Lecture 10-25: Finitely generated modules over principal ideal domains

October 25, 2024

So far I have considered modules over very general rings. Now I will change the focus to modules over very particular rings, namely principal ideal domains, and show how finitely generated modules over them admit a very simple and elegant classification. In what follows I will cite some facts about principal ideal domains which I hope are familiar to most of you; in any event I will give references to the text.

Recall first that a principal ideal domain, or PID, is an integral domain R such that every ideal is principal, that is, generated by a single element (Definition, p. 279). The two most familiar examples are $\mathbb{Z}$ and the ring $k[x]$ of polynomials inone variable over a field k. In general, one says that a nonzero $x \in R$ is prime if given any factorization $x = yz$ either y or z is a unit in R (has a multiplicative inverse); see the Definition on p. 284. I will assume that you have seen the result that every PID R is a unique factorization domain, or UFD; that is, given a nonzero nonunit $x \in R$ one can write $x = p_1 \dots p_m$ as a finite product of primes p_i and that given any two such products $x = p_1 \dots p_m = q_1 \dots q_n$ one has $m = n$ and the q_i agree with the p_i up to reordering and multiplying by units (Theorem 14, p. 287). Also any two nonzero elements z, y of R have greatest common divisor $z = ax + by$ for some $a, b \in R$, so that the ideal (z) generated by z is the same as the ideal (x, y) generated by x and y (Proposition 6, p. 280).

Before stating the main result I need to generalize a piece of terminology used earlier for free modules. Given any integral domain I (not necessarily a PID), the rank of an I-module M is the maximum number n of linearly independent elements $m_1, \ldots, m_n$ of M, so that $\sum_{j=1}^{m} i_j m_j = 0$ for $i_j \in I$ only if $i_j = 0$ for all j (p. 460). By the same argument as for vector spaces over a field, any two maximal linearly independent subsets of M have the same cardinality I showed in the lecture on October 9 that a a free module of rank n over an integral domain continues to have rank n in this new sense (Proposition 3, p. 459).

Now let R be a PID. The main classification result follows from

Theorem 4, p. 460

Given any submodule N of $M = R^n$, there is a basis $y_1, \ldots, y_n$ of M and $a_1 \ldots, a_m \in R$ such that $m \leq n$, $a_1|a_2|\cdots|a_m$, and $a_1 y_1, \ldots, a_m y_m$ is a basis of N. In particular, N is free of rank at most n.

Proof.

Following the text, I argue by induction on the rank m of N. If $m = 0$ then we must have $N = 0$ since given a nonzero $n \in N$, the only $r \in R$ with $rn = 0$ is $r = 0$; so the result is clear. In general, for every R-homomorphism $\phi \in \hom_R(M, R)$, the image $\phi(N)$ is an ideal, which must be principal; say $\phi(N) = (a_\phi)$ for $a_\phi \in R$. If $N \neq 0$, then by looking at the projections of N to the coordinates of R one sees that $a = a_\phi \neq 0$ for some ϕ. Write $a = p_1 \ldots, p_m$ with the p_i prime in R. □

Proof.

Then a has only finitely many factors in R, up to multiplication by units, namely the products of some of the p_i. Equivalently, there are only finitely many ideals of R containing (a). It follows that the set Σ of all ideals (a_ϕ) as ϕ runs through $\hom_R(M, R)$ has a nonzero element not contained in any other, say $(a_\nu) = (a_1)$; one also has $a_1 = \nu(y)$ for some $y \in N$. Next I claim that a_1 divides $\phi(y)$ for all $\phi \in \hom_R(M, R)$. Indeed, if there is ϕ with a_1 not dividing $\phi(y)$, then the greatest common divisor d of a_1 and $\phi(y)$ generates a strictly larger ideal than (a_1). Writing $d = aa_1 + b\phi(y)$, one finds that the homomorphism $\psi = a\nu + b\phi$ takes the value d at y, whence the corresponding ideal (a_ψ) is larger than (a_1), a contradiction. In particular, looking at the coordinate projections, we see that we must have $y = (c_1, \ldots, c_n) \in M$ with $c_i = a_1 b_i$ for some $b_i \in R$. Setting $y_1 = (b_1, \ldots, b_n)$ we get $\nu(y_1) = 1$. □

Proof.

Given $m \in M$, setting $a = \nu(m)$, one has $m = ay_1 + m_1$, where $m_1 \in \ker \nu$. Hence M is the sum of Ry_1 and $\ker \nu$; it is easy to see that this sum is direct. Similarly N is the direct sum of Ra_1y_1 and $N \cap \ker \nu$. Thanks to the directness of these sums, the rank of $(N \cap \ker \nu) \subset M$ is less than that of N, so by inductive hypothesis it has a basis which extends to a basis of M. Adding ay_1 to this basis, we get a basis of N. I have shown in particular that every submodule of a free R-module of finite rank n is itself free of rank at most n; in particular, $\ker \nu$ is also free over R. □

Proof.

Now the induction hypothesis applies again to $N \cap \ker \nu \subset \ker \nu$. It yields a basis $a_2 y_2, \ldots a_m y_m$ of $N \cap \ker \nu$ such that $y_2, \ldots, y_n$ is a basis of $\ker \nu$ for some $n \geq m$ and $a_2, \ldots, a_m \in R$ satisfy $a_2 | \cdots | a_m$. Then $y_1, \ldots, y_n$ is a basis of M and $a_1 y_1, \ldots, a_m y_m$ is a basis of N; it only remains to show that $a_1 | a_2$. Define a homomorphism $\phi : M \to R$ via $\phi(\sum r_i y_1) = r_1 + r_2$. Since $a_1 y_1, a_2 y_2 \in N$, we have $(a_1) \in \phi(N), a_2 \in \phi(N)$. Since (a_1) is maximal among all ideals $\psi(N)$ as ψ ranges over $\hom_R(M, R)$, we must have $\psi(N) = (a_1), a_2 \in (a_1)$, and $a_1 | a_2$, as desired. □

Now I am finally ready to state the classification theorem.

Theorem 5 (1), p. 462

Any finitely generated module M over a PID R is isomorphic to a direct sum $R^r \oplus R/(a_1) \oplus \cdot R/(a_m)$ for some nonzero $a_1, \ldots, a_m \in R$ with $a_1 | \cdots | a_m$.

This follows at once from the preceding result since if M is generated by n elements we must have $M \cong R^n/N$ for some submodule N. In particular, observe that if M is generated by n elements, then in the statement of the theorem we must have $r + m \leq n$. Note also that a finitely generated R-module M is projective if and only if it is free, or if and only if it is torsion-free in the sense that $rm = 0$ for $r \in R, m \in M$ if and only if $r = 0$ or $m = 0$.

Actually there are two versions of the classification theorem; the one just given is called the invariant factor form, with the a_i being the invariant factors. To state the other version, I need a couple of simple facts about general commutative rings R. Let I, J be comaximal ideals in such a ring, so that by definition the sum $I + J$ is all of R. Then one has the Chinese Remainder Theorem (see p. 246), which states that the intersection $IJ = I \cap J$ and the quotient $R/(I \cap J)$ is isomorphic to the direct sum $R/I \oplus R/J$.

To prove the first assertion, note first that $IJ \subset I \cap J$ by definition; conversely, if $x \in I \cap J$ and $i \in I, j \in J$ satisfy $i + j = 1$ then $x = ix + jx = ix + xj \in IJ$. To prove the second assertion define $\phi : R \to R/I \oplus R/J$ via $\phi(r) = (r + I, r + J)$. Clearly the kernel of ϕ is $I \cap J = IJ$; to see that its image is all of $R/I \oplus R/J$, again choose $i \in I, j \in J$ with $i + j = 1$. Then $i + J = 1 + J, j + I = 1 + I$, so the image of ϕ contains $(1, 0)$ and $(0, 1)$ and thus the entire direct sum.

By repeatedly applying this theorem, one sees that if $r = p_1^{n_1} \dots p_m^{n_m}$ with the p_i distinct primes in R, then $R/(r) \cong \oplus_{i=1}^m R/(p_i^{n_1})$ (see p. 464).

Theorem 6, p. 464

Any finitely generated R=module M is isomorphic to a direct sum $R^r \oplus R/(p_1^{n_1}) \oplus \dots \oplus R/(p_m^{n_m})$, where the p_i are (not necessarily distinct) primes in R.

This follows at once from the previous theorem and the Chinese Remainder Theorem, writing each invariant factor a_i as a product of prime powers in R. This called the elementary divisor form or the primary decomposition (of M); the prime powers $p_i^{n_i}$ are the elementary divisors. Notice that there is no bound on the number m of factors required, even given the number of n of generators of M, since each quotient $R/(a_i)$ in the invariant factor form might be the sum of several quotients $R/(p_i^{n_i})$.

Even more is true:

Theorem 9, p. 466

The invariant factors a_i and elementary divisors $p_i^{n_i}$ of a finitely generated module M are unique up to multiplication by units; also any two invariant factor or elementary divisor decompositions of M involve the same number r of copies of R.

Given M, denote by $\text{Tor}(M)$ its torsion submodule (p. 459), consisting of all $m \in M$ such that $rm = 0$ for some nonzero $r \in R$. It is easily checked that this is indeed a submodule and is the sum of the proper quotients of R in any invariant factor or elementary divisor decomposition of M. Thus $M/\text{Tor}(M)$ is free of rank r equal to the number of copies of R in such a decomposition; since the rank of a free R-module is uniquely determined, so too is r.

I will prove the remaining uniqueness assertion only for the primary decomposition, leaving the other case as an exercise (or you can consult the proof in the text). For each fixed prime $p \in R$ and nonnegative integer m it suffices to show that the number of factors in M of the form $R/(p^m)$ for for a fixed prime power p^m occurring in any primary decomposition depends only on M; in turn for this it suffices to show that the number k of factors $R/(p^n)$ for some $n \geq m$ depends only on M. Letting $T = \text{Tor}(M)$, we note as a simple consequence of the Chinese Remainder Theorem that for any quotient $Q = R/(q^r)$ of R with q prime is such that $p^m Q/p^{m+1}Q = 0$ if q is not a unit multiple of p, or if q is a unit multiple of p and $r \leq m$, while if q is a unit multiple of p and $r > m$, we have $p^m Q/p^{m+1}Q \cong R/(p)$. Hence $p^m T/p^{m+1}T$ is a free $R/(p)$-module of rank equal to the number k defined above, and this number is indeed determined by M. □