Algebra IV

doc. Lukáš Vokřínek, PhD.

February 13, 2024

Contents

Introduction

Introduction will be here at some point (or not).

Lukáš Vokřínek

Syllabus

Syllabus will be here at some point (or not).

1. Noetherian rings

Definition 1.1. Let A be a ring. We say that an A-module M is *noetherian*, if it satisfies the *ascending chain condition* for submodules, i.e. if there exists no strictly increasing chain

$$M_0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_n \subsetneq \cdots$$

of submodules of M. As a special case, we say that A is a noetherian ring, if it is noetherian as an A-module, i.e. if it satisfies the ascending chain condition for ideals in A.

Theorem 1.2. *An A-module M is noetherian iff every submodule of M is finitely generated.*

Proof. Assume that M is noetherian and let $L \subseteq M$ be infinitely generated. We construct inductively a stricly increasing sequence of finitely generated submodules $L_n \subseteq L$ in the following way: we start with $L_0 = 0$ and then inductively $L_n \subsetneq L$ for otherwise L would be finitely generated and we set $L_{n+1} = L_n + Rx_{n+1}$ where $x_{n+1} \in L \setminus L_n$.

For the opposite implication, assume that every submodule of M is finitely generated and that $M_0 \subseteq M_1 \subseteq \cdots$ is a sequence of submodules of M. Then $M_\infty = \cup_n M_n$ is a submodule. It is finitely generated by assumption, $M_\infty = R\{x_1, \ldots, x_k\}$ and since each x_i lies in some M_j, there exists n such that $x_1, \ldots, x_k \in M_n$. Then $M_n = M_{n+1} = \cdots$. □

Theorem 1.3. *Let $0 \to M' \xrightarrow{\alpha} M \xrightarrow{\beta} M'' \to 0$ be a short exact sequence of A-modules. Then M is noetherian iff both M' and M'' are noetherian.*

Proof. If M is noetherian then the lattice of submodules of both M' and $M'' \cong M/M'$ are sublattices of the lattice of ideals of M and as such do not contain an infinite chain.

Assume conversely that both M', M'' are noetherian and let $M_0 \subseteq M_1 \subseteq \cdots$ be a sequence of submodules. Then $M'_n = \alpha^{-1}(M_n)$ is constant for $n \gg 0$ and so is $M''_n = \beta(M_n)$. But then so must be M_n: for if $x \in M_{n+1}$, then $\beta(x) \in M''_{n+1} = M''_n$ and so $\beta(x) = \beta(y)$ for some $y \in M_n$. Analogically, $x - y = \alpha(z)$ for some $z \in M'_n$, and thus $x = y + \alpha(z) \in M_n$. (Alternatively: the inclusion $M_n \to M_{n+1}$ is an extension of inclusions $M'_n \to M'_{n+1}$ and $M''_n \to M''_{n+1}$, which are isomorphisms for $n \gg 0$ and 5-lemma gives the result.) □

Proof. Assume conversely that M', M'' are noetherian and let $L \subseteq M$ be a submodule. Then for $L' = \alpha^{-1}(L)$, $L'' = \beta(L)$ we get a short exact sequence

$$0 \to L' \to L \to L'' \to 0.$$

Since both $L' \subseteq M'$ and $L'' \subseteq M''$ are finitely generated, so is L.

Corollary 1.4. *If A is a noetherian ring, the every finitely generated module M is noetherian.*

Proof. The sum of two modules can be expressed via a (split) short exact sequence

$$0 \to M' \to M' \oplus M'' \to M'' \to 0,$$

the previous theorem thus shows that every finitely generated free module A^n is noetherian and also every quotient of it, i.e. any finitely generated module. □

In the proceeding, the commutativity assumption is crucial.

Definition 1.5. An *A-algebra* is a homomorphism of rings $\rho\colon A \to B$. Mostly, it will be a mono and we will thus think of B as a supring of A.

Example 1.6. $A[x_1, \dots, x_n]$ is an A-algebra.

Since B is canonically a B-module, by restricting scalars along ρ we may also treat it as an A-module. An alternative definition of an A-algebra is as an A-module B together with an A-bilinear mapping $B \times B \to B$ (multiplication) that, together with the addition, makes B into a ring (this means that B is a monoid object in $A-\mathsf{Mod}$).

Definition 1.7. We say that an A-algebra B is *finitely generated*, when there exist $b_1, \dots, b_n \in B$ that generate B as an A-algebra, i.e. via addition, multiplication and multiplication by scalars from A. We write $B = A[b_1, \dots, b_n]$.

We say that an A-algebra B is *finite*, when B is a finitely generated A-module (i.e. there exist $b_1, \dots, b_n \in B$ that generate B via addition and multiplication by scalars from B). We write $B = A\{b_1, \dots, b_n\}$.

We remark that finite generation is equivalent to the existence of a surjective homomorphism of A-algebras $A[x_1, \dots, x_n] \to B$ (sending x_i to the generators b_i of B; this is so because $A[x_1, \dots, x_n]$ is a free A-algebra on generators $x_1, \dots, x_n$). For any finite A-algebra there exists a surjective homomorphism of A-modules $A\{x_1, \dots, x_n\} \to B$.

Theorem 1.8. *Let A be a noetherian ring and B a finite A-algebra. Then B is also a noetherian ring.*

Proof. By the corollary, B is a noetherian A-module, so every A-submodule of B is finitely generated as an A-module. This implies easily that every ideal of B (B-submodule $\Rightarrow$ A-submodule) is a finite generated as an ideal (B-submodule). □

Example 1.9. The ring $\mathbb{Z}$ is noetherian. Then also $\mathbb{Z}[i] = \{a + bi \mid a, b \in \mathbb{Z}\}$ is noetherian.

Theorem 1.10. *Let A be a noetherian ring, $D \subseteq A$ a multiplicative subset. The also the localization $D^{-1}A$ is a noetherian ring.*

Proof. Again, the lattice of ideals of $D^{-1}A$ is a sublattice of the one for A. □

Theorem 1.11 (Hilbert basis theorem)**.** *If A is noetherian, then so is $A[x]$.*

Proof. Let $I \subseteq A[x]$ be an ideal. We define an ideal

$$J = \{a \in A \mid \exists p \in I \colon p = ax^r + \text{lot}\},$$

i.e. the ideal of leading coefficients of the polynomials from I. Let $J = (a_1, \dots, a_k)$ and pick polynomials $p_i \in I$ with leading coefficients a_i; we may assume that they all have the same degree r. The set $A_{<r}[x]$ of polynomials of degree smaller than r forms a finitely generated A-module, thus noetherian and we may write $A_{<r}[x] \cap I = A\{q_1, \dots, q_l\}$. We then claim that $I = (p_1, \dots, p_k, q_1, \dots, q_l)$: since every $p \in I$ of degree smaller than r lies in $(q_1, \dots, q_l)$, we consider $p \in I$ of degree at least r. Then $p = ax^s + \text{lot}$, with $a \in J$. Therefore

$$p = (b_1 a_1 + \dots + b_k a_k)x^s + \text{lot} = b_1 x^{s-r} p_1 + \dots + b_k x^{s-r} p_k + \text{lot},$$

where the first k terms lie in $(p_1, \dots, p_k)$ and the rest has smaller degree, so it lies in $(p_1, \dots, p_k, q_1, \dots, q_l)$ by induction. □

Corollary 1.12. *Let A be a noetherian ring. Every finitely generated A-algebra B is also a noetherian ring.*

Proof. We have $B \cong A[x_1, \dots, x_n]/I$, where $A[x_1, \dots, x_n]$ is noetherian according to the previous theorem and thus so is its quotient B: the lattice of ideals of $A[x_1, \dots, x_n]/I$ is a sublattice of the one for $A[x_1, \dots, x_n]$. □

The motivation for the Gröbner basis comes from the following problems:

- decide whether a polynomial f belongs to an ideal $I = (g_1, \dots, g_s)$
- decide whether two ideals are equal $(g_1, \dots, g_s) \stackrel{?}{=} (h_1, \dots, h_t)$

In one variable, the first problem can be solved easily by dividing f/g with a remainder. Thus, let us abstractly describe the division algorithm for polynomials f/g in one variable: we consider the leading terms of the polynomials, divide these $q = \operatorname{LT} f / \operatorname{LT} g$ and then write

$$f/g = q \cdot g + (f - q \cdot g)/g$$

where the term $f - q \cdot g$ has smaller leading term so we can proceed inductively. The situation with more variables has two problems that we will have to deal with. Firstly, the notion of a leading term is not obvious; this will be solved by considering an ordering on monomials (as an extra structure) and the result will depend on this ordering. The second problem is that ideals in $\Bbbk[x_1, \dots, x_n]$ are not principal and for applications we will then have to divide by multiple polynomials at the same time (we will see this very shortly) and the result may depend on the choice of the polynomial used for the denominator.

As explained above, we will need a monomial order and mostly we will suffice with the so called lexicographical ordering:

$$x^\alpha = {x_1}^{\alpha_1} \cdots {x_n}^{\alpha_n} > {x_1}^{\beta_1} \cdots {x_n}^{\beta_n} = x^\beta,$$

iff for some $i \geq 1$ we have $\alpha_1 = \beta_1, \dots, \alpha_{i-1} = \beta_{i-1}, \alpha_i > \beta_i$. Since this is a linear order, we may speak of the leading term of a polynomial $f \in \Bbbk[x_1, \dots, x_n]$: when

$$f = a_\alpha x^\alpha + \sum_{\beta < \alpha} a_\beta x^\beta = a_\alpha x^\alpha + \text{lot}$$

with $a_\alpha \neq 0$, we call $\operatorname{LC} f = a_\alpha$ the *leading coefficient*, $\operatorname{LM} f = x^\alpha$ the *leading monomial* and $\operatorname{LT} f = a_\alpha x^\alpha$ the *leading term*.

Somewhat more generally on monomial orders: we require that it should be a total order closed under the multiplication, i.e. $x^\alpha \leq x^\beta$ and $x^\gamma \leq x^\delta \Rightarrow x^{\alpha+\gamma} \leq x^{\beta+\delta}$ and (often) also $1 \leq x^\alpha$. The second condition is equivalent to the monomial order being a well order. To prove this, we introduce a (useful) notion of a monomial ideal, i.e. an ideal generated by a collection of monomials f_s. One can show that this is exactly

$$\{f \mid \text{each monomial contained in } f \text{ is divisible by some } f_s\}$$

(by observing that the right hand side is indeed an ideal, contains f_s and is smallest among such). Now we are ready to prove that the monomial order is a well order: given a non-empty collection of monomials, consider a monomial ideal generated by them. By Hilbert basis theorem, it is finitely generated, so generated by a finite subcollection. Clearly any of the old generators (in fact any monomial in the ideal) is divisible by one of the new generators, thus sits above that generator in the order. We have thus reduced to a finite subset that has a smallest element since we assume the order to be total.

Exercise 1.13. Show that an ideal J is monomial iff it satisfies $f \in J \Leftrightarrow$ each monomial contained in f belongs to J.

We will now give a definition that will allow us to describe a simple generalization of the membership algorithm $f \stackrel{?}{\in} (g_1, \dots, g_s)$ from the univariate to the multivariate case.

Definition 1.14. We say that elements $g_1, \dots, g_s \in I$ form a *Gröbner basis* if the leading monomial of any $g \in I$ is divisible by the leading monomial of some g_i; in other words

$$(\operatorname{LM} g \mid g \in I) = (\operatorname{LM} g_1, \dots, \operatorname{LM} g_s)$$

according to the characterization of monomial ideals above.

With a Gröbner basis one can easily test membership $f \in I$: first we verify whether $\operatorname{LM} f$ is divisible by some $\operatorname{LM} g_i$. If not we get $f \notin I$. If $\operatorname{LM} f = x^\alpha \operatorname{LM} g_i$, we replace f by the polynomial

$$f - \frac{\operatorname{LC} f}{\operatorname{LC} g_i} x^\alpha g_i$$

and we continue with testing (here we use that monomials are well ordered).

Remark. We say that a Gröbner basis of an ideal I is *reduced* if all its members g_i are monic and $\operatorname{LM} g_i$ does not divide any term of any g_j (this is in analogy with the reduced echelon form of a matrix, which is at the same time a special case).

It holds that every ideal has a *unique* reduced Gröbner basis (we will not prove this). In the proceeding we explain how to compute such a basis. Then testing equality of two ideals becomes easy – we compute the reduced Gröbner bases and compare them. Even without reduced Gröbner bases, one can simply test whether each generator of one ideal belongs to the other and the other way around.

1.1. Buchberger algorithm

The algorithm for finding a Gröbner basis of an ideal $I = (f_1, \dots, f_l)$ proceeds in steps as follows: we compute for each f_i, f_j the so called *S-polynomial* $S(f_i, f_j)$ by determining the least common multiple x^α of the monomials $\operatorname{LM} f_i$, $\operatorname{LM} f_j$ and we set

$$S(f_i, f_j) = \frac{x^\alpha}{\operatorname{LT} f_i} f_i - \frac{x^\alpha}{\operatorname{LT} f_j} f_j.$$

Then we use $f_1, \dots, f_l$ to reduce the result, i.e. we subsequently subtract monomial multiples of the f_k's to exactly cancel the actual leading term (i.e. we use division with a remainder). If we get a nonzero polynomial whose leading term is no longer divisible by the leading term of any f_k, we add the result to the set of generators, so that l increases by one and the bigger system of polynomials clearly generates I (the added polynomial may depend on the reduction, since we may choose multiple times the generator f_k whose multiple gets subtracted). Since in each step the ideal $(\operatorname{LM} f_1, \dots, \operatorname{LM} f_l)$ is enlarged and since $\Bbbk[x_1, \dots, x_n]$ is noetherian, after a finite number of steps we get to the situation where the reductions of all S-polynomials are zero. We will now show that the generating set then forms a Gröbner basis: Let $f \in I = (f_1, \dots, f_l)$, so that $f = p_1 f_1 + \dots + p_l f_l$, and we assume that in this expression the monomial

$$\max\{\operatorname{LM}(p_i f_i) \mid i = 1, \dots, l\}$$

is the smallest possible; let i_0 be an index for which the maximum above is achieved. There are two possibilities:

- the leading terms do not cancel out, i.e. $\operatorname{LM} f = \operatorname{LM}(p_{i_0} f_{i_0})$; then $\operatorname{LM} f \in (\operatorname{LM} f_1, \dots, \operatorname{LM} f_l)$;
- the leading terms do cancel out; then for indices $i \neq i_0$ with $\operatorname{LM}(p_i f_i)$ maximal, we may replace
$$p_i f_i - \frac{\operatorname{LC}(p_i f_i)}{\operatorname{LC}(p_{i_0} f_{i_0})} p_{i_0} f_{i_0} = q_i S(f_i, f_{i_0}) + \text{lot}$$
(the non-leading terms of p_i and p_{i_0} contribute to the lower order terms "lot", the leading terms yield a multiple of the S-polynomial, since this was obtained as the *smallest* monomial combination in which the leading terms cancel out). By construction, every S-polynomial $S(f_i, f_{i_0})$ can be replaced by a combination of the f_j with smaller leading terms, and the terms in "lot" already have smaller leading terms; this gives a contradiction with minimality.

Example 1.15. Compute the Gröbner basis of $I = (f_1, f_2)$, where $f_1 = x^3 - 2xy$, $f_2 = x^2 y + x - 2y^2$.

Solution. In the first step

$$S(f_1, f_2) = y f_1 - x f_2 = -x^2 \qquad f_3 = x^2$$

and no reduction is necessary. In the next step, the reduction of $S(f_1, f_2) = -f_3$ is zero, further

$$\begin{aligned} S(f_1, f_3) &= f_1 - x f_3 = -2xy & f_4 &= xy \\ S(f_2, f_3) &= f_2 - y f_3 = x - 2y^2 & f_5 &= x - 2y^2 \end{aligned}$$

and again, no reductions are necessary. In fact, it is now possible to throw out f_1, f_2, since they lie in (f_3, f_4, f_5). Let us compute

$$\begin{aligned} S(f_3, f_4) &= y f_3 - x f_4 = 0 \\ S(f_3, f_5) &= f_3 - x f_5 = 2xy^2 \equiv 0 \\ S(f_4, f_5) &= f_4 - y f_5 = 2y^3 & f_6 &= y^3 \end{aligned}$$

and now it is possible to leave out $f_3 = x f_5 + 2y f_4$ a $f_4 = y f_5 + 2 f_6$. In the last step

$$S(f_5, f_6) = y^3 f_5 - x f_6 = -2y^5 \equiv 0$$

Therefore (f_5, f_6) is a reduced Gröbner basis. ◇

Example 1.16. Compute the Gröbner basis of $I = (f_1, f_2, f_3)$, where $f_1 = x^2 + y^2 + z^2 - 1$, $f_2 = x^2 - y + z^2$, $f_3 = x - z$.

Solution. It will be convenient to write the subtraction of multiples of f_i as reductions $x^2 \equiv -y^2 - z^2 + 1$, $x^2 \equiv y - z^2$, $x \equiv z$, etc. In the first step we get

$$\begin{aligned} S(f_1, f_2) &= f_1 - f_2 = \underline{y^2} + y - 1 & f_4 &= y^2 + y - 1 \\ S(f_1, f_3) &= f_1 - x f_3 = y^2 + z^2 - 1 + \underline{xz} \equiv y^2 + 2z^2 - 1 & f_5 &= y^2 + 2z^2 - 1 \\ S(f_2, f_3) &= f_2 - x f_3 = -y + z^2 + \underline{xz} \equiv -y + 2z^2 & f_6 &= y - 2z^2 \end{aligned}$$

It is now possible to throw out $f_1 = f_2 + f_4$, $f_2 = (x+z)f_3 - f_4$, $f_4 = f_5 + f_6$ so that we have

$$\begin{aligned} S(f_3, f_5) &= y^2 f_3 - x f_5 = -y^2 z - \underline{2xz^2} + x \equiv -y^2 z - 2z^3 + \underline{x} \\ &\equiv \underline{-y^2 z} - 2z^3 + z \equiv -(1 - 2z^2)z - 2z^3 + z = 0 \\ S(f_3, f_6) &= y f_3 - x f_6 = -yz + \underline{2xz^2} \equiv \underline{-yz} + 2z^3 \\ &\equiv -2z^3 + 2z^3 \equiv 0 \\ S(f_5, f_6) &= f_5 - y f_6 = 2z^2 - 1 + \underline{2yz^2} \equiv 4z^4 + 2z^2 - 1 \qquad f_7 = z^4 + (1/2)z^2 - 1/4 \end{aligned}$$

Again we can leave out $f_5 = (y + 2z^2)f_6 + 4f_7$, so the Gröbner basis is (f_3, f_6, f_7).

As an application, we can now solve the system of equations $f_1 = f_2 = f_3 = 0$. This is equivalent to the system $f_3 = f_6 = f_7 = 0$ and similarly to linear systems we may now compute the solution "from the back": by solving $f_7 = 0$ we get $z = \frac{\sqrt{-1 \pm \sqrt{5}}}{2}$. Substituting into $f_6 = 0$ we then obtain $y = 2z^2 = -2 \pm 2\sqrt{5}$ an finally by substituting into $f_3 = 0$ gives $x = z = \frac{\sqrt{-1 \pm \sqrt{5}}}{2}$. ◇

Example 1.17. Compute the Gröbner basis of $I = (f_1, f_2)$, where $f_1 = x^2 - y$, $f_2 = x^2 + (y-1)^2 - 1$.

Solution. In the first step

$$S(f_1, f_2) = f_1 - f_2 = -y^2 + y \qquad f_3 = y^2 - y$$

and no reduction is necessary. In the next step we can leave out $f_2 = f_1 - f_3$, further

$$S(f_2, f_3) = y^2 f_2 - x^2 f_3 = \underline{x^2 y} + y^4 - 2y^3 \equiv y^2 + y^4 - 2y^3 \equiv 0$$

(any power y^k, $k \geq 1$ reduces to y just using f_3) and the reduced Gröbner basis is (f_1, f_3).

The quotiend $\Bbbk[x,y]/I$ or perhaps rather $\Bbbk[x,y]/\sqrt{I}$ has a close connection to the solution set of $f_1 = 0$, $f_2 = 0$. It consists of three points $[0,0]$, $[-1,1]$, $[1,1]$ and therefore $\dim \Bbbk[x,y]/\sqrt{I} = 3$. At the same time $\dim \Bbbk[x,y]/I = 4$, since the point $[0,0]$ should be taken "twice", concretely $x(y-1) \notin I$, but $(x(y-1))^2 \in I$, so that $x(y-1) \in \sqrt{I} \setminus I$ (the function $x(y-1)$ vanishes on the three points, but not up to a sufficiently high order). ◇

Lemma 1.18. *If* $\mathrm{LM}(f)$, $\mathrm{LM}(g)$ *are coprime, then* $S(f,g)$ *can be reduced to zero, using only* f, g.

Proof. For simplicity, we may assume f, g monic. By assumption $S(f,g) = \mathrm{LM}(g)f - \mathrm{LM}(f)g$ and in each step we will subtract a multiple of the from tf where t is a term of g or adding a multiple of the form sg where s is a term of f, in such a way that in the end the S-polynomial will reduce to $gf - fg = 0$ (the point is that every term st turns up once with a plus sign and once with a minus sign and it can only be a leading term when s is a leading term of f or t is a leading term of g). □

HW 1. Solve the following system of polynomial equations using Gröbner bases

$$\begin{aligned} x^2 + y + z &= 1 \\ x + y^2 + z &= 1 \\ x + y + z^2 &= 1 \end{aligned}$$

1.2. The confluence approach

In the case of one variable, understanding polynomials modulo g is quite simple computationally. One can always simplify any polynomial f to its remainder modulo g and two polynomials are congruent iff they give the same remainder. We will now try to outline a theory for more variables that still allows one to associate remainders (more precisely canonical forms) that are in bijection with congruence classes. The congruence relation will be expressed via a simpler relation of reduction.

Consider a monic polynomial $g = x^\beta - r$ with $x^\beta > LMr$. First, simply subtracting g from x^β yields r and we represent it as

$$x^\beta \to_g r$$

(replacement of x^β by r, like in a substitution in systems of equations). More generally, we get a similar rule by subtracting an ax^α-multiple

$$ax^\alpha \cdot x^\beta \to_g ax^\alpha \cdot r$$

and we may think of this again as replacing the copy of x^β in the product on the left by r. Yet more generally, if a polynomial f contains a term $ax^\alpha \cdot x^\beta$, we write

$$\begin{aligned} f &\to_g f' = f - ax^\alpha \cdot g \\ \cdots + ax^\alpha \cdot x^\beta + \cdots &\to_g \cdots + ax^\alpha \cdot r + \cdots \end{aligned}$$

and in effect this yet again replaces this particular appearance of x^β by r. More precisely, we may say that $f \to_g f'$ at $x^\alpha \cdot x^\beta$. Of course, if $a = 0$ we get $f' = f$. We may describe the result of $f \to_g f'$ at $x^\alpha \cdot x^\beta$ equivalently in the following way: f' is the unique polynomial that differs form f by a scalar multiple of $x^\alpha \cdot g$ and has coefficient at $x^\alpha \cdot x^\beta$ zero. Using this, if f and f' differ by a scalar multiple of $x^\alpha \cdot g$, reducing both polynomials at $x^\alpha \cdot x^\beta$ clearly yields the same f'':

$$f \to_g f''_g \leftarrow f'$$

This implies easily that two polynomials are congruent modulo (g) iff they can be joined by an (arbitrarily long) zig-zag of reductions $\to_g$.

More generally, for a set of polynomials G we denote by $\to_G$ the union of the above reductions $\to_g$ with g ranging over G. In fact, we will denote by $\to_G$ the closure under iterations, so that $f \to_G f'$ if f can be reduced to f' by a finite sequence of reductions $\to_g$ with g ranging over G. If we want to emphasize that there is exactly one reduction used we write $f \to_g^1 f'$. Generalizing the above, we see that f and f' are congruent modulo (G) iff f and f' can be joined by a zig-zag of reductions $\to_G$. Our goal for this section is to understand when zig-zags of reductions can be replaced by reductions – this will then give the canonical forms and decision for congruence, as promised.

There are two important properties that a reduction can satisfy: termination and confluence. The *termination* property asserts that any sequence $f_0 \to_G f_1 \to_G \cdots$ is eventually constant. Termination holds always as follows from the well foundedness of the monomial order (assume that the squence does not stabilize and consider, for each n, the biggest monomial in f_n that admits a G-reduction; it must exist for otherwise f_n would be reduced and the sequence would stabilize; but now the sequence of these monomials must stabilize; next, consider the sequence of the second biggest such monomials etc.; the stabilized monomials

then form an infinite decreasing sequence, giving a contradiction). The termination implies that, starting from any given f and applying reductions, at some point we arrive at a reduced polynomial h, i.e. one that does not allow any nontrivial reduction. We will say that h is a normal form of f.

Heading towards uniqueness of normal forms, the *confluence* property asserts that given any f, any two reductions of it admit a common reduction, i.e. we can complete the diagram:

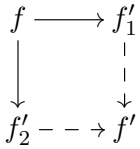

Clearly, if this is the case and both f_1' and f_2' are reduced then they must be equal to f' and thus equal to each other. Consequently, a normal form of a polynomial is unique

Proposition 1.19. *Suppose that $\to_G$ is confluent. Then two polynomials are congruent modulo (G) iff they have the same normal form. In other words, the canonical map*

$$\{\textit{normal forms}\} \to \Bbbk[\mathbb{x}]/(G)$$

is a bijection.

Proof. We have seen that f_1 and f_2 are congruent modulo (G) iff they can be joined by a zig-zag of G-reductions. Confluence implies that we may then replace this by the bottom span in

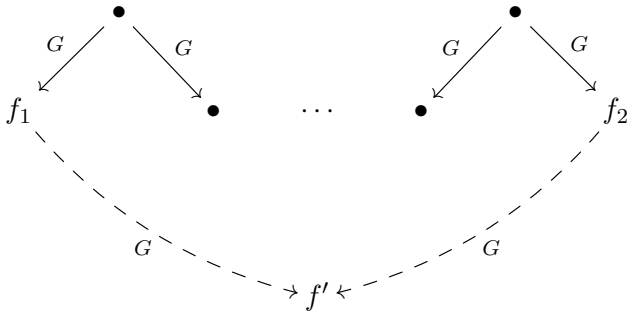

and thus the normal form of f_1 equals that of f' and symmetrically for f_2. □

In the presence of termination, we will now show that the full strength of confluence is implied by a weaker version with both reductions of f being one-step reductions. Assuming this weaker version true, we temporarily call f bad if it admits two non-confluent reductions and we split each of them into its first step and the rest, as in the solid part of the diagram below. If none of f_1'', f_2'', f'' was bad then we could complete the picture starting from the top left square and obtain that also the two given reductions of f were confluent:

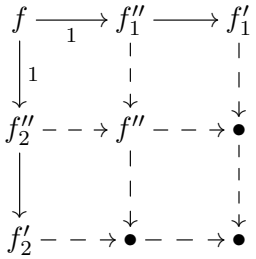

Thus, at least one of f_1'', f_2'', f'' must be bad and we can then proceed inductively and construct in this way an infinite sequence of nontrivial reductions, yielding a contradiction.

The bad news is that confluence does not always hold for $\to_G$, but we will see that it does hold if (and only if) G is a Gröbner basis. We will need a useful observation that uses the additive structure of polynomials:

Lemma 1.20. *Two reductions of f as in*

$$\begin{array}{ccc} f & \longrightarrow & f_1' \\ \downarrow & & \vdots \\ f_2' & \dashrightarrow & f' \end{array}$$

are confluent if $f_2' - f_1' \to_G 0$.

Proof. Decompose the reduction as

$$f_2' - f_1' \to_g^1 f'' \to_G 0$$

with the first step happening at x^α and apply corresponding reductions

$$f_1' \to_g^1 f_1'' \text{ at } x^\alpha, \quad f_2' \to_g^1 f_2'' \text{ at } x^\alpha.$$

It is then easy to see that $f'' = f_2'' - f_1''$. Proceeding in this way, we produce $f_1' \to_G h_1$ and $f_2' \to_G h_2$ with $0 = h_2 - h_1$, so that $h_1 = h_2$ is the required f'. □

Theorem 1.21. *The reduction $\to_G$ is confluent iff G is Gröbner.*

Proof. Assuming G Gröbner, any $h \in (G)$ may be reduced at its leading monomial to obtain $h \to_G h'$ with $h' \in (G)$ smaller (in terms of its LM) so that we obtain $h \to_G 0$ by well foundedness. Now for any two reductions

$$\begin{array}{ccc} f & \longrightarrow & f_1' \\ \downarrow & & \vdots \\ f_2' & \dashrightarrow & f' \end{array}$$

we have $f_2' - f_1' \in (G)$ and thus $f_2' - f_1' \to_G 0$, implying confluence through the previous lemma.

In the opposite direction, if $h \in (G)$ then h is congruent to 0 modulo (G) and thus they have the same normal form. Since 0 is reduced (having no term), this means in effect that $h \to_G 0$. Finally, since this reduction must eliminate the leading monomial of h at some point, that leading monomial must be divisible by one of the $\operatorname{LM} g_i$ and G is indeed Gröbner. □

We have just seen that confluence is equivalent to $\forall f \in (G) \colon f \to_G 0$. We will now show that it is enough to check this condition for some very special elements: Consider $g_1, g_2 \in G$ and denote by x^β the least common multiple of $\operatorname{LM} g_1$, $\operatorname{LM} g_2$. We define the *S-polynomial*

$$S(g_1, g_2) = \frac{x^\beta}{\operatorname{LM} g_1} g_1 - \frac{x^\beta}{\operatorname{LM} g_2} g_2.$$

Clearly, the S-polynomial belongs to (G), showing the necessity in the next theorem.

Theorem 1.22. *The reduction $\to_G$ is confluent iff $S(g_1, g_2) \to_G 0$ for each $g_1, g_2 \in G$.*

Proof. To prove sufficiency, consider two one-step reductions $f \to^1_{g_1} f_1'$ and $f \to^1_{g_2} f_2'$ as in the lemma. If they happen at different monomials then $f_1' - f_2'$ is a linear combination of g_1 and g_2 with non-cancelling leading terms, so that we can easily reduce $f_2' - f_1' \to_G 0$ using only g_1, g_2 (first use the one with the bigger leading monomial). If the reduction happens at the same monomial then this monomial must be divisible by x^β from the definition of the S-polynomial. Writing the corresponding term of f as $ax^\alpha \cdot x^\beta$, it is easy to see that

$$f_2' - f_1' = ax^\alpha \cdot \left(\frac{x^\beta}{\operatorname{LM} g_1} g_1 - \frac{x^\beta}{\operatorname{LM} g_2} g_2 \right) = ax^\alpha \cdot S(g_1, g_2).$$

Since this reduces to zero, the previous lemma gives confluence. □

This gives correctness of the Buchberger algorithm. Every time we add anything to G, the monomial ideal generated by the leading monomials of G increases, so the algortihm must terminate by Hilbert basis theorem. In that case, $S(g_1, g_2) \to_G 0$ for all $g_1, g_2 \in G$, so the reduction $\to_G$ is confluent and thus G is Groöbner.

We will now give a few applications.

- membership test: Given f and G, decide whether $f \in (G)$. The algorithm first enlarges G to a Gröbner basis and then reduces f to its normal form $f \to_G h$. Now $f \in (G)$ iff $h = 0$.
- equality test: Given G and H, decide whether $(G) = (H)$. The algorithm first enlarges both G and H to Gröbner bases and then tests whether $\forall g \in G\colon g \in (H)$ and also the symmetric version.
- elimination ideal: For the lexicographical order of monomials with $\mathtt{x} > \mathtt{y}$, let G be a Gröbner basis. Then $\Bbbk[\mathtt{y}] \cap (G) = (\Bbbk[\mathtt{y}] \cap G)$, since any $f \in \Bbbk[\mathtt{y}] \cap (G)$ reduces to zero, $f \to_G 0$ and in this process we may only use elements of G lying in $\Bbbk[\mathtt{y}]$ by our assumptions on the monomial order.
- systems of polynomial equations: Let G be a system of polynomials over an algebraically closed field. Then $G = 0$ implies $f = 0$ iff some power of f lies in (G). Thus, the system implies some equation $f = 0$ with $f \in \Bbbk[y]$ iff (G) contains some such f and by the above, this is equivalent to the Gröbner basis containing some such f. This allows one to compute solutions of systems of polynomial equations to some extent: Find such an f, find its roots, plug in one after another into the remaining polynomials and thus continue with fewer variables.
- intersection of ideals: One checks that $(G) \cap (H) = \Bbbk[\mathtt{x}] \cap ((1-t)G, tH)$, where the right hand side takes place inside $\Bbbk[t, \mathtt{x}]$. Thus, one may compute the intersection using the elimination ideal method, this time with $t > \mathtt{x}$.

2. Invariant theory

This is a nice application of the Hilbert basis theorem. We assume here that $\Bbbk$ is a field of characteristic coprime to the order of a finite group G (this condition will also be important for the representation theory later in the course). We will consider an action of G on the polynomial ring $\Bbbk[\mathtt{x}]$. The invariants (the collection of invariant polynomials in this case) is

the subset $\Bbbk[\mathbb{x}]^G = \{f \in \Bbbk[\mathbb{x}] \mid \forall a \in G\colon a \cdot f = f\}$. As an example, the symmetry group S_n acts on the variables and thus on the polynomials, e.g. $(1\ 2) \cdot x_1^2 x_2 = x_1 x_2^2$. The main theorem in this respect is that for the elementary symmetric polynomials

$$s_1 = x_1 + \cdots + x_n,\ s_2 = x_1x_2 + \cdots + x_ix_j + \cdots + x_{n-1}x_n,\ \ldots,\ s_n = x_1 \cdots x_n$$

the canonical map

$$\Bbbk[\mathbb{y}] \to \Bbbk[\mathbb{x}],\ y_i \mapsto s_i$$

is an isomorphism onto the invariants $\Bbbk[\mathbb{x}]^{S_n}$. The action of $\sigma \in S_n$ on $\Bbbk[\mathbb{x}]$ is obtained from the action on variables in two steps:

$$\begin{array}{ccc} \{\mathbb{x}\} & \xrightarrow{\mathsf{Set}} & \{\mathbb{x}\} \\ \downarrow & & \downarrow \\ \Bbbk\{\mathbb{x}\} & \xrightarrow{\mathsf{Vect}} & \Bbbk\{\mathbb{x}\} \\ \downarrow & & \downarrow \\ \Bbbk[\mathbb{x}] & \xrightarrow{\mathsf{Alg}} & \Bbbk[\mathbb{x}] \end{array}$$

The action on the set of variables induces a linear action on the vector space of linear forms (i.e. essentially on $\Bbbk^n$) and that induces an algebra action on the polynomial ring, i.e. it satisfies

$$a \cdot (f + g) = a \cdot f + a \cdot g,\ a \cdot 1 = 1,\ a \cdot (fg) = (a \cdot f)(a \cdot g).$$

It follows that the G-invariants form a $\Bbbk$-subalgebra.

Theorem 2.1 (Hilbert's on finite generation of invariants)**.** *The $\Bbbk$-algebra $\Bbbk[\mathbb{x}]^G$ is finitely generated, i.e. there exists a surjection $\Bbbk[\mathbb{y}] \twoheadrightarrow \Bbbk[\mathbb{x}]^G$.*

Proof. Denote by $i\colon \Bbbk[\mathbb{x}]^G \to \Bbbk[\mathbb{x}]$ the inclusion. In this way, we can think of $\Bbbk[\mathbb{x}]$ as an algebra over $\Bbbk[\mathbb{x}]^G$. We will now construct a retraction p in the category of $\Bbbk[\mathbb{x}]^G$-modules

$$\begin{array}{ccccc} & & \Bbbk[\mathbb{x}] & & \\ & \overset{i}{\nearrow} & & \overset{p}{\searrow} & \\ \Bbbk[\mathbb{x}]^G & & \xrightarrow[1]{} & & \Bbbk[\mathbb{x}]^G \end{array}$$

by the formula $p(f) = \frac{1}{|G|} \cdot \sum_{a \in G} a \cdot f$ (the average of the elements in the orbit of f). The compatibility of the action with the algebra structure gives for $f \in \Bbbk[\mathbb{x}]^G$:

$$a \cdot (fg) = (a \cdot f)(a \cdot g) = f(a \cdot g)$$

so that the action is indeed by $\Bbbk[\mathbb{x}]^G$-linear maps and consequently so is p. Now $\operatorname{im} p \subseteq \Bbbk[\mathbb{x}]^G$

since

$$\begin{aligned} b \cdot p(f) &= b \cdot (\frac{1}{|G|} \cdot \sum_{a \in G} a \cdot f) \\ &= \frac{1}{|G|} \cdot \sum_{a \in G} b \cdot (a \cdot f) \\ &= \frac{1}{|G|} \cdot \sum_{a \in G} (ba) \cdot f \\ &= \frac{1}{|G|} \cdot \sum_{a \in G} a \cdot f \\ &= p(f) \end{aligned}$$

since as a runs over elements of G, ba runs over the same set of elements (in a different order, i.e. $a \mapsto ba$ is a bijection). Finally $pi = 1$ since, for $f \in \Bbbk[\mathbb{x}]^G$, we have

$$p(f) = \frac{1}{|G|} \cdot \sum_{a \in G} a \cdot f = \frac{1}{|G|} \cdot \sum_{a \in G} f = f.$$

Now let $I \subseteq \Bbbk[\mathbb{x}]$ be the ideal generated by the homogeneous elements of $\Bbbk[\mathbb{x}]^G$ of positive degree, so that

$$\Bbbk^{(>0)}[\mathbb{x}]^G \subseteq I \subseteq \Bbbk[\mathbb{x}].$$

By Hilbert basis theorem, we get $I = (f_1, \ldots, f_k)$ with f_i some of the above generators, i.e. homogeneous elements of $\Bbbk[\mathbb{x}]^G$ of positive degree. We claim that $\Bbbk[\mathbb{x}]^G$ is generated as a $\Bbbk$-algebra by the same set of elements $f_1, \ldots, f_k$. Since the G-action respects the degrees of polynomials, a polynomial is G-invariant iff all its homogeneous components are G-invariant that leaves us to prove $f \in \Bbbk[\mathbb{x}]^G$ homogeneous $\Rightarrow$ $f \in \Bbbk[f_1, \ldots, f_k]$. We prove this by induction on $\deg f$. If $\deg f = 0$, there is nothing to prove, so assume that $\deg f > 0$. Now $f \in I = (f_1, \ldots, f_k)$ so we have an expression

$$f = f_1 g_1 + \cdots + f_k g_k$$

and we may assume that all the g_i are homogeneous (replace the g_i by their homogeneous components of the appropriate degrees and the equality will remain valid). Now apply the retraction p to get

$$f = f_1 p(g_1) + \cdots + f_k p(g_k)$$

(both the left hand side and the f_i are G-invariant and p is $\Bbbk[\mathbb{x}]^G$-linear). By induction, we may assume that all $p(g_i)$ already lie in $\Bbbk[f_1, \ldots, f_k]$. Thus, the same is true for f. □

3. Localization

Definition 3.1. A *local ring* is a ring (commutative with 1) with a unique maximal ideal.

Theorem 3.2. *A ring A is local iff its non-units form an ideal. In that case this ideal is the unique maximal ideal of A.*

Proof. The implication $\Rightarrow$ follows from $(a) \subseteq M$ for every non-unit a, the opposite implication is obvious. □

Definition 3.3. Let A be a ring and $D \subseteq A$ a *multiplicative subset*, i.e. a subset satisfying $1 \in D$ and $x, y \in D \Rightarrow xy \in D$. The decomposition of $A \times D$ with respect to the equivalence relation

$$(a_1, d_1) \sim (a_2, d_2) \quad \Leftrightarrow \quad \exists d \in D\colon (a_1 d_2 - a_2 d_1)d = 0$$

will be denoted $D^{-1}A$ and called the *localization* of the ring A with respect to the multiplicative subset D. Its class will be denoted $[a, d] = \frac{a}{d}$. A ring structure on $D^{-1}A$ is introduced by the formulas

$$\frac{a_1}{d_1} + \frac{a_2}{d_2} = \frac{a_1 d_2 + a_2 d_1}{d_1 d_2}, \qquad \frac{a_1}{d_1} \cdot \frac{a_2}{d_2} = \frac{a_1 a_2}{d_1 d_2}.$$

The mapping $\lambda\colon A \to D^{-1}A$, $a \mapsto \frac{a}{1}$ is a homomorphism of rings.

The localization $D^{-1}A$ has the following universal property, saying that it is the universal supring where all the elements $d \in D$ admit an inverse.

Theorem 3.4. *Let $\rho\colon A \to B$ be a ring homomorphism such that $\rho(d) \in B^\times$ is a unit for all $d \in D$. Then there exists a unique ring homomorphism $\widetilde{\rho}\colon D^{-1}A \to B$ such that $\rho = \widetilde{\rho}\lambda$.*

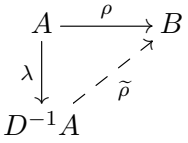

Proof. Since $\frac{a}{d} = \lambda(a)\lambda(d)^{-1}$, we are forced to set $\widetilde{\rho}(\frac{a}{d}) = \rho(a)\rho(d)^{-1}$. We will now show that this is a well defined mapping; we will leave to the reader to show that it is then a ring homomorphism. So let $\frac{a_1}{d_1} = \frac{a_2}{d_2}$, meaning that there exists $d \in D$ such that $(a_1 d_2 - a_2 d_1)d = 0$. Then also

$$(\rho(a_1)\rho(d_2) - \rho(a_2)\rho(d_1))\rho(d) = 0.$$

Since $\rho(d)$ is a unit, we also have $\rho(a_1)\rho(d_2) - \rho(a_2)\rho(d_1) = 0$, implying easily that $\rho(a_1)\rho(d_1)^{-1} = \rho(a_2)\rho(d_2)^{-1}$. □

Important special cases are:

- $D = \{1, a, a^2, \ldots\}$, then $D^{-1}A$ is obtained from A by adding an inverse to a, and we denote the result by $A[a^{-1}]$.
- $D = A \setminus \mathfrak{p}$, where $\mathfrak{p} \subseteq A$ is a prime ideal. Then D is indeed multiplicative and $D^{-1}A$ is denoted $A_\mathfrak{p}$ – it is the so called localization of A at a prime ideal $\mathfrak{p}$.
- In particular, when A is an integral domain, 0 is a prime ideal and then A_0 is the fraction field of A.

HW 2. Prove the following isomorphisms:

- $A[a^{-1}] \cong A[t]/(at - 1)$,
- $(A/I)[t] \cong A[t]/J$ and describe the ideal J,
- $A/(I+J) \cong (A/I)/J'$ and describe the ideal J' along the lines "it is essentially J, just…".

Proposition 3.5. *The localization $D^{-1}R$ is the trivial ring iff $0 \in D$.*

Proof. The triviality means $1/1 = 0/1$ and, by definition, this happens iff for some $d \in D$ we have $d = 0$, i.e. iff $0 \in D$. □

Theorem 3.6. *The localization $A_\mathfrak{p}$ at a prime ideal $\mathfrak{p}$ is a local ring.*

Proof. It is easily seen that the complement of the ideal $\mathfrak{m} = \{\frac{a}{d} \mid a \in \mathfrak{p},\, d \notin \mathfrak{p}\}$ is composed of units. □

Definition 3.7. Let A be a ring. An *A-module* is an abelian group M together with an operation

$$M \times A \to M, \quad (x, a) \mapsto xa$$

satisfying the axioms of a vector space, i.e.

$$\begin{aligned} x1 &= x, & (xa)b &= x(ab) \\ x(a+b) &= xa + xb, & (x+y)a &= xa + ya. \end{aligned}$$

An important example is an ideal – it is closed under addition and multiplication by elements of the ring.

Theorem 3.8 (Nakayama lemma)**.** *Let A be a local ring with a maximal ideal $\mathfrak{m}$. Let N be a finitely generated A-module such that $N\mathfrak{m} = N$. Then $N = 0$.*

Proof. Let $x_1, \ldots, x_n$ be a generating set of N. We may then write

$$x_j = x_1 a_{1j} + \cdots + x_n a_{nj}$$

for suitable $a_{ij} \in \mathfrak{m}$. Moving everything to the left, we obtain $(x_1, \ldots, x_n)(E - M) = 0$, where M is a matrix composed of the elements a_{ij}. Multiplying by the adjoint matrix, we get

$$(x_1, \ldots, x_n) \det(E - M) = 0,$$

i.e. $x_j \det(E - M) = 0$. This means that the multiplication by $\det(E - M)$ gives on N the zero map. However $\det(E - M) \in 1 + \mathfrak{m}$ and it is thus a unit (it does not lie in $\mathfrak{m}$). Therefore $N = 0$. □

For a multiplicative subset $D \subseteq A$ and the associated localization map $\lambda\colon A \to D^{-1}A$ we study the relationship between the ideals of A and those of $D^{-1}A$. We have maps between these sets that clearly preserve the ordering

$$\lambda_*\colon \{\text{ideals of } A\} \rightleftarrows \{\text{ideals of } D^{-1}A\} : \lambda^*$$

with

$$\lambda^*(J) = \lambda^{-1}(J) = \{a \in A \mid \tfrac{a}{1} \in J\}$$

that clearly preserves primeness (e.g. $A/\lambda^{-1}(J) \to B/J$ is clearly injective and a subring of a domain is itself a domain) and with

$$\lambda_*(I) = \underbrace{D^{-1}A \cdot \lambda(I)}_{D^{-1}I} = \{\tfrac{a}{d} \in D^{-1}A \mid a \in I\}$$

(i.e. the ideal generated by the image $\lambda(I)$.

Clearly $\lambda_*(\lambda^*(J)) = J$ and in the opposite direction

$$\lambda^*(\lambda_*(I)) = \{a \in A \mid \exists d \in D\colon da \in I\}$$

We call this the D-saturation of I and also say that I is D-saturated if it equals its saturation, i.e. if $da \in I \Rightarrow a \in I$ (division by $d \in D$). Obviously, by restriction, we get a bijection

$$\lambda_* \colon \{D\text{-saturated ideals of } A\} \cong \{\text{ideals of } D^{-1}A\} : \lambda^*$$

Further, a prime ideal P is D-saturated iff it is disjoint from D (if saturated then $d = d1 \in I \Rightarrow 1 \in I$, i.e. nonsense, so that $d \notin I$; if disjoint, one can divide by d showing D-saturatedness).

$$\lambda_* \colon \{\text{prime ideals of } A \text{ disjoint from } D\} \cong \{\text{prime ideals of } D^{-1}A\} : \lambda^*$$

Thus, if $D = R \setminus P$ the left hand side consists of prime ideals contained in P and as such contains a maximal element P, implying that $D^{-1}A = A_P$ has a unique maximal ideal, namely

$$D^{-1}P = \{\tfrac{a}{b} \mid a \in P,\, b \notin P\}$$

(alternatively, it consists exactly of the non-units of A_P). More generally, any ideal I that is maximal among those disjoint from D must be D-saturated, since its D-saturation is still a proper ideal disjoint from D, so that it is in fact a maximal D-saturated ideal and as such is a pullback of a maximal ideal of $D^{-1}A$, hence prime.

The point of the localization $D^{-1}A$ lies in having less ideals, in particular prime ideals, and this simplifies the structure theory of modules. We will see some examples of this.

The localization of a module is defined similarly by universal property

$$\begin{array}{ccc} M & \xrightarrow{\ \rho\ } & N \\ {\scriptstyle\lambda}\downarrow & \nearrow_{\tilde{\rho}} & \\ D^{-1}M & & \end{array}$$

where N is assumed to be an $D^{-1}A$ module, i.e. an A-module in which the multiplication map $d\cdot : N \to N$ is an isomorphism (look at the action map $D^{-1}A \to \mathrm{End}(N)$ and employ the universal property of the localization $D^{-1}A$). Straight from the definition we see that if the multiplication maps are isomorphisms on M then we can take $\lambda = \mathrm{id}$, i.e. $D^{-1}M = M$.

In general, since

$$\mathrm{Hom}_A(M,N) \cong \mathrm{Hom}_A(M, \mathrm{Hom}_{D^{-1}A}(D^{-1}A, N)) \cong \mathrm{Hom}_{D^{-1}A}(D^{-1}A \otimes_A M, N)$$

the so called extension of scalars gives a concrete construction $D^{-1}M = D^{-1}A \otimes_A M$. It is then important that $D^{-1}A$ is a flat A-module (see below) and thus the localization functor is exact. We will now give a second construction

$$D^{-1}M = \{\tfrac{x}{d} \mid x \in M,\, d \in D\}$$

where similarly to the case of A, it is imposed that $\frac{x}{d} = \frac{y}{e}$ iff $fex = fdy$ for some $f \in D$. To prove that this gives the previous localization, one has to prove that the maps

$$D^{-1}A \otimes_A M \rightleftarrows D^{-1}M,$$

given by $a/d \otimes x \mapsto (ax)/d$ and $1/d \otimes x \leftarrow\!\shortmid\, x/d$, are well defined (the first is the extension of the canonical inclusion $\lambda \colon M \to D^{-1}M$) and inverse to each other. This implies easily that $D^{-1}A$ is flat since for $f \colon M \to N$ injective the induced $D^{-1}f \colon D^{-1}M \to D^{-1}N$ satisfies

$D^{-1}f(x/d) = f(x)/d = 0$ iff $ef(x) = 0$, i.e. $f(ex) = 0$ and $ex = 0$ by injectivity of f; finally this gives $x/1 = 0$. Alternatively, one can express $D^{-1}A = \bigcup_{d\in D} d^{-1}A = \operatorname{colim}_{d\in D} A$ where the maps in the diagram are exactly of the form $e\cdot : A \to A$ from the copy of A with index d to the copy with index ed. It remains to show that the colimit indeed gives $D^{-1}A$ (easy) and that the diagram is filtered (very easy).

Again, for $D = R \setminus P$ we denote $M_P = D^{-1}M$.

Theorem 3.9. *For an A-module M we have: $M = 0 \Leftrightarrow \forall P$ maximal: $M_P = 0$.*

Before starting the proof we define the annihilator of $x \in M$ to be the ideal

$$\operatorname{Ann}(x) = \operatorname{Ann}_M(x) = \{a \in A \mid ax = 0\}.$$

Clearly $x = 0$ iff $\operatorname{Ann}(x) \ni 1$.

The fraction $\frac{a}{d} \in D^{-1}A$ then annihilates $\lambda(x) = \frac{x}{1}$, i.e. $\frac{ax}{d} = 0$ iff $\exists e \in D$: $eax = 0$ (i.e. $ea \in \operatorname{Ann}(x)$) iff $\frac{a}{d} \in D^{-1}\operatorname{Ann}(x)$, so that we finally get

$$\operatorname{Ann}(\tfrac{x}{1}) = D^{-1}\operatorname{Ann}(x).$$

(This implies, in particular, that $x \in \ker\lambda$ iff $D \cap \operatorname{Ann}(x) \neq \emptyset$ since these are exactly ideals giving the trivial ideal in the localization $D^{-1}A$.)

Proof. The implication $\Rightarrow$ is clear, so assume that $0 \neq x \in M$. Then $\operatorname{Ann}(x) \subsetneqq A$ is a proper ideal and there exists a maximal ideal $P \supseteq \operatorname{Ann}(x)$. Denoting $D = A \setminus P$ as usual, we obtain $D \cap \operatorname{Ann}(x) = \emptyset$ so that $D^{-1}\operatorname{Ann}(x) \not\ni 1$ is also proper. Since it equals $\operatorname{Ann}(\frac{x}{1})$, we must have $0 \neq \frac{x}{1} \in M_P$ and this module is thus also non-zero. □

Corollary 3.10. *For an A-linear map $f\colon M \to N$ we have: f is mono/epi/iso $\Leftrightarrow \forall P$ maximal: the localized map $f_P\colon M_P \to N_P$ is such.*

Proof. This follows from the chain of equivalences: f mono iff $\ker f = 0$ iff $(\ker f)_P = 0$ iff $\ker f_P = 0$ (since the localization, being exact, commutes with kernels) iff f_P mono. □

4. Primary decomposition

Let R be a (possibly graded) noetherian ring and let M be an R-module. Let us investigate when multiplication by $r \in R$ on the module M is non-injective – we may say that r is a zero divisor on M because this exactly means that there exists a nonzero $x \in M$ such that $rx = 0$. We denote

$$\operatorname{Ann}(x) = \operatorname{Ann}_M(x) = \{r \in R \mid rx = 0\},$$

the so called *annihilator* of the element x; it is easy to see that this is an ideal. The zero divisors on M are thus exactly the elements of the union of all annihilators $\operatorname{Ann}(x)$ for $x \neq 0$. Of course, it is enough to consider the maximal such and we will show that these are prime ideals.

We say that a prime ideal $\mathfrak{p}$ is an *associated prime* of the module M if $\mathfrak{p} = \operatorname{Ann}(x)$ for some $x \in M$. The set of all associated primes of M is denoted $\operatorname{Ass}(M)$.

We will now explain a useful characterization of annihilators: an R-submodule generated by x is isomorphic to $Rx \cong R/\operatorname{Ann}(x)$ by the first isomorphism theorem applied to the R-linear map $R \to M$ sending $1 \mapsto x$, whose image is obviously Rx and whose kernel is $\operatorname{Ann}(x)$. Thus, equivalently, a prime ideal $\mathfrak{p}$ is associated iff M contains a submodule isomorphic to the cyclic module $R/\mathfrak{p}$.

Example 4.1. $\mathrm{Ass}(R/\mathfrak{p}) = \{\mathfrak{p}\}$ because $R/\mathfrak{p}$ is an integral domain and thus the multiplication by any nonzero element is injective, i.e. $\mathrm{Ann}(x) = \mathfrak{p}$ for $x \neq 0$.

Lemma 4.2. *Every maximal element of* $\{\mathrm{Ann}(x) \mid x \neq 0\}$ *is an associated prime.*

In particular, for R *noetherian, every annihilator* $\mathrm{Ann}(x)$*, for* $x \neq 0$*, is contained in some associated prime.*

Remark. It is also true that, for a multiplicative subset D, a maximal element $\{\mathrm{Ann}(x) \mid \mathrm{Ann}(x) \cap D = \emptyset\}$ is an associated prime. This was proved in an earlier version and may be needed at some point...

Proof. Let $\mathrm{Ann}(x)$ be maximal and let $rs \in \mathrm{Ann}(x)$. Then either $sx = 0$ and thus $s \in \mathrm{Ann}(x)$ or $r \in \mathrm{Ann}(sx) = \mathrm{Ann}(x)$ by maximality. □

As a simple consequence, we obtain the following theorem:

Theorem 4.3. *Let* R *be a noetherian ring. Multiplication by* $r \in R$ *on an* R*-module* M *is injective iff* r *does not lie in any associated prime of* M*.* □

This theorem is useful especially because we will show that $\mathrm{Ass}(M)$ is finite for every noetherian (i.e. finitely generate) module M. The main tool here will be a so called primary decomposition. We say that a module M is $\mathfrak{p}$-*primary* if $\mathrm{Ass}(M) = \{\mathfrak{p}\}$. We also say that M is primary if it is $\mathfrak{p}$-primary for some prime ideal $\mathfrak{p}$.

In the case that M is not primary, it contains two submodules $P \cong R/\mathfrak{p}$ and $Q \cong R/\mathfrak{q}$ and in that case $\mathrm{Ass}(P \cap Q) \subseteq \mathrm{Ass}(P) \cap \mathrm{Ass}(Q) = \{\mathfrak{p}\} \cap \{\mathfrak{q}\} = \emptyset$. Since every nonzero module has some associated prime, we get $P \cap Q = 0$.

Theorem 4.4. *Let* M *be a finitely generated module over a noetherian ring* R*. Then there exists a finite collection of submodules* M_i*, for* $i = 1, \ldots, n$*, such that* $0 = \bigcap_i M_i$ *and such that each* M/M_i *is* $\mathfrak{p}_i$*-primary and the prime ideals* $\mathfrak{p}_i$ *are all distinct. If this expression is irredundant (i.e. no* M_i *can be removed) then* $\mathrm{Ass}(M) = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_n\}$*.*

Proof. Let us call an expression $M_0 = \bigcap M_i$ with all M/M_i primary a decomposition of the submodule M_0. We will show that if M_0 has no decomposition then there exists a strictly larger submodule without a decomposition and this would contradict M being noetherian. Since M_0 admits no decomposition, M/M_0 cannot be primary (for otherwise $M_0 = M_0$ would be a decomposition). As above, there exist two submodules $M_1/M_0, M_1'/M_0 \subseteq M/M_0$ with zero intersection, i.e. with $M_1 \cap M_1' = M_0$. If both M_1 and M_1' had decompositions we would obtain a decomposition for M_0 by intersecting these, so one of them does not admit a decomposition, as claimed.

We will now show $\mathrm{Ass}(M) \subseteq \{\mathfrak{p}_1, \ldots, \mathfrak{p}_n\}$. First we prove

$$\mathrm{Ann}(x) = \mathrm{Ann}_{M/M_1}(x) \cap \cdots \cap \mathrm{Ann}_{M/M_n}(x)$$

(for a nonzero $x \in M$, we have $r \in \mathrm{Ann}(x)$ iff $rx = 0$ iff $\forall i\colon rx \in M_i$ iff $\forall i\colon r \in \mathrm{Ann}_{M/M_i}(x)$). Assuming now that $\mathrm{Ann}(x) \subsetneq \mathrm{Ann}_{M/M_i}(x)$ for all i, we pick $s_i \in \mathrm{Ann}_{M/M_1}(x) \setminus \mathrm{Ann}(x)$. Their product $s_1 \cdots s_n$ then lies in the intersection, hence in $\mathrm{Ann}(x)$, but $s_i \notin \mathrm{Ann}(x)$, so $\mathrm{Ann}(x)$ is not prime. So for prime $\mathrm{Ann}(x)$ this must equal one of the $\mathrm{Ann}_{M/M_i}(x)$, and the latter can only be prime if it equals $\mathfrak{p}_i$.

For the opposite inclusion we need the irredundancy: It gives $\bigcap_{i \neq j} M_i \neq 0$ and this intersection thus contains some non-zero element, necessarily $x \notin M_j$, that has $\mathrm{Ann}_M(x) = \mathrm{Ann}_{M/M_j}(x)$ and, for some multiple $y \in Rx \subseteq \bigcap_{i \neq j} M_i$, we obtain $\mathrm{Ann}_{M/M_j}(y) = P_j$ since M/M_j is P_j-primary. □

Finally, we will study the behaviour of primary decomposition under localization, so let D be a multiplicative subset.

Lemma 4.5. *Let $x \in M$. A maximal element of $\{\operatorname{Ann}(dx) \mid d \in D\}$ equals the D-saturation of $\operatorname{Ann}(x)$.*

In particular, for R noetherian, the D-saturation of every annihilator $\operatorname{Ann}(x)$ is an annihilator $\operatorname{Ann}(d_0 x)$.

Proof. Let $\operatorname{Ann}(d_0 x)$ be maximal and let $dr \in \operatorname{Ann}(d_0 x)$. Then $r \in \operatorname{Ann}(dd_0 x) = \operatorname{Ann}(d_0 x)$ and it is D-saturated. Clearly, it has the same D-saturation as $\operatorname{Ann}(x)$. □

For the localization map $\lambda\colon M \to D^{-1}M$ we recall that $\operatorname{Ann}(x/1) = D^{-1}\operatorname{Ann}(x)$ and since the localization gives a bijection

$$\{\text{prime ideals of } A \text{ disjoint from } D\} \cong \{\text{prime ideals of } D^{-1}A\}$$

(and those intersecting D give the full ring on the right hand side) we can determine the associated primes of $D^{-1}M$:

$$\operatorname{Ass}(D^{-1}M) = \{D^{-1}P \mid P \in \operatorname{Ass}(M),\ D \cap P = \emptyset\}.$$

This takes a particularly simple form for a P-primary module M over a noetherian ring (is this necessary?): then either $D^{-1}M$ is $D^{-1}P$-primary when $D \cap P = \emptyset$ or $D^{-1}M = 0$ when $D \cap P \neq \emptyset$ (since then $D^{-1}M$ has no associated prime). Now apply this to a primary decomposition $0 = \bigcap M_i$ with M/M_i being P_i-primary. We get

$$0 = \bigcap D^{-1}M_i$$

with $D^{-1}M/D^{-1}M_i$ being $D^{-1}P_i$-primary; when some $D^{-1}M/D^{-1}M_i$ is zero, i.e. $D^{-1}M_i = D^{-1}M$, we may remove it from the decomposition. For a minimal associated prime P_j we then get only one non-zero submodule, namely

$$0 = D^{-1}M_j$$

that together with the monomorphism (since the module M/M_j is P_j-primary, we have $\operatorname{Ann}(x/1) = D^{-1}\operatorname{Ann}(x) \subseteq D^{-1}P_j$ and is thus proper, showing that $x/1 \neq 0$)

$$\begin{array}{ccc} M & \xrightarrow{\lambda_j} & D^{-1}M \\ \downarrow & & \downarrow \cong \\ M/M_j & \rightarrowtail & D^{-1}M/D^{-1}M_j \end{array}$$

gives that $M_j = \ker \lambda_j$ and as such is unique.

For completeness, still over a noetherian ring, we prove that for any prime $P \supseteq \operatorname{Ann}(x)$ there is an associated prime lying between these two: consider $\lambda\colon M \to M_P$ and observe that $\operatorname{Ann}(x/1) = \operatorname{Ann}(x)_P$ is non-trivial. It is thus contained in some associated prime $D^{-1}Q \in \operatorname{Ass}(M_P)$. As above, this means that $Q \in \operatorname{Ass}(M)$ (this is a bit circular, it seems that the general version of Lemma 4.2 is needed to conclude that there exists an annihilator maximal among those disjoint from $D = R \setminus P$ and as such is the prime Q as above). This

implies that any proper ideal I lies in prime that is minimal above it: since $\mathrm{Ass}(R/I)$ is finite, it contains a minimal element; by the above it must in fact be minimal among all primes containing $I = \mathrm{Ann}(1)$.

As a final application of this, if I is P-primary then, in particular, P is the unique minimal prime above I (so smallest) and thus Proposition 4.6 gives

$$\sqrt{I} = \bigcap_{I \subseteq Q \text{ prime}} Q = P.$$

In particular, we obtain the following consequence: $rs \in I \Rightarrow r \in \sqrt{I}$ or $s \in I$ (the first condition means that r belongs to the unique associated prime of R/I, the second that $s = 0$ in R/I). In the opposite direction, when I satisfies this condition we get that $\sqrt{I}$ is the union of all the associated primes by Theorem 4.3, and these include all minimal primes above I. Since it is also the intersection of all minimal primes above I, there must be only one associated prime and I is necessarily primary (with associated prime $\sqrt{I}$).

So for a primary ideal I, the radical $\sqrt{I}$ is a prime ideal. The converse is generally not true, since $\sqrt{I}$ being prime only means that there is a unique minimal prime over I, but some bigger prime may be associated as well. However, if $\sqrt{I}$ is a maximal ideal, then I is automatically primary. In particular, for a maximal ideal M, each power M^n is primary. It is interesting that P^n needs not be P-primary for a general prime P and its P-primary component $P^{(n)} = \lambda^*\lambda_* P^n$ is generally bigger and is called the n-th symbolic power of P.

Proposition 4.6. *The intersection of all prime ideals is the nilradical* $\sqrt{0} = \{r \in R \text{ nilpotent}\}$. *More generally* $\bigcap_{P \supseteq I} P = \sqrt{I}$.

Proof. Clearly every nilpotent element lies in every prime. Thus let $a \in R$, denote $D = \{1, a, a^2, \ldots\}$ the corresponding multiplicative subset and assume that a lies in every prime or, equivalently, every prime intersects D. Then the localization $D^{-1}R$ contains no prime and thus $1 = 0$ in $D^{-1}R$. By Proposition 3.5, this is equivalent to D containing zero, i.e. that some power of a is zero.

The second point is obtained from the first by applying to the quotient ring R/I. □

5. Chain complexes

Definition 5.1. A sequence of R-modules $A \xrightarrow{f} B \xrightarrow{g} C$ is said to be exact at B if $\operatorname{im} f = \ker g$. Similarly, one can define exactness of a longer sequence at any inner term. A sequence is exact, if it is exact at every inner term.

Exercise 5.2. Characterize exactness of $0 \to A \to B$, $A \to B \to 0$, $0 \to A \to B \to 0$, $0 \to A \to B \to C$, $A \to B \to C \to 0$ and $0 \to A \to B \to C \to 0$ (the last is referred to as a short exact sequence). In particular prove that any short exact sequence is isomorphic to an "extension" $0 \to A \hookrightarrow B \twoheadrightarrow B/A \to 0$.

In the condition $\operatorname{im} f = \ker g$, the inclusion $\subseteq$ is equivalent to $g \circ f = 0$, under which one may form the quotient $\ker g / \operatorname{im} f$ that measures the difference between the two submodules. One may thus express the exactness equivalently as $g \circ f = 0$ and $\ker g / \operatorname{im} f = 0$. These two parts are the main idea of the definition of a chain complex.

Definition 5.3. A *chain complex* C is a diagram

$$\cdots \to C_{n+1} \xrightarrow{d_{n+1}} C_n \xrightarrow{d_n} C_{n-1} \to \cdots$$

in which $d_n \circ d_{n+1} = 0$ for all n. We often abbreviate all the maps to d and call them the *boundary maps* or *differentials*. The elements of C_n are called the *n-chains*. Denoting

$$B_n = B_n(C) = \operatorname{im} d_{n+1} \subseteq C_n$$

the submodule of *n-boundaries* and

$$Z_n = Z_n(C) = \ker d_n \subseteq C_n$$

the submodule of *n-cycles*, the "square zero" condition $d \circ d = 0$ is equivalent to $B_n \subseteq Z_n$. The corresponding quotient module

$$H_n = H_n(C) = Z_n/B_n$$

is called the n-th *homology group*, or just the n-th *homology* of C.

As observed above, the corresponding sequence is exact at C_n iff $H_n(C) = 0$. If this happens for all n we say that the chain complex is *acyclic* or that the corresponding sequence is a *long exact sequence*.

Example 5.4. Simplicial homology: Let K be a simplicial complex and choose a total ordering of its vertices (this can be avoided, see below). We write the n-simplices as ordered $(n+1)$-tuples of its vertices, i.e. $\sigma = [v_0, \ldots, v_n]$ with $v_0 < \cdots < v_n$. We define the operator

$$d_i \colon K_n \to K_{n-1},\ d_i[v_0, \ldots, v_n] = [v_0, \ldots, \widehat{v_i}, \ldots, v_n].$$

The idea now is that the boundary of σ should be the collection $d_0\sigma, \ldots, d_n\sigma$. This is algebraically achieved by considering $C_n(K) = RK_n$ the free module on the set of n-simplices and by defining $d = \sum(-1)^i d_i$. The n-th homology of the chain complex $C_n(K)$ is the n-th simplicial homology of the simplicial complex K.

The problem with ordering is solved when one defines $C_n(K) = RK_n^{\mathrm{ord}}/\!\sim$, the free module on ordered n-simplices written again as $[v_0, \ldots, v_n]$, but this time without any restriction on the ordering of vertices of this n-simplex, modulo the relation

$$[v_{\sigma(0)}, \ldots, v_{\sigma(n)}] = \operatorname{sign}\sigma \cdot [v_0, \ldots, v_n].$$

Singular homology: Let X be a space and define similarly $C_n(X)$ to be the free module on the set of all continuous maps $\Delta^n \to X$, where Δ^n denotes the standard n-simplex, i.e. the convex hull of any $(n+1)$-tuple of affine independent points, e.g. $E_0, \ldots, E_n \in \mathcal{B}_n$ the standard point basis of the affine space $x_0 + \cdots + x_n = 1$. The operators d_i are now given by restricting to the faces as above. The differential on $C(X)$ is yet again $d = \sum(-1)^i d_i$.

de Rham cohomology: $\Omega^n M = \{\text{smooth } n\text{-forms on } M\}$. Here the differential points in the opposite direction $\Omega^n M \to \Omega^{n+1} M$. We will formalize this later as a cochain complex and de Rham cohomology is the cohomology of this cochain complex.

Definition 5.5. A *chain map* $f\colon C \to D$ between two chain complexes is a collection of homomorphisms f_n for which the (ladder shaped) diagram

$$\begin{array}{ccccccc} \cdots & \longrightarrow & C_n & \xrightarrow{d_n} & C_{n-1} & \longrightarrow & \cdots \\ & & \downarrow{\scriptstyle f_n} & & \downarrow{\scriptstyle f_{n-1}} & & \\ \cdots & \longrightarrow & D_n & \xrightarrow[d_n]{} & D_{n-1} & \longrightarrow & \cdots \end{array}$$

commutes, i.e. $df = fd$.

Every chain map induces maps $B_n(C) \to B_n(D)$ and $Z_n(C) \to Z_n(D)$ and thus also $H_n(C) \to H_n(D)$. We obtain

Proposition 5.6. *The n-th homology forms a functor $H_n\colon \mathsf{Ch}(\mathsf{Mod}_R) \to \mathsf{Mod}_R$.* □

Definition 5.7. We say that a chain map f is a quasi-isomorphism if the induced map on homology is an isomorphism.

As an example, a chain complex C is acyclic iff the unique map $0 \to C$ is a quasi-isomorphism iff the unique map $C \to 0$ is a quasi-isomorphism.

We will now present a special class of quasi-isomorphisms, analogous to homotopy equivalences in topology. First we need the corresponding algebraic notion of homotopy.

Definition 5.8. Let f and g be two chain maps $C \to D$. A chain homotopy from f to g is a collection of homomorphisms h_n as in the (non-commutative!) diagram

$$\begin{array}{ccccccccc} \cdots & \longrightarrow & C_{n+1} & \xrightarrow{d_{n+1}} & C_n & \xrightarrow{d_n} & C_{n-1} & \longrightarrow & \cdots \\ & & \downarrow\downarrow & \swarrow{\scriptstyle h_n} & \downarrow\downarrow & \swarrow{\scriptstyle h_{n-1}} & \downarrow\downarrow & & \\ \cdots & \longrightarrow & D_{n+1} & \xrightarrow[d_{n+1}]{} & D_n & \xrightarrow[d_n]{} & D_{n-1} & \longrightarrow & \cdots \end{array}$$

such that $dh + hd = g - f$. We write $h\colon f \sim g$ or $f \sim_h g$ or simply $f \sim g$ if the homotopy is not important.

A chain homotopy equivalence is a chain map $f\colon C \to D$ that admits a homotopy inverse, i.e. a chain map $g\colon D \to C$ together with homotopies $gf \sim 1$, $fg \sim 1$.

Remark. Any continuous map between spaces induces a chain map between their singular chain complexes and any chain homotopy induces a chain homotopy (this is not completely straightforward). The simplicial situation is a bit more straightforward, but complicated enough to be explained at this point. We will give a nice interpretation (two in fact) of chain homotopy later.

Proposition 5.9. *Chain homotopic maps induce equal maps on homology. In particular, chain homotopy equivalences are quasi-isomorphisms.*

Proof. Let $[z] \in H_n(C)$ be represented by a cycle z. Then

$$g(x) - f(x) = dh(x) + h\underbrace{d(x)}_{0}$$

so that $[g(x)] = [f(x) + dh(x)] = [f(x)]$. □

Proposition 5.10. *Chain homotopy equivalence is an equivalence relation that respects composition.*

We may thus form the homotopy category of chain complexes and chain homotopy classes of maps where chain equivalences are exactly the isomorphisms.

Proof. We prove transitivity: if $f_1 - f_0 = dh + hd$ and $f_2 - f_1 = dk + kd$ then $f_2 - f_0 = d(h+k)+(h+k)d$. Similarly if $f_1 - f_0 = dh+hd$ then $gf_1 - gf_0 = g(dh+hd) = d(gh)+(gh)d$. □

We index the modules in our chain complexes by integers, but we will be using a lot chain complexes indexed by non-negative integers only. One can extend such a chain complex by zeros and thus think of it as a chain complex in the original sense. In doing so, the non-negatively graded chain complex

$$\cdots \to C_1 \to C_0$$

will also have the zero homology $H_0(C) = C_0/B_0(C) = \operatorname{coker}(d_1)$ since every 0-chain is a cycle.

Another variation, briefly mentioned above with connection to de Rham cohomology is that of a cochain complex. Another situation where cochain complexes arise naturally is upon applying contravariant functors to chain complexes – the direction of homomorphisms changes. We will distingiush notationally by using upper indices.

Definition 5.11. A *cochain complex* C is a diagram

$$\cdots \to C^{n-1} \xrightarrow{d^{n-1}} C^n \xrightarrow{d^n} C^{n+1} \to \cdots$$

in which $d^n \circ d^{n-1} = 0$ for all n. We get notions of cochains, cocycles, coboundaries and cohomology, cochain maps and cochain homotopy in an obvious way.

Again, non-negatively graded cochain complexes will play an important role and they will look

$$C^0 \to C^1 \to \cdots$$

so that the zeroth cohomology will be $H^0(C) = Z^0(C)/0 = \ker d^0$.

Proposition 5.12. *In a pullback square*

$$\begin{array}{ccc} B & \xrightarrow{g} & C \\ \downarrow & \lrcorner & \downarrow \\ B' & \xrightarrow[g']{} & C' \end{array}$$

the induced map $\ker g \to \ker g'$ *is an iso. In addition, if* g' *is epi then so is* g.

Proof. The first point follows from simple properties of pullbacks (tutorial). The second point is verified on elements and was done in the tutorial (more formally, it follows from Mod_R being abelian, but that I have to understand first). □

Theorem 5.13 (snake lemma)**.** *Given a commutative diagram with exact rows*

$$\begin{array}{ccccccc} A & \xrightarrow{i} & B & \xrightarrow{p} & C & \longrightarrow & 0 \\ \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \\ 0 \longrightarrow A' & \xrightarrow{i'} & B' & \xrightarrow{p'} & C' & & \end{array}$$

there exists a natural exact sequence

$$\ker\alpha \to \ker\beta \to \ker\gamma \xrightarrow{\delta} \operatorname{coker}\alpha \to \operatorname{coker}\beta \to \operatorname{coker}\gamma.$$

If, in addition the map i is injective (i.e. the top exact sequence can be prolonged to the left by zero to a short exact sequence), so is the map $\ker\alpha \to \ker\beta$ and similarly for the surjectivity of the map $B' \to C'$.

Proof. One first proves the version with short exact sequences in both rows. Since limits commute with limits, starting with the square

$$\begin{array}{ccc} B & \longrightarrow & C \\ \downarrow & & \downarrow \\ B' & \longrightarrow & C' \end{array}$$

and applying kernels first in the horizontal and then in the vertical direction yields the same result as applying them in the opposite order, i.e. $\ker\alpha$ is indeed the kernel of the map $\ker\beta \to \ker\gamma$ and this proves exactness at $\ker\alpha$ and $\ker\beta$. By the dual argument, we are left to construct the "connecting homomorphism" δ and to prove exactness at its domain and codomain. We define

$$\delta(c) = (i')^{-1}\beta p^{-1}(c)$$

where we need to verify that the preimage of $\beta p^{-1}(c)$ indeed exists. By exactness, this amounts to showing that $0 = p'\beta p^{-1}(c) = \gamma p p^{-1}(c) = \gamma(c)$ and this holds since we assume $c \in \ker\gamma$. Now the preimage $p^{-1}(c)$ is not unique and we have to show that the result does not depend on the choice. However, the choice is unique up to $\operatorname{im} i$ that is mapped by β to $\operatorname{im}(i'\alpha)$ and further by $(i')^{-1}$ to $\operatorname{im}\alpha$ that is zero in $\operatorname{coker}\alpha$. Clearly, if $c = p(b)$ for some $b \in \ker\beta$ then the above prescription yields zero, so the composition

$$\ker\beta \to \ker\gamma \xrightarrow{\delta} \operatorname{coker}\alpha$$

is indeed zero. Let now $c \in \ker\gamma$ be such that $\delta(c) = 0$, i.e. $(i')^{-1}\beta p^{-1}(c) = \alpha(a)$. Now $p^{-1}(c) + i(a)$ is still a preimage of c, and lies in $\ker\beta$, by an easy inspection.

Finally, if i is not mono, replace A by $\operatorname{im} i$ (since this equals $\ker p$, it admits a map α' to $A' = \ker p'$) and apply the mono case.

$$\begin{array}{ccccccc} A & \twoheadrightarrow & \operatorname{im} i & \rightarrowtail & B & \xrightarrow{p} & C \longrightarrow 0 \\ & \searrow^{\alpha} & \downarrow{\scriptstyle\alpha'} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} \\ 0 & \longrightarrow & A' & \xrightarrow{i'} & B' & \xrightarrow{p'} & C' \end{array}$$

The mono case also easily gets that the map $\ker\alpha \to \ker\alpha'$ is epi and thus upon replacing $\ker\alpha'$ by $\ker\alpha$, the sequence remains exact everywhere except at $\ker\alpha$, as claimed. □

Remark. One can construct from a chain $A \xrightarrow{\alpha} B \xrightarrow{\beta} C$ a diagram (coming from certain map of double complexes)

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \longrightarrow & A \oplus B & \longrightarrow & B & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \alpha} & & \downarrow{\scriptstyle \beta\alpha \oplus 1} & & \downarrow{\scriptstyle \beta} & & \\ 0 & \longrightarrow & B & \longrightarrow & C \oplus B & \longrightarrow & C & \longrightarrow & 0 \end{array}$$

(the rows are not comprised of the inclusions and projections, they have to be twisted slightly), which gives the exact sequence relating kernels and cokernels of the maps α, $\beta\alpha$ and β.

Proposition 5.14 (5-lemma)**.** *In a commutative diagram with exact rows*

$$\begin{array}{ccccccccc} A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & D & \longrightarrow & E \\ \downarrow{\scriptstyle \alpha} & & \downarrow{\scriptstyle \beta} & & \downarrow{\scriptstyle \gamma} & & \downarrow{\scriptstyle \delta} & & \downarrow{\scriptstyle \varepsilon} \\ A' & \longrightarrow & B' & \longrightarrow & C' & \longrightarrow & D' & \longrightarrow & E' \end{array}$$

if α, β, δ, ε are iso, then so is γ.

More precisely, α is only required to be epi and ε to be mono.

Proof. Apply Lemma 5.20; denoting the image of the map $B \to C$ for simplicity by BC etc., we obtain short exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & BC & \longrightarrow & C & \longrightarrow & CD & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \beta\gamma} & & \downarrow{\scriptstyle \gamma} & & \downarrow{\scriptstyle \gamma\delta} & & \\ 0 & \longrightarrow & B'C' & \longrightarrow & C' & \longrightarrow & C'D' & \longrightarrow & 0 \end{array}$$

and the snake lemma gives that γ is mono provided that $\beta\gamma$ and $\gamma\delta$ are mono. The second is easier, just apply the snake lemma in

$$\begin{array}{ccccccccc} 0 & \longrightarrow & CD & \longrightarrow & D & \longrightarrow & DE & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \gamma\delta} & & \downarrow{\scriptstyle \delta} & & \downarrow{\scriptstyle \delta\varepsilon} & & \\ 0 & \longrightarrow & C'D' & \longrightarrow & D' & \longrightarrow & D'E' & \longrightarrow & 0 \end{array}$$

to get $\gamma\delta$ mono if δ is. The other condition is more complicated and involves the application of the snake lemma to

$$\begin{array}{ccccccccc} 0 & \longrightarrow & AB & \longrightarrow & B & \longrightarrow & BC & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \alpha\beta} & & \downarrow{\scriptstyle \beta} & & \downarrow{\scriptstyle \beta\gamma} & & \\ 0 & \longrightarrow & A'B' & \longrightarrow & B' & \longrightarrow & B'C' & \longrightarrow & 0 \end{array}$$

to obtain $\beta\gamma$ mono provided that β is mono and $\alpha\beta$ is epi. Finally, $\alpha\beta$ epi follows from α epi by the last application of snake lemma where one needs to prolong the sequence one step to the left, say by kernels of the maps $A \to B$ and $A' \to B'$. Altogether, γ is mono if β and δ are mono and α is epi. Dually, γ epi follows from β and δ epi and ε mono. □

The following is a converse to Proposition 5.12. I left it as an exercise, I think.

Proposition 5.15. *In a commutative square*

$$\begin{array}{ccc} B & \xrightarrow{g} & C \\ \downarrow & & \downarrow \\ B' & \xrightarrow[g']{} & C' \end{array}$$

if both g and g' are epi and the induced map $\ker g \to \ker g'$ is an iso then it is a pullback square.

Exercise 5.16. This is about self-duality of homology. Show that there exists a factorization

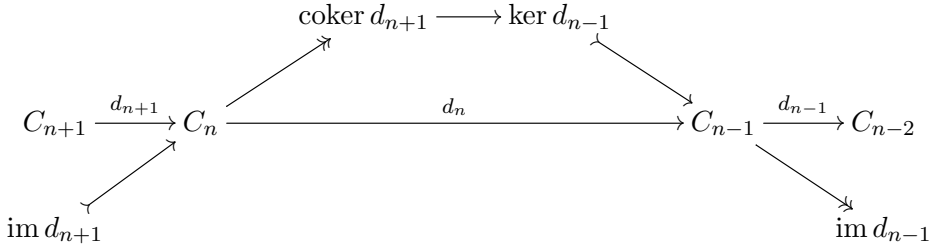

and that the map on the top has kernel $H_n(C)$ and cokernel $H_{n-1}(C)$. The diagram is self-dual, so starting with a cochain complex, interpreting it as a chain complex in the opposite category and taking homology there yields exactly the cohomology of the original cochain complex.

Theorem 5.17 (long exact sequence of homology)**.** *A short exact sequence $0 \to A \to B \to C \to 0$ of chain complexes induces a natural long exact sequence of homology*

$$\cdots \to H_{n+1}(C) \xrightarrow{\delta} H_n(A) \to H_n(B) \to H_n(C) \xrightarrow{\delta} H_{n-1}(A) \to \cdots .$$

Proof. Applying the previous exercise, we will consider the map $\operatorname{coker} d_{n+1} \to \ker d_{n-1}$ for the involved chain complexes and write them as $C_n(C)/B_n(C) \to Z_{n-1}(C)$ etc. so that we obtain a diagram

$$\begin{array}{ccccccc} & & C_n(A)/B_n(A) & \longrightarrow & C_n(B)/B_n(B) & \longrightarrow & C_n(C)/B_n(C) \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & Z_{n-1}(A) & \longrightarrow & Z_{n-1}(B) & \longrightarrow & Z_{n-1}(C) \end{array}$$

with exact rows (coker commutes with coker, similarly ker commutes with ker). Snake lemma gives a portion of the claimed long exact sequence, as required. □

Corollary 5.18. *In a short exact sequence of chain complexes as above, A is acyclic iff $B \to C$ is a q-iso. Dually, C is acyclic iff $A \to B$ is a q-iso.*

Corollary 5.19. *In a commutative diagram of chain complexes with exact rows*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \xrightarrow{i} & B & \xrightarrow{p} & C & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \alpha} & & \downarrow{\scriptstyle \beta} & & \downarrow{\scriptstyle \gamma} & & \\ 0 & \longrightarrow & A' & \xrightarrow{i'} & B' & \xrightarrow{p'} & C' & \longrightarrow & 0 \end{array}$$

if two of α, β, γ are q-iso's, so is the third.

Lemma 5.20. *A long exact sequence*

$$\cdots \to C_{n+1} \to C_n \to C_{n-1} \to \cdots$$

can be split into short exact sequences

$$0 \to B_n \to C_n \to B_{n-1} \to 0.$$

Conversely, any collection of short exact sequences as above can be spliced into a long exact sequence.

Proof.

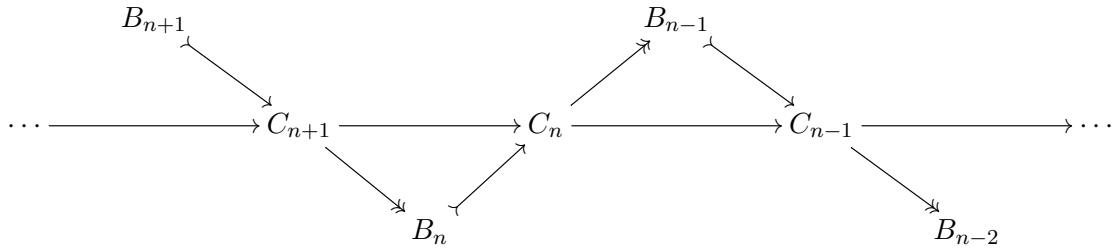

□

In a general chain complex, one has to replace the short exact sequences by

$$0 \to Z_n \to C_n \to B_{n-1} \to 0$$

and add to these the short exact sequences

$$0 \to B_n \to Z_n \to H_n \to 0$$

that define the homology as the quotient Z_n/B_n. Again, one can splice such short exact sequences into a chain complex C with homology H.

Definition 5.21. A *resolution* of a module A is a non-negatively graded chain complex C together with an "augmentation" map $\varepsilon\colon C_0 \to A$ such that

$$\cdots \to C_1 \to C_0 \xrightarrow{\varepsilon} A$$

is an acyclic chain complex (the "augmented" chain complex).

We say that C is a *projective resolution* if, in addition, C consists of projective modules.

There is a nice "global" characterization of this, using the chain map $\varepsilon\colon C \to A[0]$ where $A[0]$ denotes the chain complex whose only nonzero chains are in dimension zero and are A. Thus, the map is precisely

$$\begin{array}{ccccc} \cdots & \longrightarrow & C_1 & \longrightarrow & C_0 \\ & & \downarrow & & \downarrow{\scriptstyle \varepsilon} \\ \cdots & \longrightarrow & 0 & \longrightarrow & A \end{array}$$

Now the homology of the augmented chain complex agrees with the homology of C except in dimensions 0 and -1, where it is $\ker \varepsilon / B_0(C)$ and $\operatorname{coker} \varepsilon$. The first can be rewritten as

$$\ker(C_0/B_0(C) \xrightarrow{\varepsilon} A) = \ker(H_0(C) \xrightarrow{\varepsilon} A)$$

while the second can be rewritten as the cokernel of the same map $H_0(C) \xrightarrow{\varepsilon} A$. Since this is the induced map on homology, we are finished with the equivalence. We will give a different, more conceptual proof later.

Definition 5.22. Dually, a (co)resolution is a cochain complex C together with a coaugmentation map $A \to C^0$ such that the coaugmented cochain complex

$$A \to C^0 \to C^1 \to \cdots$$

is acyclic. An injective (co)resolution has all objects C^n injective.

Definition 5.23. A functor F is *additive* if its action on morphisms $\mathcal{C}(c', c) \to \mathcal{D}(Fc', Fc)$ is a homomorphism of groups.

Any additive functor F has an extension to a functor between chain complexes since it preserves composition and zero. We denote this extension again by F.

Example 5.24. Hom functors $\mathrm{Hom}_R(A, -)$, $\mathrm{Hom}_R(-, A)$ (the second contravariant though), tensor product functors $- \otimes_R A$, $A \otimes_R -$.

We will now discuss basic properties of functors related to exactness or, equivalently, homology. It is easy to see that $F0 = 0$ from the characterization of 0 as an object where $1 = 0$ (one may also view this as a nullary case of Lemma 5.29). We now make a couple of definitions:

Definition 5.25. An additive functor F is said to be *right exact* if the image of an exact sequence

$$A \to B \to C \to 0$$

is an exact sequence

$$FA \to FB \to FC \to 0$$

Equivalently, F pereserves cokernels. (By Lemma 5.29, it preserves finite coproducts and this is thus equivalent to preservation of finite colimits.)

A *left exact* functor is an additive functor that preserves kernels.

A functor is *exact* if it preserves short exact sequences.

Exercise 5.26. Show that a functor is exact iff it preserves all exact sequences iff it is left exact and right exact.

In particular, an exact functor preserves acyclic chain complexes. A generalization of this is the following.

Lemma 5.27. *An exact functor F commutes with homology, i.e. $H_n(FC) = FH_n(C)$. In particular, F preserves q-iso's.*

Proof. This is so since homology is defined using kernels and cokernels. □

Example 5.28. The tensor product functor $- \otimes_R A$ is right exact; it is exact iff A is flat. Similarly for the other tensor product functor. The hom functor $\mathrm{Hom}_R(A, -)$ is left exact; it is exact iff A is projective. Similarly for the other hom functor (here A should be injective for exactness), note however that this depends on writing the contravariant one as $\mathsf{Mod}_R^{\mathrm{op}} \to \mathsf{Ab}$ (and *not* as $\mathsf{Mod}_R \to \mathsf{Ab}^{\mathrm{op}}$).

Lemma 5.29. *Additive functors preserve biproducts. Equivalently, additive functors preserve exactness of split short exact sequences.*

Proof. A biproduct is a diagram

$$A \underset{p}{\overset{i}{\rightleftarrows}} C \underset{j}{\overset{q}{\rightleftarrows}} B$$

satisfying

$$\begin{pmatrix} i & j \end{pmatrix} \cdot \begin{pmatrix} p \\ q \end{pmatrix} = 1, \qquad \begin{pmatrix} p \\ q \end{pmatrix} \cdot \begin{pmatrix} i & j \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Since F preserves composition, addition, identities and zeros, the same is true for the image under F. □

Thus, any additive functor is exact on certain exact sequences. Over a field, any additive functor is exact. As we saw, this should mean that it preserves certain q-iso's. Here is a precise claim.

Lemma 5.30. *Additive functors preserve chain homotopies.*

Proof. The proof is practically the same as for the previous lemma: A chain homotopy is given by some formulas and these are presereved by additive functors. □

Example 5.31. Consider a short exact sequence

$$0 \to \mathbb{Z} \xrightarrow{2} \mathbb{Z} \to \mathbb{Z}/2 \to 0$$

and apply $- \otimes \mathbb{Z}/2$; this yields

$$0 \to \mathbb{Z}/2 \xrightarrow{0} \mathbb{Z}/2 \xrightarrow{1} \mathbb{Z}/2 \to 0$$

that is clearly not exact. Thus, $- \otimes \mathbb{Z}/2$ is not exact and does not preserve q-iso's (since $C \to 0$ is a q-iso while $FC \to 0$ is not).

In the next sections, our main aim will be to measure the non-exactness of an additive functor. We will see that in the second short exact sequence the zero on the left can be replaced by a continuation – a long exact sequence of derived functors.

6. Abelian categories

This is my personal note. The most complicated are those properties that relate limits and colimits. The definition of an abelian category is that of a finitely bicomplete Ab-enriched category where image equals coimage, i.e. where epi and mono together imply iso. The main application is that

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

is exact iff f and g form the kernel-cokernel pair (f kernel of g and g cokernel of f; by definition) iff f is mono and g is its cokernel iff g is epi and f is its kernel. We will now apply this to pullbacks and pushouts. A square

$$\begin{array}{ccc} A' & \xrightarrow{g} & A \\ {\scriptstyle f'}\downarrow & & \downarrow{\scriptstyle f} \\ B' & \xrightarrow[h]{} & B \end{array}$$

is a pullback square iff (we chose the minus sign with accordance to sign conventions for double complexes)

$$0 \longrightarrow A' \xrightarrow{\begin{pmatrix} -f' \\ g \end{pmatrix}} B' \oplus A \xrightarrow{\begin{pmatrix} h & f \end{pmatrix}} B$$

is exact (form the construction of the pullback using products and equalizers). Thus, provided that $\begin{pmatrix} h & f \end{pmatrix}$ is epi (i.e. h and f are jointly epi) this becomes a short exact sequence and by the dual argument, the square will also be a pushout! Now assume more concretely that f is epi. Since cokernels of parallel maps in any pushout agree, $\operatorname{coker} f' \cong \operatorname{coker} f = 0$ and f' is also epi!

Snake lemma: The connecting homomorphism is constructed using the pullback and pushout in the diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & A & \longrightarrow & X & \twoheadrightarrow & \ker\gamma & & \\
 & & \| & & \downarrow & \lrcorner & \downarrow & & \\
0 & \longrightarrow & A & \xrightarrow{i} & B & \xrightarrow{p} & C & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \\
0 & \longrightarrow & A' & \xrightarrow{i'} & B' & \xrightarrow{p'} & C' & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \| & & \\
 & & \operatorname{coker}\alpha & \rightarrowtail & Y & \longrightarrow & C' & \longrightarrow & 0
\end{array}$$

that then gives a map $X \to Y$. By the above, $\ker\gamma$ is the cokernel of the map $A \to X$, so in order to get a factorization of $X \to Y$ through $\ker\gamma$, we need to show that the composition $A \to X \to Y$ is zero, but this is obvious. Similarly the composition $X \to Y \to C'$ is zero and so $X \to Y$ also factors through $\operatorname{coker}\alpha$; it is a simple matter to show that it then factors through both at the same time, i.e. it induces a unique map $\delta\colon \ker\gamma \to \operatorname{coker}\alpha$. The exactness at $\ker\gamma$ needs to be checked and is probably not too difficult, but it is completely straightforward using elements.

self-duality of homology:

$$A \xrightarrow{f} B \xrightarrow{g} C$$

with $gf = 0$ induces the following diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \operatorname{im} f & \longrightarrow & B & \longrightarrow & \operatorname{coker} f & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle i} & & \| & & \downarrow{\scriptstyle p} & & \\
0 & \longrightarrow & \ker g & \longrightarrow & B & \longrightarrow & \operatorname{coim} g & \longrightarrow & 0
\end{array}$$

and the connecting homomorphism in snake lemma gives an isomorphism $\ker p \cong \operatorname{coker} i$, where $\operatorname{coker} i$ is the usual definition of homology, while $\ker p$ is the dual version.

Theorem 6.1. *Any small abelian category $\mathcal{A}$ admits an exact fully faithful embedding into ${}_R\mathsf{Mod}$ for some ring R.*

Sketch proof. The Yoneda embedding

$$y\colon \mathcal{A} \to [\mathcal{A}^{\mathrm{op}}, \mathsf{Ab}]$$

lands in the subcategory $\mathcal{L}$ of left exact functors. One shows that this is an exact localization and that y as a functor $\mathcal{A} \to \mathcal{L}$ is exact (and still fully faithful). In addition, $\mathcal{L}$ admits all small coproducts; denoting $P = \sum_{A \in \mathcal{A}} y(A)$ one shows that this is a projective object that admits an epi (joint epi would suffice) into the image of every $A \in \mathcal{A}$. Thus, the representable functor

$$\mathrm{Hom}_{\mathcal{L}}(P, -) \colon \mathcal{L} \to {}_{\mathrm{End}(P)}\mathsf{Mod}$$

is also exact (since P is projective) fully faithful (we need to show that $\mathrm{Hom}_{\mathcal{A}}(A, B) \to \mathrm{Hom}_{\mathrm{End}(P)}(\mathrm{Hom}(P, A), \mathrm{Hom}(P, B))$ is an iso; clearly if the image of $\alpha \colon A \to B$ is zero then $\alpha = 0$ by applying to any epi $p \colon P \twoheadrightarrow A$; for surjectivity, consider again an epi $p \colon P \twoheadrightarrow A$ with kernel K; further form $f \colon P \twoheadrightarrow K \rightarrowtail P$; then the image of p on the rhs is some $\bar{p} \colon P \to B$ and the image of $0 = pf$ is $0 = \overline{pf} = \bar{p}f$ so that $\bar{p}$ factors through $\mathrm{coker}\, f = A$, giving a preimage).

We return to the exact localization $\mathcal{L}$, i.e. the category of objects injective (automatically orthogonal) w.r.t.

$$\mathrm{coker}\, y(e) \to 0, \quad \mathrm{coker}\, y(m) \to y(C)$$

for any s.e.s. $0 \to A \xrightarrow{m} B \xrightarrow{e} C \to 0$ (the second is probably not necessary, see Weibel). One produces the localization functor from the small object argument. One then shows that in

$$0 \to y(A) \to y(B) \xrightarrow{y(e)} y(C) \to \mathrm{coker}\, y(e) \to 0$$

the cokernel $W = \mathrm{coker}\, y(e)$ is weakly effaceable (i.e. for any A and any $x \in WA$ there exists an epi $e \colon P \twoheadrightarrow A$ such that $y(e)(x) = 0 \in WP$) and that the localization of any weakly effaceable functor is zero. □

7. Derived functors

Derived functors of F at A are defined by taking a projective resolution $P \to A[0]$ then applying F and taking homology, i.e. $L_nF(A) = H_n(FP)$. The main technical problem to solve is showing the independence of the choice of a projective resolution (and then obviously proving basice properties). Classically, one shows that between any two projective resolutions, there exists a chain map

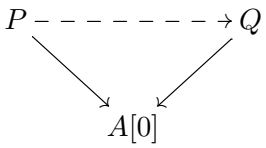

and that any two chain maps are chain homotopic (i.e. the map is unique up to chain homotopy). Applycation of F preserves this and taking homology makes the comparison map unique. There are, however, situation where projective resolutions do not exist, only some weaker version. In such situations, the existence of maps directly from P to Q cannot be expected. There is a weaker version of uniqueness (very much in the modern higher categorical sense), namely the category of (weakly) projective resolutions $P \to A[0]$, i.e. $\mathsf{Ch}(\mathsf{Mod}_R)_{\mathrm{proj}}/A[0]$ is contractible (i.e. its classifying space is) from which we will only need that any two objects can be connected by a zig-zag of morphisms and any two such zig-zags can be connected by a zig-zag of zig-zags (this is just one dimensional triviality). The proof is

not more difficult and one can find the classical approach in all books on homological algebra that I know of (I might later add a short summary).

We will be working with a collection of objects $\mathcal{P}$, called $\mathcal{P}$-projective objects or $\mathcal{P}$-projectives, that is required to satisfy

- every object admits an epi from a $\mathcal{P}$-projective and
- it is closed under kernels of epis, i.e. for a ses $0 \to A \to B \to C \to 0$ with $\mathcal{P}$-projective B and C, the same is true of A.

Typically, $\mathcal{P}$ is also assumed to be closed under finite biproducts, but we will not need this assumption.

We say that $\mathcal{P}$ is adopted to F if in addition to the above assumptions F is exact on ses's of $\mathcal{P}$-projectives. By splicing, it is then exact on bounded below les's of $\mathcal{P}$-projectives as well.

We can then construct a $\mathcal{P}$-projective resolution of any object A in the following way. Construct inductively ses's

$$0 \to K_n \to P_n \to K_{n-1} \to 0,$$

starting from $K_{-2} = 0$ and $P_{-1} = A$, so that $K_{-1} = A$ as well, in such a way that $P_n \in \mathcal{P}$ for each $n \geq 0$ (it exists by the first point). Then splice these ses's to get a les

$$\cdots \to P_1 \to P_0 \xrightarrow{\varepsilon} A \to 0.$$

We will denote this $\mathcal{P}$-projective resolution $\varepsilon\colon P \to A[0]$ (with the above augmented chain complex the mapping cone of this augmentation map ε). We will now show (only partially[1], as required for our exposition) that the category of $\mathcal{P}$-projective resolutions of a fixed A is weakly contractible. First we present a relative version of the above construction: Let $f\colon A \to B$ be a map and $\varepsilon\colon Q \to B[0]$ a resolution, not necessarily $\mathcal{P}$-projective. Then we can complete the following diagram

$$\begin{array}{ccc} P & \xrightarrow{\varepsilon} & A[0] \\ {\scriptstyle\varphi}\downarrow & & \downarrow{\scriptstyle f[0]} \\ Q & \xrightarrow[\varepsilon]{} & B[0] \end{array}$$

in such a way that if f is epi then so is φ (I guess that φ is epi from dimension 1 onwards regardless of f!!! In fact, take the pullback of ε along $f[0]$, which is an epi q-iso – see below – and thus we only need to consider the case $f = 1$). One proceeds exactly as above but using

[1] In general, let $P_i \to A[0]$ be an $\mathcal{I}$-diagram of $\mathcal{P}$-projective resolutions and replace it by a fibrant diagram P_i^{f} in the model structure with pointwise cofibrations and weak equivalences (here the fibrations of chain complexes are not necessarily surjective, but acyccic fibrations are and that should be enough); take the limit of this diagram and pullback

$$\begin{array}{ccc} \overline{P} & \longrightarrow & \lim P_i^{\mathrm{f}} \\ \downarrow & \lrcorner & \downarrow{\scriptstyle\sim} \\ A[0] & \longrightarrow & \lim A[0] \end{array}$$

This then forms a resolution $\overline{P} \to A[0]$ and one can then replace it by a $\mathcal{P}$-resolution with the map $P \to \lim P_i^{\mathrm{f}}$ corresponding to a cone $P \Rightarrow P_i^{\mathrm{f}} \Leftarrow P_i$ that together with the natural transformation shows that the category of $\mathcal{P}$-resolutions is weakly contractible – thus, one should in fact assume that P_i^{f} is itself a $\mathcal{P}$-resolution, i.e. that $\mathcal{P}$ should be closed under products (finite should be sufficient as we can assume $\mathcal{I}$ to be finite in the sense that $N\mathcal{I}$ is (locally) finite).

a pullback square

$$\begin{array}{ccccccccc}
0 & \longrightarrow & K_n & \longrightarrow & P_n & \longrightarrow & K_{n-1} & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \| & & \\
0 & \longrightarrow & L_n & \longrightarrow & \overline{P}_n & \longrightarrow & K_{n-1} & \longrightarrow & 0 \\
 & & \| & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & L_n & \longrightarrow & Q_n & \longrightarrow & L_{n-1} & \longrightarrow & 0
\end{array}$$

(this requires observation that kernels and pullbacks commute and also that a pullback of epi is epi).

This applies in particular to the following situation: given two resolutions $P' \to A[0]$ and $P'' \to A[0]$ we can form their pullback

$$\begin{array}{ccc}
\overline{P} & \longrightarrow & P'' \\
\downarrow & & \downarrow \\
P' & \longrightarrow & A[0]
\end{array}$$

and since epi q-iso's are closed under pullbacks (the epi part we know, then one takes kernels, which agree and are acyclic) we see that all the maps in the square are such. Replacing $\overline{P}$ by a $\mathcal{P}$-projective resolution P as above, we thus get a span of epis between $\mathcal{P}$-projective resolutions $P' \leftarrow P \to P''$. We will need a further level of dimension: Given two spans $\widehat{P}$ and $\widetilde{P}$ between P' and P'' as above (i.e. two epis $\widehat{P} \to \overline{P}$ and $\widetilde{P} \to \overline{P}$) we may form their pullback over $\overline{P}$ and get

$$\begin{array}{ccccc}
\underline{P} & \longrightarrow & \widetilde{P} & & \\
\downarrow & & \downarrow & & \\
\widehat{P} & \longrightarrow & \overline{P} & \longrightarrow & P'' \\
 & & \downarrow & & \downarrow \\
 & & P' & \longrightarrow & A[0]
\end{array}$$

and finally resolve $\underline{P}$ by a $\mathcal{P}$-projective P to get a span between spans:

$$\begin{array}{ccccc}
 & & P' & & \\
 & \nearrow & \uparrow & \nwarrow & \\
\widehat{P} & \longleftarrow & P & \longrightarrow & \widetilde{P} \\
 & \searrow & \downarrow & \swarrow & \\
 & & P'' & &
\end{array}$$

Now we finally utilize the above to the definition and properties of derived functors. We assume that $\mathcal{P}$ is adopted to F. Given an epi $\varphi\colon P \to P'$ between $\mathcal{P}$-projective resolutions of A, the kernel $\ker\varphi$ is then an acyclic chain complex of $\mathcal{P}$-projectives and as such remains acyclic upon applying F. Thus, the map $FP \to FP'$ is still an (epi) q-iso. Applying homology

then yields an iso $H_*FP \to H_*FP'$. We thus get a diagram of isomorphisms

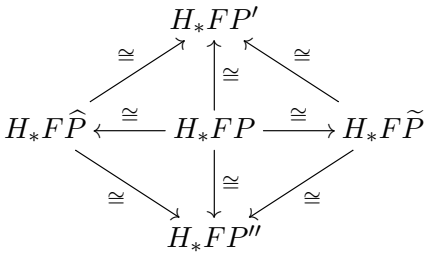

so that we get a well defined (i.e. unique) comparison isomorphism $H_*FP' \cong H_*FP''$. We may thus define $L_*F(A) = H_*FP$ where $P \to A[0]$ is any $\mathcal{P}$-projective resolution and we get that any two possible such definitions are isomorphic in a canonical way, allowing us to talk e.g. about individual elements of $L_*F(A)$.

Now given a ses

$$0 \to A \to B \to C \to 0$$

we take an arbitrary $\mathcal{P}$-projective resolution $R \to C[0]$, then construct a $\mathcal{P}$-projective resolution $Q \to B[0]$ together with an epi $Q \to R$ and finally take kernels to get

$$\begin{array}{ccccccccc} 0 & \longrightarrow & P & \longrightarrow & Q & \longrightarrow & R & \longrightarrow & 0 \\ & & \downarrow & & \downarrow\simeq & & \downarrow\simeq & & \\ 0 & \longrightarrow & A[0] & \longrightarrow & B[0] & \longrightarrow & C[0] & \longrightarrow & 0 \end{array}$$

By the properties of $\mathcal{P}$ we conclude that P consists of $\mathcal{P}$-projectives and, by 5-lemma, the left vertical map is a q-iso, making it a $\mathcal{P}$-projective resolution. Thus, upon applying H_*F to the top row we get a les consisting of the left derived functors $L_*F(A)$, $L_*F(B)$, $L_*F(C)$.

Finally, if A is itself $\mathcal{P}$-projective then the augmented chain complex $P \to A[0]$ remains exact upon applying F and, thus, $L_0F(A) = FA$ and $L_nF(A) = 0$ for $n > 0$. In particular, we get that $\mathcal{P}$ is contained in the collection $\overline{\mathcal{P}} = \{A \mid L_nF(A) = 0 \text{ for all } n > 0\}$. Since this class satisfies the properties, it is thus the maximal such class for a given functor F.

Remark. We will now show independence of L_nF of the class $\mathcal{P}$. Thus, let $\mathcal{Q}$ be another class and consider a $\mathcal{Q}$-resolution $Q \to A[0]$, further a $\mathcal{P}$-resolution $P \to A[0]$ etc. as in

$$\begin{array}{ccc} P' & \longrightarrow & A[0] \\ \sim\downarrow & & \| \\ Q' & \longrightarrow & A[0] \\ \sim\downarrow & & \| \\ P & \longrightarrow & A[0] \\ \sim\downarrow & & \| \\ Q & \longrightarrow & A[0] \end{array}$$

Since both composites $P' \to P$ and $Q' \to Q$ are epi q-iso's between complexes of $\mathcal{P}$-projectives or $\mathcal{Q}$-projectives, they remain q-iso's upon applying F so that the middle map $FQ' \to FP$ is a q-iso by the 2-out-of-6 property, proving $L^{\mathcal{Q}}_*F \cong L^{\mathcal{P}}_*F$. The uniqueness of this isomorphism follows by comparing to the maximal class above

$$\mathcal{P} \subseteq \overline{\mathcal{P}} = \overline{\mathcal{Q}} \supseteq \mathcal{Q}$$

that is independent by the mere existence of an isomorphism showing that the above comparison maps can be thought of as comparison maps for the class $\overline{\mathcal{P}} = \overline{\mathcal{Q}}$ and are thus unique.

Remark. One should also prove that each L_nF is additive and I thought that this would require the class $\mathcal{P}$ to be closed under finite biproducts, and it seems so.

Proposition 7.1. *A right exact functor F is exact iff $\forall n > 0\colon L_nF = 0$ iff $L_1F = 0$.*

8. Balancing Tor and Ext

We define $\operatorname{Tor}_n^R(A,B) = L_n(-\otimes_R B)(A)$. There is a second candidate, namely $L_n(A\otimes_R -)(B)$. One we show that these are the same, we will know that these can be defined using flat resolutions (since flat modules are acyclic).

A similar situation arises for $\operatorname{Ext}^n_R(A,B) = R^n(\operatorname{Hom}_R(A,-))(B)$ and the symmetric version obtained from the contravariant hom functor. We will show that in both cases, the two derived functors are canonically isomorphic. We will concentrate on the tensor products since these are both covariant and thus easier. The two derived functors are obtained as homology of chain complexes $P\otimes_R B$ and $A\otimes_R Q$ where $P\to A[0]$ and $Q\to B[0]$ are projective resolutions. It thus seems more than logical to compare these using a span

$$P\otimes_R B \leftarrow P\otimes_R Q \to A\otimes Q.$$

The question is what this $P\otimes_R Q$ should be. We can draw a diagram where we write only $\otimes$ for simplicity

$$\begin{array}{ccccc}
\vdots & & \vdots & & \\
\downarrow & & \downarrow & & \\
\cdots \longleftarrow & P_{p-1}\otimes Q_q & \xleftarrow{d\otimes 1} & P_p\otimes Q_q & \longleftarrow \cdots \\
\downarrow{\scriptstyle 1\otimes d} & & \downarrow{\scriptstyle 1\otimes d} & & \\
\cdots \longleftarrow & P_{p-1}\otimes Q_{q-1} & \xleftarrow{d\otimes 1} & P_p\otimes Q_{q-1} & \longleftarrow \cdots \\
\downarrow & & \downarrow & & \\
\vdots & & \vdots & &
\end{array}$$

where $d^h = d\otimes 1$ and $d^v = 1\otimes d$ will not be exactly one's first guess, resulting in the squares anti-commuting rather than commuting, i.e. $d^hd^v = -d^vd^h$. We will now make this kind of structure formal.

Definition 8.1. A double (chain) complex is a diagram of modules $D_{p,q}$ and homomorphisms $d^h\colon D_{p,q}\to D_{p-1,q}$, $d^v\colon D_{p,q}\to D_{p,q-1}$ that satisfy $(d^h+d^v)^2 = 0$, i.e.

$$d^hd^h = 0, \quad d^vd^v = 0, \quad d^hd^v + d^vd^h = 0.$$

This way of presenting the axioms suggests the following definition.

Definition 8.2. For a double complex D, the total complex $\operatorname{Tot}^+ D$ is a chain complex with

$$(\operatorname{Tot}^+ D)_n = \sum_{p+q=n} D_{p,q}$$

and with differential $d = d^h + d^v$.

There is another version of the total complex $\operatorname{Tot}^\times D$ where the sum is replaced by the product.

The two versions are useful for different applications, the first is related to left derived functors and the second for right derived functors. Since we concentrate on the tensor product case, we will stick to the sum version and will denote it simply by $\operatorname{Tot} D$.

In order to make the tensor product into an example, we have to introduce signs. We define

$$(f \otimes g)(x \otimes y) = (-1)^{|g|\cdot|x|} fx \otimes gy$$

where $|g| = |gx| - |x|$ is the degree of g (we will treat this more formally later). Thus the identity has degree $|1| = 0$, while the differential has degree $|d| = -1$. This gives as particular cases

$$d^h(x \otimes y) = (d \otimes 1)(x \otimes y) = dx \otimes y, \quad d^v(x \otimes y) = (1 \otimes d)(x \otimes y) = (-1)^{|x|} x \otimes dy$$

and the squares clearly anti-commute with this notation (one can also prove this formally by first verifying $(f \otimes g)(h \otimes k) = (-1)^{|g|\cdot|h|} fh \otimes gk$ and then using this to compute $(d \otimes 1 + 1 \otimes d)^2 = 0$). We speak of Koszul sign convention.

Later, when we will deal with chain complexes, we will be using $P \otimes Q$ to denote the total complex $\operatorname{Tot} P \otimes Q$ of this double complex.

We will now introduce two very useful special cases of the tensor product construction. The motivation comes from topology, where $C(X \times Y) = C(X) \otimes C(Y)$ (also for Koszul sign convention), at least when one deals with cellular complexes where products of cells are cells. In this way one obtains the cylinder of C by tensoring with the chain complex of the interval:

$$\operatorname{cyl} C = \operatorname{Tot}(\operatorname{cyl} R[0] \otimes C)$$

since $R[0]$ is interpreted as a point and then the cylinder on the point is the interval; it remains to specify this interval:

$$\operatorname{cyl} R[0] = \cdots \to 0 \to R \xrightarrow{d} R \oplus R$$

Denoting the 1-dimensional generator by e (edge) and the 0-dimensional generators by v_-, v_+ (initial and terminal vertices), we define $de = v_+ - v_-$.

Exercise 8.3. Prove that chain maps $\operatorname{cyl} C \to D$ are in bijection with triples (f, g, h) where f and g are chain maps $C \to D$ and h is a chain homotopy $f \sim g$. In topology, one can recover the two involved maps from a homotopy (as restrictions to the two ends of the cylinder), while in homological algebra, this is not the case – the best one can get is the difference $g - f = dh + hd$.

Another example that we will need is the cone. We define similarly

$$\operatorname{cone} C = \operatorname{Tot}(\operatorname{cone} R[0] \otimes C)$$

where again the chain complex

$$\operatorname{cone} R[0] = \cdots \to 0 \to R \xrightarrow{d} R$$

has differential $de = v$, i.e. $d = 1$. We will now draw a picture of the double complex

cone $R[0] \otimes C$ and a simpler realization thereof:

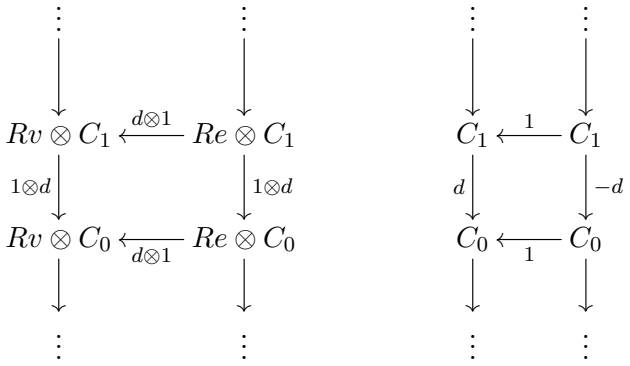

where the minus sign comes from $(1 \otimes d)(e \otimes x) = (-1)^{|e|} e \otimes dx = e \otimes (-dx)$, since $|e| = 1$. Clearly the zeroth column form a subcomplex (since $R[0] \subseteq \operatorname{cone} R[0]$ and then one applies the tensor product), so $C \hookrightarrow \operatorname{cone} C$. The quotient is the first column, i.e. the chain complex $C[1]$ – called the suspension of C – is just C shifted by one dimension up and with opposite differential (again, the quotient of $R[0] \hookrightarrow \operatorname{cone} R[0]$ is $R[1]$ and $C[1] = \operatorname{Tot} R[1] \otimes C$).

What comes now is a concrete description of the pushout

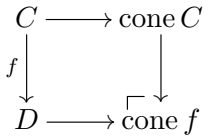

(with horizontal cokernels $C[1]$, see below) that results from replacing the subcomplex $C \subseteq \operatorname{cone} C$ by D via f. It is the total complex of the double complex

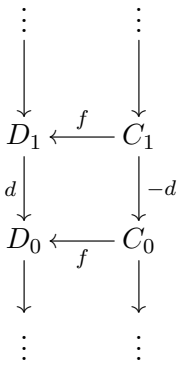

Similarly to the case of cone C, we get a subcomplex and a quotient, forming a short exact sequence

$$0 \to D \to \operatorname{cone} f \to C[1] \to 0$$

Exercise 8.4. Verify that the connecting homomorphism in the homology long exact sequence is the map $H_{n+1}(C[1]) = H_n(C) \to H_n(D)$ induced by f. Conclude that f is a q-iso iff cone f is acyclic. Apply to the augmentation map.

Proposition 8.5. *Let D be a first quadrant double complex, i.e. such that $D_{p,q} = 0$ whenever $p < 0$ or $q < 0$. If D has exact columns, i.e. if for each p the chain complex $(D_{p,\bullet}, d^v)$ is acyclic, then* $\operatorname{Tot} D$ *is acyclic.*

Dually, the same conclusion holds for first quadrant double complexes with exact rows.

Remark. This version works for right halfplane double complexes (or upper halfplane complexes in the second case), but the $\operatorname{Tot}^\times$-version requires this stronger assumption, I think.

Proof. Denote $D^{(0)} = D$. As above, the zeroth column forms a subcomplex $D_{0,\bullet}$ with quotient $D^{(1)}$, obtained by removing the zeroth column. Continuing this way, we obtain short exact sequences

$$0 \to D_{p,\bullet} \to \operatorname{Tot} D^{(p)} \to \operatorname{Tot} D^{(p+1)} \to 0$$

which shows that the natural projection maps

$$\operatorname{Tot} D = \operatorname{Tot} D^{(0)} \to \operatorname{Tot} D^{(1)} \to \cdots$$

are all q-iso's. Since $\operatorname{Tot} D^{(n+1)}$ is concentrated in dimensions $\geq n+1$, it has zero H_n and thus the same is true for $\operatorname{Tot} D$. □

Now consider the double complex obtained as a tensor product of the augmented chain complex P and the chain complex Q, i.e.

$$\begin{array}{ccccccc}
\vdots & & \vdots & & \vdots & & \\
\downarrow & & \downarrow & & \downarrow & & \\
A \otimes Q_1 & \longleftarrow & P_0 \otimes Q_1 & \longleftarrow & P_1 \otimes Q_1 & \longleftarrow & \cdots \\
\downarrow & & \downarrow & & \downarrow & & \\
A \otimes Q_0 & \longleftarrow & P_0 \otimes Q_0 & \longleftarrow & P_1 \otimes Q_0 & \longleftarrow & \cdots
\end{array}$$

This has exact rows since these are obtained by tensoring the augmented chain complex P with a projective Q_q. Thus, the total complex is acyclic. Since it is (up to suspension) the cone of the map $\varepsilon \otimes 1 \colon P \otimes Q \to A \otimes Q$, this map is a q-iso. Symetrically, the map $1 \otimes \varepsilon \colon P \otimes Q \to P \otimes B$ is also a q-iso and we obtain the following theorem:

Theorem 8.6 (balancing of Tor)**.** *There exists a natural isomorphism*

$$L_n(- \otimes B)(A) \cong L_n(A \otimes -)(B)$$

between the derived functors of the tensor product functor.

In fact, one can see easily that for a right exact bifunctor F, we only need that $F(P, -)$ should be exact for any projective P as well as $F(-, Q)$ for any projective Q and we obtain

$$L_n F(A, -)(B) \cong L_n F(-, B)(A)$$

We will now shortly comment on the derived functors of the hom functor. The covariant one is easier, so we start with this:

$$R^n(\operatorname{Hom}(A, -))(B) = H^n(\operatorname{Hom}(A, I))$$

where $B[0] \to I$ is an injective resolution. The situation of the other hom functor is exactly the same when interpreted in the opposite category; translating to the ordinary category of modules, we get

$$R^n(\mathrm{Hom}(-,B))(A) = H^n(\mathrm{Hom}(P,B))$$

where $P \to A[0]$ is a projective resolution. Again, we can form a double cochain complex $\mathrm{Hom}(P,I)$ and its total complex $\mathrm{Tot}^\times \mathrm{Hom}(P,I)$ that admits a cospan

$$\mathrm{Hom}(P,B) \to \mathrm{Tot}^\times \mathrm{Hom}(P,I) \leftarrow \mathrm{Hom}(A,I)$$

with both maps q-iso's by an analogous argument.

Theorem 8.7 (balancing of Ext)**.** *There exists a natural isomorphism*

$$R^n(\mathrm{Hom}(A,-)(B) \cong R^n(\mathrm{Hom}(-,B)(A)$$

between the derived functors of the hom functor.

We will now study hom complexes from a different perspective – related, but it may be easier to forget about what we did up to now. So let C and D be chain complexes and construct a chain complex $\mathrm{Hom}(C,D)$ with the aim of giving the category of chain complexes the closed symmetric monoidal structure. Symmetry is perhaps worth mentioning first, since it is given by a (not so much now) surprising isomorphism

$$B \otimes C \xrightarrow{\cong} C \otimes B, \quad x \otimes y \mapsto (-1)^{|x|\cdot|y|} y \otimes x.$$

Now our goal is the adjointness

$$\frac{B \otimes C \to D}{B \to \mathrm{Hom}(C,D)}$$

We will first study this on the level of the underlying graded modules (i.e. ignore the differentials). This becomes

$$\frac{\sum_{n+k=\ell} B_n \otimes C_k \to D_\ell}{B_n \to \prod_k \mathrm{Hom}(C_k, D_{n+k})}$$

so we want to endow the graded module $\mathrm{Hom}(C,D)_n = \prod_k \mathrm{Hom}(C_k, D_{n+k})$ with a differential so that the map at the top is a chain map iff the map at the bottom is. The differential will be derived from the requirement that the counit is a chain map, by observing that the counit is (as usual) the evaluation map

$$\mathrm{ev}\colon \mathrm{Hom}(C,D) \otimes C \to D, \quad f \otimes c \mapsto fc.$$

We will denote the differential on the hom complex by D and we thus require

$$d\,\mathrm{ev} = \mathrm{ev}(D \otimes 1 + 1 \otimes d),$$

that by applying to $f \otimes c$ amounts to

$$d(fc) = (Df)c + (-1)^{|f|} f(dc).$$

We can thus write $Df = df - (-1)^{|f|} fd = [d,f]$ (the graded commutator). It is rather straightforward that this indeed makes $\mathsf{Ch}(\mathsf{Mod}_R)$ into a closed symmetric monoidal category.

The 0-chains of $\mathrm{Hom}(C, D)$ are by the construction not-necessarily-chain maps $f\colon C \to D$. The 0-cycles are those that satisfy $Df = 0$, i.e.

$$df - fd = 0$$

and these are exactly the chain maps.[2] (Chain maps of degree n are defined as n-cycles, i.e. they are required to satisfy $df = (-1)^n fd$.) For two chain maps f, g, i.e. 0-cycles, we have $[f] = [g]$ in $H_0(\mathrm{Hom}(C, D))$ iff $g - f \in B_0(\mathrm{Hom}(C, D))$ iff there exists $h \in \mathrm{Hom}(C, D)_1$ with $Dh = g - f$; this means

$$dh + hd = g - f$$

and this h is a chain homotopy from f to g. As a result

$$H_0(\mathrm{Hom}(C, D)) = [C, D]$$

the group of chain homotopy classes of chain maps. This is another explanation of chain homotopy (we had definition, then as a map from the cylinder, now as a homology relation in the hom complex).

9. Ext and extensions

We will now apply the derived functor Ext^1 to study extensions of modules, i.e. short exact sequences

$$\xi\colon 0 \to B \to X \to A \to 0$$

We start with a simple question: When does the sequence split? By applying $\mathrm{Hom}(A, -)$ we obtain an exact sequence

$$0 \to \mathrm{Hom}(A, B) \to \mathrm{Hom}(A, X) \to \mathrm{Hom}(A, A) \xrightarrow{\partial} \mathrm{Ext}^1(A, B) \to \mathrm{Ext}^1(A, X) \to \cdots$$

Clearly, ξ admits a splitting iff $1 \in \mathrm{Hom}(A, A)$ lies in the image (more precisely, any preimage is such a splitting $A \to X$) iff $\partial(1) = 0$. We define

$$\theta(\xi) = \partial(1) \in \mathrm{Ext}^1(A, B)$$

and we just observed that this is the (unique) obstruction to the existence of a splitting.

Lemma 9.1. *ξ splits iff $\theta(\xi) = 0$. In particular, it splits when $\mathrm{Ext}^1(A, B) = 0$.* □

Naturality of the class θ with respect to maps of ses's should be rather clear: we need

$$\begin{array}{lccccccc} \xi\colon & 0 & \longrightarrow & B & \longrightarrow & X & \longrightarrow & A & \longrightarrow & 0 \\ & & & \| & & \downarrow & & \| & & \\ \zeta\colon & 0 & \longrightarrow & B & \longrightarrow & Y & \longrightarrow & A & \longrightarrow & 0 \end{array}$$

[2] This corresponds to the fact that the unit is $R[0]$ and maps out of $R[0]$ are exactly the 0-cycles.

(the map $X \to Y$ is then necessarily an iso by 5-lemma) so that the portions of long exact sequences of derive functors

$$\begin{array}{ccc} \mathrm{Hom}(A,A) & \xrightarrow{\partial} & \mathrm{Ext}^1(A,B) \\ {\scriptstyle 1}\downarrow & & \downarrow{\scriptstyle 1} \\ \mathrm{Hom}(A,A) & \xrightarrow{\partial} & \mathrm{Ext}^1(A,B) \end{array}$$

have both vertical maps identities – they are induced by maps in the transformation above so we require these to be identities. We will then say that the extensions ξ and ζ are isomorphic and we see that then $\theta(\xi) = \theta(\zeta)$. We have just defined, for fixed A and B, a map

$$\theta\colon \{\text{extensions } 0 \to B \to X \to A \to 0\}/\text{iso} \to \mathrm{Ext}^1(A,B).$$

Theorem 9.2. *The above map θ is bijective.*

Proof. We first produce a map in the opposite direction. Let $B \to I$ be an embedding of B into an injective module and let C be the cokernel so that we have a ses

$$\zeta\colon 0 \to B \xrightarrow{i} I \xrightarrow{p} C \to 0.$$

The les of $\mathrm{Ext}(A,-)$ then gives

$$\cdots \to \mathrm{Hom}(A,I) \xrightarrow{p_*} \mathrm{Hom}(A,C) \to \mathrm{Ext}^1(A,B) \to \underbrace{\mathrm{Ext}^1(A,I)}_{0} \to \cdots$$

(since I is injective). Thus, for any $\alpha \in \mathrm{Ext}^1(A,B)$ there exists a preimage $f\colon A \to C$ and any other preimage is of the form $f + pg$ for $g\colon A \to I$. We form a pullback of the ses above along f and obtain

$$\begin{array}{lccccccccc} \xi_f\colon & 0 & \longrightarrow & B & \longrightarrow & X_f & \longrightarrow & A & \longrightarrow & 0 \\ & & & \| & & \downarrow & \lrcorner & \downarrow{\scriptstyle f} & & \\ \zeta\colon & 0 & \longrightarrow & B & \xrightarrow{i} & I & \xrightarrow{p} & C & \longrightarrow & 0 \end{array}$$

Concretely $X_f = \{(x,a) \mid p(x) = f(a)\}$ and we thus have an isomorphism

$$X_f \xrightarrow{\cong} X_{f+pg}, \quad (x,a) \mapsto (x + g(a), a)$$

that respects the inclusion of B and projection onto A so this is in fact an isomorphism of extensions $\xi_f \cong \xi_{f+pg}$. This finishes the construction of the inverse mapping, we need to verify that these are indeed inverse to each other.

We thus study the obstruction $\theta(\xi_f)$ of the obstruction above. Again, the transformation of ses's gives a transformation between the sequences of derived functors

$$\begin{array}{ccc} \mathrm{Hom}(A,A) & \xrightarrow{\partial} & \mathrm{Ext}^1(A,B) \\ {\scriptstyle f_*}\downarrow & & \downarrow{\scriptstyle 1} \\ \mathrm{Hom}(A,C) & \xrightarrow{\partial} & \mathrm{Ext}^1(A,B) \end{array}$$

and this means precisely $\theta(\xi_f) = \partial(1) = \partial f_*(1) = \partial(f) = \alpha$ by construction (f was chosen as a preimage of α). It remains to show that the constructed inverse mapping is surjective, i.e. that every extension is obtained as a pullback from ζ:

$$\begin{array}{lccccccc} \xi\colon & 0 & \longrightarrow & B & \longrightarrow & X & \longrightarrow & A & \longrightarrow & 0 \\ & & & \| & & \downarrow & & \downarrow & & \\ \zeta\colon & 0 & \longrightarrow & B & \xrightarrow{i} & I & \xrightarrow{p} & C & \longrightarrow & 0 \end{array}$$

Start by extending the map i into the injective I along the inclusion $B \to X$, as suggested in the diagram. We then complete the diagram by a map $f\colon A \to C$ that is just the induced maps on cokernels. Since both $X \to A$ and p are epi and the induced map on kernels is an iso, Proposition 5.15 yields that the square is indeed a pullback. □

Example 9.3. Since $\mathrm{Ext}^1(\mathbb{Z}/m, \mathbb{Z}/n) = 0$ if $\gcd(m, n) = 1$, every ses

$$0 \to \mathbb{Z}/n \to X \to \mathbb{Z}/m \to 0$$

splits. We will prove later a more general result for non-commutative groups.

Since Ext^1 as a derived functor is additive in both variables, we get $\mathrm{Ext}^1(B, A) = 0$ for any finite abelian groups of coprime orders (split into a direct sum and apply the above).

Interestingly, since $\mathrm{Ext}^1(A, B)$ is an abelian group, the same must be true for the set of extensions up to isomorphism. It is instructive to transport the addition along the above isomorphism. The result looks as follows. Take two extensions and consider their biproduct

$$0 \to B \oplus B \to X \oplus Y \to A \oplus A \to 0.$$

Now take the pullback as in the above proof along the diagonal $A \to A \oplus A$ to obtain an extension of A by $B \oplus B$. Perform the dual construction, i.e. form the pushout along the codiagonal $B \oplus B \to B$ (i.e. the addition) and finally obtain an extension of A by B; this is the sum of the original extensions.

Remark. The higher groups $\mathrm{Ext}^n(A, B)$ are in bijection with classes of longer extensions

$$0 \to B \to X_n \to \cdots \to X_1 \to A \to 0$$

modulo an equivalence relation generated by not necessarily invertible transformations

$$\begin{array}{ccccccccccc} 0 & \longrightarrow & B & \longrightarrow & X_n & \longrightarrow & \cdots & \longrightarrow & X_1 & \longrightarrow & A & \longrightarrow & 0 \\ & & \| & & \downarrow & & & & \downarrow & & \| & & \\ 0 & \longrightarrow & B & \longrightarrow & Y_n & \longrightarrow & \cdots & \longrightarrow & Y_1 & \longrightarrow & A & \longrightarrow & 0 \end{array}$$

There is a way of explaining this in a more natural way as follows. Consider the middle part as a chain complex X concentrated in dimension 1 through n and rewrite the exact sequence as

$$0 \to B[n] \to X \to A[1] \to 0,$$

a ses of chain complexes. The natural transformation above becomes

$$\begin{array}{ccccccccc} 0 & \longrightarrow & B[n] & \longrightarrow & X & \longrightarrow & A[1] & \longrightarrow & 0 \\ & & \| & & \downarrow & & \| & & \\ 0 & \longrightarrow & B[n] & \longrightarrow & Y & \longrightarrow & A[1] & \longrightarrow & 0 \end{array}$$

and the middle map is a q-iso by 5-lemma. To make this comparison complete, one should show that chain complexes that are not concentrated in dimensions 1 through n can be truncated to the latter (easy).

10. Homological dimension

We have just shown that $\mathrm{Ext}^1(-,-) = 0$ iff every ses splits and we know that vanishing of $\mathrm{Ext}^1(A,-)$ is equivalent to A being projective so $\mathrm{Ext}^1(-,-) = 0$ is also equivalent to every module being projective and dually also to every module being injective. We will now study higher dimensional analogues of such statements.

Definition 10.1. A *projective dimension* of a module A, denoted $\mathrm{pd}(A)$, is defined to be the length of the shortest projective resolution of A, i.e. $\mathrm{pd}(A) \leq n$ iff A admits a projective resolution

$$\begin{array}{ccccccccc} \cdots & \longrightarrow & 0 & \longrightarrow & P_n & \longrightarrow & \cdots & \longrightarrow & P_0 \\ & & & & & & & & \downarrow \\ & & & & & & & & A \end{array}$$

There are similar notions of a flat dimension and injective dimension (the second using injective coresolutions):

$$\begin{array}{ccccccccc} A & & & & & & & & \\ \downarrow & & & & & & & & \\ I^0 & \longrightarrow & \cdots & \longrightarrow & I^n & \longrightarrow & 0 & \longrightarrow & \cdots \end{array}$$

Lemma 10.2. *TFAE*

1. $\mathrm{pd}(A) \leq n$,
2. *in any exact sequence*

$$0 \to M_n \to P_{n-1} \to \cdots \to P_0 \to A \to 0$$

with P_i projective, also M_n is projective,

3. $\mathrm{Ext}^{n+1}(A,-) = 0$.

Proof. The implications $2 \Rightarrow 1 \Rightarrow 3$ are trivial. Thus, let $\mathrm{Ext}^{n+1}(A,-) = 0$ and consider an exact sequence

$$0 \to M_n \to P_{n-1} \to \cdots \to P_0 \to A \to 0.$$

Denoting $M_0 = A$, we split it into ses's

$$0 \to M_{k+1} \to P_k \to M_k \to 0.$$

Applying $\mathrm{Ext}(-,B)$ yields, by projectivity of P_k the following isomorphisms for $i \geq 1$:

$$\underbrace{\mathrm{Ext}^{i+1}(P_k,B)}_{0} \leftarrow \mathrm{Ext}^{i+1}(M_k,B) \xleftarrow{\cong} \mathrm{Ext}^{i}(M_{k+1},B) \leftarrow \underbrace{\mathrm{Ext}^{i}(P_k,B)}_{0}$$

We thus obtain

$$\mathrm{Ext}^1(M_n,B) = \cdots = \mathrm{Ext}^n(M_1,B) = \mathrm{Ext}^{n+1}(M_0,B) = 0$$

and since this holds for any B, the module M_n is projective. □

A dual statement then shows that the injective dimension $\mathrm{id}(B) \leq n$ is also equivalent to $\mathrm{Ext}^{n+1}(-, B) = 0$. Together, these results yield.

Corollary 10.3. $\sup\{\mathrm{pd}(A) \mid A \in \mathsf{Mod}_R\} = \inf\{n \mid \mathrm{Ext}^{n+1} = 0\} = \sup\{\mathrm{id}(A) \mid A \in \mathsf{Mod}_R\}$. □

Remark. It is also true that this equals $d = \sup\{\mathrm{pd}(R/J) \mid J \subseteq R \text{ ideal}\}$. For an arbitrary module A, consider

$$0 \to B \to I^0 \to \cdots \to I^{d-1} \to M^d \to 0$$

and conclude $0 = \mathrm{Ext}^{d+1}(R/J, B) \cong \mathrm{Ext}^1(R/J, M^n)$, i.e. $\mathrm{Hom}(R, M^n) \twoheadrightarrow \mathrm{Hom}(J, M^n)$ and M^n is injective by Baer criterion.

Definition 10.4. The number from the previous corollary is called the global dimension of R and denoted $\mathrm{gl.dim}(R)$.

Similarly, one can prove that the supremum of flat dimensions of modules does not depend on the side and equals the smallest n for which $\mathrm{Tor}^{n+1} = 0$, called $\mathrm{Tor.dim}(R)$.

Example 10.5. A ring R has global dimension 0 iff $\mathrm{Ext}^1 = 0$ iff every module is projective iff every module is injective.

A ring R has global dimension 1 iff $\mathrm{Ext}^2 = 0$ iff in every ses $0 \to M \to P \to A \to 0$, the module M is projective; since A could be arbitrary, e.g. the cokernel of an arbitrary inclusion $M \to P$, this is equivalent to a submodule of a projective module being projective. Dually, this is equivalent to a quotient of an injective module being injective.

Any PID has global dimension 1: It can be proved by induction on n that a submodule of R^n is free of rank $\leq n$, starting from $n = 1$ where this is just the definition of a PID.

Theorem 10.6 (Hilbert on szyzygies)**.** $\mathrm{gl.dim}\, \Bbbk[x_1, \ldots, x_n] = n$. *More generally, if* $\mathrm{gl.dim}\, R = d$ *then* $\mathrm{gl.dim}\, R[x] = d + 1$.

The full strength of the Hilbert theorem on szyzygies gives part 2 of Lemma 10.2 with projective replaced by free.

Theorem 10.7 (Künneth)**.** *Assume that* $\mathrm{Tor.dim}\, R \leq 1$*, i.e. that every submodule of a flat module is flat. Let C be a chain complex of flat modules and A and module. Then there exists a natural ses (unnaturally split)*

$$0 \to H_n(C) \otimes A \to H_n(C \otimes A) \to \mathrm{Tor}_1(H_{n-1}(C), A) \to 0.$$

Proof. Apply $\mathrm{Tor}(-, A)$ to the ses

$$0 \to Z_n \to C_n \to B_{n-1} \to 0$$

of flat modules to obtain a ses of chain complexes

$$0 \to Z \otimes A \to C \otimes A \to B[1] \otimes A \to 0$$

where the outer chain complexes are endowed with zero differential. Now apply the les of homology:

$$\cdots \to (B[1] \otimes A)_{n+1} \xrightarrow{\partial_n} (Z \otimes A)_n \to H_n(C \otimes A) \to (B[1] \otimes A)_n \xrightarrow{\partial_{n-1}} (Z \otimes A)_{n-1} \to \cdots .$$

The connecting homomorphism ∂_n is the canonical map

$$i \otimes 1 \colon B_n \otimes A \to Z_n \otimes A$$

that fits into a les of $\operatorname{Tor}(-, A)$ applied to $0 \to B_n \to Z_n \to H_n \to 0$, yielding (recalling that Z_n must be flat)

$$0 \to \operatorname{Tor}(H_n, A) \to B_n \otimes A \xrightarrow{\partial_n} Z_n \otimes A \to H_n \otimes A \to 0.$$

In other words, $\operatorname{coker} \partial_n = H_n \otimes A$ and $\ker \partial_{n-1} = \operatorname{Tor}(H_{n-1}, A)$. Thus, one may replace the les above by a ses with $H_n(C \otimes A)$ in the middle, surrounded by the cokernel and the kernel. □

11. Group cohomology

This is a particular derived functor for modules over the group ring $\mathbb{Z}G$ for a group G. It is a free abelian group on the set G, i.e. its elements are formal $\mathbb{Z}$-linear combinations of elements of the group G, say $\sum a_g \cdot g$ with only finitely many nonzero coefficients $a_g \in \mathbb{Z}$. The multiplication is extended $\mathbb{Z}$-linearly from the multiplication in G, i.e.

$$(\sum a_h \cdot h) \cdot (\sum b_k \cdot k) = \sum \sum (a_h b_k) \cdot hk = \sum (\sum_{hk=g} a_h b_k) \cdot g.$$

More abstractly, the free abelian group functor turns finite products into finite tensor products, $\mathbb{Z}(X \times Y) \cong \mathbb{Z}X \otimes \mathbb{Z}Y$, i.e. it is strongly monoidal. Thus, the multiplication in G induces

$$\mathbb{Z}G \otimes \mathbb{Z}G \cong \mathbb{Z}(G \times G) \longrightarrow \mathbb{Z}G$$

and similarly for the unit (which is then just the element $1 \in G$ interpreted as an element of $\mathbb{Z}G$).

A $\mathbb{Z}G$-module is then equivalently an abelian group M together with an action of G via homomorphisms of groups, i.e. $a \cdot (x + y) = a \cdot x + a \cdot y$. This is easily seen to be so by interpreting the module structure as a ring homomorphism $\mathbb{Z}G \to \operatorname{End}_{\mathbb{Z}}(M)$ and by the freeness of $\mathbb{Z}G$, this is induced uniquely by a group homomorphism $G \to \operatorname{Aut}_{\mathbb{Z}}(M)$. Another point of view is that this is a functor $G \to \mathsf{Ab}$ hitting M (or in the first interpretation an Ab-enriched functor $\mathbb{Z}G \to \mathsf{Ab}$).

Example 11.1. The symmetry group S_n acts on $V^{\otimes n}$ by permuting the vectors in the tensor product. An important construction is that of the invariants $(V^{\otimes n})^{S_n}$, i.e. the submodule of tensors that are invariant under the action, i.e. such that $t \cdot \sigma = t$ (since it is naturally a right action, i.e. a right $\mathbb{Z}G$-module). Another related construction are the coinvariants $(V^{\otimes n})_{S_n}$, i.e. the quotient by the congruence generated by $t \cdot \sigma \sim t$. When $\operatorname{char} \Bbbk = 0$, these are two equivalent definitions of the n-th symmetric power $S^n V$.

Definition 11.2. The invariants of a $\mathbb{Z}G$-module M is the submodule

$$M^G = \{x \in M \mid \forall a \in G \colon a \cdot x = x\}.$$

The coinvariants of a $\mathbb{Z}G$-module M is the quotient module

$$M_G = M/(a \cdot x \sim x \mid a \in G,\, x \in M).$$

The first is the limit of the diagram

$$G \circlearrowright M$$

the second is the colimit of the same diagram. There is another intrepretation of the same, using the trivial $\mathbb{Z}G$-module $\mathbb{Z}$. In general any abelian group admits a trivial action where $a \cdot x = x$ for any $a \in G$.

Lemma 11.3. *$M^G = \mathrm{Hom}_{\mathbb{Z}G}(\mathbb{Z}, M)$ and $M_G = M \otimes_{\mathbb{Z}G} \mathbb{Z}$ (for a right $\mathbb{Z}G$-module M).*

Proof. The point is that $\mathbb{Z} = (\mathbb{Z}G)_G$, since the congruence identifies exactly the generators of $\mathbb{Z}G$. Thus,

$$\frac{f\colon \mathbb{Z} \to M}{\frac{f\colon (\mathbb{Z}G)/(a \cdot x \sim x) \to M}{\frac{f\colon \mathbb{Z}G \to M \text{ such that } f(a \cdot x) = f(x)}{\frac{m \in M \text{ such that } a \cdot m = m}{m \in M^G}}}}$$

and similarly

$$M \otimes_{\mathbb{Z}G} \mathbb{Z} \cong M \otimes_{\mathbb{Z}G} \mathbb{Z}G/(a \cdot x \sim x) \cong (M \otimes_{\mathbb{Z}G} \mathbb{Z}G)/(m \otimes a \cdot x \sim m \otimes x) \cong M/(a \cdot m \sim m) = M_G.$$

Perhaps, it is better to relate it to $\mathrm{Hom}_{\mathbb{Z}}$ and $\otimes_{\mathbb{Z}}$. □

Definition 11.4. The n-th group homology with coefficients in a $\mathbb{Z}G$-module M is

$$H_n(G; M) = L_n(-)_G(M) = \mathrm{Tor}_n^{\mathbb{Z}G}(M, \mathbb{Z}) \text{ or } \mathrm{Tor}_n^{\mathbb{Z}G}(\mathbb{Z}, M).$$

The n-th group cohomology with coefficients in a $\mathbb{Z}G$-module M is

$$H^n(G; M) = R^n(-)^G(M) = \mathrm{Ext}^n_{\mathbb{Z}G}(\mathbb{Z}, M).$$

We will study these via a projective resolution of $\mathbb{Z} \in \mathsf{Mod}_{\mathbb{Z}G}$.

Example 11.5. Denote by C_k the cyclic group of order k written multiplicatively (i.e. it is $\mathbb{Z}/k$ but that is usually written additively), with elements $1, t, \cdots, t^{k-1}$ and with $t^k = 1$. A projective resolution was constructed in the tutorial, where the norm $N = t^{k-1} + \cdots + 1$ denotes the sum of all the elements of the group

$$\cdots \xrightarrow{t-1} \mathbb{Z}C_k \xrightarrow{N} \mathbb{Z}C_k \xrightarrow{t-1} \mathbb{Z}C_k \xrightarrow{\mathrm{ev}_1} \mathbb{Z}$$

Now compute the homology of C_k with coefficients $\mathbb{Z}$, i.e. apply $- \otimes_{\mathbb{Z}C_k} \mathbb{Z}$

$$\cdots \xrightarrow{0} \mathbb{Z} \xrightarrow{k} \mathbb{Z} \xrightarrow{0} \mathbb{Z}$$

and take homology to obtain

$$H_n(C_k; \mathbb{Z}) = \begin{cases} \mathbb{Z} & n = 0 \\ \mathbb{Z}/k & n = 1, 3, 5, \ldots \\ 0 & n = 2, 4, 6, \ldots \end{cases}$$

Example 11.6. Denote by C_∞ the infinite cyclic group with elements powers t^k of the generator t. The group ring $\mathbb{Z}C_\infty$ is then the ring of Laurent polynomials. A projective resolution was constructed in the tutorial

$$\begin{array}{ccccccc} \cdots \longrightarrow 0 \longrightarrow \mathbb{Z}C_\infty & \xrightarrow{t-1} & \mathbb{Z}C_\infty \\ & & \downarrow{\scriptstyle \mathrm{ev}_1} \\ & & \mathbb{Z} \end{array}$$

Now compute the homology of C_∞ with coefficients $\mathbb{Z}$, i.e. apply $-\otimes_{\mathbb{Z}C_\infty}\mathbb{Z}$

$$\cdots \longrightarrow 0 \longrightarrow \mathbb{Z} \xrightarrow{0} \mathbb{Z}$$

and take homology to obtain

$$H_n(C_\infty;\mathbb{Z}) = \begin{cases} \mathbb{Z} & n = 0,1 \\ 0 & n = 2,3,4,\dots \end{cases}$$

Remark. It can be shown that $H_n(G;\mathbb{Z}) = H_n(BG;\mathbb{Z})$ equals the singular homology of the classifying space $BG = K(G,1)$. Thus, despite $G = C_k$ being finite, the classifying space BC_k is infinite dimensional (and also $\mathbb{Z}C_k$ has infinite global dimension). On the other hand $BC_\infty \simeq S^1$ is homotopy equivalent to a circle.

We will now construct a general projective resolution of $\mathbb{Z} \in \mathsf{Mod}_{\mathbb{Z}G}$, the so called bar resolution. It has two versions – reduced and unreduced. We start with the second.

Definition 11.7. The *unreduced bar resolution* is the chain complex B^u with chains

$$B^u_n = \mathbb{Z}G(G \times \cdots \times G) = \mathbb{Z}(G \times (G \times \cdots \times G))$$

where we denote the $\mathbb{Z}G$-generators as $[g_1 \otimes \cdots \otimes g_n]$ and thus the $\mathbb{Z}$-generators as $g[g_1 \otimes \cdots \otimes g_n]$. The differential in this complex is $d = \sum (-1)^i d_i$ for $\mathbb{Z}G$-linear operators

$$\begin{aligned} d_0[g_1 \otimes \cdots \otimes g_n] &= g_1 \cdot [g_2 \otimes \cdots \otimes g_n] \\ d_i[g_1 \otimes \cdots \otimes g_n] &= [g_1 \otimes \cdots \otimes g_i g_{i+1} \otimes \cdots \otimes g_n] \\ d_n[g_1 \otimes \cdots \otimes g_n] &= [g_1 \otimes \cdots \otimes g_{n-1}] \end{aligned}$$

and with augmentation $\varepsilon\colon B^u_0 \to \mathbb{Z}$, $\varepsilon[] = 1$.

The (unreduced) *bar resolution* is the quotient B of B^u by the subcomplex spanned by $[g_1 \otimes \cdots \otimes 1 \otimes \cdots \otimes g_n]$, where 1 denotes the unit of the group G. The classes are denoted $[g_1 \mid \cdots \mid g_n]$ and it is thus understood that thys symbol is zero when some $g_i = 1$.

In the formula for d_n, one could imagine the rhs multiplied from the right by g_n to get a more symmetrical version that will also be correct if we equip B^u with *trivial* right $\mathbb{Z}G$-module structure, see Hochschild (co)homology.

Theorem 11.8. *Both B^u and B are free resolutions of $\mathbb{Z} \in \mathsf{Mod}_{\mathbb{Z}G}$.*

Proof. The proof that $B^u \to \mathbb{Z}[0]$ is indeed an augmented chain complex was done in the tutorial. We will now show that the generators that contain 1 somewhere span a subcomplex. In the expression for

$$d[g_1 \otimes \cdots \otimes 1 \otimes \cdots \otimes g_n]$$

with 1 at position i, all terms contain this very same 1 except for the contributions d_{i-1} and d_i that give

$$(-1)^{i-1} \cdot [g_1 \otimes \cdots \otimes g_{i-1}1 \otimes \cdots \otimes g_n] + (-1)^i \cdot [g_1 \otimes \cdots \otimes 1g_i \otimes \cdots \otimes g_n] = 0$$

(the same generator with opposite signs).

We will now show that $B^u \to \mathbb{Z}[0]$ is a q-iso in $\mathsf{Ch}(\mathsf{Mod}_{\mathbb{Z}G})$ or equivalently in $\mathsf{Ch}(\mathsf{Ab})$, since the homology is computed the same way in $\mathsf{Mod}_{\mathbb{Z}G}$ and Ab. We will prove this by showing that the augmented chain complex is chain homotopy equivalent to the zero complex in $\mathsf{Ch}(\mathsf{Ab})$, i.e. that it admits a contraction $h\colon 0 \sim 1$, $dh + hd = 1$. Since the chain homotopy will only be $\mathbb{Z}$-linear, we define it on the $\mathbb{Z}$-generators by setting

$$h(g \cdot [g_1 \otimes \cdots \otimes g_n]) = [g \otimes g_1 \otimes \cdots \otimes g_n].$$

Easily $d_{i+1}h = hd_i$ and, thus, all terms in $dh + hd$ cancel out with the exception of $d_0h = 1$. We need to treat separately the cases involving the augmentation

$$(dh + h\varepsilon)(g \cdot []) = d[g] + h1 = g \cdot [] - [] + h1$$

so that we need to set $h1 = []$. Finally

$$(\varepsilon h + h0)1 = \varepsilon[] = 1.$$

The same formula works for the reduced version. □

Now we study examples. Obviously $H_0(G;\mathbb{Z}) = \mathbb{Z}_G = \mathbb{Z}$ and this corresponds to $H_0(BG;\mathbb{Z}) = \mathbb{Z}$ since BG is always connected. We proceed to H_1 so we write out explicitly the lower dimensions of the unreduced bar resolution

$$B^u = \cdots \to \mathbb{Z}G\{[g \otimes h]\} \to \mathbb{Z}G\{[g]\} \to \mathbb{Z}G\{[]\}.$$

The coinvariants $(-)_G = \mathbb{Z} \otimes_{\mathbb{Z}G} -$ then replace the free $\mathbb{Z}G$-modules by the corresponding free $\mathbb{Z}$-modules, i.e.

$$\mathbb{Z} \otimes_{\mathbb{Z}G} B^u = \cdots \to \mathbb{Z}\{[g \otimes h]\} \to \mathbb{Z}\{[g]\} \to \mathbb{Z}\{[]\}.$$

We will now compute the first homology of this complex, i.e. $H_1(G;\mathbb{Z})$. The differential in the original bar resolution takes $d[g] = g[] - []$ and after quotienting out the action this becomes zero. Going up by one dimension $d[g \otimes h] = g[h] - [gh] + [g]$ becomes on coinvariants

$$\cdots \longrightarrow \mathbb{Z}\{[g \otimes h]\} \longrightarrow \mathbb{Z}\{[g]\} \xrightarrow{0} \mathbb{Z}\{[]\}$$

$$[g \otimes h] \longmapsto [h] - [gh] + [g]$$

Altogether we get

$$H_1(G;\mathbb{Z}) = \mathbb{Z}\{[g]\}/([gh] \sim [g] + [h])$$

the free abelian group generated by the elements of the group with addition forced to equal the original group multiplication. This is easily seen to give the abelianization G_{ab} of G, e.g. by its universal property

$$\begin{array}{ccc} g & & G \longrightarrow A \\ \downarrow & & \downarrow \nearrow \\ [g] & & G_{\mathrm{ab}} \end{array}$$

This corresponds to the fact that for a path connected space X we have $H_1(X;\mathbb{Z}) = \pi_1(X)_{\mathrm{ab}}$ and $\pi_1(BG) = \pi_1(K(G,1)) = G$.

Exercise 11.9. Show that $H^1(G;M) = \mathrm{Der}(\mathbb{Z}G;M)/\mathrm{PDer}(\mathbb{Z}G;M)$ using the concrete description of the cochain complex $\mathrm{Hom}_{\mathbb{Z}G}(\mathbb{Z},M)$ below. Here a derivation of a ring G with coefficients in an R-R-bimodule M is a group homomorphism $D\colon R \to M$ satisfying the Leibniz rule

$$D(r \cdot s) = Dr \cdot s + r \cdot Ds.$$

A principal derivation is one of the form $D_x(r) = rx - xr$ for some $x \in M$. In the case $R = \mathbb{Z}G$ and a left $\mathbb{Z}G$-module M made into a right $\mathbb{Z}G$-module trivially (as above), the formulas become

$$D(g \cdot h) = Dg + g \cdot Dh, \quad D_x(g) = gx - x.$$

We will now study certain extensions of groups very closely related to $H^2(G;M)$. We will restrict to certain extensions (to be specified in a minute via a certain action)

$$1 \to M \xrightarrow{i} X \xrightarrow{p} G \to 1$$

where we write all groups multiplicatively and assume M commutative (later on, we will rewrite M additively, but at this point it would seem rather confusing). There is an action of X on M by conjugation (since it is the kernel of p and thus a normal subgroup:

$$X \to \mathrm{Aut}(M), \quad x \mapsto (m \mapsto xmx^{-1} = {}^x m)$$

and by commutativity, the restriction to M is trivial and thus this action factors through $X/M \cong G$ and we denote the action by the power on the left as above, i.e. ${}^a m$.

Definition 11.10. Let M be a $\mathbb{Z}G$-module. An extension

$$1 \to M \xrightarrow{i} X \xrightarrow{p} G \to 1$$

is to be understood as an extension of groups with M commutative and such that the given G-action agrees with the conjugation action coming from the extension.

Our aim will be to classify the extensions for a fixed $G \in \mathsf{Grp}$ and $M \in {}_{\mathbb{Z}G}\mathsf{Mod}$. We will now choose a *based section* of p, i.e. a mapping $\sigma\colon G \to X$ satisfying $p(\sigma(a)) = a$ and $p(1) = 1$ (i.e. thinking of $G = X/M$ it is a mapping that picks a representative in each class and picks $1 \in 1M$). Now we may rewrite the conjugation action as

$${}^a m = \sigma(a) \cdot m \cdot \sigma(a)^{-1}.$$

If σ happens to be a homomorphism then the extension is split and we will see that it is then isomorphic to the so called semidirect product $M \rtimes G$. We will now construct the so called factor set that is an obstruction to σ being a homomorphism:

$$[a, b] = \sigma(a) \cdot \sigma(b) \cdot \sigma(ab)^{-1}.$$

By the based property, we get $[a, 1] = 1 = [1, b]$ and we say again that the factor set is based.

We will now explain the importance of the factor set: Given an extension and a based section, the mapping

$$M \times G \to X, \quad (m, a) \mapsto m \cdot \sigma(a)$$

is a bijection with inverse $x \mapsto (x \cdot \sigma(p(x))^{-1}, p(x))$. We may thus transport the group structure from X to $M \times G$ and obtain an isomorphic group in a somewhat "canonical form" that will allow us to compare two extensions: We first compute the product of the images of (m, a) and (n, b) inside X:

$$m \cdot \sigma(a) \cdot n \cdot \sigma(b) = m \cdot {}^a n \cdot \sigma(a) \cdot \sigma(b) = \underbrace{m \cdot {}^a n \cdot [a, b]}_{\in M} \cdot \underbrace{\sigma(ab)}_{\in \sigma(G)}.$$

Thus this corresponds to the pair with components as indicated and the transported group structure is

$$(m, a) \cdot (n, b) = (m \cdot {}^a n \cdot [a, b], ab).$$

Assuming now that $\varphi \colon G \times G \to M$ is a based mapping, we observe that the identity for this product is always $(1, 1)$. We will now study when this product is associative (it will then have inverses as well), in which case we denote the resulting group $M \times_\varphi G$ with multiplication

$$(m, a) \cdot (n, b) = (m \cdot {}^a n \cdot \varphi(a, b), ab).$$

By construction, $M \times_{[-,-]} G$ is always a group.

Lemma 11.11. *The product is associative iff* ${}^a\varphi(b, c) \cdot \varphi(a, bc) = \varphi(a, b) \cdot \varphi(ab, c)$. *If this is the case, inverses exist.*

Proof. This is a straightforward computation:

$$\begin{aligned}
(m, a) \cdot ((n, b) \cdot (p, c)) &= (m, a) \cdot (n \cdot {}^b p \cdot \varphi(b, c), bc) \\
&= (m \cdot {}^a(n \cdot {}^b p \cdot \varphi(b, c)) \cdot \varphi(a, bc), abc) \\
&= (m \cdot {}^a n \cdot {}^{ab} p \cdot {}^a\varphi(b, c) \cdot \varphi(a, bc), abc) \\
((m, a) \cdot (n, b)) \cdot (p, c) &= (m \cdot {}^a n \cdot \varphi(a, b), ab) \cdot (p, c) \\
&= (m \cdot {}^a n \cdot \varphi(a, b) \cdot {}^{ab} p \cdot \varphi(ab, c), abc)
\end{aligned}$$

The first claim thus follows from commutativity of M. The second claim follows from the observation that the equation

$$(m, a) \cdot (n, b) = (m \cdot {}^a n \cdot \varphi(a, b), ab) = (1, 1)$$

with parameter (m, a) has a unique solution $b = a^{-1}$ and then one can solve for n from the first component, giving a right inverse. Symmetrically, the unique solution has $a = b^{-1}$ and one can solve for m, giving a left inverse. When the product is associative, these have to be equal. □

We will now show that the factor set $[-,-]\colon G\times G\to M$ can be interpreted as a 2-cocycle in the cochain complex $\mathrm{Hom}_{\mathbb{Z}G}(B,M)$. We write out the terms of this cochain complex in low dimensions

$$B=\cdots\to\mathbb{Z}G\{[a\mid b\mid c]\}\to\mathbb{Z}G\{[a\mid b]\}\to\mathbb{Z}G\{[a]\}\to\mathbb{Z}G\{[]\}$$
$$\mathrm{Hom}_{\mathbb{Z}G}(B,M)=\cdots\leftarrow\mathrm{Map}_{\mathrm{b}}(G^3,M)\leftarrow\mathrm{Map}_{\mathrm{b}}(G^2,M)\leftarrow\mathrm{Map}_{\mathrm{b}}(G^1,M)\leftarrow\mathrm{Map}_{\mathrm{b}}(G^0,M)$$

where $\mathrm{Map}_{\mathrm{b}}$ denotes the set of all "based" mappings, i.e. those whose value is 1 whenever one of the arguments is 1. Thus, the factor set is a 2-cochain. The differential of a 2-cochain is (remember that we write the group M multiplicatively and the action as a power)

$$(\delta\varphi)(a,b,c)={}^a\varphi(b,c)\cdot\varphi(ab,c)^{-1}\cdot\varphi(a,bc)\cdot\varphi(a,b)^{-1}.$$

The equation from the lemma claims exactly this. Any other based section differs by $\sigma'(a)=\beta(a)\cdot\sigma(a)$ for some based mapping $\beta\colon G\to M$ and we compute the corresponding factor set:

Lemma 11.12. *$[a,b]'=[a,b]\cdot(\delta\beta)(a,b)$, i.e. the two factor sets differ by a 2-coboundary.*

Proof. This is again a simple computation:

$$\begin{aligned}[a,b]'&=\sigma'(a)\sigma'(b)\sigma'(ab)^{-1}=\beta(a)\sigma(a)\beta(b)\sigma(b)\sigma(ab)^{-1}\beta(ab)^{-1}\\&=\beta(a)\cdot{}^a\beta(b)\cdot[a,b]\cdot\beta(ab)^{-1}\end{aligned}$$

with all factors in the commutative group M and with $(\delta\beta)(a,b)={}^a\beta(b)\cdot\beta(ab)^{-1}\cdot\beta(a)$. □

We may thus summarize this technical part by stating that there is a well defined 2-cohomology class associated with an extension, called the factor set $[-,-]\in H^2(G;M)$.

Theorem 11.13. *The mapping*

$$\{\text{extensions } 0\to M\to X\to A\to 0\}/\text{iso}\to H^2(G;M),$$

associating to an extension its factor set, is bijective.

Proof. We first show that the mapping is injective. Given two extensions X and X' with factor sets $[-,-]$ and $[-,-]'$ that are cohomologous, i.e. differ by a coboundary $\delta\beta$, one can change the based section of X by β to obtain a new based section with corresponding factor set $[-,-]'$. Now the construction above gives isomorphisms

$$X\cong M\times_{[-,-]'}G\cong X'.$$

To prove surjectivity, let φ be a 2-cocycle and consider the group $M\times_\varphi G$. If we equip it with the obvious based section $\sigma(a)=(1,a)$ then the corresponding factor set will be

$$\begin{aligned}[a,b]&=\sigma(a)\sigma(b)\sigma(ab)^{-1}\\&=(1,a)\cdot(1,b)\cdot(1,ab)^{-1}\\&=(\varphi(a,b),ab)\cdot(1,ab)^{-1}\\&=(\varphi(a,b),1)\cdot(1,ab)\cdot(1,ab)^{-1}\\&=(\varphi(a,b),1)\end{aligned}$$

that equals $\varphi(a,b)$ as an element of $M\subseteq M\times_\varphi G$, as required. □

Theorem 11.14. *Let G be a finite group of order k. Then the multiplication by k is zero on $H_n(G;M)$ and $H^n(G;M)$ for $n>0$, i.e. the order of any element divides k.*

Proof. We will show that the multiplication by k map on B is homotopic to the map that is zero in all dimensions except dimension 0 where it is multiplication by $N=\sum_{g\in G} g$.

$$\begin{array}{ccccc} B_2 & \longrightarrow & B_1 & \longrightarrow & B_0 \\ {\scriptstyle 0}\downarrow\uparrow{\scriptstyle k} & & {\scriptstyle 0}\downarrow\uparrow{\scriptstyle k} & & {\scriptstyle N}\downarrow\uparrow{\scriptstyle k} \\ B_2 & \longrightarrow & B_1 & \longrightarrow & B_0 \end{array}$$

We define

$$h[g_1 \mid \cdots \mid g_n] = (-1)^{n+1} \cdot \sum_{g\in G} [g_1 \mid \cdots \mid g_n \mid g].$$

Clearly $d_i h = -h d_i$ (thanks to the above alternating sign) so that everything in $dh+hd$ cancels out except

$$d_{n+1} h[g_1 \mid \cdots \mid g_n] = \sum_{g\in G} [g_1 \mid \cdots \mid g_n] = k \cdot [g_1 \mid \cdots \mid g_n]$$

so that $dh+hd=k$ as claimed except in dimension 0 where

$$(dh+hd)[] = d(-\sum_{g\in G}[g]) = \sum_{g\in G} [] - g[] = k[] - N[].$$

Now apply either $M\otimes_{\mathbb{Z}G} -$ or $\mathrm{Hom}_{\mathbb{Z}G}(-,M)$ to obtain a chain homotopy between the corresponding maps on the resulting chain complexes. In (co)homology the maps become equal. □

Corollary 11.15. *Let G and M be finite with $\gcd(|G|,|M|)=1$. Then $H_n(G;M)=0$ and $H^n(G;M)=0$ for $n>0$. Consequently, any extension $0\to M\to X\to G\to 0$ splits, i.e. is isomorphic to the semidirect product $M\rtimes G$.*

Proof. The multiplication by $k=|G|$ is both zero by the previous theorem and an isomorphism, since it is induced by an isomorphism $M \xrightarrow{k} M$ (let $l=|M|$ and $ak+bl=1$; then the inverse is clearly the multiplication by a). □

Remark. This is a generalization of Example 9.3 to the case of nonabelian G. The theorem holds even for nonabelian M and is proved from the above abelian case by "group theoretic induction" (descreasing order be quotienting out the centre, I think).

12. Flatness is stalkwise

We use this opportunity to talk about various special instances of flat modules. The main goal is to prove the theorem

Theorem 12.1. *Let R be a commutative ring. An R-module A is flat iff for every maximal ideal $P\subseteq R$ the localization A_P is a flat R_P-module.*

The main ingredient of the proof is the so called flat base change for Tor. Let S be an R-algerba that is flat as an R-module. Then for an R-module A and S-module B we have

$$\mathrm{Tor}_n^R(A,B) \cong \mathrm{Tor}_n^S(A \otimes_R S, B)$$

Proof of the claim. Consider a projective resolution $P \to A[0]$. Extending the scalars via the exact functor $S \otimes_R -$ we thus obtain a resolution

$$P \otimes_R S \to A \otimes_R S[0]$$

that is easily seen to be projective again (the extension takes $R \mapsto S$, thus free to free and thus projective to projective). We may thus use this to compute

$$\mathrm{Tor}_n^S(A \otimes_R S, B) = H_n(P \otimes_R S \otimes_S B) = H_n(P \otimes_R B) = \mathrm{Tor}_n^R(A,B). \qquad \square$$

Now we are ready to prove the theorem.

Proof of the theorem. Assuming A flat, we get

$$\mathrm{Tor}_n^{R_P}(A_P, B) = \mathrm{Tor}_n^{R_P}(A \otimes_R R_P, B) = \mathrm{Tor}_n^R(A,B) = 0$$

and A_P is flat.

In the opposite direction, assuming A_P flat over R_P, we need to show $\mathrm{Tor}_n^R(A,B) = 0$ and this is equivalent to $\mathrm{Tor}_n^R(A,B)_P = 0$ for all P maximal. Since Tor is some homology group and localization is exact, we may view this as

$$H_n(Q \otimes_R B)_P \cong H_n((Q \otimes_R B)_P) \cong H_n(Q_P \otimes_{R_P} B_P) \cong \mathrm{Tor}_n^{R_P}(A_P, B_P) = 0. \qquad \square$$

We will now show that over local rings, flat modules are very close to free modules. First a general result.

Definition 12.2. An R-module A is *finitely presentable* if there exists a ses

$$R^s \to R^t \to A \to 0$$

for some finite s and t.

Exercise 12.3. Show that for a finitely presentable A and any ses

$$0 \to K \to L \to A \to 0$$

with L finitely generated, also K is finitely generated.

Consider the following map

$$B \otimes_R A^* \cong \mathrm{Hom}_R(R,B) \otimes_R \mathrm{Hom}_R(A,R) \xrightarrow{\circ} \mathrm{Hom}_R(A,B).$$

Proposition 12.4. *If A is finitely presentable and B is flat then this map is an isomorphism.*

Proof. We proved this at the tutorial. The idea is to prove this for $A = R$, then for finite coproducts, then for cokernels (using B flat). $\square$

Corollary 12.5. *If A is finitely presentable and flat, it is projective.*

Proof. In the map $A \otimes A^* \to \mathrm{Hom}_R(A, A)$ the simple tensor $a_i \otimes \eta^i$ maps to the composition $A \xrightarrow{\eta^i} R \xrightarrow{a^i} A$ and a finite sum of such clearly maps to

$$A \xrightarrow{\begin{pmatrix}\eta^1\\ \vdots \\ \eta^n\end{pmatrix}} R^n \xrightarrow{\begin{pmatrix}a_1 & \cdots & a_n\end{pmatrix}} A.$$

If this happens to be the preimage of the identity then A is a direct summand of R^n and is thus projective. □

Remark. This implies that over a noetherian ring, $\mathrm{gl.dim}(R) = \mathrm{Tor.dim}(R)$: This is because we can translate this claim to equality in

$$\sup\{\mathrm{pd}(R/J) \mid J \subseteq R\} \geq \sup\{\mathrm{fd}(R/J) \mid J \subseteq R\} = d.$$

Now over a noetherian ring the cyclic modules R/J admit a resolution by f.g. free modules, so we consider an exact sequence of f.g. modules

$$0 \to M_d \to P_{d-1} \to \cdots \to P_0 \to R/J \to 0$$

and conclude that M_d is flat; as it is also f.p., it must be projective and $\mathrm{pd}(R/J) \leq d$, as claimed.

Over a local ring, projective modules are exactly free modules (Kaplansky theorem). We will prove a simpler version.

Theorem 12.6. *A finitely generated projective module over a commutative local ring is free.*

Proof. Let A be a f.g. projective module. Then $A/MA \cong R/M \otimes_R A$ is a f.d. vector space over the residue field R/M. Let $a_1, \ldots, a_n \in A$ be such that their images in A/MA form a basis. These give an R-linear map

$$p\colon R^n \to A,\ e_i \mapsto a_i.$$

Since A is projective, the s.e.s.

$$0 \to \ker p \to R^n \to A \to 0$$

splits and is thus preserved by $R/M \otimes_R -$. Since p yields an iso by assumption, we must have $\ker p/M \ker p = 0$, i.e. $\ker p = M \ker p$ and Nakayama lemma yields $\ker p$. □

13. Simplicial resolutions

Let $\mathcal{C}$ be a category. A *monad* on $\mathcal{C}$ is a monoid in the strictly monoidal category $([\mathcal{C}, \mathcal{C}], \circ)$ of endofunctors of $\mathcal{C}$. Thus, it is an endofunctor T equipped with two natural transformations

$$\mu\colon T \circ T \to T, \quad \eta\colon 1 \to T,$$

the multiplication and the unit, satisfying the associativity and unitality axioms

$$\mu \circ (\mu \circ 1) = \mu \circ (1 \circ \mu), \quad \mu \circ (1 \circ \eta) = 1 = \mu \circ (\eta \circ 1).$$

More concisely, one requires natural transformations $\mu_k \colon T^k \to T$ that are closed under compositions. This means that $1 = \mu_1$ (as a composition of zero μ_k's) and e.g.

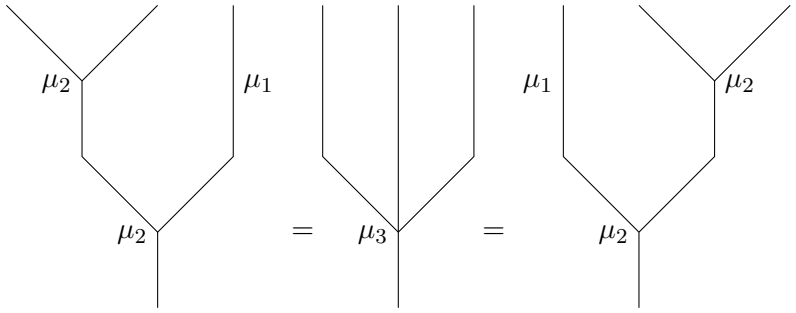

and similarly for the unary

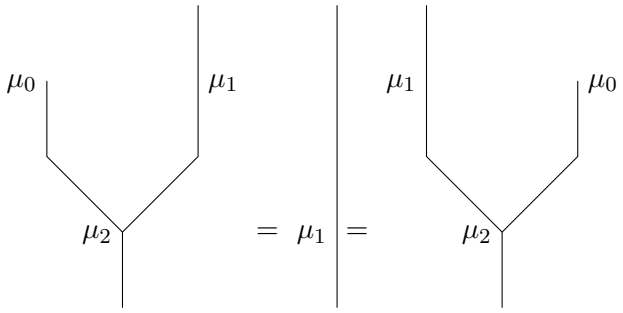

There is a universal strictly monoidal category with a monoid. We will give a concrete description and then, instead of showing the universal property, give the unique instance of it that we are interested in, i.e. to monads. It is the category Δ of all finite ordinals (topologists would only consider non-empty ordinals; this would correspond to a non-unital version)

$$[n] = \{0 < 1 < \cdots < n\}$$

with $[0] = \{0\}$ and $[-1] = \emptyset$. Morphisms in Δ are the order preserving maps and the monoidal product is the "join" $[m] * [n] = [m + 1 + n]$ or more intuitively

$$\{0 < \cdots < m\} * \{0 < \cdots < n\} = \{\underbrace{0 < \cdots m}_{0<\cdots<m} < \underbrace{m + 1 < \cdots < m + 1 + n}_{0<\cdots<n}\}$$

Geometrically, this is the join construction for simplices. The monoidal unit is clearly $[-1]$ and $[0]$ is a monoid with $\mu_n \colon [n - 1] \to [0]$ the unique map. We will now outline why Δ is a universal strictly monoidal category with a monoid: any map in Δ can be decomposed canonically as

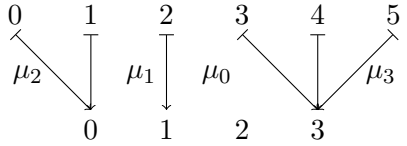

a join of the μ_k's.

Proposition 13.1. *For any monad T, there is a unique stricti monoidal functor $\Delta \to [\mathcal{C}, \mathcal{C}]$ sending $[0]$ to T.*

Proof. We must send $[n] \mapsto T^{n+1}$ and a map decomposed, as above, into a join of μ_k's to the composition of the corresponding transformations $\mu_k \colon T^k \to T$. □

Example 13.2. Any adjunction $F \colon \mathcal{C} \underset{\longleftarrow}{\overset{\longrightarrow}{\perp}} \mathcal{D} : G$ induces a monad on $\mathcal{C}$ with $T = GF$ and $\eta \colon 1 \to GF$ the unit and $\mu = G\varepsilon F \colon GFGF \to GF$ the multiplication.

Dually, it gives a comonad on $\mathcal{D}$ with $\perp = FG$ and $\varepsilon \colon FG \to 1$ the counit and $\delta = F\eta G \colon FG \to FGFG$ the comultiplication.

Now dually to the above proposition, a comonad gives a strict monoidal functor $\Delta^{\mathrm{op}} \to [\mathcal{D}, \mathcal{D}]$, sending $[0]$ to $\perp$. By composing with the evaluation at $A \in \mathcal{D}$, this gives a functor $\Delta^{\mathrm{op}} \to \mathcal{D}$, $[n] \mapsto \perp^{n+1} A$; functors of this shape are called *(augmented) simplicial objects* in $\mathcal{D}$.

Example 13.3. There is an adjunction $F \colon \mathsf{Ab} \underset{\longleftarrow}{\overset{\longrightarrow}{\perp}} {}_{\mathbb{Z}G}\mathsf{Mod} : U$, where U is the forgetful functor and U is the extension of scalars, $F(A) = \mathbb{Z}G \otimes A$. The induced comonad on ${}_{\mathbb{Z}G}\mathsf{Mod}$ is then again $\perp A = \mathbb{Z}G \otimes A$ with counit $\perp A \to A$, $r \otimes a \mapsto ra$ the $\mathbb{Z}G$-multiplication in A and comultiplication $\perp A \to \perp^2 A$, $r \otimes a \mapsto r \otimes 1 \otimes a$. The induced augmented simplicial object for $A \in {}_{\mathbb{Z}G}\mathsf{Mod}$ looks

$$\cdots \underset{\longleftarrow}{\overset{\longrightarrow}{\overset{d_0}{\underset{d_2}{\overset{\longleftarrow}{\underset{\longleftarrow}{\overset{\longrightarrow}{\underset{d_1}{\longrightarrow}}}}}}}} \mathbb{Z}G \otimes \mathbb{Z}G \otimes A \underset{\longleftarrow}{\overset{\longrightarrow}{\overset{d_0}{\underset{d_1}{\longrightarrow}}}} \mathbb{Z}G \otimes A \xrightarrow{d_0} A$$

(with the right most d_0 the augmentation). The maps d_i are of the form $1 * \cdots * 1 * \varepsilon * 1 * \cdots * 1$, i.e. they are all induced by the counit and are all given by mulitiplication of a pair of neighbours in the tensor product. The unnamed maps are the so called degeneracy maps and are induced by the comultiplication (i.e. 1 is inserted at various points).

In general, an (augmented) simplicial object $X_n = X[n]$ in an abelian category, such as the category ${}_{\mathbb{Z}G}\mathsf{Mod}$ above, gives an (augmented) chain complex, the so called Moore chain complex of the simplicial object:

$$\cdots \longrightarrow X_2 \xrightarrow{d} X_1 \xrightarrow{d} X_0 \xrightarrow{d} X_{-1}$$

with the last $d \colon X_0 \to X_{-1}$ the augmentation and with all $d = \sum (-1)^i d_i$.

In this way, applying the general machinery to $\mathbb{Z} \in {}_{\mathbb{Z}G}\mathsf{Mod}$ produces the standard bar resolution.

Example 13.4 (towards Hochschild (co)homology)**.** Let R be a non-commutative k-algebra over a commutative ring k, typically a field. We then get an adjunction $F \colon {}_k\mathsf{Mod} \underset{\longleftarrow}{\overset{\longrightarrow}{\perp}} {}_R\mathsf{Mod} : U$ with U the forgetful functor and F the extension of scalars $FA = R \otimes A$ where all the tensor products will be taken over the ground ring k. The general machinery, applied to the R-module R gives the bar resolution B_R as follows

$$\cdots \longrightarrow R \otimes R \otimes R \otimes R \xrightarrow{d} R \otimes R \otimes R \xrightarrow{d} R \otimes R \underbrace{\xrightarrow{d} R}_{\text{augm}}.$$

Here again $d = \sum (-1)^i d_i$ and

$$d_i(r_0 \otimes \cdots \otimes r_n) = r_0 \otimes \cdots \otimes r_i r_{i+1} \otimes \cdots \otimes r_n.$$

Hochschild cohomology of R with coefficients in an R-R-bimodule A is

$$H^n(R;A) = H^n(\mathrm{Hom}_{R-R}(B_R, A)).$$

Dually the Hochschild homology is

$$H_n(R;A) = H^n(B_R \otimes_{R-R} A)$$

but some care has to be taken with the tensor product (it coequalizes right action and a left action but also the other way around, i.e. $xr \otimes a = x \otimes ra$ but also $rx \otimes a = x \otimes ar$).

One can show that again $H^1(R;A) = \mathrm{Der}(R;A)/\mathrm{PDer}(R;A)$ and that $H^2(R;A)$ corresponds to the so called square zero extensions, i.e. $A \subseteq X$ is a square zero ideal $A^2 = 0$ such that $X/A \cong R$.

At the tutorial we discussed operations on the Hochschild cohomology $H^*(R;R)$ and Deligne conjecture.

14. Representation theory

The lectures were following the text by John Bourke, but simplified some parts considerably. We will thus give an exposition that concetrates on these parts were the lectures departed from the text.

We will concentrate on representations of *finite* groups, so all our groups will be assumed to be finite.

Definition 14.1. A representation of a group G over a field $\Bbbk$ is a $\Bbbk G$-module V.

Equivalently, this is a ring homomorphism $\Bbbk G \to \mathrm{End}_{\mathbb{Z}}(V)$. By restriction of scalars along the inclusion $\Bbbk \subseteq \Bbbk G$, the abelian group V becomes a vector space over $\Bbbk$. Since the elements of the field and elements of the group G commute inside $\Bbbk G$, we obtain a group homomorphism

$$\begin{array}{ccc} G & \longrightarrow & \mathrm{End}_{\Bbbk}(V) \\ & \searrow & \uparrow \\ & & \mathrm{Aut}_{\Bbbk}(V) = \mathrm{GL}(V) \end{array}$$

Thus, equivalently a representation is a vector space V over $\Bbbk$ together with a homomorphism of groups $G \to \mathrm{GL}(V)$.

Example 14.2. Dihedral group D_8 with 8 elements, i.e. the group of symmetries of a square, has a canonical action on $\mathbb{R}^2$ (via these symmetries).

As we explained above, $\Bbbk G$ is a Hopf algebra with comultiplication induced by the diagonal

$$\delta\colon \Bbbk G \to \Bbbk(G \times G) \cong \Bbbk G \otimes \Bbbk G$$

and counit by the constant map

$$\Bbbk G \to \Bbbk * \cong \Bbbk$$

This allows us to make a tensor product of two representations into a representation:

$$a \cdot (v \otimes w) = av \cdot aw$$

for $a \in G$, but not for more general elements of $\Bbbk G$. Formally, the multiplication is

$$\Bbbk G \otimes V \otimes W \xrightarrow{\delta\otimes 1\otimes 1} \Bbbk G \otimes \Bbbk G \otimes V \otimes W \xrightarrow{1\otimes\rho\otimes 1} \Bbbk G \otimes V \otimes \Bbbk G \otimes W \xrightarrow{\mu\otimes\mu} V \otimes W.$$

We stress that the tensor product here is over $\Bbbk$ and not over $\Bbbk G$.

Most importantly, this monoidal structure is closed, i.e. there exists an internal hom that we constructed at the tutorial

$$\frac{U \otimes V \to W}{U \to [V, W]}$$

On the level of vector spaces, we must have $[V, W] = \mathrm{Hom}_{\Bbbk}(V, W)$ and it remains to come up with a G-action so that the bottom map is $\Bbbk G$-linear iff the top map is. We concluded

$${}^g\varphi = g\varphi g^{-1}, \quad {}^g\varphi(v) = g \cdot \varphi(g^{-1} \cdot v).$$

In particular, a fixed point for the action is φ such that ${}^g\varphi = \varphi$ or, equivalently, $g\varphi = \varphi g$, i.e. φ is $\Bbbk G$-linear (this corresponds to the fact that the unit is $\Bbbk$ and maps from the unit are exactly the fixed points).

$$[V, W]^G = \mathrm{Hom}_{\Bbbk G}(V, W).$$

We will now present an important tool – the projection $\pi\colon U \to U^G$ onto the fixed points. Here we have to assume that $\operatorname{char}\Bbbk \nmid |G|$, typically $\operatorname{char}\Bbbk = 0$.

$$\pi(u) = \frac{1}{|G|} \cdot \sum_{g\in G} gu.$$

This has two properties: $\mathrm{im}(\pi) \subseteq U^G$ and $\pi|_{U^G} = \mathrm{id}$, both easily verified.

Definition 14.3. A $\Bbbk G$-module U is said to be *irreducible* (simple) if its only quotients (equivalently submodules) are U and 0 (equivalently 0 and U). In other words there exist only two extensions

$$0 \longrightarrow 0 \longrightarrow U \longrightarrow U \longrightarrow 0$$

$$0 \longrightarrow U \longrightarrow U \longrightarrow 0 \longrightarrow 0$$

The module 0 is not considered irreducible.

Definition 14.4. A $\Bbbk G$-module U is said to be *indecomposable* if $U = V \oplus W$ only for $V = 0$, $W = U$ and $V = U$, $W = 0$. In other words, the only split extensions are as above.

Obviously every irreducible module is indecomposable.

Theorem 14.5 (Maschke)**.** *If* $\operatorname{char}\Bbbk \nmid |G|$ *then every short exact sequence in* ${}_{\Bbbk G}\mathsf{Mod}$ *splits. Consequently, every indecomposable* $\Bbbk G$*-module is irreducible.*

Proof. Let $V \subseteq U$ be a submodule. Since the inclusion splits over $\Bbbk$, we get a projection $p\colon U \to V$. So $p \in [U, V]$ and we may apply the projection $\pi\colon [U, V] \to [U, V]^G = \mathrm{Hom}_{\Bbbk G}(U, V)$ to it to obtain

$$\pi(p)(u) = \frac{1}{|G|} \cdot \sum_{g\in G} gp(g^{-1}u).$$

It remains to check that it is still a projection onto V, i.e. that $\pi(p)(v) = v$ for $v \in V$. But in this case $g^{-1}v \in V$ as well and in the formula above we may ignore p (being identity on V). □

In the tutorial, we showed that even if $\operatorname{char}\Bbbk \mid |G|$, the module $\Bbbk G$ is still injective (the group algebra $\Bbbk G$ is self-injective). Together with $\Bbbk G$ being noetherian, it follows that projective and injective modules coincide (we only proved $\Rightarrow$).

Corollary 14.6. *Every (finite dimensional) representation splits into a direct sum of irreducible representations.*

Proof. By induction on $\dim_\Bbbk U$. □

Remark (Jordan-Hölder theorem). A composition series for U is a finite filtration

$$0 \subseteq U_1 \subseteq \cdots \subseteq U_n = U$$

with filtration quotients $Q_i = U_i/U_{i-1}$ irreducible. In our case $U \cong \bigoplus U_i$. The theorem says that the collection of the Q_i's is independent of the fultration. So consider

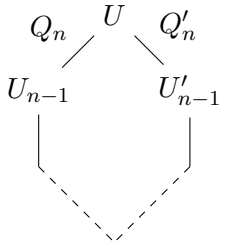

Now consider the sum $U_{n-1}+U'_{n-1}$. Since this lies between U_{n-1} and U and the quotient is the irreducible module Q_n, this must equal either U_{n-1} or U and similarly for U'_{n-1}. Out of the four possibilities, only two make sense. One possibility is that $U_{n-1} = U_{n-1} + U'_{n-1} = U'_{n-1}$ in which case we may apply induction on this common submodule. The other possibility is that $U_{n-1} + U'_{n-1} = U$ and we get

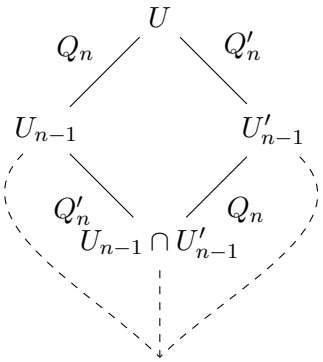

with filtration quotients equal on the opposite sides since the square is a pushout. Now apply induction to the smaller modules U_{n-1} and U'_{n-1}: The left most path has filtration quotients Q_n and those of U_{n-1}, i.e. Q_n and Q'_n and the filtration quotients of $U_{n-1} \cap U'_{n-1}$. Symmetrically, the same is true for the right path and thus these are equal. (In more general contexts, one has to prove alongside that $U_{n-1} \cap U'_{n-1}$ admits a composition series, for finite dimensional representations this is clear.)

Corollary 14.7. *The decomposition of a finite dimensional representation into a direct sum of irreducible representations is unique up to the order of submodules.* □

Corollary 14.8. *Let $\Bbbk G \cong U_1 \oplus \cdots \oplus U_n$ be a decomposition of $\Bbbk G$ into a direct sum of irreducible representations. Then any irreducible representation is isomorphic to one of the U_i.*

Proof. Let U be an irreducible representation. Any $0 \neq x \in U$ gives a $\Bbbk G$-linear map

$$\Bbbk G \to U,\ 1 \mapsto x$$

whose image must be equal to U by irreducibility. This epi splits by Maschke theorem, so

$$\Bbbk G \cong U \oplus V \cong U \oplus V_1 \oplus \cdots \oplus V_k.$$

By uniqueness, U must be one of the U_i's. □

Theorem 14.9 (Schur)**.** *Let $\varphi\colon U \to V$ be a $\Bbbk G$-linear map between irreducible $\Bbbk G$-modules. Then either $\varphi = 0$ or φ is an isomorphism.*

If $\Bbbk$ is algebraically closed then any $\varphi\colon U \to U$ is a multiplication by some $\lambda \in \Bbbk$, i.e. $\varphi(u) = \lambda u$, i.e.

$$\operatorname{Hom}_{\Bbbk G}(U, U) = \Bbbk.$$

Proof. Since $\ker\varphi \subseteq U$ is a submodule, either φ is mono or zero. Similarly $\operatorname{im}\varphi \subseteq V$ is a submodule, so either φ is zero or epi.

The second part is similar: $\varphi\colon U \to U$ is $\Bbbk$-linear so has some eigenvalue λ. Then $\ker(\varphi - \lambda \cdot 1) \subseteq U$ is a nonzero submodule, so this eigenspace equals U. □

Thus, for $U \ncong V$ we have $\operatorname{Hom}_{\Bbbk G}(U, V) = 0$ and for $U \cong V$ we have $\operatorname{Hom}_{\Bbbk G}(U, V) \cong \Bbbk$ (noncanonically; in this respect Schur lemma is better).

Theorem 14.10. *Assume $\Bbbk$ algebraically closed. The number of times U appears in the direct sum decomposition $\Bbbk G \cong U_1 \oplus \cdots \oplus U_n$ equals $\dim U$.*

Proof. This follows from the computation

$$U \cong \operatorname{Hom}_{\Bbbk G}(\Bbbk G, U) \cong \operatorname{Hom}_{\Bbbk G}(U_1 \oplus \cdots \oplus U_n, U) \cong \bigoplus_i \operatorname{Hom}_{\Bbbk G}(U_i, U) \cong \bigoplus_{i \text{ s.t. } U_i \cong U} \Bbbk. \qquad \square$$

Let $V_1, \ldots, V_r$ be a complete set of irreducible representations, i.e. containing a single representative of each isomorphism class of irreducible representations.

Corollary 14.11. $|G| = \sum_i \dim(V_i)^2$.

Proof. $|G| = \dim \Bbbk G = \sum_i \dim U_i = \sum_i \dim V_i \cdot \dim V_i$. □

Proposition 14.12. *G is abelian iff all its complex irreducible representations are one-dimensional.*

Proof. Assume G abelian. Then $\Bbbk G$ is commutative, so $g\cdot\colon U \to U$ is $\Bbbk G$-linear. For U irreducible, it must be multiplication by some $\lambda \in \Bbbk$. Since this holds for any g, all subspaces of U are $\Bbbk G$-submodules and U must be one-dimensional.

If every irreducible representation is one-dimensional then the action

$$G \to \mathrm{GL}(V) \cong \Bbbk^\times$$

lands in a commutative group so the multiplication by g and by h on V commute. Thus, the same is true in a direct sum of irreducible representations, i.e. in any representation and, in particular, in $\Bbbk G$. Thus $g \cdot h \cdot 1 = h \cdot g \cdot 1$. □

Example 14.13. We studied the two-dimensional representation of D_8 and we tried to show that it is irreducible over $\mathbb{C}$. The reflection across some line ℓ has invariant subspaces 0, ℓ, $\ell^\perp$ and $\mathbb{C}^2$. Since D_8 contains reflections across the lines $x = 0$ and $x = y$, the only common invariant subspaces are 0 and $\mathbb{C}^2$. Thus, we have

$$8 = |G| = 2^2 + 1^2 + 1^2 + 1^2 + 1^2$$

since there always exists a trivial one-dimensional representation. In the tutorial, we described all four one-dimensional representations.

15. Characters of groups

Definition 15.1. Let U be a representation. The function

$$\chi = \chi_U \colon G \to \Bbbk,\ g \mapsto \operatorname{tr}(g\times \colon U \to U)$$

is called a *character* of G (associated to the representation U). It is said to be an *irreducible character* if U is irreducible.

The basic property is that isomorphic representations give equal characters and that $\chi(gh) = \chi(hg)$ or $\chi(ghg^{-1}) = \chi(h)$. A function $G \to \Bbbk$ is a *class function* if it is constant along each conjugacy class. We write $C(G)$ for vector space of class functions. We thus have $\chi \in C(G)$.

Lemma 15.2. $\dim C(G) = |G/\mathrm{conj}|$.

Proof. This is rather obvious since $C(G) = \Bbbk^{G/\mathrm{conj}}$. □

We will now restrict to $\Bbbk = \mathbb{C}$. Our goal now will be to show that the irreducible characters form an orthonormal basis of $C(G)$, for which we have to introduce an inner product on $C(G)$. We could do so right now, but we will get to the definition naturally by studying characters of induced representations.

- $\chi_{U\oplus V} = \chi_U + \chi_V$. □
- $\chi_{U\otimes V} = \chi_U \cdot \chi_V$.
  This follow from writing $g \cdot e_j = \sum_i a_j^i e_i$ and $g \cdot \overline{e_l} = \sum_k b_l^k \overline{e_k}$ so that

  $$g \cdot (e_j \otimes e_l) = ge_j \otimes ge_l = \sum_i a_j^i e_i \otimes \sum_k b_l^k \overline{e_k} = \sum_{i,k} a_j^i b_l^k e_i \otimes \overline{e_k}$$

  and the sum across the diagonal equals

  $$\sum_{i,k} a_i^i b_k^k = \sum_i a_i^i \cdot \sum_k b_k^k.$$

  A more conceptual proof uses string diagrams and the corresponding definition of trace (equivalently the contraction of $\varphi \in T_1^1$).
- $\chi_{U^*} = \chi_{\bar{U}} = \overline{\chi_U}$.
  We have shown in the tutorial that $U^* \cong \bar{U}$ as representations, where $U^* = [U, \Bbbk]$ is the dual vector space with action $g \cdot \eta = \eta g^{-1}$ and $\bar{U}$ is U with complex multiplication $z * u = \bar{z} \cdot u$. The operator $g\times$ remains the same but its matrix in $\bar{U}$ is complex conjugate of the matrix in U.

- $\chi_{[U,V]} = \chi_{V \otimes U^*} = \chi_V \cdot \overline{\chi_U}$.
- $\dim U^G = \frac{1}{|G|} \cdot \sum_{g \in G} \chi_U(g)$.

  The trace of every projection equals the dimension of its image (just write the matrix of the projection in a basis formed by vectors from the image and vectors from the kernel). Applying this to the projection $\pi\colon U \to U$ with image U^G gives the left hand side. The right hand side is obtained from the concrete formula for π.

The last two points then give the following theorem.

Theorem 15.3. $\dim \mathrm{Hom}_{\Bbbk G}(U, V) = \dim [U, V]^G = \frac{1}{|G|} \cdot \sum_{g \in G} \chi_V(g)\overline{\chi_U(g)}$. □

For class functions $f_1, f_2 \in C(G)$ we define their inner product

$$\langle f_1, f_2 \rangle = \frac{1}{|G|} \cdot \sum_{g \in G} f_1(g) \cdot \overline{f_2(g)}.$$

Up to the factor $1/|G|$, this is the standard inner product (so in particular it indeed is an inner product).

Corollary 15.4. *For irreducible representations U, V we have*

$$\langle \chi_V, \chi_U \rangle = \begin{cases} 1 & \text{if } U \cong V \\ 0 & \text{if } U \ncong V \end{cases}$$

so that irreducible characters form an orthonormal system in $C(G)$. □

Corollary 15.5. *Two finite dimensional representations U, V are isomorphic iff $\chi_U = \chi_V$.*

Proof. Decomposing both into a direct sum of irreducible representations, with a_i, b_i the multiplicities of the irreducible V_i, we get $\langle \chi_U, \chi_{V_i} \rangle = \langle \sum_j a_j \chi_{V_j} \rangle = a_i$ and similarly for V. Assuming that $\chi_U = \chi_V$ thus gives $a_i = b_i$ and thus $U \cong V$. □

Corollary 15.6. *A representation U is irreducible iff $\langle \chi_U, \chi_U \rangle = 1$.*

Proof. Continuing the notation from the last proof, $\langle \chi_U, \chi_U \rangle = \sum a_i^2$. □

Example 15.7. We computed the character of the two-dimensional representation of D_8. All the reflections have eigenvalues 1 and -1 and thus have zero trace. Among the four rotations, two have zero trace, identity has trace 2 and the rotation by 180° has trace -2. Thus,

$$\langle \chi_{\mathbb{C}^2}, \chi_{\mathbb{C}^2} \rangle = \frac{1}{8} \cdot \sum_{g \in D_8} |\chi_{\mathbb{C}^2}(g)|^2 = 1$$

and we have another proof that the representation is irreducible.

Going back to the proof that irreducible characters form a basis of $C(G)$, it remains to compute the number of irreducible representations.

Lemma 15.8. *The number of irreducible representations is at least* $\dim Z(\mathbb{C}G)$.

Proof. Decompose $\mathbb{C}G$ into a direct sum of subrepresentations

$$\mathbb{C}G = W_1 \oplus \cdots \oplus W_r,$$

where each W_i is a direct sum of the $\dim V_i$ copies of the irreducible representation V_i, as in Theorem 14.10. Now express

$$1 = e_1 + \cdots + e_r$$

with each $e_i \in W_i$. We will now show that the centre $Z(\mathbb{C}G) \subseteq [e_1, \ldots, e_r]$. Let $z \in Z(\mathbb{C}G)$. Then multiplication by z is $\mathbb{C}G$-linear on every representation. By Schur's lemma, it coincides with multiplication by some $\lambda_i \in \mathbb{C}$ on V_i and thus also on W_i. Therefore

$$z = z \cdot 1 = z \cdot (e_1 + \cdots + e_r) = \lambda_1 e_1 + \cdots + \lambda_r e_r \in [e_1, \ldots, e_r]. \qquad \square$$

Lemma 15.9. $\dim Z(\mathbb{C}G) = |G/\text{conj}|$.

Proof. Consider $r = \sum a_g \cdot g \in Z(\mathbb{C}G)$. Then $hrh^{-1} = r$ and clearly the left hand side contains g with coefficient $a_{h^{-1}gh}$ since the corresponding term of r is $a_{h^{-1}gh} \cdot h^{-1}gh$ and gets conjugated to $a_{h^{-1}gh}g$. The equality then gives $a_{h^{-1}gh} = a_g$ and the coefficients are constant across each conjugacy class. Denoting the conjugacy classes by C_i, and the sum of elements within the conjugacy class by $\bar{C}_i$, which easily lies in the centre, we may write

$$r = \sum_i a_i \bar{C}_i$$

where $a_i = a_g$ for any $g \in C_i$. We have just shown that the $\bar{C}_i$ generate $Z(\mathbb{C}G)$ and clearly they are linearly independent. $\square$

Putting everything together, we see that the number of irreducible representations $\geq$ the number of conjugacy classes, i.e. $\dim C(G)$. Since the irreducible characters form an orthonormal, hence linearly independent, system in $C(G)$ we must have equality and they must generate $C(G)$. We have thus proved:

Theorem 15.10. *The irreducible characters form an orthonormal basis of the space $C(G)$ of class functions.* $\square$

16. Representations of symmetry groups S_n

This was rather informative and I followed very closely John's notes.

17. Integrally closed rings, valuation rings, Dedekind domains

In this section all rings will be commutative with 1 as usual, but additionally also domains.

The motivation for Dedekind domains is the existence and uniquness of factorization of ideals into a product of prime ideals. This clearly holds for PID's since this is then just the UFD property. However, many important examples are Dedekind domains but not PID's. For example, the coordinate ring $\Bbbk[V]$ of an irreducible smooth curve over an algebraically closed field (or in fact any field, I think) is such an example. The smoothness is a local property and we will introduce and study these rings in terms of their localizations at (nonzero) prime ideals.

Dedekind domains will be rings whose localizations at nonzero primes are discrete valuation rings; these are closely related to integrally closed rings so we start with them.

Definition 17.1. Let R be a domain and let K be its fraction field. We say that an element of K is integral over R if it is a root of a monic polynomial from $R[x]$.

We say that R is integrally closed (or normal) if every element of K that is integral over R lies in R. (One may prove that the collection of integral elements forms an intermideate ring $R \subseteq \overline{R} \subseteq K$ and the condition says $R = \overline{R}$.)

Proposition 17.2. *Every UFD is integrally closed.*

Proof. Let $a/b \in K$ be integral over R and we may assume that a and b are coprime. Then it satisfies

$$(a/b)^n + r_{n-1}(a/b)^{n-1} + \cdots + r_0 = 0.$$

Clearing the denominators and expressing a^n from this we get

$$a^n = -b \cdot (r_{n-1}a^{n-1} + \cdot + r_0 b^{n-1}).$$

This means that $b \mid a^n$ and by coprimality we get that b is a unit, so that $a/b \in R$. □

We will now show that the property of being integrally closed is local, or in fact stalkwise (here R_P are the stalks of the affine scheme $\operatorname{Spec} R$):

Theorem 17.3. *A domain R is integrally closed iff for every prime/maximal ideal $P \subseteq R$ the localization R_P is integrally closed.*

Proof. We interpret the ring R and all its localizations R_P as subrings of the fraction field K that is clearly also the fraction field of R_P. In the $\Rightarrow$ direction, let $a \in K$ be a root of a monic polynomial from $R_P[x]$. We may thus write

$$a^n + r_{n-1}/d_{n-1} \cdot a^{n-1} + \cdots + r_0/d_0 = 0.$$

Denote $d = d_{n-1} \cdots d_0 \notin P$, we multiply this equation by d^n and get

$$(ad)^n + r_{n-1}d/d_{n-1}(ad)^{n-1} + \cdots + r_0 d^n/d_0 = 0.$$

This shows ad integral over R and by the assumption $ad \in R$, implying $a \in R_P$.

In the opposite direction $\Leftarrow$, any element of K that is integral over R is, in particular, integral over each R_P and thus belongs to $\bigcap_{P \text{ maximal}} R_P$. We claim that this equals R (this is a general fact, that might be proved independently). Thus, let $a \in \bigcap_{P \text{ maximal}} R_P$ and form the denominator ideal

$$D = \{d \in R \mid d \cdot a \in R\}$$

and the membership $a \in R_P$ means that D contains an element from the complement of P, i.e. $D \not\subseteq P$. Since this holds for every maximal ideal, we must have $D = R$, giving $a \in D$ and finally $a \in R$. □

Definition 17.4. We say that a domain R is a discrete valuation ring if there exists a discrete valuation v on its fraction field K so that $R = \{0\} \cup v^{-1}(\mathbb{N}_0)$; in detail, $v \colon K^\times \to \mathbb{Z}$ is required to satisfy

- v is surjective,
- $v(a \cdot b) = v(a) + v(b)$,
- $v(a + b) \geq \min\{v(a), v(b)\}$.

Remark. It may be advantegeous to extend v to all of K by declaring $v(0) = \infty$.

For any discrete valuation on a field K, the inverse image $\{0\} \cup v^{-1}(\mathbb{N}_0)$ is closed under addition and multiplication, by the above properties, and contains 1 by virtue of the easily checked $v(1) = 0$.

The structure of DVRs is rather rigid, as we will now explore. Any element t with $v(t) = 1$ will be called a local parameter for R and we will fix an arbitrary choice of such.

Proposition 17.5. • *Units in a DVR R are exactly the elements $u \in R$ with $v(u) = 0$. Every nonzero element $r \in R$ can be written uniquely as $r = u \cdot t^n$ with $u \in R^\times$ (where clearly $n = v(r)$).*

- *A domain R is a DVR iff it is a UFD with a unique irreducible element, up to associatedness.*
- *A DVR is a local ring with the unique maximal ideal*

$$M = \{r \in R \mid v(r) > 0\} = (t).$$

Every nonzero ideal is of the form $M^n = (t^n)$ (so that R is in fact a PID and noetherian). Conversely, if R has nonzero ideals exactly $M^n = (t^n)$, it is a DVR.

- *The prime ideals of a DVR R are exactly 0 and M so that R has Krull dimension 1, i.e. the longest chain of primes consists of one inclusion – in this case $0 \subseteq M$.*

Proof. This is all fairly straightforward. For the first point, use $v(u^{-1}) = -v(u)$; further, since $r/t^n \in K$ has valuation 0, it is a unit of the ring.

For the second point, the implication $\Rightarrow$ is exactly the first point. For the implication $\Leftarrow$, observe that every nonzero element of the fraction field K can be written uniquely as $k = u \cdot t^n$ with $t \in \mathbb{Z}$ and we may thus introduce $v(k) = n$.

For the third point, observe that $r \mid s$ iff $v(r) \leq v(s)$ so that every nonzero ideal I is generated by any nonzero element of minimal valuation. By the first part, we may write it as $r = u \cdot t^n$ and thus $I = (r) = (t^n)$. In the opposite direction, R must then be a PID, hence UFD. Since irreducible elements of a PID, up to associatedness, correspond precisely to nonzero prime ideals and the only such is (t), there is a unique irreducible and the second point applies.

The last point is clear. □

Theorem 17.6. *For a domain R, the following conditions are equivalent*

- *R is a DVR,*
- *R is a noetherian local ring whose unique maximal ideal is nonzero and principal,*
- *R is a noetherian local ring of Krull dimension 1 that is also integrally closed.*

Proof. We have proved that the first point implies the other (except we did not mention explicitly DVR $\Rightarrow$ UFD $\Rightarrow$ integrally closed).

It remains to prove that any of the other conditions imply that R is a DVR. Start with the second point. Let $M = (t)$ be the maximal ideal of R. We will show that all nonzero proper ideals are of the form M^n. Clearly $I \subseteq M$ and we claim that there exists the largest n for which $I \subseteq M^n$. Otherwise, I would lie in the intersection $M^\infty \stackrel{\text{def}}{=} \bigcap_n M^n$ that we will show to be zero by Nakayama lemma: M^∞ is finitely generated since R is noetherian and clearly satisfies $M \cdot M^\infty = M$, thus $M^\infty = 0$. Thus let $a \in I \subseteq M^n = (t^n)$ with $a \notin M^{n+1}$.

We may write $a = u \cdot t^n$ and by assumption $u \notin M \Rightarrow u \in R^\times$. Thus a is associate to t^n and thus already a alone generates M^n; we get $I = M^n = (t^n)$, as claimed.

Now assume the third set of conditions. We will prove all conditions in the second point, where $M = 0$ would imply that R has Krull dimension 0, so it remains to show that M is principal. By Nakayama lemma $M^2 \subsetneq M$, for equality would give $M = 0$. Let $t \in M \setminus M^2$. We claim that $M = (t)$. Clearly $I = (t)$ is a proper nonzero ideal, thus contained in a unique prime ideal M. Proposition 4.6 gives $\sqrt{I} = M$ and this implies $M^n \subseteq I$ for some n by finite generation of M. Starting from this, we will show inductively $M^n \subseteq I \Rightarrow M^{n-1} \subseteq I$, finishing with $M \subseteq I \subseteq M$, as claimed. Thus let $x \in M^{n-1}$. Since we want to show that $x \in I = (t)$, we consider the element $x/t \in K$ of the fraction field and we want $x/t \in R$, which we prove by exploiting the fact that R is integrally closed. Consider the multiplication by x/t:

$$x/t \cdot : M \to R$$

Clearly $x/t \cdot M \subseteq 1/t \cdot M^n \subseteq R$. Now the image must be a submodule, i.e. an ideal, and we claim that it cannot be the trivial ideal R: for otherwise there would exist $m \in M$ such that $x/t \cdot m = 1$, i.e. $t = xm \in M^n$ and we obviously assume $n \geq 2$ and $t \notin M^2$. Thus the image of the multiplication map must be contained in M:

$$x/t \cdot : M \to M$$

Now the Cayley-Hamilton-Nakayama-like argument below gives a monic polynomial $F \in R[x]$, such that the multiplication by $F(x/t)$ is a zero map. Since K is a field and $M \neq 0$, this implies that $F(x/t) = 0$ in K, as required. □

Theorem 17.7. *Let S be a (commutative) R-algebra and let M be an S-module that is finitely generated over R. Then for every $s \in S$ there exists a monic polynomial $F \in R[x]$ such that $F(s) \cdot x = 0$ for all $x \in M$, i.e. $F(s)$ lies in the kernel of $S \to \mathrm{End}(M)$.*

Remark. Keeping R commutative, we may replace a non-commutative S by its commutative subalgebra $R[s] \subseteq S$ and apply the theorem to it, getting the same conclusion even for S non-commutative.

Proof. Write $M = R\{x_1, \ldots, x_n\}$ and express the action of $s \in S$ on M in two ways with respect to this generating set:

$$(x_1, \ldots, x_n) \cdot sE = (s \cdot x_1, \ldots, s \cdot x_n) = (x_1, \ldots, x_n) \cdot \begin{pmatrix} r_{11} & \cdots & r_{1n} \\ \vdots & & \vdots \\ r_{n1} & \cdots & r_{nn} \end{pmatrix} = (x_1, \ldots, x_n) \cdot A$$

where A denotes the n-by-n matrix in the formula, with elements in R. One can write this concisely as

$$(x_1, \ldots, x_n) \cdot (sE - A) = (0, \ldots, 0).$$

Multiplying by the adjoint matrix gives

$$(x_1, \ldots, x_n) \cdot \det(sE - A) = (0, \ldots, 0),$$

i.e. the multiplication by $\det(sE - A)$ annihilates the generators $x_1, \ldots, x_n$ and thus M. We may set $F(x) = \det(xE - A) \in R[x]$. □

We may now apply this characterization of DVRs to introduce Dedekind domains. Importantly, the last condition localizes well, so we define:

Definition 17.8. A *Dedekind domain* is a noetherian domain of Krull dimension 1 that is integrally closed.

Theorem 17.9. *Let R be a domain. TFAE*

- *R is a Dedekind domain,*
- *R is noetherian and for all nonzero primes P, the localization R_P is a DVR.*

Proof. We have proved that R is integrally closed iff R_P is integrally closed and it remains to show the same for the Krull dimension, but this is easy, since localization at P picks out of the prime ideals of R those that are contained in P. The point is that the primes in a DD and in a DVR form the following posets:

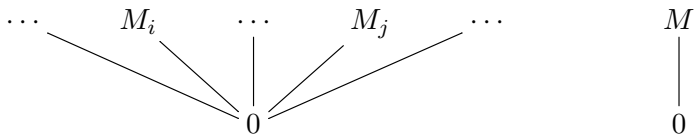

These clearly correspond to one another. □

Now we want to show an interpretation of a DD in terms of fractional ideals.

Definition 17.10. Let R be a domain with a fraction field K. A fractional ideal is an R-submodule $A \subseteq K$ of the form $1/d \cdot I$ for an ideal $I \subseteq R$.

Remark. Over a noetherian domain, this is equivalent to A being a finitely generated R-submodule of K.

We introduce a product of fractional ideals similarly to that of ideals, i.e. AB is the ideal generated by the products ab, for $a \in A$ and $b \in B$. Clearly $(1/d \cdot I)(1/e \cdot J) = 1/(de) \cdot IJ$ so that this product is indeed a fractional ideal. Clearly the unit is R so that we get an induced notion of an *invertible* fractional ideal A as that for which there exists a fractional ideal B such that $AB = R$.

Example 17.11. A principal fractional ideal is one of the form $(k) = (r/d) = 1/d \cdot (r)$ for $k = r/d \in K$. Clearly, this has inverse (k^{-1}). In a principal ideal domain, these are all examples.

Consider, for a nonzero fractional ideal A, the following fractional ideal

$$A' = \{k \in K \mid kA \subseteq R\}$$

(since A contains some element $d \in R$, we have $A'd \subseteq R$ so that $A' = 1/d \cdot I$ for the ideal $I = A'd$). By definition, $A'A \subseteq R$ and we will prove that the equality holds iff A is invertible, in which case $A^{-1} = A'$. The implication $\Rightarrow$ is obvious, so assume that A is invertible. Then $A^{-1} \subseteq A'$ and consequently

$$R = A^{-1}A \subseteq A'A \subseteq R$$

and so we must get equality everywhere and A' is also an inverse. But inverses are unique in monoids.

Theorem 17.12. *In a Dedekind domain, every fractional ideal is invertible. In addition, every nonzero proper ideal $I \subseteq R$ admits a unique decomposition $I = P_1 \cdots P_r$ into a product of prime ideals.*

Proof. Let A be a fractional ideal and consider the fractional ideal A' as above. We need to show that $A'A = R$. This means that the inclusion $A'A \to R$ is an isomorphism and we know that this may be checked on localizations. These are[3]

$$(A_P)'A_P = (A')_P A_P = (A'A)_P \to R_P$$

(we think of the localization A_P as the R_P-submodule generated by A). These will be isomorphisms provided that A_P is invertible. This follows from R_P being a PID.

Now let $I \subseteq R$ be an ideal. The primary decomposition of I is

$$I = I_1 \cap \cdots \cap I_s$$

with each I_i primary, say $\operatorname{Ass} R/I_i = \{P_i\}$. Thus, I_i is contained in a unique maximal ideal P_i and consequently these are pairwise comaximal, i.e. $I_i + I_j = R$, giving

$$I = I_1 \cdots I_s$$

by the following proposition. Now the P_i-primary component I_i of I is uniquely determined since P_i is minimal over I and in fact I_i/I is the kernel of the localization map $R/I \to R_{P_i}/I_{P_i}$ or, slightly better, I_i is the preimage of I_{P_i} under the localization map $\lambda_i \colon R \to R_{P_i}$. Now since R_{P_i} is a DVR, the ideal I_{P_i} is a power $M_i^{k_i}$ of the maximal ideal and thus pulls back to the corresponding power $P_i^{k_i}$, since this power is P_i-primary $P_i^{k_i}$, thus $(R \setminus P_i)$-saturated, and clearly maps to $M_i^{k_i}$.

The uniqueness follows easily from all primes being invertible: for if $P_1 \cdots P_r = Q_1 \cdots Q_s$ then for any prime Q_j we have $P_1 \cdots P_r \subseteq Q_j$ so $P_i \subseteq Q_j$. Symmetrically $Q_{j'} \subseteq P_i \subseteq Q_j$ and appying this for Q_j minimal must give equality. We may this common prime in the group of invertible fractional ideals and proceed by induction. □

In fact, these are both equivalent conditions, i.e. a domain where every nonzero fractional ideal is invertible is a Dedekind domain (or a field). Also, a domain where every nonzero proper ideal factors uniquely into a product of prime ideals is a Dedekind domain (or a field). The first claim is not difficult: One shows that all invertible fractional ideals must be finitely generated ($AB = R$ implies $\widehat{A}B = R$ for some finitely generated fractional subideal $\hat{A} \subseteq A$ by looking at $1 \in R$; by uniqueness of inverses $A = \widehat{A}$), hence R is noetherian. Every localization R_P at a nonzero prime P will also have all nonzero fractional ideals invertible (the fractional ideals are of the form A_P and as such admit an inverse A'_P). Thus, it remains to show that every noetherian local ring with all fractional ideals invertible must be a DVR. Let $t \in M \setminus M^2$ and consider the fractional ideal $t^{-1}M$ with inverse $(t)M^{-1} \subseteq MM^{-1} = R$. Now $(t)M^{-1} \not\subseteq M$ since otherwise $t \in (t) \subseteq M^2$, so $(t)M^{-1} = R$ giving $(t) = M$ as required. The second claim is much more complicated.

Proposition 17.13. *Assume that $I + J = R$. Then $IJ = I \cap J$. More generally if I_i are pairwise comaximal then $I_1 \cdots I_r = I_1 \cap \cdots \cap I_r$.*

Proof. The containment $IJ \subseteq I \cap J$ holds always, so let $z \in I \cap J$. Write $x + y = 1$, giving $z = (x + y)z = xz + zy$ with both terms in IJ.

[3]The first equality requires finite generation $A = R\{a_1, \ldots, a_r\}$. Clearly $(A')_P \subseteq (A_P)'$ (since e.g. $(A')_P A_P = (A'A)_P \subseteq R_P$) and the righ hand side consists of $k \in K$ such that $ka_i \in R_P$, implying the existence of $d \notin P$ such that $dka_i \in R$, i.e. $dk \in A'$, and thus $k = (dk)/d \in (A')_P$.

The general case is obtained by application to $I_1 \cdots I_{r-1}$ and I_r once we show that these are comaximal which is a bit tricky. So let $x_i + y_i = 1$ with $x_i \in I_i$ and $y_i \in I_r$. Then we have

$$1 = (x_1 + y_1) \cdots (x_{r-1} + y_{r-1}) = x_1 \cdots x_{r-1} + \underbrace{\text{terms containing some } y_i}_{\in I_r} \in I_1 \cdots I_{r-1} + I_r. \ \square$$

Remark. I would say that I_i are comaximal if $I_j + \bigcap_{i \neq j} I_i = R$ and the second part shows that pairwise comaximal implies comaximal, which I find a bit surprising.

18. Some interesting exercises

A left adjoint is right exact. A left adjoint is exact iff its right adjoint preserves injectives (for the reverse implication, it should be useful that $\mathrm{Hom}(-, I)$ preserves and jointly reflects exactness – here $0 = H^n(\mathrm{Hom}(C, I)) \cong \mathrm{Hom}(H_n C, I)$ so it remains to show that it jointly reflects zero; then use $\mathrm{Hom}(F-, I) \cong \mathrm{Hom}(-, GI)$ for I and thus also GI injective). Both adjoints are exact iff $\mathrm{Ext}^*(Fx, y) \cong \mathrm{Ext}^*(x, Gy)$ (apply the previous to a projective resolution of F and/or to an injective resolution of y).

A square with vertical maps mono and horizontal maps epi is a pullback iff it is a pushout (the respective maps are jointly epi and jointly mono).

Prove that a square is a pushout square iff the induced map on (say vertical) cokernels is iso and on kernels is epi. (Make the square into a double complex, and form the long exact sequence of homology groups for the columns and the total space.) Dually, it is a pullback square iff the induced map on kernels is iso and on cokernels is mono.

Let $0 \to A \to B \to C \to 0$ be a ses and $g\colon Y \to B$ an arbitrary map. By factoring $Y \to B \to C$ through its image Z, construct a diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & X & \longrightarrow & Y & \twoheadrightarrow & Z & \longrightarrow & 0 \\ & & \downarrow & & \downarrow{\scriptstyle g} & & \downarrow & & \\ 0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \end{array}$$

Prove that the left square is a pullback square (not so much interesting I guess, but it gives a concrete construction of $g^{-1}(A)$; it would be more challenging to start from the pullback, take the cokernel and show that the induced map on cokernels is mono). More interestingly, reprove that noetherian modules are closed under extensions: Assume that A and C are noetherian and that $Y \subseteq B$ is a submodule and apply the above to this inclusion. You will need to show that an extension of f.g. modules is f.g.

Prove that in a ses

$$0 \to A \to B \to C \to 0$$

if C is f.p. and B is f.g. then also A is f.g. (write C as a cokernel $R^s \to R^t \to C \to 0$ and lift $R^t \to C$ to B, getting a diagram

$$\begin{array}{ccccccccc} & & R^s & \longrightarrow & R^t & \longrightarrow & C & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & \| & & \\ 0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \end{array}$$

Observe that $\operatorname{coker} f \cong \operatorname{coker} g$ with the latter f.g., so that we get

$$R^s \to A \to \operatorname{coker} g \to 0$$

and conclude that A is f.g.).

over a noetherian ring, every f.g. module has a projective resolution consisting of f.g. free modules.

truncations and homology.

injective generators, e.g. $\mathrm{Hom}(R, \mathbb{Q}/\mathbb{Z})$; related to the exercise about exact left adjoints.

derived functors are universal δ-functors (i.e. $\mathrm{Hom}(T_*, L_*F) \cong \mathrm{Hom}(T_0, F)$ it is a right adjoint to the 0-component functor; homology is such a functor on non-negatively graded chain complexes); if T_* is a δ-functor, define $\mathcal{P} = \{A \mid \forall n > 0 \colon T_n A = 0\}$. If there is enough $\mathcal{P}$-projectives then T_* is universal.

a functor is additive iff it preserves biproducts (binary, but maybe zero is also needed). in particular, any left or right adjoint is automatically additive!

19. Possible essay topics

COMMUTATIVE ALGEBRA:

flatness, faithful flatness, going up/down – Matsumara: Commutative algebra

Grobner bases and primary decomposition, radicals etc. – Robbiano etal.: Computational aspects of commutative algebra

combinatorics and commutative algebra, face ring of a simplicial complex – Stanley: Combinatorics and commutative algebra

noncommutative localization

Hilbert functions, Hilbert polynomials, Koszul resolutions (overlap to homological algebra)

symbolic powers of an ideal

local properties of commutative rings (e.g. flatness)

Morita equivalence, Morita invariance

—

HOMOLOGICAL ALGEBRA:

derived categories/model categories point of view

cohomology of associative/commutative/Lie algebras (Hochschield, André–Quillen, Chevalley–Eilenberg)

A_∞-algebras (and C_∞ and L_∞, possibly E_∞-algebra)

spectral sequences

Galois cohomology, Tate cohomology

sheaf theory, introduction of Ext and Tor

abelian categories

simplicial methods, Dold-Kan correspondence

derived Morita equivalence

homotopy limits and colimits

general Künneth theorem

Eilenberg–Zilber theorem for simplicial abelian groups

differential graded algebras, general Ext vs extensions

satellites, δ-functors, universal δ-functors (universal property of derived functors)

—

REPRESENTATION THEORY:

representation theory of Lie groups (including Weyl group)

representation theory of Lie algebras (including Weyl group)

modular representation theory (when characterstics divides the order of the group)
bialgebras, Hopf algebras, Frobenius algebras
representation ring and equivariant stable homotopy theory