Finitely Generated Modules over a PID, I

A will throughout be a fixed PID. We will develop the structure theory for finitely generated A-modules.

Lemma 1 *Any submodule $M \subset F$ of a free A-module is itself free, with $rank(M) \leq rank(F)$.* □

PROOF We prove the finite rank case $M \subset A^n$. For free modules of infinite rank, some set theoretic tool, like well-ordering a basis, is required. In the finite rank case, if $n = 1$ then any (nonzero) submodule of A is a principal ideal $(a) \cong A$. By induction, look at the exact sequence

$$0 \to A^{n-1} \xrightarrow{i} A^n \xrightarrow{p} A \to 0$$

where the map p is projection on the last coordinate. Restricting to $M \subset A^n$, we get

$$0 \to M \cap A^{n-1} \to M \to p(M) \to 0.$$

Then either $p(M) \cong A$ or $p(M) = 0$. By induction $M \cap A^{n-1}$ is free of rank $\leq n-1$. If $p(M) \cong A$, the second sequence splits, so M is free of rank $\leq n$. If $p(M) = 0$ then $M = M \cap A^{n-1}$. ■

If M is an A module, we define its *torsion submodule* to be $T(M) = T = \{t \in M \mid bt = 0, \text{ some } b \neq 0 \in A\}$. ($T$ is defined for left modules over any ring, and $t \in T$ just says $ann(t) \neq 0$. But to prove $T \subset M$ is a submodule you want the ring to be commutative and have no zero divisors.) We say M is *torsion free* if $T(M) = 0$.

Lemma 2 *A torsion-free, finitely generated A-module is isomorphic to a submodule of a free module, hence is free.* □

PROOF Say M is generated by $\{x_1, x_2, \cdots, x_s\}$. Choose a maximal linearly independent subset, which by reordering we may assume is $\{x_1, \cdots, x_d\}$. Then $M \supset N =< x_1, \cdots, x_d > \cong A^d$. By maximality of the linearly independent set, for $1 \leq i \leq s-d$ there exists $0 \neq a_i \in A$ with $a_i x_{d+i} \in N$. Let $a = \prod a_i$. The A module map $a : M \to M$ is injective, since M is torsion free. But obviously $aM \subset N$. ■

Lemma 3 *If M is a finitely generated A-module and $T \subset M$ is its torsion submodule, then $M/T = F$ is a finitely generated torsion-free, hence free, A-module and one has a direct sum decomposition $M \cong T \oplus F \cong T \oplus A^n$. Moreover, T is a finitely generated A-module, and the integer $n = rank(M/T)$ is uniquely determined by M.* □

PROOF If $t \in M$ and $\bar{t} \in M/T$ is a torsion element, then $bt \in T$ for some $b \neq 0$. This means $abt = 0 \in M$ for some $a \neq 0$. But $ab \neq 0$, so $t \in T$ hence $\bar{t} = 0 \in M/T$ and M/T is torsion free. Obviously this part of the argument works for any integral domain A. Always if M is finitely generated so is M/T. If A is a PID then from Lemma 2 we conclude that $F = M/T$ is free, and hence the sequence

$$0 \to T \to M \to F = M/T \to 0$$

splits. Then T is also a quotient of M, hence finitely generated. (Or, T is finitely generated since A is Noetherian.) ■

We now need to analyze a finitely generated, torsion module, T, over a PID A. We will do this by first using a decomposition of T into p-torsion summands, for primes $p \in A$, and then an analysis of a p-torsion module. In general, if $p \in A$ is a prime and T is a torsion A-module, set $T_p = \{x \in T \mid p^n x = 0, \text{ for some } n \geq 0\}$. If $x \in T$, set $ord(x) = (a)$ if $ann(x) = (a) \subset A$. Thus $\langle x \rangle \cong A/(a)$. If $x \in T_p$ then $ord(x) = (p^n)$ where n is least so that $p^n x = 0$.

Lemma 4 *For all torsion modules, $T \cong \oplus_{p\ prime} T_p$.*

PROOF If $x \in T$ and $dx = 0$, factor $d = \prod p_j^{n_j}$. Let $d_j = d/p_j^{n_j}$. Then the d_j have no common factors, so $1 = \sum a_j d_j$ for some $a_j \in A$. Then $t = \sum a_j d_j t = \sum t_j$ with $p_j^{n_j} t_j = 0$. This shows every element of T is a sum of elements of prime power order. The direct sum decomposition statement is equivalent to the statement that $T_p \cap \sum_{q \neq p} T_q = (0)$. But this follows easily from the fact that p^n is relatively prime to a product r of any other prime powers. Namely, write $1 = ap^n + br$. Then $p^n x = 0$ and $rx = 0$ implies $x = 1x = 0$. ∎

Remark 1 Lemma 4 is pretty much a module version of the Chinese Remainder Theorem. In fact, if $x \in T$ has $ord(x) = (d)$ then $\langle x \rangle \cong A/(d)$. If $d = \prod p_j^{n_j}$ then $A/(d) \cong \prod A/p_j^{n_j}$, as rings, so certainly as A modules. □

Now suppose T is a finitely generated p torsion module. Choose a finite set generators. Each one has order some power of p, so we see easily that for some $y \in T$ we have $ord(y) = (p^e)$ and $p^e T = 0$. Consider the exact sequence

$$0 \to \langle y \rangle \to T \to \overline{T} = T/\langle y \rangle \to 0.$$

Lemma 5 *If $\overline{x} \in \overline{T}$ has $ord(\overline{x}) = (p^f)$ then there exists $x \in T$ projecting to $\overline{x}$ with $ord(x) = (p^f)$.* □

PROOF Certainly $f \leq e$ since $p^e T = 0$. Choose any $z \in T$ projecting to $\overline{x}$. Then $p^f z = ay$. But also $0 = p^e z = p^{e-f} ay$. Since $ord(y) = (p^e)$, we must have $a = p^f b$. Set $x = z - by$. Then x projects to $\overline{x}$, so x has order at least p^f. But $p^f x = p^f z - p^f by = ay - ay = 0$. ∎

Theorem 1 *If T is a finitely generated nonzero p-primary torsion module, then*

$$T \cong A/(p^{e_1}) \oplus A/(p^{e_2}) \oplus \cdots \oplus A/(p^{e_m}),$$

where $0 < e_1 \leq e_2 \leq \cdots \leq e_m$. The exponents e_j are uniquely determined by T. The integer m is the least number of generators of T.

PROOF Say $T = \langle y_1, \cdots, y_m \rangle$ with m least. If $m = 1$ there is nothing to prove, T is cyclic. In general, assume e is least with $p^e T = 0$. We may as well assume y_m has order exactly (p^e), that is, $\langle y_m \rangle \cong (p^e)$. Consider the exact sequence

$$0 \to \langle y_m \rangle \to T \to \overline{T} = T/\langle y_m \rangle \to 0.$$

$\overline{T}$ can be generated by $m-1$ elements, (but no fewer). By induction, we know $\overline{T}$ has a decomposition as stated in the Theorem, with $m-1$ summands (but the $\overline{y_j},\ j < m$, are not necessarily independent cyclic generators),

$$\overline{T} \cong A/(p^{e_1}) \oplus A/(p^{e_2}) \oplus \cdots \oplus A/(p^{e_{m-1}}).$$

Certainly $e_j \leq e$ for $1 \leq j \leq m-1$ since $p^e T = 0$. We set $e_m = e$. Applying Lemma 5, we see the exact sequence above splits, and this establishes the decomposition statement of the Theorem for T.

Given such a decomposition, obviously m is the dimension of T/pT as vector space over $A/(p)$. Furthermore, the dimension of pT/p^2T over $A/(p)$ is the number of e_j which are greater than 1. In general, the dimension of $p^i T/p^{i+1}T$ is the number of e_j which are greater than i. These dimensions are invariant, and determine the e_j, hence the e_j are uniquely determined by T. ∎

Theorem 1, combined with Lemma 4, presents one normal form for a finitely generated torsion A-module T. The powers (p^{e_j}) which occur in the formula for T_p in Theorem 1, including the number of times each occurs, as p varies over prime divisors of $\text{ann}(T) = (d)$, are called **elementary divisors** of T. Here is a second normal form.

Theorem 2 *If T is a finitely generated torsion module over a PID A, then*

$$T \cong A/(d_1) \oplus A/(d_2) \oplus \cdots \oplus A/(d_m),$$

where the $d_j \in A$ are nonzero nonunits, and where $d_1 \mid d_2 \mid \cdots \mid d_m$. The integer m is the least number of generators of T. The ideals $(d_j), 1 \leq j \leq m$, are uniquely determined by T, with $(d_m) = ann(T)$. □

The ideals $(d_j) \subset A$ of Theorem 2 that successively divide each other are called the **invariant factors** of T.

Remark 2 It is quite easy to go back and forth between the invariant factor form and the elementary divisor form. Here is how the translation goes. Use the Chinese Remainder Theorem to convert an invariant factor formula for T, as in Theorem 2, to elementary divisor form. That is, if $d_i = \prod_j p_{ij}^{f_{ij}}$ is the factorization of d_i into distinct prime powers, then $A/(d_i) \cong \oplus_j A/(p_{ij}^{f_{ij}})$. Conversely, given the elementary divisor form of Theorem 1 for each T_p, p prime, reconstruct the invariant factors d_i as follows. The last $(d_m) = ann(T)$ must be the product of the prime powers (p^e), where $(p^e) = ord(T_p)$ is the highest prime power seen in the formula for T_p in Theorem 1. Then remove one cyclic summand of T corresponding to each of these highest prime powers, and look at the remaining summands. Apply the same recipe to these summands to construct (d_{m-1}). Namely, (d_{m-1}) must be the product of one each of the highest remaining prime powers. Namely, any other (d) would contribute the wrong prime power factors to ever recover the unique normal form of Lemma 4 and Theorem 1, using divisors of d as later d_j. Continue this algorithm to find all the d_i. Examining these two translations, recovering each normal form from the other, reveals that a uniqueness result for either normal form implies uniqueness for the other normal form. Thus, uniqueness of the invariant factors follows from the rather clean proof of uniqueness of the elementary divisors. □

Finitely Generated Modules over a PID, II

If M is any finitely generated module over a Noetherian ring R, there exist exact sequences

$$R^m \xrightarrow{\alpha} R^n \to M \to 0.$$

In terms of standard bases, we can represent the map α by an $n \times m$ matrix A over R. If we change bases in R^m and R^n then A is replaced by $B = PAQ$, where P is an invertible $n \times n$ matrix and Q is an invertible $m \times m$ matrix. But of course $M = R^n/\alpha(R^m)$ doesn't change. So a strategy for the existence part of the structure theorem when R is a PID is to find P and Q so that the matrix B is especially simple. In fact, this strategy proves more than just the structure theorem.

We first consider the ideals in R generated by the matrix entries a_{ij} of A and b_{ij} of $B = PAQ$.

Lemma 6 *$(a_{ij}) = (b_{ij})$ as ideals of R.* □

PROOF The formula for matrix multiplication shows $(b_{ij}) \subset (a_{ij})$. But P and Q are invertible, so the argument reverses. ■

Now we specialize to a PID R. Among the principal ideals generated by the upper left corner entries of all $B = PAQ$, as P and Q vary, choose a maximal member (b_{11}) of this family of ideals. The following is the key result.

Lemma 7 *$(b_{11}) = gcd(b_{ij}) = gcd(a_{ij})$.* □

PROOF We first claim that $b_{11} \mid b_{1j}$ and $b_{11} \mid b_{k1}$ for all j, k. To see this with $j = 2$, let $gcd(b_{11}, b_{12}) = rb_{11} + sb_{12}$ with $gcd(r, s) = 1$. So there exists a 2×2 matrix of determinant 1, $\begin{pmatrix} r & u \\ s & v \end{pmatrix}$. Let Q_0 be the block $m \times m$ matrix with this upper left 2×2 corner and with the $m - 2 \times m - 2$ identity matrix as lower right block. Then the $(1, 1)$ entry of BQ_0 is $rb_{11} + sb_{12}$, which generates a larger ideal than b_{11} unless $b_{11} \mid b_{12}$. For $j > 2$, we can first multiply B by a permutation matrix Q to switch the second and j^{th} columns, then repeat the argument here.

Operating on the left of B with similar block P_0 and with row permutation matrices, we get $b_{11} \mid b_{k1}$.

Next, we can use elementary row and column operations to replace B by some $C = PBQ$, with all first row and first column entries 0, except for b_{11}. Suppose the matrix C has entry $c = c_{jk}$ with $j, k > 1$ By either an elementary row or elementary column operation, we can put this entry on the first row or the first column, without changing the $(1, 1)$ entry b_{11}. Repeating the first part of the argument, we conclude $b_{11} \mid c$. Thus b_{11} is the *gcd* of all entries of C, which is the same as the *gcd* of all entries of our original A. ■

Now we can apply induction to the $n - 1 \times m - 1$ matrix obtained by removing the first row and column of the matrix C in the previous paragraph. Of course this first row and column only have the single non-zero entry b_{11}. The conclusion is the strong result known as "The Good Bases Theorem".

Theorem 3 *Suppose R is a PID and suppose $\alpha : R^m \to R^n$ is a (nonzero) R linear map. Then there is an integer $1 \leq s \leq min(m, n)$ and nonzero elements $d_1 \mid d_2 \mid \cdots \mid d_s \in R$ and bases $\{e_i\}$ of R^m and $\{f_j\}$ of R^n such that $\alpha(e_i) = d_i f_i$ if $1 \leq i \leq s$ and $\alpha(e_i) = 0$ if $i > s$. Consequently, the finitely generated R module $M = R^n/\alpha(R^m)$ is isomorphic to $R/(d_1) \oplus \cdots \oplus R/(d_s) \oplus R^{n-s}$.* □

PROOF The induction alluded to above the statement of the Theorem produces $B = PAQ$ with successively dividing main diagonal entries d_i, $1 \leq i \leq s$ and all other entries 0. Specifically, d_1 is $gcd(a_{ij})$ and then d_2 is the *gcd* of entries of an $n-1 \times m-1$ matrix, each of which belongs to the ideal (a_{ij}). Hence $d_1 \mid d_2$, and so on. With respect to the original coordinates R^m and R^n, the vectors e_i are the columns of Q and the vectors f_j are the columns of P^{-1}. But this specific detail is hardly necessary to see the truth of the Theorem from the diagonal form of PAQ. ∎

Remark 3 It is certainly possible that some $d_i = 1$. One just ignores these in the consequence about the structure of $M = R^n/\alpha(R^m)$. If one begins with an M then one theoretically could start with a minimal set of n generators, which implies no d_i is a unit. Also one can exploit the fact that $ker(R^n \to M)$ is actually free of rank $m \leq n$. So we could assume $\alpha : R^m \to R^n$ is injective, which implies $s = m$. But the theorem applies to all matrices, that is, to all finite presentations of R modules, hence is more powerful than the structure theorem by itself.

Remark 4 If R is a Euclidean Domain, like $\mathbb{Z}$ or $k[X]$, the general 2×2 blocks used in the proof of Lemma 7 to obtain *gcd*'s are not needed. One can just perform elementary row and column operations to a matrix A to bring it to the desired diagonal form. This is then very constructive, using repeatedly the division algorithm with quotients and remainders and the elementary row and column operations of Gauss-Jordan elimination, to arrive at the *gcd*'s in the final 'diagonal' matrix PAQ. This is one reason the proof given here of the Theorem for arbitrary PID's has some advantages over more abstract proofs. The method applied to Euclidean Domains is extremely explicit and constructive. □

Remark 5 If one thinks about minors in PAQ, one can see the following nice fact. With the d_j as in the statement of the Theorem, the product $d_1 d_2 \cdots d_k$ for each k is the *gcd* of all $k \times k$ minor determinants of the original matrix A. Specifically, one sees that just like the *gcd* of matrix entries, this minor determinant *gcd* does not change passing from A to PAQ. □

Remark 6 In order to get a uniqueness statement for the structure of M from Theorem 3, first note $n - s = rank(M/T)$ is uniquely determined. Let's assume we have arranged that no d_i are units. Then the integer n is the least number of generators of M and the integer s is the least number of generators of T. To see these two statements, choose any prime divisor p of d_1. Then M/pM is a vector space of dimension n over the field $R/(p)$, hence M cannot be generated by fewer than n elements. Similarly, s is the dimension of T/pT over $R/(p)$, so T cannot be generated by fewer than s elements.

The uniqueness of the ideals (d_i) can be proved in different ways. The uniqueness of the (d_i) based on the translation process between elementary divisors and invariant factors, given in Remark 2 above, is fairly easy to understand.

Here is a nice direct characterization of (d_i). Of course, $(d_s) = ann(T)$. For $i < s$ and $e \in R$, the module eT can be generated by $s - i$ elements if and only if d_i divides e. Thus, d_i is the gcd of all such elements e. The idea is, multiply all the summands of one fixed decomposition by e. You get something isomorphic to another sum of cyclic modules

$$eT \cong R/(e_1) \oplus R/(e_2) \oplus \cdots \oplus R/(e_s)$$

with $e_i \mid e_{i+1}$. Precisely, $e_i = d_i/\gcd(e, d_i)$. The number of *non-trivial* summands is then the least number of generators of eT. But a summand disappears, $(e_i) = (1)$, if and only if d_i divides e. □