be a P'-submodule of $(P')^r$, and let M be the P-submodule of P^r generated by the elements of M'.*

a) *A σ-Gröbner basis of M' is also a σ-Gröbner basis of M. In particular, we have $\mathrm{LT}_\sigma\{M'\} = \mathrm{LT}_\sigma\{M\}$.*

b) *The reduced σ-Gröbner basis of M' is also the reduced σ-Gröbner basis of M.*

Proof. Let $G = \{g_1, \ldots, g_s\} \subseteq (P')^r \setminus \{0\}$ be a σ-Gröbner basis of M'. Since the set G generates the P'-module M' and the set M' generates the P-module M, the set G generates the P-module M.

Let $\mathcal{G}$ be the tuple $(g_1, \ldots, g_s)$. Using Theorem 2.3.7, we see that $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) = \langle \sigma_{ij} \mid 1 \leq i < j \leq s, \gamma_i = \gamma_j \rangle$, where $\sigma_{ij} \in (P')^s$ is given by $\sigma_{ij} = \frac{1}{c_i} t_{ij} \varepsilon_i - \frac{1}{c_j} t_{ji} \varepsilon_j$. By Condition D_1) of Theorem 2.4.1, the elements σ_{ij} have liftings in $(P')^s$. These liftings are also liftings in P^s of the elements σ_{ij} if we consider those as elements of P^s. Using Condition D_3) of Theorem 2.4.1, we deduce that G is in fact a σ-Gröbner basis of M. This proves a).

To prove b), we observe that the extra conditions required in Definition 2.4.12 are independent of the base field. $\qquad\square$

Theorem 2.4.17. (Existence and Uniqueness of the Field of Definition)
Let M be a non-zero P-submodule of P^r.

a) *There exists a unique field of definition of M.*
b) *Given any module term ordering σ, let G be the corresponding reduced σ-Gröbner basis of M. Then the field of definition of M is the field generated over the prime field of K by the coefficients of the terms in the support of the vectors in G.*

Proof. Let σ be a module term ordering, and let G be the reduced σ-Gröbner basis of M. Moreover, let k be the field generated over the prime field of K by the coefficients of the elements of G. Since the set G generates M, the module M is defined over k.

Suppose now that $K' \subseteq K$ is a subfield over which M is defined, i.e. suppose there exists a system of generators $\{m_1, \ldots, m_t\}$ of the P-module M which is contained in $K'[x_1, \ldots, x_n]^r \setminus \{0\}$. Let $G' = \{g_1', \ldots, g_s'\} \subseteq K'[x_1, \ldots, x_n]^r$ be the reduced σ-Gröbner basis of the $K'[x_1, \ldots, x_n]$-module $\langle m_1, \ldots, m_t \rangle \subseteq K'[x_1, \ldots, x_n]^r$. Since the reduced σ-Gröbner basis of a module is unique, Lemma 2.4.16.b implies $G = G'$. From this we infer that $k \subseteq K'$.

The facts that M is defined over k, and that every other field over which M is defined contains k, together imply both claims of the theorem. $\qquad\square$

Exercise 1. Let $I = (g)$ with $g \in P \setminus \{0\}$ be a principal ideal in P. Show that $G = \{g\}$ is a Gröbner basis of I with respect to every term ordering.

Exercise 2. Let $m_1, \ldots, m_s \in P^r$ be terms, and let $M = \langle m_1, \ldots, m_s \rangle$. Show that $\{m_1, \ldots, m_s\}$ is a Gröbner basis of M with respect to every term ordering.

Exercise 3. Let $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$ be a σ-Gröbner basis of the P-module $M = \langle g_1, \ldots, g_s \rangle$, and let $m \in M$. Show that $G \cup \{m\}$ is a σ-Gröbner basis of M.

Exercise 4. Let $g_1 = x_2 - x_1^2$ and $g_2 = x_3 - x_1^3$ be polynomials in $K[x_1, x_2, x_3]$. Find a term ordering σ on $\mathbb{T}^3$ such that $G = \{g_1, g_2\}$ is a σ-Gröbner basis of the ideal $I = (g_1, g_2)$, and a term ordering τ such that it is not.

Exercise 5. Let $P = K[x_1, \ldots, x_n]$, let $m \leq n$, let $G = \{f_1, f_2, \ldots, f_m\}$, where $f_i \in K[x_i]$ for $i = 1, \ldots, m$, and let $I \subseteq P$ be the ideal generated by G.

a) Use Condition C_3) of Theorem 2.4.1 to show that G is a σ-Gröbner basis of I with respect to every term ordering σ.

b) If, moreover, the polynomials f_i are monic, show that G is the reduced σ-Gröbner basis of I with respect to every term ordering σ.

Exercise 6. Let R be a Noetherian integral domain. Show that the following conditions are equivalent.

a) For all $a, b \in R \setminus \{0\}$, the ideal $(a) \cap (b)$ is principal.

b) The ring R is factorial.

Hint: Use Exercise 6 in Section 1.2.

Exercise 7. Using Corollary 2.4.11 and CoCoA, find a set of terms whose residue classes form a basis of $\mathbb{Z}/(5)[x, y, z]/(x^2 - yz, y^3 + z^3, z^5 - x^2 y^2)$ as a $\mathbb{Z}/(5)$-vector space. (*Hint:* You may use the CoCoA function `GBasis(...)`.)

Exercise 8. A system of generators $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$ of a P-module $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$ is called a **minimal σ-Gröbner basis** of M if $\{\mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s)\}$ is a minimal system of generators of $\mathrm{LT}_\sigma(M)$.

a) Prove that any two minimal σ-Gröbner bases of M have the same number of elements.

b) Give an example of a module M which has two different minimal σ-Gröbner bases, all of whose elements g_i have leading coefficients $\mathrm{LC}_\sigma(g_i) = 1$.

Exercise 9. Let σ be a term ordering on $\mathbb{T}^n \langle e_1, \ldots, e_r \rangle$. We set $\mathrm{LT}_\sigma(0) = \infty$. In particular, we are assuming $\mathrm{LT}_\sigma(g) <_\sigma \mathrm{LT}_\sigma(0)$ for every $g \in P^r$. Given a tuple $(g_1, \ldots, g_s) \in (P^r)^s$, we identify it with the tuple $(g_1, \ldots, g_s, 0) \in (P^r)^{s+1}$, hence with $(g_1, \ldots, g_s, 0, 0) \in (P^r)^{s+2}$, and so on.

For two tuples $\mathcal{G} = (g_1, \ldots, g_s) \in (P^r)^s$ and $\mathcal{G}' = (g_1', \ldots, g_{s'}') \in (P^r)^{s'}$, we define $\mathcal{G} \preceq \mathcal{G}'$ if and only if $\mathrm{LT}_\sigma(\mathcal{G}) \leq_{\mathrm{Lex}} \mathrm{LT}_\sigma(\mathcal{G}')$. This means that either there exists an index $i \geq 1$ such that $\mathrm{LT}_\sigma(g_i) <_\sigma \mathrm{LT}_\sigma(g_i')$ and $\mathrm{LT}_\sigma(g_j) = \mathrm{LT}_\sigma(g_j')$ for $1 \leq j < i$, or we have $\mathcal{G} = \mathcal{G}'$.

A tuple $\mathcal{G} = (g_1, \ldots, g_s)$ of elements in P^r is said to be **increasingly ordered** with respect to σ if $\mathrm{LT}_\sigma(g_1) \leq_\sigma \cdots \leq_\sigma \mathrm{LT}_\sigma(g_s)$. It is said to be **interreduced** if $g_i \neq 0$ for $i = 1, \ldots, s$ and $\mathrm{LT}_\sigma(g_i)$ does not divide any term in $\mathrm{Supp}(g_j)$ for $i, j \in \{1, \ldots, s\}$ such that $i \neq j$. Finally, the tuple $\mathcal{G}$ it is called **monic** if all its components are monic.

a) For $g_1, \ldots, g_s, g_{s+1}, \ldots, g_t \in P^r$, show $(g_1, \ldots, g_s, g_{s+1}, \ldots, g_t) \preceq (g_1, \ldots, g_s)$.

b) Prove that the relation $\preceq$ is reflexive and transitive, but not a total ordering on the set of the increasingly ordered tuples of elements of P^r.

c) Let M be a non-zero submodule of P^r, and let $\mathcal{G}$ be an increasingly ordered, interreduced tuple of elements of M. Show that the following conditions are equivalent.

1) With respect to $\preceq$, the tuple $\mathcal{G}$ is minimal among all increasingly ordered, interreduced, monic tuples of elements of M.

2) The tuple $\mathcal{G}$ is obtained by increasingly ordering the reduced σ-Gröbner basis of M.

Exercise 10. Let I be an ideal of $P = K[x_1, \ldots, x_n]$, let σ be a term ordering on $\mathbb{T}^n$, and let Γ be a group of K-algebra automorphisms of P. Show by example that if I is Γ-stable (i.e. if $\gamma(I) \subseteq I$ for all $\gamma \in \Gamma$), then the reduced σ-Gröbner basis of I need not be Γ-stable.

Tutorial 21: Linear Algebra

The purpose of this tutorial is to show how Gaußian Elimination in Linear Algebra relates to the theory of Gröbner bases. Let K be a field, let $m, n > 0$, and let $\mathcal{A} = (a_{ij})$ be an $m \times n$-matrix with coefficients in K. We equip the ring $P = K[x_1, \ldots, x_n]$ with the lexicographic term ordering `Lex`.

a) Write a CoCoA program `RowReduce(...)` which uses row operations to bring the matrix $\mathcal{A}$ into row echelon form and then returns the matrix $\mathcal{B} = (b_{ij})$ obtained in this way.

b) For $i = 1, \ldots, m$, let $f_i = a_{i1}x_1 + \cdots + a_{in}x_n$ and $g_i = b_{i1}x_1 + \cdots + b_{in}x_n$. Show that $G = \{g_i \mid 1 \le i \le m, \ g_i \ne 0\}$ is a `Lex`-Gröbner basis of the ideal $I = (f_1, \ldots, f_m)$.

c) Find and prove an algorithm which computes the `Lex`-Gröbner basis of an ideal I of P which is generated by polynomials of degree ≤ 1.

d) Implement your algorithm in a CoCoA function `LinearGB(...)` which takes a list of polynomials of degree ≤ 1 generating I and returns the `Lex`-Gröbner basis of I.

e) Use `LinearGB(...)` to compute the `Lex`-Gröbner bases of the following ideals.

1) $I_1 = (3x_1 - 6x_2 - 2x_3, \ 2x_1 - 4x_2 + 4x_4, \ x_1 - 2x_2 - x_3 - x_4) \subseteq \mathbb{Q}[x_1, x_2, x_3, x_4]$

2) $I_2 = (x_1 + x_2 + x_3, \ x_1 - x_2, \ x_1 - x_3) \subseteq \mathbb{Q}[x_1, x_2, x_3]$

3) $I_3 = (x_1 + 1, \ x_2 + x_3 + 1, \ x_4 + x_5 + 1, \ x_1 + x_4 - 1) \subseteq \mathbb{Q}[x_1, \ldots, x_5]$

Tutorial 22: Reduced Gröbner Bases

In this tutorial we shall implement an algorithm to find the reduced Gröbner
basis from an arbitrary one, and we shall study various particular cases of
reduced Gröbner bases. So let K be a field, $n \geq 1$, $P = K[x_1, \ldots, x_n]$, $r \geq 1$,
σ a module term ordering on $\mathbb{T}^n \langle e_1, \ldots, e_r \rangle$, and $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$
a σ-Gröbner basis of the P-submodule $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$.

a) Implement the method described in the proof of Theorem 2.4.13. Write
   a CoCoA function `ReduceGB(...)` which takes any σ-Gröbner basis of M
   and computes the reduced σ-Gröbner basis from it.
b) Apply your function `ReduceGB(...)` in the following cases, assuming each
   time that the given sets are `Lex`-Gröbner bases of the ideals they generate.
   1) $G_1 = \{x^2 + y^2 + 1,\ x^2y + 2xy + x,\ -2xy - x + y^3 + y,\ -y^5 - 2y^4 + y^2 + y - 2\} \subseteq \mathbb{Z}/(5)[x, y]$.
   2) $G_2 = \{xz^3 - x - 3y^6 - 18y^4 - 12y^3 - 18y^2 - 12y - 3,\ 15x - y^6 - 12y^5 - 79y^3 - 24y^2 - 67y + z^3 - 26,\ y^6 + 6y^4 + 4y^3 + 6y^2 + 4y - z^3 + 2,\ z^3 - 1\} \subseteq \mathbb{Q}[x, y, z]$.
   3) $G_3 = \{x^2 + y - 1,\ xy - 2y^2 + 2y,\ 4y^3 - 7y^2 + 3y,\ 1/2x^2 + 1/2xy - y^2 + 3/2y - 1/2\} \subseteq \mathbb{Q}[x, y]$.
c) Now we equip the polynomial ring P with its standard grading (see
   Example 1.7.2). Prove that an ideal $I \subseteq P$ is homogeneous if and only
   if its reduced σ-Gröbner basis consists of homogeneous polynomials.
   Hint: First show that any homogeneous ideal has a σ-Gröbner basis
   consisting of homogeneous polynomials.
d) Let $m \geq 1$, let $\mathcal{A} = (a_{ij})$ be an $m \times n$-matrix with coefficients in K,
   and let $f_i = a_{i1}x_1 + \cdots + a_{in}x_n$ for $i = 1, \ldots, m$. Using row operations
   only, we bring $\mathcal{A}$ to reduced row echelon form $\mathcal{B} = (b_{ij})$, i.e. in the row
   echelon form we clear out everything starting from the bottom. For the
   non-zero rows numbered $i = 1, \ldots, t$ of $\mathcal{B}$, we form the linear polynomials
   $g_i = b_{i1}x_1 + \cdots + b_{in}x_n$. Prove that $\{g_1/\mathrm{LC}_{\mathtt{Lex}}(g_1), \ldots, g_t/\mathrm{LC}_{\mathtt{Lex}}(g_t)\}$ is
   the reduced `Lex`-Gröbner basis of the ideal $I = (f_1, \ldots, f_m)$ of P.
e) Write a CoCoA program `LinRedGB(...)` which computes the reduced `Lex`-
   Gröbner basis of an ideal $I = (f_1, \ldots, f_m)$ as in d) using the method
   described there.
f) Apply your function `LinRedGB(...)` to the ideals I_1 and I_2 of Tutorial 21.e.
   Check your results by comparing them to the results of `LinearGB(...)` and
   `ReduceGB(...)`.
g) Suppose that $\{m_1, \ldots, m_t\} \subseteq P^r \setminus \{0\}$ is any system of generators of
   the P-module M, and that $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$ is the reduced
   σ-Gröbner basis of M. Then there are matrices $\mathcal{A} = (a_{ij})$ and $\mathcal{B} = (b_{ij})$
   with coefficients in P such that $m_i = a_{i1}g_1 + \cdots + a_{is}g_s$ for $i = 1, \ldots, t$
   and $g_j = b_{j1}m_1 + \cdots + b_{jt}m_t$ for $j = 1, \ldots, s$. Give an example in which
   $\mathcal{A}\mathcal{B}$ is not the identity matrix.

2.5 Buchberger's Algorithm

Knowing + and × is good enough,
understanding their interaction is ideal*.*
(Bruno Buchberger)

In the last section we saw some theoretical applications of Gröbner bases, especially of reduced Gröbner bases. But Gröbner bases would be hardly more than a small side subject in commutative algebra if we did not have the possibility of computing them. The key to almost all applications of Gröbner bases in Computational Commutative Algebra, and therefore to the remainder of these volumes, is the algorithm developed by Bruno Buchberger in his doctoral thesis [Bu65].

Knowing that every ideal has a unique reduced Gröbner basis is good enough, but actually computing it is *ideal*. This computation is based on the characterization of Gröbner bases via lifting of syzygies. The idea is that we need to check whether the set of generators satisfies Condition D_3). If a syzygy of the leading terms is found which does not lift to a syzygy of the generators, we can find an element of the module which has a *new* leading term. By adding it to the set of generators, we can achieve the desired lifting. Then the termination of the algorithm is guaranteed by Dickson's Lemma (more precisely, by Corollary 1.3.10), and its correctness follows from the fact that lifting of syzygies characterizes Gröbner bases (see Theorem 2.4.1).

Since Buchberger's Algorithm is the basic tool underlying most calculations in Computational Commutative Algebra, it is very important to study possibilities for optimizing it. First indications on how to avoid some unnecessary steps in the execution of the algorithm are given in Remark 2.5.6 and Proposition 2.5.8. Some additional possibilities are contained in Tutorial 25. For the case of systems of generators consisting of homogeneous polynomials or vectors of polynomials, an efficient version of Buchberger's Algorithm will be explained in Volume 2.

At the end of this section we discuss the Extended Buchberger Algorithm. Besides a Gröbner basis, it also yields the change of basis matrix from the given system of generators to the Gröbner basis (see Proposition 2.5.11).

As usual, let K be a field, let $n \geq 1$, let $P = K[x_1, \ldots, x_n]$ be a polynomial ring, let $r \geq 1$, and let σ be a module term ordering on $\mathbb{T}^n \langle e_1, \ldots, e_r \rangle$. Our goal is to compute a σ-Gröbner basis of a P-submodule $M \subseteq P^r$ which is explicitly given by a system of generators $G = \{g_1, \ldots, g_s\} \subseteq P^r \setminus \{0\}$. Let $\mathcal{G}$ be the tuple $(g_1, \ldots, g_s)$. We start by writing $\mathrm{LM}_\sigma(g_i) = c_i t_i e_{\gamma_i}$ with $c_i \in K \setminus \{0\}$, $t_i \in \mathbb{T}^n$, and $\gamma_i \in \{1, \ldots, r\}$ for $i = 1, \ldots, s$, and by recalling the fundamental diagram

$$
\begin{array}{ccccccccc}
0 & \longrightarrow & \mathrm{Syz}(\mathcal{G}) & \longrightarrow & P^s & \xrightarrow{\lambda} & P^r & \longrightarrow & P^r/M & \longrightarrow & 0 \\
& & & & \downarrow{\scriptstyle \mathrm{LF}} & & \downarrow{\scriptstyle \mathrm{LM}} & & & & \\
0 & \longrightarrow & \mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G})) & \longrightarrow & P^s & \xrightarrow{\Lambda} & P^r & \longrightarrow & P^r/N & \longrightarrow & 0
\end{array}
$$

studied in Section 2.3. Then we introduce or recall the following abbreviations.

Definition 2.5.1. Let $\mathbb{B}$ be the set $\mathbb{B} = \{(i,j) \mid 1 \leq i < j \leq s, \gamma_i = \gamma_j\}$. The elements of $\mathbb{B}$ are called the **critical pairs** of $\mathcal{G}$. Moreover, we let $t_{ij} = \frac{\mathrm{lcm}(t_i,t_j)}{t_i} = \frac{t_j}{\gcd(t_i,t_j)} \in \mathbb{T}^n$ and $\sigma_{ij} = \frac{1}{c_i}t_{ij}\varepsilon_i - \frac{1}{c_j}t_{ji}\varepsilon_j \in P^s$ for all $i,j \in \{1,\ldots,s\}$. For every critical pair $(i,j) \in \mathbb{B}$, we call

$$S_{ij} = \lambda(\sigma_{ij}) = \frac{1}{c_i}\,t_{ij}\,g_i - \frac{1}{c_j}\,t_{ji}\,g_j \in M$$

the **S-vector** of g_i and g_j. If $r = 1$, we call $S_{ij} \in P$ also the **S-polynomial** of g_i and g_j.

We can rephrase Theorem 2.3.7 by saying that if $(i,j) \in \mathbb{B}$ is a critical pair, then the fundamental syzygy σ_{ij} is a homogeneous element of P^s with $\deg_{\sigma,G}(\sigma_{ij}) = \mathrm{lcm}(t_i,t_j)e_{\gamma_i}$ and that the set $\Sigma = \{\sigma_{ij} \mid (i,j) \in \mathbb{B}\}$ is a homogeneous system of generators of the P-module $\mathrm{Syz}(\mathrm{LM}_\sigma(\mathcal{G}))$. Furthermore, we know by Theorem 2.4.1 that G is a σ-Gröbner basis of M if and only if all fundamental syzygies σ_{ij} have liftings in $\mathrm{Syz}(\mathcal{G})$. For some of them, this is always the case.

Proposition 2.5.2. *Let $(i,j) \in \mathbb{B}$ be such that $S_{ij} \xrightarrow{G} 0$. Then σ_{ij} has a lifting in $\mathrm{Syz}(\mathcal{G})$.*

Proof. If $S_{ij} = 0$, there is nothing to show, since σ_{ij} is a lifting of itself. Thus we may assume $S_{ij} \neq 0$. In view of Lemma 2.2.6, we can use $S_{ij} \xrightarrow{G} 0$ to obtain a representation $S_{ij} = \sum_{k=1}^s f_k g_k$ with $f_1,\ldots,f_s \in P$ such that $\mathrm{LT}_\sigma(S_{ij}) = \max_\sigma\{\mathrm{LT}_\sigma(f_k g_k) \mid 1 \leq k \leq s, \ f_k g_k \neq 0\}$. Since σ_{ij} is homogeneous, we have $\Lambda(\mathrm{LF}(\sigma_{ij})) = \Lambda(\sigma_{ij}) = 0$, and Proposition 2.3.6.b yields $\deg_{\sigma,G}(\sigma_{ij}) >_\sigma \mathrm{LT}_\sigma(S_{ij})$. Now we consider the element $\tau_{ij} = \sigma_{ij} - \sum_{k=1}^s f_k\varepsilon_k \in P^s$. From $\deg_{\sigma,G}(\sum_{k=1}^s f_k\varepsilon_k) = \mathrm{LT}_\sigma(S_{ij}) <_\sigma \deg_{\sigma,G}(\sigma_{ij})$ we deduce that $\mathrm{LF}_{\sigma,G}(\tau_{ij}) = \sigma_{ij}$. From $\lambda(\tau_{ij}) = \lambda(\sigma_{ij}) - S_{ij} = 0$ and $\mathrm{LF}(\tau_{ij}) = \sigma_{ij}$ we conclude that τ_{ij} is a lifting of σ_{ij} in $\mathrm{Syz}(\mathcal{G})$. $\qquad\square$

Corollary 2.5.3. (Buchberger's Criterion)
Let $M \subseteq P^r$ be a P-submodule generated by $G = \{g_1,\ldots,g_s\} \subseteq P^r \setminus \{0\}$, and let $\mathcal{G} = (g_1,\ldots,g_s)$. Then the following conditions are equivalent.

a) The set G is a σ-Gröbner basis of M.
b) For all critical pairs $(i,j) \in \mathbb{B}$, we have $\mathrm{NR}_{\sigma,\mathcal{G}}(S_{ij}) = 0$.

Proof. If G is a σ-Gröbner basis of M, then $S_{ij} \in M$ yields $\mathrm{NR}_{\sigma,\mathcal{G}}(S_{ij}) = 0$ by Corollary 2.4.9.a and Proposition 2.4.10.a. Conversely, if condition b) holds, then $S_{ij} \xrightarrow{G} 0$. Using Proposition 2.5.2 we see that, for every pair $(i,j) \in \mathbb{B}$, the element σ_{ij} has a lifting in $\mathrm{Syz}(\mathcal{G})$. Thus Condition $D_3)$ of Theorem 2.4.1 holds. $\qquad\square$

Let us see how this criterion applies in practice. The following example also shows that *the leading term ideal of the square of an ideal is, in general,* NOT *the square of the leading term ideal.*

Example 2.5.4. Let $P = \mathbb{Q}[x, y, z]$, let $\sigma = \mathtt{DegRevLex}$, and let I be the ideal of P generated by $g_1 = x^2 - y^2$, $g_2 = xy^2 - z^3$, and $g_3 = y^4 - xz^3 = -y^2 g_1 + x g_2$. Successively, we compute

$$S_{12} = y^2 g_1 - x g_2 = -y^4 + xz^3 \xrightarrow{g_3} 0$$
$$S_{13} = y^4 g_1 - x^2 g_3 = -y^6 + x^3 z^3 \xrightarrow{g_3} x^3 z^3 - xy^2 z^3 \xrightarrow{g_1} 0$$
$$S_{23} = y^2 g_2 - x g_3 = -y^2 z^3 + x^2 z^3 \xrightarrow{g_1} 0$$

Thus Buchberger's Criterion applies and says that $\{g_1, g_2, g_3\}$ is a σ-Gröbner basis of I. In particular, the leading term ideal of I is $\mathrm{LT}_\sigma(I) = (x^2, xy^2, y^4)$.

By the way, in this example the obvious inclusion $\mathrm{LT}_\sigma(I)^2 \subseteq \mathrm{LT}_\sigma(I^2)$ is a strict one, disproving a claim in [CLS92], p. 443. More precisely, the element $f = g_2^2 - g_1 g_3 = y^6 + x^3 z^3 - 3xy^2 z^3 + z^6 \in I^2$ has a leading term $\mathrm{LT}_\sigma(f) = y^6$ which is not in $\mathrm{LT}_\sigma(I)^2$.

The idea of Buchberger's Algorithm is to enlarge G in such a way that eventually all elements σ_{ij} with $(i, j) \in \mathbb{B}$ have a lifting in $\mathrm{Syz}(\mathcal{G})$. By Theorem 2.4.1, this ensures that the enlarged set is a σ-Gröbner basis of M.

Theorem 2.5.5. (Buchberger's Algorithm)
Let $(g_1, \ldots, g_s) \in (P^r)^s$ be a tuple of non-zero elements which generate a submodule $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$. For $i = 1, \ldots, s$, let $\mathrm{LM}_\sigma(g_i) = c_i t_i e_{\gamma_i}$ with $c_i \in K \setminus \{0\}$, $t_i \in \mathbb{T}^n$, and $\gamma_i \in \{1, \ldots, r\}$. Consider the following sequence of instructions.

1) *Let $\mathcal{G} = (g_1, \ldots, g_s)$, $s' = s$ and $B = \{(i, j) \mid 1 \le i < j \le s', \gamma_i = \gamma_j\}$.*
2) *If $B = \emptyset$, return the result $\mathcal{G}$. Otherwise, choose a critical pair $(i, j) \in B$ and delete it from B.*
3) *Compute $S_{ij} = \frac{t_j}{c_i \gcd(t_i, t_j)} g_i - \frac{t_i}{c_j \gcd(t_i, t_j)} g_j$ and $\mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij})$. If the result is $\mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij}) = 0$, continue with step 2).*
4) *Increase s' by one. Append $g_{s'} = \mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij})$ to $\mathcal{G}$ and the set of critical pairs $\{(i, s') \mid 1 \le i < s', \gamma_i = \gamma_{s'}\}$ to B. Then continue with step 2).*

This is an algorithm, i.e. it stops after finitely many steps. It returns a tuple $\mathcal{G}$ of vectors which form a σ-Gröbner basis of M.

Proof. Every time step 2) is executed, one critical pair is cancelled from B. The set B is enlarged only in step 4). When this happens, an element is appended to $\mathcal{G}$ which has a leading term with respect to σ which is not in the monomodule generated by the leading terms of the previous elements of $\mathcal{G}$. Corollary 1.3.10 shows that P^r cannot contain an infinite chain

$$\langle \mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_s) \rangle \subset \langle \mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_{s+1}) \rangle \subset \cdots$$

Therefore step 4) can be executed only a finite number of times, i.e. the procedure stops after finitely many steps.

It remains to show that when the algorithm stops, the vectors in the resulting tuple $\mathcal{G}$ form a σ-Gröbner basis of M. Let s'' be the number of elements of the resulting tuple $\mathcal{G}$. During the execution of the procedure all pairs (i, j) such that $1 \leq i < j \leq s''$ and $\gamma_i = \gamma_j$ are considered, since whenever s' is increased in step 4), all necessary new pairs (i, s') are added to B. By Corollary 2.5.3, it suffices to show that, for every such pair (i, j), we have $\mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij}) = 0$. If at a certain step $S_{ij} = 0$ or $\mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij}) = 0$, there is nothing to prove. If at a certain step $\mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij}) \neq 0$, then $\mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij})$ is added to the tuple $\mathcal{G}$. Hence $\mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij})$ reduces to 0 via the rewrite rule defined by the vectors in the new tuple. $\qquad \square$

A closer look at this proof shows that a number of variants and optimizations of Buchberger's Algorithm are possible. Some of the most effective ones will be discussed in Tutorial 25 and in Volume 2. Here we limit ourselves to pointing out some obvious opportunities for improvement.

Remark 2.5.6. (First Optimizations of Buchberger's Algorithm)

a) In Buchberger's Algorithm, one can substitute the computation of the normal remainder $\mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij})$ by any procedure producing an element $m \in P^r$ which satisfies $S_{ij} \xrightarrow{\ \mathcal{G}\ } m$, and $\mathrm{LT}_\sigma(m) \notin \langle \mathrm{LT}_\sigma(g_1), \ldots, \mathrm{LT}_\sigma(g_{s'}) \rangle$ if $m \neq 0$.

b) If $\mathbb{B}' \subseteq \mathbb{B}$ is a subset with the property that also the set $\{\sigma_{ij} \mid (i, j) \in \mathbb{B}'\}$ generates $\mathrm{Syz}(\mathrm{LM}_\sigma(g_1), \ldots, \mathrm{LM}_\sigma(g_s))$, it suffices to start with $B = \mathbb{B}'$ in step 1) of Buchberger's Algorithm. This follows from Proposition 2.3.11.

c) In step 2) of the theorem we did not specify which critical pair $(i, j) \in B$ we should choose. One possibility is to take the pair (i, j) for which $\mathrm{lcm}(t_i, t_j)$ is minimal with respect to σ. This is called the **normal selection strategy**. It works well in practice if the term ordering σ is degree-compatible. Another possibility which avoids sorting the terms $\mathrm{lcm}(t_i, t_j)$ with respect to σ is to take any pair (i, j) for which the degree of $\mathrm{lcm}(t_i, t_j)$ is minimal.

To help the reader understand Theorem 2.5.5 better, we now apply Buchberger's Algorithm in a concrete case.

Example 2.5.7. Let $n = 2$, let $r = 1$, let $M \subseteq P = K[x, y]$ be the ideal generated by $g_1 = x^2$ and $g_2 = xy + y^2$. We compute a Gröbner basis of M with respect to $\sigma = \mathtt{Lex}$ and follow the steps of Buchberger's Algorithm.

1) Let $\mathcal{G} = (g_1, g_2)$, $s' = 2$ and $B = \{(1, 2)\}$.
2) Choose $(1, 2) \in B$ and set $B = \emptyset$.
3) We compute $S_{12} = yg_1 - xg_2 = -xy^2 \xrightarrow{\ g_2\ } y^3 = \mathrm{NR}_{\sigma, \mathcal{G}}(S_{12}) \neq 0$.

4) Let $s' = 3$, let $\mathcal{G} = (g_1, g_2, g_3)$ with $g_3 = y^3$, and let $B = \{(1,3), (2,3)\}$. Then return to step 2).

2) Choose $(1,3) \in B$ and set $B = \{(2,3)\}$.

3) We compute $S_{13} = y^3 g_1 - x^2 g_3 = 0$ and return to step 2).

2) Choose $(2,3) \in B$ and set $B = \emptyset$.

3) We compute $S_{23} = y^2 g_2 - x g_3 = y^4$. Then we calculate $S_{23} \xrightarrow{g_3} 0 = \mathrm{NR}_{\sigma, \mathcal{G}}(S_{23})$ and return to step 2).

2) Since $B = \emptyset$, we return the result $\mathcal{G} = (g_1, g_2, g_3)$.

If $r = 1$, i.e. if M is an ideal in P, there is another optimization of Buchberger's Algorithm which turns out to be useful in practise.

Proposition 2.5.8. *Let $\mathcal{G} = (g_1, \ldots, g_s)$ be a tuple of non-zero polynomials, let $I = (g_1, \ldots, g_s) \subseteq P$, and let $t_i = \mathrm{LT}_\sigma(g_i)$ for $i = 1, \ldots, s$. Suppose that $\gcd(t_i, t_j) = 1$ for some pair $(i,j) \in \mathbb{B}$. Then σ_{ij} has a lifting in $\mathrm{Syz}(\mathcal{G})$.*

Proof. This follows from the observations that $\sigma_{ij} = \frac{1}{c_i} t_j \varepsilon_i - \frac{1}{c_j} t_i \varepsilon_j$ and that $\tau_{ij} = \frac{1}{c_i c_j} g_j \varepsilon_i - \frac{1}{c_i c_j} g_i \varepsilon_j$ is a lifting of σ_{ij} in $\mathrm{Syz}(\mathcal{G})$. $\qquad\square$

Remark 2.5.9. For $f, g \in P$, the pair $(-g, f)$ is called the **trivial syzygy** of (f, g). Therefore Proposition 2.5.8 can be rephrased by saying that if $\gcd(t_i, t_j) = 1$, then the trivial syzygy of $(\mathrm{LM}_\sigma(g_i), \mathrm{LM}_\sigma(g_j))$ can be lifted to the trivial syzygy of (g_i, g_j).

The above result can be used to detect some special Gröbner bases.

Corollary 2.5.10. *Let $G = \{g_1, \ldots, g_s\} \subseteq P \setminus \{0\}$, and let $I = (g_1, \ldots, g_s)$. Assume that the leading terms of the elements $g_1, \ldots, g_s$ are pairwise coprime. Then G is a σ-Gröbner basis of I.*

Proof. Let $\mathcal{G} = (g_1, \ldots, g_s)$. By Proposition 2.5.8, every element σ_{ij} has a lifting in $\mathrm{Syz}(\mathcal{G})$. Thus G satisfies Condition D_3) of Theorem 2.4.1. $\qquad\square$

Finally, we can extend Buchberger's Algorithm in such a way that it not only computes a Gröbner basis of a submodule $M \subseteq P^r$, but also a matrix of polynomials which describes how the Gröbner basis can be expressed in terms of the original system of generators of M.

Proposition 2.5.11. (The Extended Buchberger Algorithm)
Let $(g_1, \ldots, g_s) \in (P^r)^s$ be a tuple of non-zero vectors in P^r which generate a submodule $M = \langle g_1, \ldots, g_s \rangle \subseteq P^r$. We write $\mathrm{LM}_\sigma(g_i) = c_i t_i e_{\gamma_i}$ with $c_i \in K \setminus \{0\}$, $t_i \in \mathbb{T}^n$, and $\gamma_i \in \{1, \ldots, r\}$ for $i = 1, \ldots, s$. Consider the following sequence of instructions.

1) *Let $\mathcal{G} = (g_1, \ldots, g_s)$, let $s' = s$, let $\mathcal{A}$ be the $s \times s$ identity matrix, and let $B = \{(i,j) \mid 1 \le i < j \le s, \gamma_i = \gamma_j\}$.*

2) *If $B = \emptyset$, return the result $(\mathcal{G}, \mathcal{A})$. Otherwise, choose a critical pair $(i,j) \in B$ and delete it from B.*

3) *Using the Division Algorithm, find* $q_1, \ldots, q_{s'} \in P$ *and* $p \in P^r$ *such that* $S_{ij} = q_1 g_1 + \cdots + q_{s'} g_{s'} + p$ *and the conditions of Theorem 1.6.4 hold. If* $p = 0$, *continue with step 2).*

4) *If* $p \neq 0$ *in step 3), then increase* s' *by one, append* $g_{s'} = p$ *to* $\mathcal{G}$, *add* $\{(i, s') \mid 1 \leq i < s', \gamma_i = \gamma_{s'}\}$ *to* B, *and append the column vector* $\frac{t_j}{c_i \gcd(t_i, t_j)} a_i - \frac{t_i}{c_j \gcd(t_i, t_j)} a_j - q_1 a_1 - \cdots - q_{s'-1} a_{s'-1}$ *to* $\mathcal{A}$, *where* $a_1, \ldots, a_{s'-1}$ *denote the previous columns of* $\mathcal{A}$. *Continue with step 2).*

This is an algorithm, i.e. it stops after finitely many steps. It returns a tuple $\mathcal{G} = (g_1, \ldots, g_{s'})$ *of vectors which form a* σ-*Gröbner basis of* M, *where* $s' \geq s$, *together with an* $s \times s'$-*matrix* $\mathcal{A} = (a_{ij})$ *of polynomials such that* $g_j = a_{1j} g_1 + \cdots + a_{sj} g_s$ *for* $j = 1, \ldots, s'$.

Proof. In view of Theorem 2.5.5, it suffices to prove the last claim. Each time a new column is appended to $\mathcal{A}$ in step 4), we have $g_j = a_{1j} g_1 + \cdots + a_{sj} g_s$ for $j < s'$, where s' is the current number of columns of $\mathcal{A}$. Now the calculation

$$
\begin{aligned}
g_{s'} = p &= S_{ij} - q_1 g_1 - \cdots - q_{s'-1} g_{s'-1} \\
&= \tfrac{t_{ij}}{c_i}(a_{1i} g_1 + \cdots + a_{si} g_s) - \tfrac{t_{ji}}{c_j}(a_{1j} g_1 + \cdots + a_{sj} g_s) \\
&\quad - \textstyle\sum_{k=1}^{s'-1} q_k (a_{1k} g_1 + \cdots + a_{sk} g_s) \\
&= (g_1, \ldots, g_s) \cdot (\tfrac{t_j}{c_i \gcd(t_i, t_j)} a_i - \tfrac{t_i}{c_j \gcd(t_i, t_j)} a_j - q_1 a_1 - \cdots - q_{s'-1} a_{s'-1}) \\
&= (g_1, \ldots, g_s) \cdot (a_{1s'}, \ldots, a_{ss'})^{\mathrm{tr}} = a_{1s'} g_1 + \cdots + a_{ss'} g_s
\end{aligned}
$$

finishes the proof. $\qquad\square$

To show how this extended algorithm works in practice, let us apply it in the situation of Example 2.5.7.

Example 2.5.12. Let $n = 2$, let $r = 1$, let $M \subseteq P = K[x, y]$ be the ideal generated by $g_1 = x^2$ and $g_2 = xy + y^2$. As in Example 2.5.7, we follow the steps of the Buchberger Algorithm, except that we now use the extended version above.

1) Let $\mathcal{G} = (g_1, g_2)$, let $s' = 2$, let $\mathcal{A} = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$, and let $B = \{(1, 2)\}$.

2) Choose $(1, 2) \in B$ and set $B = \emptyset$.

3) We compute $S_{12} = -xy^2 = 0 \cdot g_1 + (-y) \cdot g_2 + y^3$ and let $q_1 = 0$, $q_2 = -y$, and $p = y^3$.

4) Let $s' = 3$, let $\mathcal{G} = (g_1, g_2, g_3)$ with $g_3 = y^3$, and let $B = \{(1, 3), (2, 3)\}$. We append the column vector $ya_1 - xa_2 - 0 \cdot a_1 + ya_2$ to the matrix $\mathcal{A}$ and get $\mathcal{A} = \left(\begin{smallmatrix} 1 & 0 & y \\ 0 & 1 & -x+y \end{smallmatrix}\right)$. Then we return to step 2).

2) Choose $(1, 3) \in B$ and set $B = \{(2, 3)\}$.

3) We compute $S_{13} = y^3 g_1 - x^2 g_3 = 0$ and return to step 2).

2) Choose $(2, 3) \in B$ and set $B = \emptyset$.

3) We compute $S_{23} = y^4 = 0 \cdot g_1 + 0 \cdot g_2 + y g_3$. Then we return to step 2).

2) Since $B = \emptyset$, we return the result $(\mathcal{G}, \mathcal{A})$, where $\mathcal{G} = (g_1, g_2, g_3)$ and $\mathcal{A} = \left(\begin{smallmatrix} 1 & 0 & y \\ 0 & 1 & -x+y \end{smallmatrix}\right)$.

Exercise 1. Let $P = K[x, y, z]$, let $\mathcal{F} = (x^2 - y, xy - z) \in P^2$, and let $\sigma = \mathtt{DegRevLex}$. Perform all steps of Buchberger's Algorithm applied to $\mathcal{F}$. Then find a term ordering σ such that $\mathcal{F}$ is a σ-Gröbner basis of the ideal $(x^2 - y, xy - z)$.

Exercise 2. Apply Buchberger's Algorithm as in Example 2.5.7 to compute a $\mathtt{DegLexPos}$-Gröbner basis of the submodule $M = \langle g_1, g_2, g_3, g_4 \rangle$ of $\mathbb{Q}[x, y]^3$ in the following cases.

a) $g_1 = (x^2, xy, y^2)$, $g_2 = (y, 0, x)$, $g_3 = (0, x, y)$, $g_4 = (y, 1, 0)$
b) $g_1 = (y - x, y, y)$, $g_2 = (xy, x, x)$, $g_3 = (x, y, y)$, $g_4 = (x, y, 0)$
c) $g_1 = (0, y, x)$, $g_2 = (0, x, xy - x)$, $g_3 = (y, x, 0)$, $g_4 = (y^2, y, 0)$

Exercise 3. In the cases of Exercise 2, determine representatives for a K-basis of $\mathbb{Q}[x, y]^3 / M$.

Exercise 4. Find out which module $M \subseteq \mathbb{Q}[x, y]^3$ in Exercise 2 contains the vector

$$m = (x^2 y - y^2 + xy^2,\ xy^2 - y^2 + x^2 + 2xy - x - y,\ x^2 y + xy^2 - 3xy + x)$$

Exercise 5. A polynomial $f \in P = K[x_1, \ldots, x_n]$ is called a **binomial** if it is of the form $f = at + a't'$ with $a, a' \in K \setminus \{0\}$ and $t, t' \in \mathbb{T}^n$. Let σ be a term ordering on $\mathbb{T}^n$ and I a **binomial ideal**, i.e. an ideal generated by binomials.

a) Prove that every element of the reduced σ-Gröbner basis of I is a binomial or a monomial.
b) Given a term $t \in \mathbb{T}^n$, show that $\mathrm{NF}_{\sigma, I}(t)$ is a scalar multiple of a term.

Exercise 6. Consider the polynomial ring $P = \mathbb{Q}[x, y]$, the P-submodule $M = \langle g_1, g_2, g_3, g_4 \rangle \subseteq P^3$ such that $g_1 = (xy, x, y)$, $g_2 = (y^2 + y, x + y^2, x)$, $g_3 = (-x, y, x)$, $g_4 = (y^2, y, x)$, and the module term ordering $\sigma = \mathtt{LexPos}$.

a) Using the algorithm given in Proposition 2.5.11, compute a σ-Gröbner basis $\{g_1, \ldots, g_{s'}\}$ of M, where $s' \geq 4$, and a matrix $\mathcal{A}$ such that $(g_1, \ldots, g_{s'}) = (g_1, \ldots, g_4) \cdot \mathcal{A}$.
b) Now use the method described in the proof of Proposition 2.4.13 to compute the reduced σ-Gröbner basis $\{g_1', \ldots, g_6'\}$ of M. Then find a matrix $\mathcal{A}'$ such that $(g_1', \ldots, g_6') = (g_1, \ldots, g_4) \cdot \mathcal{A}'$.
c) For the following elements of P^3, check whether they lie in M, and if they do, find their representations in terms of both $\{g_1', \ldots, g_6'\}$ and $\{g_1, \ldots, g_4\}$.
    1) $m_1 = (-2y, y - 1, xy + y)$
    2) $m_2 = (xy^5 - xy + y, xy^4 + x + 2y^2 - y, y^5 + xy)$
d) For the following pairs of elements of P^3, check whether $m_1 + M$ agrees with $m_2 + M$ in the residue class module P^3 / M.
    1) $m_1 = (2y, x^2 y + x^2 + xy + 2x - 3y, -x + y)$, $m_2 = (-x^2 + y - x, x^3 + 2x^2, x^2 - y)$
    2) $m_1 = (x^3 + x^2 + y - x, x^2 + x, x + y)$, $m_2 = (y, x^3 + 2x^2 - xy - y, 0)$

Tutorial 23: Buchberger's Criterion

In this tutorial we shall implement Buchberger's Criterion 2.5.3 and use it to decide whether certain sets of polynomials are Gröbner bases of the ideals they generate. As in the whole section, we let K be a field, we let $P = K[x_1, \ldots, x_n]$ be a polynomial ring over K, we let σ be a module term ordering on $\mathbb{T}^n\langle e_1, \ldots, e_r \rangle$, where $r \geq 1$, we let $\mathcal{G} = (g_1, \ldots, g_s)$ be a tuple of non-zero vectors, and we let $M \subseteq P^r$ be the P-submodule generated by the vectors in $\mathcal{G}$.

a) Write a CoCoA function `CheckGB(...)` which takes $\mathcal{G}$ and uses Buchberger's Criterion 2.5.3 to check whether it forms a σ-Gröbner basis of M. (*Hint:* You may want to use the function `NormalRemainder(...)` from Tutorial 15 or the built-in CoCoA function `NR(...)`.)

b) Let $G = \{x_2 - x_1^2, x_3 - x_1^3\} \subseteq \mathbb{Q}[x_1, x_2, x_3]$. Use the function `CheckGB(...)` to check whether G is a σ-Gröbner basis of the ideal it generates, where σ is one of the following term orderings: `Lex`, `DegLex`, `Ord(V)` where $V = \left(\begin{smallmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{smallmatrix} \right)$ or $V = \left(\begin{smallmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{smallmatrix} \right)$.

c) Use the function `CheckGB(...)` to determine which of the following systems of generators are Gröbner bases with respect to the stated term orderings of the ideals and modules they generate. In the first three cases, try to find a term ordering and a system of generators containing G such that Corollary 2.5.10 can be applied.

1) $G = \{x_1 x_2^2 - x_1 x_3 + x_2, x_1 x_2 - x_3^2, x_1 - x_2 x_3^4\} \subseteq \mathbb{Q}[x_1, x_2, x_3]$ with respect to `Lex`

2) $G = \{x_1^4 x_2^2 - x_3^5, x_1^3 x_2^3 - 1, x_1^2 x_2^4 - 2x_3\} \subseteq \mathbb{Q}[x_1, x_2, x_3]$ with respect to `DegLex`

3) $G = \{x_1 x_3 - x_2^2, x_1 x_4 - x_2 x_3, x_2 x_4 - x_3^2\} \subseteq \mathbb{Q}[x_1, x_2, x_3, x_4]$ with respect to `DegRevLex`

4) $G = \{(x_1^2 - x_2 x_3)(e_1 + e_2), (x_1 x_3 - x_2 x_4)(e_1 - e_2), (x_3^2 - x_1 x_4)e_1, (x_3^2 - x_1 x_4)e_2\} \subseteq \mathbb{Q}[x_1, x_2, x_3, x_4]^2$ with respect to `PosDegRevLex` and `DegRevLexPos`

5) $G = \{(x_1 - x_2^2)e_1, (x_1 - x_3^3)e_1, (x_2 - x_1^2)e_2, (x_2 - x_3^3)e_2, (x_3 - x_1^2)e_3, (x_3 - x_2^3)e_3\} \subseteq \mathbb{Q}[x_1, x_2, x_3]^3$ with respect to `PosDegRevLex` and `DegRevLexPos`

d) Let $n > 1$, and let σ be the lexicographic term ordering on $K[x_1, \ldots, x_n, y_1, \ldots, y_n]$ such that $x_1 >_\sigma \cdots >_\sigma x_n >_\sigma y_1 >_\sigma \cdots >_\sigma y_n$. Moreover, for $i = 1, \ldots, n$, let $s_i = \sum_{1 \leq j_1 < \cdots < j_i \leq n} x_{j_1} \cdots x_{j_i}$ be the i^{th} elementary symmetric polynomial in $x_1, \ldots, x_n$ (see also Tutorial 12), and let $h_{i,j} = \sum_{\alpha_j + \cdots + \alpha_n = i} x_j^{\alpha_j} \cdots x_n^{\alpha_n}$ for $i, j = 1, \ldots, n$. Use Buchberger's Criterion to prove that the polynomials

$$g_i = (-1)^i(y_i - s_i) + \sum_{j=1}^{i-1}(-1)^j h_{i-j,i}(y_j - s_j)$$

such that $i = 1, \ldots, n$ form a σ-Gröbner basis of the polynomial ideal $I = (y_1 - s_1, \ldots, y_n - s_n)$.

e) Verify the result of d) for $n = 1, \ldots, 5$ by applying your function $\texttt{CheckGB}(\ldots)$. Can you compute this for larger n? How far can you go?

Tutorial 24: Computing Some Gröbner Bases

The purpose of this tutorial is to implement a first version of Buchberger's Algorithm in the case of polynomial ideals, and to use it to study some particular examples. For instance, we will see that the elements of the reduced Gröbner basis of an ideal can have very high degree, even if the generators of the ideals have low degrees.

Then, for the specific ideal $I = (yz - z^2, xz - z^2, xy - z^2)$, you will be guided to find *all* possible reduced Gröbner bases of I, and to give a meaning to the picture on the cover of this book. As usual, we let $P = K[x_1, \ldots, x_n]$ be a polynomial ring over a field K.

a) Write a CoCoA function $\texttt{SPoly}(\ldots)$ which takes a tuple of non-zero polynomials $(g_1, \ldots, g_s)$ and indices $i, j \in \{1, \ldots, s\}$ with $i \neq j$ as arguments and returns the S-polynomial S_{ij} of g_i and g_j with respect to the current term ordering.

b) Implement Buchberger's Algorithm 2.5.5 in the case of polynomial ideals. To this end, write a CoCoA program $\texttt{FirstGB}(\ldots)$ which takes a tuple of non-zero polynomials generating the ideal and computes a Gröbner basis with respect to the current term ordering. (*Hint:* For step 3), use the built-in function $\texttt{NR}(\ldots)$ or $\texttt{NormalRemainder}(\ldots)$ of Tutorial 15.)

c) Using $\texttt{FirstGB}(\ldots)$, calculate the Gröbner bases of the following ideals with respect to the stated term orderings.

  1) $I = (x_2^2, x_1 x_2 x_3 + x_3^3) \subseteq \mathbb{Q}[x_1, x_2, x_3]$ with respect to $\texttt{DegRevLex}$
  2) $I = (x_1^2 x_2 - 1, x_1 x_2^2 - x_1) \subseteq \mathbb{Q}[x_1, x_2]$ with respect to $\texttt{Lex}$ and $\texttt{DegLex}$
  3) $I = (x_1 - x_3^4, x_2 - x_3^5) \subseteq \mathbb{Q}[x_1, x_2, x_3]$ with respect to $\texttt{Lex}$ and $\texttt{DegRevLex}$

d) Prove that for every number $m \geq 1$, the reduced Gröbner basis of

$$I_m = (x_1^{m+1} - x_2 x_3^{m-1} x_4, \; x_1 x_2^{m-1} - x_3^m, \; x_1^m x_3 - x_2^m x_4) \subseteq K[x_1, x_2, x_3, x_4]$$

with respect to $\texttt{DegRevLex}$ contains $f_m = x_3^{m^2+1} - x_2^{m^2} x_4$. Note that the degree $m^2 + 1$ of this polynomial is much higher than the degrees of the generators of I_m. Can you write down the whole reduced Gröbner basis of I_m with respect to $\texttt{DegRevLex}$? (Guess it or prove it!)

e) If you couldn't do the second part of d), calculate the reduced Gröbner basis of the ideal I_m with respect to $\texttt{DegRevLex}$ using $\texttt{FirstGB}(\ldots)$ for $m = 1, \ldots, 100$ and determine its length.

f) Prove that the ideal I_3 of part d) has (up to sign) the same reduced Gröbner bases with respect to `Lex` and `DegRevLex`. Does this hold for all $m \geq 1$?

In the remainder of this tutorial, we want to study the polynomial ideal $I = (xy - z^2, xz - z^2, yz - z^2)$ in $P = K[x, y, z]$. Although we are not going to use it, we mention that I is the ideal of all polynomials which vanish at three lines in $\mathbb{A}_K^3$ passing through the origin, or, equivalently, at three points in $\mathbb{P}_K^2$ (see Tutorials 27 and 35).

g) Let σ be any term ordering such that $x >_\sigma z$ and $y >_\sigma z$. Show that the reduced σ-Gröbner basis of I is $\{xz - z^2, yz - z^2, xy - z^2\}$.

h) Let σ be any term ordering such that $x >_\sigma z$ and $z >_\sigma y$. Show that the reduced σ-Gröbner basis of I is $\{xy - yz, xz - yz, z^2 - yz\}$.

i) Let σ be any term ordering such that $y >_\sigma z$ and $z >_\sigma x$. Show that the reduced σ-Gröbner basis of I is $\{z^2 - xz, yz - xz, xy - xz\}$.

j) Consider the situation where σ is a term ordering such that $z >_\sigma x$ and $z >_\sigma y$. Show that there are only two possible reduced Gröbner bases of I, according as $x >_\sigma y$ or $y >_\sigma x$. Observe that in both cases the number of elements in the reduced Gröbner basis is four.

k) Prove there are exactly five reduced Gröbner bases of I.

l) Group the term orderings in five classes, depending on the inequalities considered before. Then find five term orderings which give rise to the five reduced Gröbner bases you found above.

m) Prove that for each of the five reduced Gröbner bases, there is an infinite set of term orderings σ such that it is the reduced σ-Gröbner basis of I.

n) Consider the description of term orderings by matrices explained in Section 1.4. Try to use it to interpret the following picture.
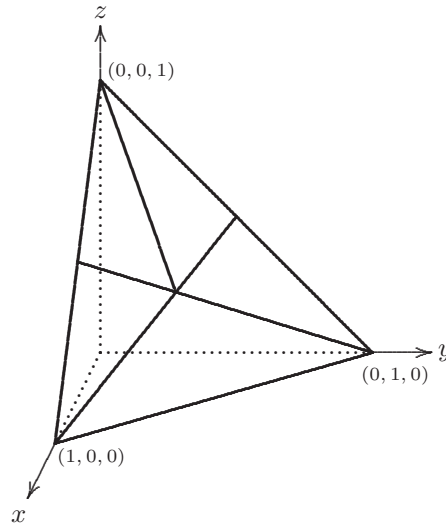
Tutorial 25: Some Optimizations of Buchberger's Algorithm

The purpose of this tutorial is to find and to implement optimized versions of Buchberger's Algorithm in the case of polynomial ideals. The amount of time consumed by a certain Gröbner basis computation depends largely on the number of critical pairs which have to be dealt with, and on the number of reduction steps which have to be performed in order to treat each critical pair. Therefore we will ask you to implement *counters* in your programs which measure these quantities, and we will judge our progress towards our goal of optimizing Buchberger's Algorithm by looking at the numbers returned by those counters.

Let $P = K[x_1, \ldots, x_n]$ be a polynomial ring over a field K, let $I \subseteq P$ be an ideal, and let $\mathcal{G} = (g_1, \ldots, g_s)$ be a tuple of non-zero polynomials which generate I. Furthermore, let σ be a term ordering, and let the elements $t_i, t_{ij} \in \mathbb{T}^n$, $\sigma_{ij} \in P^s$, and $S_{ij} \in P$ be defined as at the beginning of this section.

a) Update your CoCoA function `FirstGB`($\ldots$) from Tutorial 24 such that it returns not only a σ-Gröbner basis of I, but also the number of critical pairs (i, j) such that $S_{ij} \neq 0$, i.e. such that the normal remainder had to be computed, and the total number of reduction steps which were necessary to compute all those normal remainders.
   Hint: You will have to modify the function `NormalRemainder`($\ldots$) from Tutorial 15 suitably.

b) Apply your new function `FirstGB`($\ldots$) in the following five cases. Each time, compute a Gröbner basis with respect to `DegRevLex` and one with respect to `Lex`.

   1) $I = (x_1^2 - 2x_2^2 + 3x_1, \ x_1^3 - 2x_1x_2)$ in $\mathbb{Q}[x_1, x_2]$
   2) $I = (x_1 - 2x_3^4, \ x_2 - 3x_3^5)$ in $\mathbb{Q}[x_1, x_2, x_3]$
   3) $I = (x_1^2 - 2x_2^2, \ x_1^3 - 3x_3^3, \ x_1^4 - x_4^4)$ in $\mathbb{Q}[x_1, x_2, x_3, x_4]$
   4) $I = (x_1^3 - 4x_2^3, \ x_1^5 - 7x_3^5, \ x_1^7 - 11x_4^7)$ in $\mathbb{Q}[x_1, x_2, x_3, x_4]$
   5) $I = (x_1^2 + x_2^3 + x_3^3 - 1, \ x_1^3 + x_2^4 + x_3^5 - 1)$ in $\mathbb{Q}[x_1, x_2, x_3]$

c) Implement a CoCoA function `SecondGB`($\ldots$) which takes the list $\mathcal{G}$ and computes a σ-Gröbner basis of I via Buchberger's Algorithm 2.5.5, where the pair $(i, j) \in B$ is chosen in step 2) according to the normal selection strategy (see Remark 2.5.6.c), and where the optimization which follows from Proposition 2.5.8 is used.

d) Apply your function `SecondGB`($\ldots$) in the cases of b) and compare the results of your counters with those returned by the function `FirstGB`($\ldots$).

e) Given $1 \leq i < j < k \leq s$, find three terms $t, t', t'' \in \mathbb{T}^n$ such that $t\sigma_{ij} + t'\sigma_{jk} - t''\sigma_{ik} = 0$. Prove that one can choose $t = 1$ if and only if t_k divides $\mathrm{lcm}(t_i, t_j)$. Give similar criteria for $t' = 1$ and $t'' = 1$. The triple (i, j, k) is called a **Buchberger triple** if one can choose $t = 1$ or $t' = 1$ or $t'' = 1$.

f) Prove that one can drop a critical pair (i, j) in the execution of Buch-berger's Algorithm if it is contained in a Buchberger triple and if the other two pairs have been treated already. Write a CoCoA function $\texttt{ThirdGB}(\dots)$ which is based on $\texttt{SecondGB}(\dots)$ and adds this new optimization. To make sure that you do not drop more than one pair from a Buchberger triple, implement a list T which keeps track of the pairs which have been treated already.

g) Apply your function $\texttt{ThirdGB}(\dots)$ in the cases of b) and determine the improvement which has been achieved.

h) Start again with your implementation $\texttt{SecondGB}(\dots)$ of Buchberger's Algorithm, and replace step 4) by the following sequence of instructions.

4a) Increase s' by one. Append $g_i = \mathrm{NR}_{\sigma, \mathcal{G}}(S_{ij})$ to $\mathcal{G}$, and form the set $C = \{(i, s') \mid 1 \leq i < s',\ \gamma_i = \gamma_{s'}\}$.

4b) Delete in C all pairs (j, s') such that there exists an index i in $\{1, \dots, s'-1\}$ with the properties that $i < j$ and $t_{s'i}$ divides $t_{s'j}$.

4c) Delete in C all pairs (i, s') such that there exists an index j in $\{1, \dots, s'-1\}$ with the properties that $i < j$ and $t_{s'j}$ properly divides $t_{s'i}$.

4d) Delete in B all pairs (i, j) such that no divisibility occurs between $t_{s'i}$ and $t_{s'j}$ (hence both (i, s') and (j, s') survived the preceding two steps) and we have $\gcd(t_{is'}, t_{js'}) = 1$.

4e) Replace B by $B \cup C$ and continue with step 2).

The fact that this modified algorithm still computes a σ-Gröbner basis of M in finitely many steps will be studied in Volume 2. Implement it in a CoCoA function $\texttt{GoodGB}(\dots)$, apply this function in the cases of b), and compare the values returned by your counters with the earlier results.

2.6 Hilbert's Nullstellensatz

The art of doing mathematics
consists in finding that special case
which contains all the germs of generality.
(David Hilbert)

As in the first chapter, this closing section deviates from the main line of development. It is both a bridge to many applications of Computational Commutative Algebra and a foundation for numerous theoretical advances in later chapters. In the introduction of this book we mentioned that one of the most common areas where Computational Commutative Algebra is applied is algebraic geometry. The fundamental tool to translate statements from algebraic geometry into the language of commutative algebra and back is Hilbert's Nullstellensatz.

So, what is the relation between geometry and polynomials? Polynomial rings were introduced right at the beginning of this book. Since then, we kept trying to extract information from their intrinsic algebraic structure. But there is another way of looking at polynomials: they can be seen as *functions*. More precisely, given a polynomial f in the polynomial ring $K[x_1, \ldots, x_n]$ over a field K and an extension field $L \supseteq K$, we can evaluate f at each point of L^n and obtain a function from L^n to L.

Of special importance is then the set of *zeros* of f, i.e. the set of points $(a_1, \ldots, a_n) \in L^n$ such that $f(a_1, \ldots, a_n) = 0$. More generally, we can extend the setting to many polynomial equations and look for their common zeros. These are the geometric counterparts of polynomial ideals, and Hilbert's Nullstellensatz, in its different versions, provides the connection between both kinds of objects.

Since we are trying to be as self-contained as possible, we present a proof of Hilbert's Nullstellensatz in the current section. At several key points the theory of Gröbner bases will prove very useful. On the way, we shall also obtain a clearer picture of how the set of solutions of a system of polynomial equations depends on the field over which those equations are defined, and on the field where we look for the coordinates of the solution points. For instance, the polynomial $x^4 + 2x^2 + 1 \in \mathbb{R}[x]$ has no zeros in $\mathbb{R}$, but the two zeros i and $-i$ in $\mathbb{C}$. As we shall see, this simple special case already contains the germ of many more general phenomena.

The section begins with the proofs of some algebraic facts which lead to the field theoretic version of Hilbert's Nullstellensatz (see Theorem 2.6.6). This theorem can also be viewed as a structure theorem for maximal ideals in polynomial rings over algebraically closed fields (see Corollary 2.6.9). As a consequence, we are able to interpret the zeros of an ideal I in such a polynomial ring as the set of maximal ideals containing I (see Proposition 2.6.11). For instance, the zeros of $x^4 + 2x^2 + 1 \in \mathbb{C}[x]$ correspond to the maximal ideals $(x + i)$ and $(x - i)$ containing this polynomial.

Given a field extension $K \subset L$, an important result is proved about the behaviour of ideals under extension from $K[x_1, \ldots, x_n]$ to $L[x_1, \ldots, x_n]$. This result is the key to the weak form of Hilbert's Nullstellensatz (see Theorem 2.6.13) which provides us with an effective way to check whether a given polynomial ideal has zeros in $\overline{K}^n$, where $\overline{K}$ is the algebraic closure of K. For our ideal $(x^4 + 2x^2 + 1) \subseteq \mathbb{R}[x]$, the easy observation $1 \notin (x^4 + 2x^2 + 1)$ suffices to conclude that it has zeros in the algebraic closure $\mathbb{C}$ of $\mathbb{R}$.

Finally, we prove the Nullstellensatz in its full generality (see Theorem 2.6.16). It says that the operation of forming the vanishing ideal of a subset of $\overline{K}^n$ is an inverse to the operation of taking the set of zeros of a polynomial ideal if one considers radical ideals only. This highlights the importance of the ideal theoretic operation of forming the radical of an ideal. In the case of the principal ideal $(x^4 + 2x^2 + 1)$ in $\mathbb{C}[x]$, it says that the vanishing ideal of its set of zeros $\{i, -i\}$ is its radical ideal $(x^2 + 1)$.

2.6.A   The Field-Theoretic Version

Let K be an arbitrary field. Many algebraic geometers use the following terminology.

Definition 2.6.1. A finitely generated K-algebra is also called an **affine K-algebra**.

According to Corollary 1.1.14, such algebras are of the form P/I for some polynomial ring $P = K[x_1, \ldots, x_n]$ and some ideal $I \subseteq P$. Now we present three lemmas leading up to the first theorem of this section which is also called the field-theoretic version of Hilbert's Nullstellensatz. Recall that the field of fractions of a polynomial ring $P = K[x_1, \ldots, x_n]$ is usually denoted by $Q(P) = K(x_1, \ldots, x_n)$.

Lemma 2.6.2. *Let x be an indeterminate over our field K. Then $K(x)$ is not an affine K-algebra.*

Proof. Suppose $K(x) = K[\frac{f_1}{g_1}, \ldots, \frac{f_s}{g_s}]$ for some $f_1, \ldots, f_s, g_1, \ldots, g_s \in K[x]$ such that $g_1 \cdot g_2 \cdots g_s \neq 0$. Since $\frac{1}{x} \notin K[x]$, we may assume $g_1 \cdot g_2 \cdots g_s \notin K$. Then the fraction $\frac{1}{1+g_1 \cdot g_2 \cdots g_s}$ can be written as a polynomial expression in $\frac{f_1}{g_1}, \ldots, \frac{f_s}{g_s}$. Clearing denominators, we get $(g_1 \cdot g_2 \cdots g_s)^i = (1 + g_1 \cdot g_2 \cdots g_s) \cdot h$ for suitable $i > 0$ and $h \in K[x]$. Now $K[x]$ is a factorial domain (see Theorem 1.2.13), but clearly no irreducible factor of the non-constant polynomial $1 + g_1 \cdot g_2 \cdots g_s$ can divide one of the polynomials $g_1, \ldots, g_s$. This contradiction finishes the proof. $\square$

Lemma 2.6.3. *Let $A \subseteq B \subseteq C$ be three rings.*

  a) *If B is a finitely generated A-module, then it is also a finitely generated A-algebra.*

b) If B is a finitely generated A-algebra and if C is a finitely generated B-algebra, then C is a finitely generated A-algebra.

Proof. Let $\{b_1, \ldots, b_s\}$ be a set of generators of B as an A-module. Then $B = Ab_1 + \cdots + Ab_s \subseteq A[b_1, \ldots, b_s] \subseteq B$ implies claim a). For the proof of b), we use Corollary 1.1.14 to write $B = A[x_1, \ldots, x_n]/I$ with an ideal $I \subseteq A[x_1, \ldots, x_n]$ and $C = B[y_1, \ldots, y_m]/J$ with an ideal $J \subseteq B[y_1, \ldots, y_m]$. Then the claim follows from

$$C \cong A[x_1, \ldots, x_n, y_1, \ldots, y_m]/\big(I \cdot A[x_1, \ldots, x_n, y_1, \ldots, y_m] + \pi^{-1}(J)\big)$$

where $\pi : A[x_1, \ldots, x_n, y_1, \ldots, y_m] \longrightarrow B[y_1, \ldots, y_m]$ is the canonical homomorphism. $\qquad\square$

The next result is deeper. We want to show that under certain circumstances a K-subalgebra of an affine K-algebra is an affine K-algebra. The following example shows that this is not always the case.

Example 2.6.4. The K-subalgebra $K[x, xy, xy^2, xy^3, \ldots]$ of $K[x, y]$ is not finitely generated. Namely, for every finite set of elements of this subalgebra, the finitely many terms in the support of those polynomials can be written as polynomials in finitely many terms $x, xy, \ldots, xy^i$. But since we have $xy^i \notin K[x, xy, \ldots, xy^{i-1}]$ for $i \geq 2$, those polynomials do not generate the subalgebra.

Lemma 2.6.5. *Let A and B be two K-algebras such that $A \subseteq B$. Assume that B is an affine K-algebra and a finitely generated A-module. Then A is an affine K-algebra.*

Proof. Let $\{b_1, \ldots, b_s\}$ be a set of generators of B as a K-algebra and $\{\beta_1, \ldots, \beta_t\}$ a set of generators of B as an A-module. Then there are elements $a_{ij}, a'_{ijk} \in A$ such that we have expressions

$$b_i = \sum_{j=1}^{t} a_{ij}\beta_j \text{ for } i = 1, \ldots, s \quad \text{and} \quad \beta_i\beta_j = \sum_{k=1}^{t} a'_{ijk}\beta_k \text{ for } i, j = 1, \ldots, t.$$

Let A_0 be the K-subalgebra of A generated by all elements a_{ij} and a'_{ijk}. It is an affine K-algebra, hence Noetherian by Theorem 2.4.6. Assume for a moment that we know that B is a finitely generated A_0-module. Then B is a Noetherian A_0-module by Theorem 2.4.6 again. Thus A, an A_0-submodule of B, is a finitely generated A_0-module. Therefore A is a finitely generated A_0-algebra by Lemma 2.6.3.a. Since A_0 is an affine K-algebra, also A is an affine K-algebra by Lemma 2.6.3.b.

Consequently, to finish the proof it suffices to show that B is a finitely generated A_0-module. To this end we observe that every element of B is a polynomial expression in $\beta_1, \ldots, \beta_t$ with coefficients in A_0 because of the first set of expressions above. Using the second set of expressions, we can replace

every product $\beta_i \beta_j$ by an element of $A_0 \beta_1 + \cdots + A_0 \beta_t$. If we iterate those substitutions, we see that every element of B is in $A_0 \beta_1 + \cdots + A_0 \beta_t + A_0$, i.e. the A_0-module B is generated by $\{1, \beta_1, \ldots, \beta_t\}$. $\square$

Theorem 2.6.6. (Field-Theoretic Version of Hilbert's Nullstellensatz)

Let $P = K[x_1, \ldots, x_n]$ and $\mathfrak{m}$ a maximal ideal of P.

a) For every $i \in \{1, \ldots, n\}$, the intersection $\mathfrak{m} \cap K[x_i]$ is a non-zero ideal.

b) The affine K-algebra $P/\mathfrak{m}$ is a finitely generated K-vector space.

Proof. First we show that a) implies b). By assumption, for $i = 1, \ldots, n$, the intersection $\mathfrak{m} \cap K[x_i]$ is a non-zero principal ideal generated by some non-constant polynomial $f_i \in K[x_i]$. Then the ideal $\mathfrak{n} = (f_1, \ldots, f_n) \subseteq P$ is contained in $\mathfrak{m}$ and we have a surjective homomorphism $P/\mathfrak{n} \longrightarrow P/\mathfrak{m}$. Therefore it suffices to show that $P/\mathfrak{n}$ is a finitely generated K-vector space. Now $\{f_1, \ldots, f_n\}$ is a Gröbner basis of $\mathfrak{n}$ with respect to every term ordering σ by Corollary 2.5.10. If we write $\mathrm{LT}_\sigma(f_i) = x_i^{d_i}$ with $d_i \geq 1$ for $i = 1, \ldots, n$, we may deduce from Corollary 2.4.11 that a K-vector space basis of $P/\mathfrak{n}$ is given by the finite set of residue classes of the terms $x_1^{\alpha_1} \cdots x_n^{\alpha_n}$ such that $0 \leq \alpha_i < d_i$ for $i = 1, \ldots, n$.

Now we prove a) by induction on n. If $n = 1$, the ideal $\mathfrak{m}$ is a principal ideal generated by an irreducible polynomial, and the claim holds. If $n > 1$, we denote the affine K-algebra $P/\mathfrak{m}$ by B. Let $i \in \{1, \ldots, n\}$, let $\overline{x}_i$ be the residue class of x_i in B, and let A be the field of fractions of the integral domain $K[\overline{x}_i]$ contained in the field B. Clearly, considered as an A-algebra, B is generated by $\{\overline{x}_1, \ldots, \overline{x}_{i-1}, \overline{x}_{i+1}, \ldots, \overline{x}_n\}$. Hence B is isomorphic to $A[x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n]/\mathfrak{n}$ for some maximal ideal $\mathfrak{n}$. By induction and the implication shown above, B is a finitely generated A-vector space. Therefore, by Lemma 2.6.5, the field A is a finitely generated K-algebra. Then $\overline{x}_i$ is not an indeterminate over K by Lemma 2.6.2, and hence $\mathfrak{m} \cap K[x_i]$ is different from (0). $\square$

Condition a) of the above theorem does not hold for more general rings, as the following example shows. (In this example we shall use some facts about power series and Laurent series rings which will be discussed more thoroughly in Chapter V. The inexperienced reader may safely skip it.)

Example 2.6.7. Let $R = K[[x]][y]$ be the polynomial ring in the indeterminate y over the univariate power series ring $K[[x]]$ over a field K. Then the principal ideal $\mathfrak{m} = (xy - 1)$ is maximal, because $R/(xy - 1) \cong K[[x]]_x$ is a field. But we have $\mathfrak{m} \cap K[y] = (0)$.

Definition 2.6.8. A field K is called **algebraically closed** if every irreducible polynomial in $K[x]$ is linear. This implies that every polynomial $f \in K[x]$ of degree d can be written as

$$f(x) = c \, (x - a_1)^{\alpha_1} (x - a_2)^{\alpha_2} \cdots (x - a_s)^{\alpha_s}$$

where $c, a_1, \ldots, a_s \in K$ and $\alpha_1, \ldots, \alpha_s \in \mathbb{N}$ are such that $a_1, \ldots, a_s$ are pairwise distinct and $\alpha_1 + \cdots + \alpha_s = d$.

For instance, the **Fundamental Theorem of Algebra** says that the field of complex numbers $\mathbb{C}$ is algebraically closed. The fields $\mathbb{R}$ and $\mathbb{Q}$ are not algebraically closed, since the quadratic polynomial $x^2 + 1$ is irreducible over them.

An important result which you should know is that, for every field K, there exists an algebraic extension field $\overline{K}$ which is an algebraically closed field. (And if you do not know it, you can look for instance at [La70], Ch. 7.) The field $\overline{K}$ is unique up to a K-algebra isomorphism and is called the **algebraic closure** of K. For example, the field $\mathbb{C}$ is the algebraic closure of $\mathbb{R}$, since it is algebraically closed and an algebraic extension of $\mathbb{R}$. The algebraic closure of $\mathbb{Q}$ is the field of algebraic numbers $\overline{\mathbb{Q}}$ discussed in Tutorial 18.

The field-theoretic version of Hilbert's Nullstellensatz can also be interpreted as a structure theorem for maximal ideals in polynomial rings over algebraically closed fields.

Corollary 2.6.9. *Let K be an algebraically closed field, and let $\mathfrak{m}$ be a maximal ideal in $K[x_1, \ldots, x_n]$. Then there exist elements $a_1, \ldots, a_n$ in K such that*

$$\mathfrak{m} = (x_1 - a_1, \ldots, x_n - a_n)$$

Proof. Theorem 2.6.6 yields non-zero polynomials $f_1, \ldots, f_n \in \mathfrak{m}$ such that $f_i \in K[x_i]$ for $i = 1, \ldots, n$. Every polynomial f_i factorizes completely into linear factors, since K is algebraically closed. Moreover, the ideal $\mathfrak{m}$ is maximal, hence prime. This implies that it contains one of the linear factors of each polynomial f_i, say $x_i - a_i$. Then $\mathfrak{m}$ contains the ideal $(x_1 - a_1, \ldots, x_n - a_n)$ which, on the other hand, is a maximal ideal. Thus they must be equal and the proof is complete. $\square$

2.6.B   The Geometric Version

In the remainder of this section we want to explain the geometric versions of Hilbert's Nullstellensatz. The German word "Nullstellensatz" literally means "zero-places-proposition". Let us define what this refers to.

Definition 2.6.10. Let $K \subseteq L$ be a field extension, let $\overline{K}$ be the algebraic closure of K, and let $P = K[x_1, \ldots, x_n]$.

a) An element $(a_1, \ldots, a_n) \in L^n$ (which we shall also call a **point** of L^n) is said to be a **zero** of a polynomial $f \in P$ in L^n if $f(a_1, \ldots, a_n) = 0$, i.e. if the evaluation of f at the point $(a_1, \ldots, a_n)$ is zero. The set of all zeros of f in L^n will be denoted by $\mathcal{Z}_L(f)$. If we simply say that $(a_1, \ldots, a_n)$ is a zero of f, we mean $(a_1, \ldots, a_n) \in \overline{K}^n$ and $f(a_1, \ldots, a_n) = 0$.

b) For an ideal $I \subseteq P$, the **set of zeros** of I in L^n is defined as

$$\mathcal{Z}_L(I) = \{(a_1, \ldots, a_n) \in L^n \mid f(a_1, \ldots, a_n) = 0 \text{ for all } f \in I\}$$

Again we call the set of zeros of I in $\overline{K}^n$ simply the set of zeros of I and denote it by $\mathcal{Z}(I)$. Later we shall also call $\mathcal{Z}(I)$ the **affine variety** defined by I.

It is easy to see that the set of zeros $\mathcal{Z}_L(f)$ of a polynomial $f \in P$ in L^n agrees with the set of zeros $\mathcal{Z}_L((f))$ of the principal ideal it generates. Moreover, if an ideal $I \subseteq P$ is generated by a set of polynomials $\{f_1, \ldots, f_s\}$, then we have $\mathcal{Z}_L(I) = \cap_{i=1}^s \mathcal{Z}_L(f_i)$.

Algebraically, the set of zeros of an ideal corresponds to a set of maximal ideals in the polynomial ring, as our next proposition shows.

Proposition 2.6.11. *Let K be an algebraically closed field, let I be a proper ideal in $P = K[x_1, \ldots, x_n]$, and let Σ be the set of maximal ideals in P which contain I. Then the map*

$$\varphi : \mathcal{Z}(I) \longrightarrow \Sigma$$

defined by $\varphi(a_1, \ldots, a_n) = (x_1 - a_1, \ldots, x_n - a_n)$ is bijective.

Proof. For $p = (a_1, \ldots, a_n) \in K^n$, we denote by $\mathfrak{m}_p = (x_1 - a_1, \ldots, x_n - a_n)$ the corresponding maximal ideal in P. Then the map φ can be described by $\varphi(p) = \mathfrak{m}_p$. First we prove that φ is well-defined. For a point $p \in \mathcal{Z}(I)$, all polynomials in I vanish at p. Using Theorem 1.6.4, we get the representation $f = q_1(x_1 - a_1) + \cdots + q_n(x_n - a_n) + p$, and we see that those polynomials belong to $\mathfrak{m}_p$.

The map φ is clearly injective. Hence it suffices to show that it is surjective. We choose a maximal ideal $\mathfrak{m} \in \Sigma$. By Corollary 2.6.9, there exists a point $p = (a_1, \ldots, a_n) \in K^n$ such that $\mathfrak{m} = \mathfrak{m}_p$. By the definition of Σ, we have $\mathfrak{m}_p \supseteq I$. It follows that $p \in \mathcal{Z}(I)$, and the proof is complete. $\square$

Our next result is useful for comparing the set of zeros of I in L^n for different extension fields L of K. Remember that if $K \subseteq L$ is a field extension and I is an ideal of $P = K[x_1, \ldots, x_n]$, we use the notation $IL[x_1, \ldots x_n]$ to denote the ideal of $L[x_1, \ldots x_n]$ generated by the set I.

Proposition 2.6.12. *Let $K \subseteq L$ be a field extension and I an ideal of $K[x_1, \ldots, x_n]$. Then*

$$IL[x_1, \ldots, x_n] \cap K[x_1, \ldots, x_n] = I$$

In particular, we have $IL[x_1, \ldots, x_n] = L[x_1, \ldots, x_n]$ if and only if we have $I = K[x_1, \ldots, x_n]$.

Proof. Obviously we only need to prove that the left-hand side is contained in I. We choose a term ordering σ on $\mathbb{T}^n$ and let $G = \{g_1, \ldots, g_s\}$ be a σ-Gröbner basis of I. From Lemma 2.4.16 it follows that the set G is also a σ-Gröbner basis of the ideal $IL[x_1, \ldots, x_n]$. Now let f be a polynomial in $IL[x_1, \ldots, x_n] \cap K[x_1, \ldots, x_n]$. If we compute the normal form $\mathrm{NF}_\sigma(f)$ using the Division Algorithm 1.6.4, we only perform operations inside $K[x_1, \ldots, x_n]$, and therefore $f - \mathrm{NF}_\sigma(f)$ is in the ideal generated in $K[x_1, \ldots, x_n]$ by the set of polynomials $\{g_1, \ldots, g_s\}$, which is I. But $f \in IL[x_1, \ldots, x_n]$ implies $\mathrm{NF}_\sigma(f) = 0$, hence we have $f \in I$. $\square$

The questions which ideals have zeros and how one can check that are now answered by the following theorem and its corollary.

Theorem 2.6.13. (Weak Nullstellensatz)
Let K be a field, and let I be a proper ideal of $P = K[x_1, \ldots, x_n]$, i.e. let $I \subset P$. Then $\mathcal{Z}(I) \neq \emptyset$.

Proof. Let $\overline{K}$ be the algebraic closure of K, and let $\overline{P} = \overline{K}[x_1, \ldots, x_n]$. Then $I\overline{P}$ is a proper ideal of $\overline{P}$ by Proposition 2.6.12. Since we know that $\overline{P}$ is Noetherian, the ideal $I\overline{P}$ is contained in a maximal ideal $\mathfrak{m}$ of $\overline{P}$ by Proposition 2.4.5.c. Now Corollary 2.6.9 says that there is a point $(a_1, \ldots, a_n) \in \overline{K}^n$ such that $\mathfrak{m} = (x_1 - a_1, \ldots, x_n - a_n)$. Hence $(a_1, \ldots, a_n)$ is a zero of $\mathfrak{m}$, and therefore also of $I \subseteq I\overline{P} \subseteq \mathfrak{m}$. $\square$

Of course, one cannot hope to get $\mathcal{Z}_K(I) \neq \emptyset$ if K is not algebraically closed, since for instance $\mathcal{Z}_\mathbb{Q}(x^2 + 1) = \emptyset$. Moreover, although the question of whether $\mathcal{Z}_L(I) = \emptyset$ or not does not depend on which algebraically closed field $L \supseteq K$ we choose, the set $\mathcal{Z}_L(I)$ itself clearly does. For instance, if $I = (y - x^2) \subseteq \mathbb{Q}[x, y]$, then $(\pi, \pi^2) \in \mathcal{Z}_\mathbb{C}(I)$, but $(\pi, \pi^2) \notin \mathcal{Z}_{\overline{\mathbb{Q}}}(I)$.

Corollary 2.6.14. *Let L be a field which contains the algebraic closure of K, and let I be an ideal of $K[x_1, \ldots, x_n]$. Then the following conditions are equivalent.*

a) $\mathcal{Z}_L(I) = \emptyset$
b) $1 \in I$

In particular, this result holds if K is the field of definition of I.

Proof. Clearly b) implies a). For the converse, we observe that $\mathcal{Z}_L(I) = \emptyset$ implies $\mathcal{Z}(I) = \emptyset$, and then $1 \in I$ by the Weak Nullstellensatz. $\square$

Recall that, for a ring R and an ideal I in R, the set $\{r \in R \mid r^i \in I$ for some $i \geq 0\}$ is again an ideal of R which is called the **radical** of I and denoted by $\sqrt{I}$. An ideal I such that $I = \sqrt{I}$ is called a **radical ideal**. Equivalently, an ideal I is a radical ideal in R if the residue class ring R/I has no non-zero nilpotent elements. Thus, for instance, prime ideals are radical ideals.

In the case of an ideal I of $K[x_1, \ldots, x_n]$, it is easy to see that I and $\sqrt{I}$ have the same set of zeros. Thus the operation of assigning the set of zeros to an ideal $I \subseteq K[x_1, \ldots, x_n]$ is not one-to-one. In order to study this operation more closely, let us define an operation going in the other direction.

Definition 2.6.15. Let $K \subseteq L$ be a field extension, and let $S \subseteq L^n$. Then the set of all polynomials $f \in K[x_1, \ldots, x_n]$ such that $f(a_1, \ldots, a_n) = 0$ for all points $(a_1, \ldots, a_n) \in S$ forms an ideal of the polynomial ring $K[x_1, \ldots, x_n]$. This ideal is called the **vanishing ideal** of S in $K[x_1, \ldots, x_n]$ and denoted by $\mathcal{I}(S)$.

Using this notation, the strong version of Hilbert's Nullstellensatz says that the operation $\mathcal{I}(\ldots)$ is an inverse to $\mathcal{Z}(\ldots)$ if one considers only radical ideals in polynomial rings over algebraically closed fields.

Theorem 2.6.16. (Hilbert's Nullstellensatz)
Let K be an algebraically closed field, and let I be a proper ideal of $K[x_1, \ldots, x_n]$. Then

$$\mathcal{I}(\mathcal{Z}(I)) = \sqrt{I}$$

Proof. To show the inclusion $\mathcal{I}(\mathcal{Z}(I)) \supseteq \sqrt{I}$, suppose that a polynomial $f \in P = K[x_1, \ldots, x_n]$ satisfies $f^i \in I$ for some $i \geq 0$. Then we have $f^i(a_1, \ldots, a_n) = 0$ for every point $(a_1, \ldots, a_n) \in \mathcal{Z}(I)$. Thus we also have $f(a_1, \ldots, a_n) = 0$ for every point $(a_1, \ldots, a_n) \in \mathcal{Z}(I)$, i.e. $f \in \mathcal{I}(\mathcal{Z}(I))$.

To prove the other inclusion, we may assume that $\mathcal{I}(\mathcal{Z}(I)) \neq (0)$. We choose $f \in \mathcal{I}(\mathcal{Z}(I)) \setminus \{0\}$ and generators $\{g_1, \ldots, g_s\}$ of I. Let x_{n+1} be a new indeterminate, and consider the ideal $I' = IP[x_{n+1}] + (x_{n+1}f - 1)$ in the polynomial ring $P[x_{n+1}]$. For every point $(a_1, \ldots, a_{n+1}) \in \mathcal{Z}(I')$ we have $a_{n+1}f(a_1, \ldots, a_n) = 1$ and $g_i(a_1, \ldots, a_n) = 0$ for $i = 1, \ldots, s$. But then $(a_1, \ldots, a_n) \in \mathcal{Z}(I)$ and $f(a_1, \ldots, a_n) \neq 0$ contradict the choice of f. Consequently, such a point does not exist, i.e. $\mathcal{Z}(I') = \emptyset$, and the Weak Nullstellensatz 2.6.13 yields $1 \in I'$.

Therefore there are $s + 1$ polynomials $h, h_1, \ldots, h_s \in P[x_{n+1}]$ such that $1 = \sum_{i=1}^{s} h_i \cdot g_i + h \cdot (x_{n+1}f - 1)$. In the field $K(x_1, \ldots, x_n, x_{n+1})$ we may substitute $\frac{1}{f}$ for x_{n+1}. We get the equality

$$1 = \sum_{i=1}^{s} h_i(x_1, \ldots, x_n, \tfrac{1}{f}) \cdot g_i$$

By clearing the denominators, we find $f^m = \sum_{i=1}^{s} \tilde{h}_i \cdot g_i$ for some $m \geq 0$ and suitable polynomials $\tilde{h}_1, \ldots, \tilde{h}_s \in P$, which means that $f \in \sqrt{I}$. $\square$

Our final result in this section provides a reformulation of Hilbert's Nullstellensatz which will prove useful in the final section of this book.

Corollary 2.6.17. *Let K be a field, let I be a proper ideal in the polynomial ring $P = K[x_1, \ldots, x_n]$, let $\overline{K}$ be the algebraic closure of K, let*

$\overline{P} = \overline{K}[x_1, \ldots, x_n]$, *and let* f *be a polynomial in* P. *If* f *belongs to all maximal ideals containing* $I\overline{P}$, *then* $f \in \sqrt{I}$.

Proof. First we observe that Proposition 2.6.11 implies $f \in \mathcal{I}(\mathcal{Z}(I\overline{P}))$. This ideal equals $\sqrt{I\overline{P}}$ by Hilbert's Nullstellensatz 2.6.16. Therefore there exists a number $i \in \mathbb{N}$ such that $f^i \in I\overline{P}$. The claim now follows from Proposition 2.6.12. $\qquad\square$

Exercise 1. Give a direct proof for the fact that condition b) of Theorem 2.6.6 implies condition a) of that theorem.

Exercise 2. Let K be a field and $f(x) \in K[x]$ an irreducible polynomial. Find a maximal ideal $\mathfrak{m}$ in $K[x, y]$ which contains the ideal $I = (f(x), f(y))$, and compute the intersection of $\mathfrak{m}$ with $K[x]$ and $K[y]$.

Exercise 3. (Structure of Maximal Ideals in $\mathbb{R}[x_1, \ldots, x_n]$)
Let $\mathfrak{m}$ be a maximal ideal in $\mathbb{R}[x_1, \ldots, x_n]$.
- a) Let $n = 1$. Show that $\mathfrak{m}$ is either generated by a polynomial of type $x_1 - a$ with $a \in \mathbb{R}$, or a polynomial of type $x_1^2 + ax_1 + b$ with $a, b \in \mathbb{R}$ and $a^2 - 4b < 0$.
  Hint: Use the fact that if $a + ib$ is a complex zero of a polynomial in $\mathbb{R}[x]$, then also $a - ib$ is a zero, to characterize irreducible polynomials in $\mathbb{R}[x]$.
- b) Let $n = 2$ and $f_1 = x_1^2 + a_1 x_1 + b_1$, $f_2 = x_2^2 + a_2 x_2 + b_2$ with $a_1, b_1, a_2, b_2 \in \mathbb{R}$ and $a_1^2 - 4b_1 < 0$, $a_2^2 - 4b_2 < 0$. Show that the ideal $I = (f_1, f_2)$ is not maximal in $\mathbb{R}[x_1, x_2]$.
  Hint: Use the fact that $\mathbb{R}[x_1, x_2]/(f_1)$ is isomorphic to $\mathbb{C}[x_2]$.
- c) Let $n = 2$ and assume that $x_1^2 + a_1 x_1 + b_1 \in \mathfrak{m}$ with $a_1, b_1 \in \mathbb{R}$ and $a_1^2 - 4b_1 < 0$. Show that there exist $a_2, b_2 \in \mathbb{R}$ such that we have $\mathfrak{m} = (x_1^2 + a_1 x_1 + b_1, x_2 - a_2 x_1 - b_2)$.
- d) In the general case prove the following fact: either there exist numbers $a_1, \ldots, a_n \in \mathbb{R}$ such that $\mathfrak{m} = (x_1 - a_1, \ldots, x_n - a_n)$ or, up to a permutation of the indeterminates, there exist $a_1, b_1, a_2, b_2 \ldots, a_n, b_n \in \mathbb{R}$ such that $a_1^2 - 4b_1 < 0$ and $\mathfrak{m} = (x_1^2 + a_1 x_1 + b_1, x_2 - a_2 x_1 - b_2, \ldots, x_n - a_n x_1 - b_n)$.

Exercise 4. Let $f \in \mathbb{Q}[x, y]$ be a non-constant polynomial. Prove that $\mathcal{Z}_{\mathbb{Q}}(f) \subset \mathcal{Z}(f)$.

Exercise 5. Let $K \subseteq L$ be a field extension, let $P = K[x_1, \ldots, x_n]$, let $f, f_1, \ldots, f_s \in P$, and let $I = (f_1, \ldots, f_s)$.
- a) Show that $\mathcal{Z}_L(f) = \mathcal{Z}_L((f))$.
- b) Show that $\mathcal{Z}_L(I) = \cap_{i=1}^{s} \mathcal{Z}_L(f_i)$.
- c) Show that $\mathcal{Z}_L(I) = \mathcal{Z}_L(\sqrt{I})$.

Exercise 6. Let $K \subseteq L$ be a field extension, let I be an ideal in $K[x_1, \ldots, x_n]$, and let S be a subset of L^n.
- a) Show that $\mathcal{I}(\mathcal{Z}_L(I)) \supseteq I$.
- b) Show that $\mathcal{Z}_L(\mathcal{I}(S)) \supseteq S$.

Exercise 7. Let R be a ring, and let I and J be ideals in R. Prove the following rules.

a) $\sqrt{\sqrt{I}} = \sqrt{I}$

b) $\sqrt{I \cap J} = \sqrt{IJ} = \sqrt{I} \cap \sqrt{J}$

c) $\sqrt{I^i} = \sqrt{I}$ for all $i \geq 1$.

d) If I is an intersection of prime ideals, then $\sqrt{I} = I$.

e) $\sqrt{\sqrt{I} + \sqrt{J}} = \sqrt{I + J}$

Exercise 8. Let K be an algebraically closed field and I a proper ideal of $P = K[x_1, \ldots, x_n]$. In this exercise we use the Zariski topology on K^n defined in Tutorial 27.

The ideal I is said to be **reducible** if it is the intersection of two strictly bigger ideals. A closed set of a topological space is said to be **reducible** if it is the union of two properly contained closed subsets.

a) Show that if $\mathcal{Z}(I)$ is reducible, then I is reducible.

b) Give an example which shows that the converse is not true.

c) Show that the converse of a) is true if I is a radical ideal.

d) Let I be a radical ideal. Prove that the following conditions are equivalent.

  1) There exist two ideals $I_1, I_2 \subset P$ such that $I = I_1 \cap I_2$ and $I_1 + I_2 = P$.

  2) $\mathcal{Z}(I)$ is disconnected, i.e. it is the union of two disjoint closed sets.

Exercise 9. Let K be an algebraically closed field, and let I be a proper radical ideal of $P = K[x_1, \ldots, x_n]$.

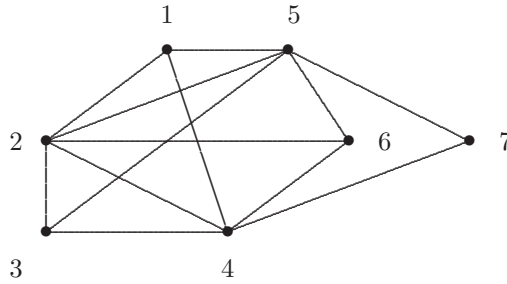a) Prove that $\mathcal{Z}(I)$ is finite if and only if I is of the form

$$I = \bigcap_{i=1}^s (x_1 - a_{i1}, \ldots, x_n - a_{in})$$

with pairwise different points $(a_{11}, \ldots, a_{1n}), \ldots, (a_{s1}, \ldots, a_{sn}) \in K^n$.

b) Show that if a) holds, then $\mathcal{Z}_L(I) = \mathcal{Z}(I)$ for every extension field $L \supseteq K$.

Tutorial 26: Graph Colourings

Suppose we are given 3 different colours and a graph Γ having n nodes and at most one arch between any two nodes, e.g.

Our goal is to find out if the nodes can be coloured in such a way that no arch connects two nodes of the same colour. In order to use the theory of Gröbner bases to solve this problem, we introduce the following notation.

The colours will be called -1, 0, and 1. They will be identified with the elements of the field $\mathbb{F}_3 = \mathbb{Z}/(3)$. For $i = 1, \ldots, n$, we choose an indeterminate x_i and form the polynomial ring $P = \mathbb{F}_3[x_1, \ldots, x_n]$. We shall identify a colouring of the graph with a point of $\mathbb{F}_3^n$ such that the i^{th} coordinate of the point corresponds to the colour of the i^{th} node.

a) Show that the set of zeros of the ideal $(x_1^3 - x_1, \ldots, x_n^3 - x_n)$ is precisely the set of all colourings.

b) Prove that the i^{th} and j^{th} node of the graph have different colours if and only if the colouring is a zero of the polynomial $x_i^2 + x_i x_j + x_j^2 - 1$.

c) In addition, we may assume that the first and second nodes are connected, that the first node has colour "0", and that the second node has colour "1". What polynomial equations does this imply for the colourings under consideration?

d) Write a CoCoA program $\texttt{Colouring}(\ldots)$ which takes a list of pairs from $\{1, \ldots, n\}^2$ representing the arches and computes an ideal $I \subseteq P$ whose zeros are precisely the colourings of the graph represented by those pairs which satisfy our additional conditions.

e) Apply your function $\texttt{Colouring}(\ldots)$ to the graph above. Then use CoCoA to compute the reduced $\texttt{Lex}$-Gröbner basis of this ideal. Does the graph have a colouring of the desired kind? If yes, how many different ones? (*Hint:* Use Hilbert's Nullstellensatz to interpret the answer of your calculation.)

f) Consider the graph formed by connecting the center of a regular 7-gon to its vertices. (It has 8 nodes and 14 arches.) Use CoCoA and the Weak Nullstellensatz to show that this graph cannot be coloured as required above.

Tutorial 27: Affine Varieties

Let $K \subseteq L$ be a field extension. For every ideal I in $K[x_1, \ldots, x_n]$ we consider the set $\mathcal{Z}_L(I) \subseteq L^n$ as given in Definition 2.6.10. For the moment, let us call a subset of L^n a **zero-set** if it is of the form $\mathcal{Z}_L(I) \subseteq L^n$ for some ideal $I \subseteq K[x_1, \ldots, x_n]$.

a) Prove the following claims.

    1) $\emptyset$ is a zero-set.
    2) L^n is a zero-set.
    3) If $E_1, \ldots, E_s$ are zero-sets, then $\cup_{i=1}^s E_i$ is a zero-set.
    4) If J is a set of indices and $\{E_j\}_{j \in J}$ a set of zero-sets indexed by J, then $\cap_{j \in J} E_j$ is a zero-set.

    Deduce that the zero-sets of L^n can be taken as the closed sets of a topology, which we denote by $\mathrm{Top}_{K,L}$. If $K = L$, then $\mathrm{Top}_{K,K}$ is called the **Zariski topology** on K^n. Moreover, K^n with the Zariski topology is called the n-**dimensional affine space** over K and denoted by $\mathbb{A}_K^n$. Zero-sets in $\overline{K}^n$ are called **affine varieties** (or **affine sets**).

b) Let p_1 and p_2 be two distinct points in K^n. Then the set of points $\overline{p_1 p_2} = \{p_1 + \lambda(p_2 - p_1) \mid \lambda \in K\}$ is called the **line** passing through p_1 and p_2. Show that $\overline{p_1 p_2}$ is a closed set in the Zariski topology.

c) Let $K \subset K' \subseteq L$ be field extensions and $\mathrm{Top}_{K,L}$, $\mathrm{Top}_{K',L}$ the corresponding topologies on L^n. Show that $\mathrm{Top}_{K',L}$ is finer than $\mathrm{Top}_{K,L}$, i.e. that every closed set with respect to $\mathrm{Top}_{K,L}$ is also closed with respect to $\mathrm{Top}_{K',L}$.

d) Let K be algebraically closed and consider the Zariski topology on K^n. Show that there is a bijection between the set of radical ideals in $K[x_1, \ldots, x_n]$ and the closed sets in $\mathbb{A}_K^n$. Then show that the statement is false if K is not algebraically closed.

e) Let K be algebraically closed, let I be an ideal in $P = K[x_1, \ldots, x_n]$, and assume that the dimension of P/I as a vector space over K is finite. Then show that $\mathcal{Z}(I)$ is a finite set of points.
Hint: For $i = 1, \ldots, n$, show that $I \cap K[x_i] = (f_i) \neq (0)$ and conclude that $\mathcal{Z}(I) \subseteq \mathcal{Z}((f_1, \ldots, f_n))$.

f) Let $P = \mathbb{Q}[x_1, \ldots, x_n]$, and let I be an ideal in P. Write two CoCoA programs which perform the following tasks.

    1) Given I and a point $p \in \mathbb{Q}^n$, check if $p \in \mathcal{Z}_{\mathbb{Q}}(I)$ and return the corresponding Boolean value.
    2) If $I \neq (0)$, find a point p which is not in $\mathcal{Z}_{\mathbb{Q}}(I)$ and return it.

g) Let S be a subset of $\mathbb{A}_K^n$. Show that the set of all Zariski-closed subsets of $\mathbb{A}_K^n$ containing S has a unique minimal element (with respect to inclusion). This zero-set is called the **Zariski closure** of S.

h) Let K be an algebraically closed field, and let $S \subseteq \mathbb{A}_K^n$. Prove that the Zariski closure of S is given by $\mathcal{Z}(\mathcal{I}(S))$.