20 Modules and Presentation Matrices

20.1 Review — Definition of Modules

Last class, we defined modules over commutative rings — we've also seen a few examples over noncommutative rings, but from now we'll stick to commutative ones.

Definition 20.1
A **module** M over a ring R is an abelian group together with an action of R — each $r \in R$ acts by a map $r : m \mapsto r(m)$ (we often denote $r(m)$ by rm), satisfying certain axioms.

In some sense, modules can be similar to vector spaces:

Example 20.2
The free module of rank n is $M = R^n$, consisting of n-tuples of elements in R (where addition and the R-action are performed componentwise).

But there are many other examples of modules over most rings, and we'll now discuss how to describe more general modules.

20.2 Generators and Relations

If M is a module, the elements $a_1, \ldots, a_n \in M$ form a **system of generators** if every $x \in M$ has the form

$$x = \sum r_i a_i$$

for some $r_i \in R$.

Choosing any n elements $a_1, \ldots, a_n \in M$ provides a homomorphism of modules $\varphi : R^n \to M$ (where we send $(r_1, \ldots, r_n) \mapsto r_1 a_1 + \cdots + r_n a_n$). Then a_1, ..., a_n are generators if and only if φ is onto.

In the analogy with vector spaces, a system of generators is a set of vectors that *span* the vector space. But in vector spaces, if we have a set of vectors which span the space, we can always find a subset which forms a basis — we can drop some of the a_i to make the presentation $x = \sum r_i a_i$ be *unique*. This is not true in general.

Stating that the presentation $x = \sum r_i a_i$ is unique is equivalent to stating that $\ker \varphi = 0$ (given two presentations, we could subtract them to get an element of the kernel). Then $\varphi : R^n \to M$ is actually an isomorphism, which means M is free. But this is often *not* the case — there are usually many R-modules which are not free.

Student Question. *Does the system of generators have to be finite?*

Answer. *Not necessarily; as a silly example, we could take* all *elements of M as our system of generators. In later classes, we'll discuss ways to sometimes show that we* can *always find a finite system of generators; but for today, we'll just* assume *that we can.*

Definition 20.3
If a finite set of generators exists, then M is called **finitely generated.**

20.3 Presentation Matrices

Most of the time, we won't be lucky enough that $\ker \varphi = 0$. But in general, by the homomorphism theorem, we can write

$$M \cong R^n / \ker \varphi.$$

Now $\ker \varphi$ is *also* a module, so we can describe it further.

Definition 20.4
If $\ker \varphi$ is also finitely generated (for some surjective homomorphism $\varphi : R^n \to M$), then we say M is **finitely presented**.

We'll see later that for a large class of rings, *every* module which is finitely generated is also finitely presented; but for now, we'll assume that M is finitely presented as well.

Then we can choose a set of generators $b_1, \ldots, b_m$ for $\ker\varphi$. Each b_i can be thought of as a column vector (since it's an element of R^n). So we can write down the $n \times m$ marix

$$B = \begin{bmatrix} | & | & & | \\ b_1 & b_2 & \cdots & b_m \\ | & | & & | \end{bmatrix},$$

called the **presentation matrix**. Now since we've fixed generators for $\ker\varphi$, we have an onto map $\psi : R^m \to \ker\varphi$. Equivalently, we can think of ψ as a map $\psi : R^m \to R^n$ whose image is exactly $\ker\varphi$. As in the case of vector spaces in linear algebra, this map can explicitly described as $\psi : x \mapsto Bx$ (where $x \in R^m$), and $\operatorname{im}\psi = BR^m$. This means we can write

$$M \cong R^n/BR^m$$

(where BR^m is the span of the column vectors of the presentation matrix B).

20.4 Classifying Modules

The rest of the lecture focuses on classifying finitely generated modules over a Euclidean domain. Since abelian groups are $\mathbb{Z}$-modules, this will also give us a classification of finitely generated groups.

For a given module M, the presentation is not at all unique.

Example 20.5
Let $R = \mathbb{Z}$ and $M = \mathbb{Z}/5\mathbb{Z}$. An obvious presentation of M is to choose one generator 1, and the generator 5 for the kernel. In this case, $B = \begin{bmatrix} 5 \end{bmatrix}$. But there are other presentations as well. For example, we can choose two generators for M, and the presentation matrix

$$B = \begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix}.$$

To see that $\mathbb{Z}^2/B\mathbb{Z}^2$ is again $\mathbb{Z}/5\mathbb{Z}$, note that the sublattice $L \subset \mathbb{Z}^2$ spanned by $(2,1)^t$ and $(1,3)^t$ has index $|\det(B)| = 5$, which means the quotient $\mathbb{Z}^2/L$ has five elements. But the only group of five elements is $\mathbb{Z}/5\mathbb{Z}$, so this construction indeed gives another presentation of $\mathbb{Z}/5\mathbb{Z}$.

Our goal is to more systematically understand how to understand the module given a presentation. Here, the analogy between vector spaces and modules will be useful — similarly to in linear algebra, we'll start with a matrix and try to perform elementary operations on the matrix which don't change the module. We'll then use these operations to write the matrix in a simpler form, from where it's easy to understand the module.

20.4.1 Elementary Row and Column Operations

We'll use the notation M_B to refer to the module M produced by a presentation matrix B.

In linear algebra, we saw the elementary column operations for matrices over a field. We can define elementary column operations on matrices over a *ring* in a similar way:

1. Multiplying a column by a unit (an invertible element) — note that in the case of a field, we could multiply by *any* nonzero element (because any nonzero element has an inverse), but it was important that we were able to invert the multiplication; so here we restrict the definition to units.
2. Adding an arbitrary multiple of one column to another column.
3. For convenience, we can also include swapping two columns; but in fact, this can be obtained as a combination of the first two operations.

These operations *do not* change the span of the columns. (This can be verified the same way as for matrices over a field.) So if B' is obtained from B by elementary column operations, then we have $BR^m = B'R^m$.

Another useful way to think of this is in terms of matrix multiplication. In other words, if B' is obtained from B by elementary column operations, then we have $B' = B \cdot C$, where C is an $m \times m$ *invertible* matrix (it's easy to see that all the elementary column operations are invertible). This means $C \in \mathrm{GL}_m(R)$ (the group of invertible matrices wth entries in R). Note that over a ring, for a matrix C to be invertible, $\det(C)$ must be a *unit* — it's not enough to require the determinant to be nonzero. (In fact, the converse is also true, but requires more work.)

Then we have $B'R^m = BCR^m$. But since C is invertible, the map defined by C is an isomorphism, so $CR^m = R^m$. So this is an alternate way of making it clear that $B'R^m = BR^m$.

The elementary *row* operations are defined analogously.

When we apply elementary row operations to a matrix B, we produce a matrix B' which is a presentation matrix for an *isomorphic* module — in order to explain why, we'll again think in terms of matrix multiplication. We can write $B' = CB$, where $C \in \mathrm{GL}_n(R)$. We then want to produce an isomorphism between R^n/BR^m and R^n/CBR^m. But that isomorphism is just given by multiplication by C. More explicitly, we can draw a commutative diagram:

$$\begin{tikzcd}
R^n/BR^m \arrow[r] & R^n/CB \cdot R^m \\
R^n \arrow[u] \arrow[r, "C"'] & R^n \arrow[u]
\end{tikzcd}$$

The map $x \mapsto Cx$ is an isomorphism $R^n \to R^n$. But by definition, $x \in BR^n$ if and only if $Cx \in CBR^m$. So if we restrict our isomorphism to BR^m, then it restricts to an isomorphism between BR^m and CBR^m; and therefore to an isomorphism between the quotients R^n/BR^m and $R^n/CB \cdot R^m$.

So the row and column operations don't change our module (up to isomorphism).

20.4.2 Smith Normal Form

Using the row and column operations, it's possible to write any presentation matrix in a much simpler form. (We'll primarily focus on the case $R = \mathbb{Z}$, but this holds for any Euclidean domain.)

Theorem 20.6

Every $n \times m$ matrix over a Euclidean domain R can be reduced by elementary row and column operations to a matrix in *Smith normal form* — if we let $B = (b_{ij})$, then we have $b_{ij} = 0$ for all $i \neq j$, and $b_{11} \mid b_{22} \mid b_{33} \mid \cdots$.

So a matrix in Smith normal form looks like

$$\begin{bmatrix} d_1 & & & & & \\ & d_2 & & & & \\ & & d_3 & & & \\ & & & \ddots & & \\ & & & & d_k & \end{bmatrix}$$

(where all entries not shown are 0's).

We'll discuss the proof next class. It combines two ideas — the method of using Gaussian elimination to solve systems of equations over a *field* (which involves reducing matrices to a simpler form using row and column operations as well), and the Euclidean algorithm. (We've previously discussed Euclidean domains in the context of them being PIDs, but it turns out that here, having an effective way of computing the gcd using division with remainder will be very useful. The theorem is also true for PIDs, with a bit of modification (and a different proof); but all the examples we'll be interested in are Euclidean domains.

Corollary 20.7

Every finitely presented module over a Euclidean domain is isomorphic to a direct sum of *cyclic* modules (modules which are generated by one element) — we can write

$$M \cong R^a \oplus R/(d_1) \oplus R/(d_2) \oplus \cdots \oplus R/(d_k),$$

where we additionally have $d_1 \mid d_2 \mid \cdots \mid d_k$.

Later we'll see that any finitely generated module over a Euclidean domain is also finitely presented; this will mean that our statement actually holds for all finitely *generated* modules.

This corollary is clear from the theorem:

Proof. Using Theorem 20.6, we can rewrite the presentation matrix B to be diagonal, with diagonal $d_1 \mid d_2 \mid \cdots \mid d_k$, where $k = \min(m, n)$. In this case, when we take a column vector in R^m and multiply by B, we simply scale each coordinate by the corresponding d_i — so when we quotient out by BR^m, this coordinate becomes $R/(d_i)$, and we get

$$M_B = \bigoplus R/(d_i) \oplus R^a.$$

(The extra free factor comes from rows with no entries — either because $d_i = 0$ or because $n > m$ — since such rows correspond to coordinates in R^n where we're not quotienting out by anything.) □

In fact, this classification can be used to understand the subgroups of lattices as well (which came up when we studied factorization in quadratic number fields). This theorem implies that to describe a subgroup, we can always choose a basis for the lattice, and simply scale these basis vectors by numbers to get a basis for the subgroup.

Student Question. *What does it mean for a module to be cyclic — is the free module cyclic?*

Answer. *A module is cyclic if it's generated by one element. The free module is cyclic — it's generated by one element with* no *relations. Meanwhile, R/d is generated by one element x, with the relation that $dx = 0$.*

Resource: Algebra II Student Notes
Spring 2022
Instructor: Roman Bezrukavnikov
Notes taken by Sanjana Das and Jakin Ng