3. Modules over PIDs

Finitely Generated Modules over Principal Ideal Domains

Main Theorem

Our goal is to prove the classification theorem for finitely generated modules over PID's, which asserts that every finitely generated module over a PID is the direct sum of a free module and a finite set of cyclic modules. Depending on how you describe the cyclic modules you get different uniqueness statements.

Theorem: Let R be a principal ideal domain and let M be a finitely generated R module. Then there is an integer k and elements $\pi_1, \ldots, \pi_m$ in R such that $\pi_1 | \pi_2 | \cdots | \pi_m$ such that

$$M = R^k \oplus R/\pi_1 R \oplus \cdots \oplus R/\pi_m R.$$

Further, the integer k and the ideals $\pi_i R$ are uniquely determined by M. The ideals $\pi_i R$ are called the invariant factors of M, and the integer k is its rank.

Notice that if $R = \mathbb{Z}$ and M is finite then this is the fundamental theorem of finite abelian groups with the π_i being the invariant factors.

Alternative formulation

Theorem: Let R be a PID and let M be a finitely generated R module. Then there is an integer k and elements $\pi_i \in R$ such that π_i is a prime power and

$$M = R^k \oplus R/\pi_1 R \oplus \cdots \oplus R/\pi_m R.$$

Again, the rank k and the prime power factors π_i are unique (up to ordering in this case).

The prime powers π_i are called the elementary divisors of M.

If $R = \mathbb{Z}$ this is the fundamental theorem of finite abelian groups, asserting that every such group is a finite product of cyclic groups of prime power order, and that the prime powers are unique up to ordering.

Strategy

Our strategy is to use ideas from linear algebra and approach the problem algorithmically.

Suppose that M is generated by n elements $e_1, \ldots, e_n$ over the PID R. Then there is a surjective map

$$\pi : R^n \to M$$

defined by $\pi((r_1, \ldots, r_n)) = \sum_{i=1}^n r_i e_i$.

If $f = (r_1, \ldots, r_n)$ is in the kernel of π, then

$$\sum_{i=1}^n r_i e_i = 0.$$

Relations

Because of this, elements of the kernel N of π are called *relations* for the generators e_i, and N is called the module of relations for M.

Since the relation module N of this map is a submodule of R^n, we know from our discussion of finite generation is generated by (at most) n elements $f_1, \ldots, f_n$.

Let's assume that our relation module has n generators $f_1, \ldots, f_n$, some of which might be zero.

The relation matrix

Expressing f_j in terms of the e_i yields an $n \times n$ matrix $A = (a_{ij})$ defined by:

$$f_j = \sum a_{ji} e_i$$

The columns of the matrix A express the generators f_j of the kernel of π in terms of the basis e_i for R^n.

A is called a relation matrix for M.

The kernel as column space of the relation matrix

If, as we do in linear algebra, we express elements of R^n as column vectors with R entries, we have a map

$$a : R^n \to R^n$$

defined by $a(v) = Av$ (matrix multiplication by A on a column vector v with entries in R).
If the entries of v are $(r_1, \dots, r_n)$ then $a(v) = \sum_{i=1}^{n} r_i f_i$ and therefore the image of the R-linear map a is N.

Standard form

We've reached a point where our module M is isomorphic to R^n/N where N is generated by the columns of our matrix A.

We will show the following:

- N is free of rank m where $m \leq n$.
- M has a basis $y_1, \ldots, y_m$ with the property that there are elements $b_1, \ldots, b_m \in R$ such that $b_1|b_2|\cdots|b_m$ and $b_1 y_1, b_2 y_2, \ldots, b_m y_m$ are a basis for N.

In terms of the relation matrix, we are saying that if we choose our basis $e_1, \ldots, e_n$ and $f_1, \ldots, f_n$ properly, then the corresponding matrix A is diagonal with entries $b_1, b_2, \ldots, b_m, 0, 0 \ldots 0$ and $b_1|b_2|\cdots|b_m$.

We will do this by modifying the set of generators f_j and e_i so that, at each stage, they continue to be sets of generators, but eventually they have the desired relation.

The result from standard form

If we achieve the standard form, then we have the picture

$$R^n \to M$$

where

$$(r_1, \ldots, r_n) \mapsto \sum r_i y_i$$

and the kernel of this map is

$$N = b_1 y_1 \oplus b_2 y_2 \oplus \cdots \oplus b_m y_m.$$

Therefore $R^n/N = R/b_1 R \oplus \cdots R/b_m R \oplus R^{n-m}$ which is the structure we are trying to establish.

Alternatively, we can think of M as having generators $e_1, \ldots, e_n$ and relations $b_i e_i = 0$

Reduction Operations

Modifying the generators of M

Lemma: Suppose $1 \leq t, s \leq n$ with $i \neq j$. If we let elements $e_i^* = e_i$ for $i \neq t, s$, and also

$$e_t^* = xe_t + ye_s$$
$$e_s^* = ze_t + we_s$$

Then $e_1^*, \ldots, e_n^*$ are also generators of M.

Proof: Write

Since $e_i = e_i^*$ for $i \neq t, s$ and

$$e_t = we_t^* - ye_s^*$$
$$e_s = -ze_t^* + xe_s^*.$$

wee see that all of the e_i are in the submodule of M generated by the e_i^*, and vice versa, so the e_i^* are again a set of generators of M.

Row operations

Let's examine the effect of this change on the relation matrix A. If

$$m = r_1 e_1 + \cdots + r_n e_n.$$

then

$$m = \sum_{i \neq t,s} r_i e_i^* + (r_t w - r_s z) e_r^* + (-y r_t + x r_s) e_s^*.$$

This means that if we construct the relation matrix A^* by writing

$$f_j = \sum a_{ji}^* e_i^*$$

we see that A^* is obtained from A by modifying rows t and s. If we use subscripts to describe rows of matrices then

$$A_t^* = w A_t - z A_s$$
$$A_s^* = -y A_t + x A_s$$

Column Operations

More generally, we see that, given any relation matrix A, and x, y, z, w such that $xw - yz = 1$, modifying A by changing rows t and s according to these formulas yields a new relation matrix giving rise to an isomorphic module M.

A similar line of argument shows that if we make the same type of modification to the generators f_j for the relations, then we modify the relation matrix A by column operations of the same type.

Outline of proof of standard form

Initial remarks

Now suppose we are given an $n \times n$ matrix A with entries in a PID R. There is a sequence of row and column operations that reduces it to standard form, so that the reduced matrix is diagonal, the first k diagonal elements are nonzero and the remaining $n-k$ are zero, and the nonzero diagonal elements satisfy

$$a_{11} | a_{22} | \cdots | a_{kk}$$

Main Steps

1. If $A = 0$, we're done, otherwise swap rows and columns so a_{11} is not zero.

Clear out the first row

2. If all a_{1i} for $i > 1$ are divisible by a_{11}, replace each column A^j where a_{1j} is not zero by $A^j - a_{11}/a_{1j}A^1$. Otherwise, for each column $j = 2, \ldots, n$ where a_{1j} is not zero, use the fact that R is a PID to find a generator d for the ideal (a_{11}, a_{1j}) for each column and write $a_{11}x - a_{1j}y = d$. Then make a column operation using this x and y with $w = a_{11}/d$ and $z = a_{1j}/d$ to obtain a matrix with $a_{11} = d$ and $a_{1j} = 0$. At the end of this step, the only nonzero entry in the first row is a_{11}.

Clear out the first column

3. If all a_{i1} for $i > 1$ are divisible by a_{11}, replace each row A_j with $A_j - a_{j1}/a_{11}A_1$. Now you've got a matrix so that the first row and column are all zero, except for a_{11}. Go to step 4. Otherwise, use the fact that R is a PID to find a generator $d = a_{11}x - a_{j1}y$ and make a row operation using this x and y with $w = a_{11}/d$ and $z = a_{j1}/d$ to obtain a matrix with $a_{11} = d$ and $a_{j1} = 0$. At the end of this process, you've got a matrix so that a_{11} is the only nonzero entry in the first column; but you may have messed up the first row. So go back to step 2.

Check divisibility; descend to submatrix

4. At this point the first row and column of A are zero except for a_{11}. If a_{11} divides every entry in the lower right $(n-1) \times (n-1)$ submatrix, then apply this algorithm to that submatrix and continue. If a_{11} does NOT divide every entry in lower submatrix, find a row A_j containing an element not divisible by a_{11} and replace the first row A_1 by $A_1 + A_j$. Now go back to step 2 and continue.

Remarks on the algorithm

There are two things to consider in this algorithm.

First, the loop through steps 2 and 3 must eventually terminate because each time you go through it, you replace a_{11} by a divisor of a_{11}. This cannot continue indefinitely, so eventually you will reach step 4.

Second, if a_{11} divides everything in the lower submatrix, then by induction, once that matrix is in standard form, the whole matrix will be in standard form. If a_{11} does *not* divide everything in the lower submatrix, then the return to step 2 will replace a_{11} by a proper divisor of a_{11} and again, that can't continue indefinitely.

Constructive for Euclidean rings

The only non-constructive part of this “algorithm” is that we invoke the PID property of R so that, given a, b we can find $ax + by = d$ where d is the gcd of a and b. If R is Euclidean, this can be done constructively, and so this algorithm can be carried out in practice.

Uniqueness

Uniqueness in DF

Proof of uniqueness is given in DF, Section 12.1 Theorem 9.