Math 154. Modules over a PID

In this handout we give prove a special case of the structure theorem for modules over a principal ideal domain (PID), sufficient for our needs at the outset in this course (the general structure theorem is needed for more advanced applications in algebraic number theory and beyond).

1. Context

Let A be a commutative ring and let M be a finitely generated A-module. (For example, a $\mathbf{Z}$-module is just an abelian group by another name – why? – and so for $A = \mathbf{Z}$ such an M is just a finitely generated abelian group.) We say that M is *free of rank* n if $M \simeq A^{\oplus n}$ as an A-module; i.e., there exists a set of n elements $x_1, \ldots, x_n \in M$ that is a linearly independent spanning set (in other words, every $m \in M$ has the form $m = \sum a_i x_i$ for unique $a_1, \ldots, a_n \in A$). Such $\{x_i\}$ is called a *basis* of M. It is true that if $A \neq 0$ then such an n is unique when it exists (so the notion of "rank" of a free finitely generated module is well-defined), but we do not need this except for PID's, in which case we will prove such well-definedness below.

Over a general commutative ring, *finitely generated modules generally do not have bases*. For instance, the $\mathbf{Z}$-module $(\mathbf{Z}/6\mathbf{Z}) \times (\mathbf{Z}/49\mathbf{Z})$ has no basis since every nonzero element of the module is killed by a nonzero integer (e.g., killed by 6×49). Although we can define the abstract notion of an *A-linear map* between two A-modules, in general such maps cannot be described via matrices (in the absence of a basis).

Here is an example that looks innocuous but is crucial:

Example 1.1. Let A be a commutative ring and M an A-module that is "cyclic"; i.e., is spanned over A by a single element. If $m \in M$ is such an element then the linear map $A \to M$ defined by $a \mapsto am$ is surjective, and the kernel J is the *annihilator* of m (i.e., the set of $a \in A$ such that $am = 0$). This kernel is an ideal of A (why?), and the resulting map $A/J \to M$ is an isomorphism. Conversely, for any ideal J of A the quotient module A/J is spanned by the element 1 mod J. So the cyclic A-modules are, up to A-linear isomorphism, given exactly by the quotients A/J modulo ideals of A.

If A is a field then the only such J's are (0) and (1), but *every* A that is not a field admits a nonzero proper ideal (why?), whence admits the A-module A/J with no A-basis (as every nonzero element is killed by the nonzero ideal J).

Here is an important source of modules:

Example 1.2. An A-submodule of A is just an ideal of A by another name. (Why?) So a special case of the question as to whether submodules of finitely generated A-modules are finitely generated is this: are all ideals of A finitely generated? Later in the course we will study the important class of rings for which this finiteness property of ideals holds, but for now we note that it certainly holds for a PID (all ideals are even cyclic, by definition of "PID"!).

Beware that if a domain A is *not* a PID (of which we shall see many examples later in the course) then the free module A of rank 1 contains non-principal ideals and hence has submodules that are *not* cyclic. So in contrast with the theory of vector spaces (i.e., linear algebra), over more general domains A it can happen that a free module admits submodules which require *more* generators than the ambient A-module!

In general, the above example shows that it can be hard to control information about generators of a submodule in terms of generators of an ambient module. But over a PID the situation is much nicer, as encoded in the following special case of the structure theorem for finitely generated modules over a PID.

Theorem 1.3. *Let A be a PID with field of fractions K, and M an A-module generated by r elements with $r \geq 1$.*

(1) *Every submodule $M' \subset M$ is generated by r elements.*
(2) *If M is nonzero and torsion-free (i.e., for nonzero $a \in A$, multiplication on M by a is injective: $am = 0 \Rightarrow m = 0$) then M is free, say $M \simeq A^{\oplus n}$, and such n is unique (called the* rank *of M). Every nonzero submodule $M' \subset M$ is free of rank $n' \leq n$.*

In case (2), if M is contained in a K-vector space V and M spans V over K then $n = \dim_K V$.

Note that in (1) we are *not* claiming that M' cannot admit a generating set of size smaller than r (even if r is minimal for M); consider the case $M' = 0$.

As an example, if we take A to be $\mathbf{Z}$ then for a field F of finite degree n over $\mathbf{Q}$, we will show later that the torsion-free $\mathbf{Z}$-module M given by the ring of integers $\mathscr{O}_F$ in F (to be defined!) is finitely generated, so it is free of some finite rank n by (2), and necessarily $n = [F : \mathbf{Q}]$ by the final assertion in Theorem 1.3 since we will see that M spans F over $\mathbf{Q}$ (every $\alpha \in F$ even admits a nonzero rational multiple that lies in $\mathscr{O}_F$).

2. Proof of Theorem 1.3

Say M is generated by elements $x_1, \ldots, x_r$. We shall induct on r first to prove (1), and then to prove (2). If $r = 1$ then $M \simeq A/J$ for an ideal J of A by Example 1.1, so submodules $M' \subset M$ are exactly I/J for ideals $I \subset A$ containing J (i.e., I is a submodule of A containing J). But A is a PID (!!), so such I admits a single generator, hence likewise the quotient $M' = I/J$ is spanned by the image of such a generator of I. (The end of Example 1.2 indicates that already this cyclic case completely breaks down beyond PID's for (1).) This settles (1) when $r = 1$.

Now suppose $r > 1$. Let N be the A-span of $\{x_1, \ldots, x_{r-1}\}$, so N admits $r-1$ generators. Since $N' := N \cap M'$ is a submodule of N, by induction it is generated by $r-1$ elements $y_1, \ldots, y_{r-1}$. Consider the quotient $M'/N' \subset M/N$. Since M/N admits a single generator (namely, the image of x_r), by the settled cyclic case it follows that M'/N' is cyclic. Letting $y_r \in M'$ represent a generator of M'/N', every $y \in M'$ satisfies $y \equiv a_r y_r \bmod N'$ for some $a_r \in A$, so $y - a_r y_r \in N' = \sum_{1 \leq i < r} A y_i$; i.e., $y - a_r y_r = a_1 y_1 + \cdots + a_{r-1} y_{r-1}$ for some $a_1, \ldots, a_{r-1} \in A$. This says exactly that $y_1, \ldots, y_r$ span M', so M' admits a spanning set of size r. This completes the proof of (1).

Now we take up (2) (so $M \neq 0$), again arguing by induction on r that $M \simeq A^{\oplus n}$ for some $n \leq r$ and that if $M \simeq A^{\oplus n'}$ then necessarily $n' = n$. (Once this is proved in general, if $M' \subset M$ is a nonzero submodule then M' inherits torsion-freeness from M and by (1) we know that M' is generated by n elements since $M \simeq A^{\oplus n}$ is generated by n elements. Thus, applying for M' the aspect of (2) we are temporarily taking as known it would follow that $M' \simeq A^{\oplus n'}$ for some $n' \leq n$.)

We will first prove existence of $n \leq r$ such that $M \simeq A^{\oplus n}$. Suppose $r = 1$, so every element of M is an A-multiple of some common element $x_1 \in M$. Necessarily $x_1 \neq 0$, so by the torsion-free hypothesis when we write a general $m \in M$ as ax_1 for some $a \in A$ then such a is *unique*. (Indeed, if $ax_1 = a'x_1$ for some $a, a' \in A$ then $(a - a')x_1 = 0$ yet $x_1 \neq 0$, so $a - a' = 0$ by torsion-freeness.) This says $M \simeq A$ as A-modules (using the basis $\{x_1\}$).

Suppose next that $r > 1$ and that it is known that every torsion-free A-module generated by $< r$ elements is isomorphic to $A^{\oplus e}$ with $e < r$. Pick a spanning set $\{x_1, \ldots, x_r\}$ of M. If the x_i's are linearly independent then $M \simeq A^{\oplus r}$ and we are done. Suppose instead that there is a nontrivial linear dependence relation: $\sum a_i x_i = 0$ with some $a_{i_0} \neq 0$. By relabeling the indices, we may assume $a_r \neq 0$. Let M' be the A-span of $x_1, \ldots, x_{r-1}$, so $a_r x_r \in M'$ and clearly $a_r x_i \in M'$ for all $i < r$, so a_r multiples every x_j into M' and thus multiples *every* element of M into M' (as everything in M is an A-linear combination of the x_j's). In other words, $x \mapsto ax$ is an A-linear

map $M \to M'$ and it is *injective* since $a \neq 0$ and M is torsion-free. Thus, this identifies M with an A-submodule of M'. But M' admits $r-1$ generators, so by (1) applied to M' and its submodule that is a copy of M we conclude that M admits $r-1$ generators! The inductive hypothesis then does the job, so we have shown that in general $M \simeq A^{\oplus n}$ for some $n \leq r$.

We have established the freeness asserted in (2), so for (2) it remains to show the *rank* of a free finitely generated A-module is well-defined (when A is a PID): if there exists an A-linear isomorphism $f : A^{\oplus n} \simeq A^{\oplus n'}$ for some $n, n' \geq 1$ then $n = n'$. If A is a field then we may use dimension considerations, so assume A is not a field. Thus, this PID contains a nonzero nonunit, so by the UFD property it contains an irreducible element π. The quotient ring $k := A/\pi A$ is then a field (why?). Reducing the A-linear isomorphism f modulo π gives a k-linear isomorphism $k^n \simeq k^{n'}$, so dimension considerations in linear algebra over k then force $n = n'$. This completes the proof of (2).

Finally, if a nonzero finitely generated torsion-free (hence free!) A-module M is contained in a d-dimensional K-vector space V with $d > 0$ and M spans V over K then we claim that M has rank equal to d. We know $M \simeq A^{\oplus r}$ for some $r \geq 1$, so M is spanned by r elements $x_1, \ldots, x_r$ linearly independent over A. These r elements then span V over K (as M spans V over K by hypothesis), so $d \leq r$. If $d < r$ then $x_1, \ldots, x_r$ must be linearly dependent over K inside V. But a nontrivial K-linear dependence relation can be scaled by a common denominator in the coefficients of such a dependence relation so that the coefficients are all in A. This is then a nontrivial A-linear dependence relation on the x_i's in M, but the x_i's were chosen to be linearly independent over A. Thus, $r = d$ as desired.