Kiran Vodrahalli February 29, 2016

Contents

Today I will be discussing how to classify finitely generated R-modules for a special class of R: principal ideal domains. First we will present the form of these R-modules and a proof demonstrating how they may be classified. Then, we will introduce Smith Normal Form as an algorithmic approach to computing these decompositions. Finally, we will see some examples of applying this theorem to linear algebraic questions.

1 Preliminaries

First, let us define the two assumptions we make on the R-modules.

Definition 1.1. Principal Ideal Domain (PID).
A **principal ideal domain** is a **commutative** ring R in which every proper ideal is generated by a single element. Note that PIDs are Unique Factorization Domains, which are Commutative Rings. Also note that every Euclidean Domain is a PID, and thus every Field is a PID. Also note that every ideal in a PID is finitely generated.

Let us also recall the definition of finitely generated for R-modules:

Definition 1.2. Finitely Generated Modules.
An R-module is **finitely generated** if there is some n for which there is a surjection $R^n \to M$ for commutative R. In other words, there exist $m_1, \cdots, m_n \in M$ such that we can write $M = Rm_1 + \cdots + Rm_n$. Some examples:

- Finitely generated modules over $\mathbb{Z}$ are finite abelian groups.
- Finitely generated modules over a field $\mathbb{F}$ are finite-dimensional vector spaces. The elements of R are the "scalars" which send elements of M to other elements of M and the elements of M are the "vectors".

We will highlight the places these two necessary assumptions come into play.

Now we give a few more definitions for background.

We can extend the definition of finitely generated in the case where $R^n \to M$ is an isomorphism.

Definition 1.3. Free module.
The R-module is **free** if the map $R^n \to M$ is an isomorphism. A perhaps simpler way of thinking about free modules is that they are essentially a module which has a **basis**, in the usual sense from linear algebra. That is, a generating set consisting of linearly independent elements.

Let us prove that if there is an isomorphism from $R^n \to M$ then the there are generating elements which are a basis. We use the natural isomorphism, where we map tuples to a linear combination of the elementary generating elements of R^n. Consider that if $x = \lambda_1 x_1 + \cdots + \lambda_n x_n = 0$ for $\lambda_i \in R$, then $\lambda_i = 0$ for all i, otherwise, it's possible for two elements in R^n to be isomorphic to 0, which is impossible. Second, since we have an isomorphism, this map spans M. Both of these gives that $x_1, \cdots, x_n$ are a basis. Therefore, we can think of R^n and as a direct sum of isomorphic copies of R.

We also need the definition of prime element:

Definition 1.4. Prime element.
A **prime element** of R is a nonzero nonunit element p such that if $p|ab$ for $a, b \in R$, then $p|a$ or $p|b$. Note that an **irreducible element** of R is a nonzero nonunit element of R with no proper divisors. For the case where R is a PID, these definitions are equivalent, which we proved when we discussed UFDs. We say that two prime elements $p_1, p_2 \in R$ are **associates** if $p_1 = up_2$, where $u \in R$ is a unit. Note that being associates is an equivalence relation.

2 Classification Theorem

In a manner completely analgous to the Classification of Finite Abelian Groups, it turns out to be possible to break down finitely generated modules over PIDS into a simple direct sum format, which makes it easy to tell if two R-modules satisfying these properties are isomorphic.

Theorem 2.1. *Classification of Finitely-Generated R-Modules over PIDs.*
Let M be a finitely generated R-module. Then

$$M \cong R^b \oplus \bigoplus_{i=1}^{k} \left[\bigoplus_{j=1}^{l_i} R/\left(p_i^{t(i)_j}\right) \right]$$

*a direct sum of cyclic modules where p_i are primes, $t(i)_j$ are positive integers. b is called the **Betti number** and the powers of p_i are the **torsion coefficients**. Think of this as $t(i)$ is a vector of length l_i.*

Note that the Classification of Finite Abelian Groups is a direct consequence of our Classification Theorem:

Corollary 2.2. *Classification of Finite Abelian Groups. Let G be a finite abelian group. Then*

$$G \cong \bigoplus_{i=1}^{l} \mathbb{Z}/p_i^{t_i}\mathbb{Z}$$

where p_i are primes. We can equivalently write

$$G \cong \bigoplus_{j=1}^{k} \mathbb{Z}/d_{jj}\mathbb{Z}$$

where $d_{11}|d_{22}|\cdots|d_{kk}$ are uniquely determined by G.

We obtain this result by applying our theorem to the ring of integers, $R = \mathbb{Z}$, where modules are groups. Note that everything is in the torsion component. We will see the second statement in our discussion of Smith Normal Form.

2.1 Strategy

We give the following general strategy for classification. The classification proceeds by decomposing our R-module in various ways. First, we need some definitions to understand the steps.

Last lecture, we learned about torsion:

Definition 2.3. Torsion.
A module M is **torsion** when $M = M^{\text{tor}}$, where

$$M^{\text{tor}} = \{m \in M : \exists r \in R \text{ s. t. } rm = 0\}$$

A module is **torsion-free** when $M^{\text{tor}} = \{0\}$.

We now introduce a new concept which will be helpful in decomposing submodules which are torsion:

Definition 2.4. p-primary submodule.
We denote the **p-primary submodule** as $M[p^\infty]$ for prime element $p \in R$:

$$M[p^\infty] = \{m \in M : p^n m = 0, \text{ for some } n \in \mathbb{N}\}$$

Furthermore, each p-primary submodule is finitely generated. Essentially, we look at the set of all elements of M which get killed by p^i for any postive integral i. Note that in the case where $R = \mathbb{Z}$, the p-primary submodules correspond directly to the Sylow p-subgroups.

We now prove that $M[p^\infty]$ is a submodule:

Proof. First we show that $M[p^\infty]$ is an abelian subgroup. Since M is abelian, all elements commute. Then, consider $y_1, y_2 \in M[p^\infty]$. We need to show that $y_1 - y_2 \in M[p^\infty]$. Consider that by definition $p^{n_1}y_1 = 0$, $p^{n_2}y_2 = 0$. Then let $n = \max(n_1, n_2)$. We have $p^n y_1 = p^n y_2 = 0$, and thus $p^n(y_1 - y_2) = 0$, thus $y_1 - y_2 \in M[p^\infty]$. Then we show that $M[p^\infty]$ is closed under the ring action for any $r \in R$. Consider $y \in M[p^\infty]$, then $p^n y = 0$ for some n. Now consider ry. Noting that $p^n(ry) = r(p^n y) = 0$ since R is commutative, we get that $ry \in M[p^\infty]$. Thus, both submodule conditions are satisfied. □

The fact that $M[p^\infty]$ is finitely generated follows from the fact that R is a PID and M, our R-module, is finitely generated by assumption. We previously proved that every submodule of R can be generated by at most the same number of elements as M. Therefore, $M[p^\infty]$ is also finitely generated. Note that submodules in general are not necessarily finitely generated if their parent module is. We needed to use the fact that R was a PID and therefore Noetherian.

Then, our strategy is as follows:

1. **Torsion decomposition**: First, we will mod by M^{tor}, the torsion of the group, leaving just the free part of the group. Using short exact sequences we can decompose our R-module into a direct product of the free and torsion parts of the group, and we show that the ring structure is preserved. This is the R^b term.

2. **Primary decomposition**: We want to further break down the torsion part of the group in a manner analagous to applying the Chinese Remainder Theorem. We essentially want to factor the torsion into a direct sum of **primary submodules**. This step corresponds to the $\bigoplus_{i=1}^{k}$ direct sum, when there are k primary modules.

3. **Cyclic decomposition**: Now, since we can split up the torsion into primary submodules, we would like to break down primary submodules even further. Essentially, we can decompose our primary p-submodule into a direct sum of submodules each of which corresponds to a power of p. This final step corresponds to the $\bigoplus_{j=1}^{l}$ direct sum, where $p_i^{t_1}$ is the minimal annihilator for the p_i-primary module.

2.2 Torsion Decomposition

There turns out to be a nice result regarding free modules and torsion when R is a PID, which was proved already:

Lemma 2.5. *Torsion-free $\equiv$ free over PIDs.*
For modules over PIDs, M/M^{tor} is torsion-free and therefore a free submodule of M.

We also need the following lemma:

Lemma 2.6. *Splitting lemma.*
Let $\varphi : N \to R^n$ be a surjective free R-module homomorphism; then, there exists an R-module homomorphism $\phi : R^n \to N$ such that $\varphi \circ \phi = id$.

Proof. Let $e_i \in R^n$ be $(1, 0, 0, \cdots, 0), (0, 1, 0, 0, \cdots), \cdots, (0, 0, \cdots, 0, 1)$, where i denotes the position where there is a 1. Then these e_i generate R^n, so to construct ϕ it suffices to define the images of e_i (this is where we need that the R-module is free) and there are no restrictions on the image which need to be enforced. Since φ is surjective we are guaranteed that $e_i \in \text{Im}(\varphi)$. Take a representative element of $\varphi^{-1}(e_i)$, f_i. Then we define a homomorphism by $\phi(e_i) = f_i$, which is homomorphic since φ was homomorphic. □

Remark 2.7. Projective R-modules.
For any commutative ring R, a **projective** R-module has the property that every surjection φ to it admits a way to go back: I.e., the existence of ϕ.

Then, we have the following one-time decomposition

Theorem 2.8.

$$M \cong M^{\text{tor}} \times M/M^{\text{tor}}$$

where we split our module into torsion and free components.

As a brief example, we can think of the finite abelian group case as a situation where everying is in the torsion, and the vector space case as an example of no torsion. It turns out that groups over elliptic curves result in modules where we have both free and torsion parts. This situation can also arize for homology groups and fundamental groups in algebraic topology.

Let us prove that we can do this:

Proof. Consider the canonical surjection $\varphi : M \to M/M^{\text{tor}}$, which by the Splitting lemma admits $\phi : M/M^{\text{tor}} \to M$ since M/M^{tor} is a free module by the fact that torsion-free modules are free modules over PIDs. Now define $\alpha : M^{\text{tor}} \times M/M^{\text{tor}} \to M$ by

$$\alpha(m_1, m_2) = m_1 + \phi(m_2)$$

We want to prove that α is the desired isomorphism. We show α is an R-module homomorphism since ϕ is a homomorphism:

$$\alpha(m_1+n_1, m_2+n_2) = (m_1+n_1)+\phi(m_2+n_2) = (m_1+\phi(m_2))+(n_1+\phi(n_2)) = \alpha(m_1, m_2)+\alpha(n_1, n_2)$$

We also verify that the R-action works properly: Take $r \in R$:

$$r \cdot \alpha(m_1, m_2) = r \cdot (m_1 + \phi(m_2)) = r \cdot m_1 + r \cdot \phi(m_2) = r \cdot m_1 + \phi(r \cdot m_2) = \alpha(r \cdot m_1, r \cdot m_2)$$

as desired. Next we show that α is surjective. Take any $m \in M$. Consider that $\alpha(m - \phi(\varphi(m)), \varphi(m)) = m - \phi(\varphi(m)) + \phi(\varphi(m)) = m$, and since $m - \phi(\varphi(m)) \in M^{\text{tor}}$ (because $\varphi(m - \phi(\varphi(m))) = 0$ implies that $m - \phi(\varphi(m))$ is in the kernel of φ which is precisely M^{tor}) this is well defined. Thus $\text{Im}(\alpha) = M$. Now we prove that α is also injective. To do this, we simply need to show $\text{Ker}(\alpha) = \{(0, 0)\}$. We have $\text{Ker}(\alpha) = \{(m_1, m_2) : m_1 + \phi(m_2) = 0\}$. By applying φ, we get $\varphi(m_1) + \varphi(\phi(m_2)) = \varphi(m_1) + m_2 = 0$. But then note that m_1 is in the kernel of φ since m_1 is torsion! Thus we have that $m_2 = 0$. Then we also have $\phi(m_2) = 0$ so necessarily, $m_2 = 0$ since ϕ is a homomorphism. Thus $\text{Ker}(\alpha) = \{(0, 0)\}$ as desired and α must be injective.

Thus α is a bijective homomorphism, and is therefore an isomorphism. □

As we proved before, since M is a finitely-generated R-module over a PID, all its submodules are also finitely generated. Futhermore since M/M^{tor} is in addition a free module, it is isomorphic to R^b for some unique positive integer b where $b \leq n$ by the definition of free. Thus, we can now write

$$M \cong R^b \times \left[M^{\text{tor}}\right]$$

It remains to decompose the torsion portion in brackets.

Remark 2.9. Exact Sequences and Short Exact Sequences.
A lot of the material in this section is highly connected with homological algebra, which is heavily used in algebraic topology and category theory. Maps are the main objects of study in this area of math. For instance, the notion of projective module is actually generalizable to the notion of **covering spaces** in algebraic topology, and ϕ, as we used it, is analagous to a lift.

Exact sequences are sequences of homomorphisms

$$M_1 \to M_2 \to M_3 \to \cdots \to M_n$$

where the image of one map is the kernel of the next. In laymen terms, we keep sending the results of one map to die under the next map, which is pretty cruel.

Short exact sequences are exact sequences of the form

$$0 \to A \to B \to C \to 0$$

Let's write out each of the homomorphisms, their images, and kernels: $(0 \to A)$ has image 0 in A. $(A \to B)$ has kernel 0, and some $\mathrm{Im}(A) \subseteq B$. $(B \to C)$ has kernel $\mathrm{Im}(A)$, and image $\mathrm{Im}(B/\mathrm{Im}(A)) \subseteq C$. Finally, $(C \to 0)$ has kernel $\mathrm{Im}(B/\mathrm{Im}(A))$, and image 0, which implies that

$$C/\mathrm{Im}\left(B/\mathrm{Im}(A)\right) \cong 0$$

which implies

$$C \cong \mathrm{Im}\left(B/\mathrm{Im}(A)\right)$$

Since the image is contained in C, we must have

$$C \cong B/\mathrm{Im}(A)$$

We can explain the relationship between M, M^{tor}, and M/M^{tor} in terms of a short exact sequence. Namely,

$$0 \to M^{\mathrm{tor}} \to M \to M/M^{\mathrm{tor}} \to 0$$

2.3 Primary Decomposition

We want to be able to further decompose the torsion part of M, M^{tor}. Fortunately, we can break up any torsion submodule into a direct sum of p-primary components. We can think of this as essentially "factoring into primes" in the manner of the Chinese Remainder Theorem (Recall: $\mathbb{Z}/n\mathbb{Z} \cong \bigoplus_{i=1}^{k} \mathbb{Z}/p_i^{m_i}\mathbb{Z}$ where $n = p_1^{m_1} \cdots p_k^{m_k}$). If we think to the case $R = \mathbb{Z}$, what we're saying is we can decompose any finite abelian group into a direct sum of Sylow p-subgroups. Now that we have some intuition, let's proceed to write the main theorem of this section:

Theorem 2.10. *For R a PID, every torsion R-module M^{tor} is a $\bigoplus$ (direct sum) of its p-primary submodules. When R is also finitely-generated, this direct sum is finite:*

$$M^{\mathrm{tor}} \cong \bigoplus_{i=1}^{k} M[p_i^{\infty}]$$

Proof. For convenience, denote $\tilde{M} := M^{\text{tor}}$ for this section.

Essentially, since M is Noetherian, the submodules of M are finitely generated (this follows since M is over a PID). Note that M being finitely generated does not imply M is Noetherian; consider $R = \mathbb{C}[x_1, \cdots]$, which when we choose $M = R$ as well, we see that it's a free module! But, it's not Noetherian. We also know that if we can express $M = \bigoplus M_p$ then since M is Noetherian, this is necessarily a finite sum. Suppose this were not the case and there were infinitely many p: Then, $M[p_1^\infty] \subsetneq M[p_1^\infty] \oplus M[p_2^\infty] \subsetneq M[p_1^\infty] \oplus M[p_2^\infty] \oplus M[p_3^\infty] \oplus \cdots$ and so on. By using the direct sum notation, we're indicating that each term in the sum is unique, as we will later on see. This construction gives an infinite ascending chain, contradicting the Noetherianness of M. Finally, we also want to note that expression M as a direct sum of p-primary modules induces the notion that the p-primary submodules are distinct, which we prove by the finite intersection property of the components of direct sums later on. Along with our ability to uniquely factor elements of R (since it is a PID which is a UFD), we can directly correspond the p-primary submodules with the prime powers of the factorization, as we will soon see.

Let us enumerate the equivalence classes of primes via association as Λ. Then choose p_λ for each class $\lambda \in \Lambda$ (using Axiom of Choice, which isn't necessary for finitely generated modules). Then we take the direct sum $\bigoplus_{\lambda\in\Lambda} \tilde{M}[p_\lambda^\infty]$ (equivalently, the set of elements $m_{\lambda_1} + \cdots + m_{\lambda_k}$ with $m_{\lambda_i} \in \tilde{M}[p_{\lambda_i}^\infty]$). We claim that

$$\tilde{M} \cong \bigoplus_{\lambda\in\Lambda} \tilde{M}[p_\lambda^\infty]$$

The definition of direct sum gives that we must show that first, $\tilde{M}[p_\lambda^\infty]$ generate all of $\tilde{M}$ ($\tilde{M} = \sum_{\lambda\in\Lambda} \tilde{M}[p_\lambda^\infty]$, which is possible to write down with finite terms by the definition of finitely generated) and second, that $\tilde{M}[p_{\lambda_0}^\infty] \bigcap \sum_{i=1}^k \tilde{M}[p_{\lambda_i}^\infty] = 0$ for any $\lambda_0, \cdots, \lambda_k \in \Lambda$ (the $\tilde{M}[p_\lambda^\infty]$ are linearly disjoint).

Consider any $m \in \tilde{M}$. Since M is finitely generated, as we have said, we can generate M with a finite set $\{m_1, \cdots, m_k\}$. Since $\tilde{M}$ is torsion, we have $\exists r_0 \in R$ such that $r_0 m = 0$ for every m, where we can take $r_0 = \text{lcm}(r_1, \cdots, r_k)$ with $r_i m_i = 0$. Then since PIDs are UFDs, we can uniquely factor

$$r_0 = u p_{\lambda_1}^{\alpha_{\lambda_1}} p_{\lambda_2}^{\alpha_{\lambda_2}} \cdots p_{\lambda_k}^{\alpha_{\lambda_k}}$$

where u is a unit and α_{λ_i} are positive integers. These primes are the only ones necessary for the primary decomposition, and we see that we only have a finite number of them.

We associate $p_{\lambda_i}^{\alpha_{\lambda_i}}$ as the power of p_{λ_i} which kills everything in $\tilde{M}[p_{\lambda_i}^\infty]$. Then, define $q_i = r_0/p_{\lambda_i}^{\alpha_{\lambda_i}}$. Note that the ideal $I = (q_1, \cdots, q_k) = R$ since R is a PID, implying that $I = (a)$ is generated by one element. Then, $a|q_i$ for all i for $(q_i) \subsetneq (a)$. By unique factorization, we must have $a = 1$ since $\gcd(q_1, \cdots, q_k) = 1$ due to primality and irreducibility being identical. Thus, $I = (1) = R$. Therefore, we can take $r_1, \cdots, r_k \in R$ such that

$$q_1 r_1 + \cdots + q_k r_k = 1$$

Letting $m_{\lambda_i} = q_i r_i m$, we first verify that $p_{\lambda_i}^{\alpha_i} m_{\lambda_i} = r_0 r_i m = r_i(r_0 m) = 0$ and thus $m_{\lambda_i} \in \tilde{M}[p_{\lambda_i}^\infty]$. Then we see

$$m(q_1 r_1 + \cdots + q_k r_k) = m q_1 r_1 + \cdots + m q_k r_k = m$$

and thus

$$m = \sum_{i=1}^{k} m_{\lambda_i}$$

Now we show linear disjointness to finish the proof that $\tilde{M}$ is expressible as a direct sum of p-primary modules. Denote the intersection of interest as $\bigcup$. Suppose there were $x \in \bigcup$. Then, $g_1 x = p_{\lambda_0}^{\lambda_0} x = 0$, $g_2 x = \left(\prod_{i=1}^{k} p_{\lambda_k}^{\alpha_k}\right) x = 0$, by definition of the p_{λ_i}-primary modules. However, these two elements generate all of R as we previously saw. Thus, some R-linear combination $Rg_1 + Rg_2 = R$ and thus $r_1 g_1 + r_2 g_2 = 1$. But $1 \cdot x = (r_1 g_1 + r_2 g_2)x = r_1 g_1 x + r_2 g_2 x = 0 + 0 = 0$. Thus $x = 0$ only. □

It remains to decompose the p-primary submodules even further.

2.4 Cyclic Decomposition

In the last phase, we decompose the p-primary submodules further. Denote M_p as a p-primary submodule for convenience. Recall from earlier that M_p is a finitely generated submodule since M is finitely generated and R is a PID.

We state our main theorem for this section directly:

Theorem 2.11. *For R a PID and $p \in R$ prime, let M_p be a finitely-generated p-primary R-submodule. Let t be the smallest integer such that p^t annihilates all $m \in M_p$. Choose $a_t \in M_p$ so that $p^t a_t = 0, p^{t-1} a_t \neq 0$. Then, there is a submodule $N_t \subsetneq M$ such that*

$$M_p \cong N_t \oplus Ra_t$$

Proof. Again, we need to show that the terms N_t, Ra_t are linearly disjoint and second that they generate M_p. That is, we must show $N_t \cap Ra_t = 0$ and $N_t + Ra_t = M_p$.

Take a maximal submodule N of M_p satisfying linear disjointness. Note that picking a maximal submodule is possible since R is a PID and therefore finitely generated; for any one of its ideal chains simply pick a maximal ideal which satisfies disjointness (this exists as long as Ra_t does not generate all of M_p). Now suppose for the sake of contradiction that $N + Ra_t \subsetneq M_p$ (is a strict subset). Take $b \in M_p \setminus (N + Ra_t)$. Then for some $s > 0$ an integer, $p^s b \in N + Ra_t$ since b is still in a p-primary submodule. Note that for $s = t$, $p^s b = 0 \in N + Ra_t$. Also note that if $p^s b \in N + Ra_t$, then $p^{s'} b \in N + Ra_t$ for all $s' \geq s$ since N and Ra_t are both closed under multiplication by external ring elements. Therefore, let us rdefine b to be the first $p^k b$ such that $p^k b \notin N + Ra_t$, which must exist since there is a transition. Choosing $k = s - 1$, we write $b := p^{s-1} b$; then, $pb \in N + Ra$ but $b \notin N + Ra$. Let $pb = n + ra$ for $n \in N$ and $r \in R$. Now, multiplying by p^{t-1} gives $0 = p^{t-1} n + p^{t-1} ra$. Thus we see $p^{t-1} ra = -p^{t-1} n$, which shows that $p^{t-1} ra = -p^{t-1} n \in N$ since N is a submodule and is therefore closed under multiplication by ring elements. We also see that $p^{t-1} ra \in Ra$ for the same reason. Thus $p^{t-1} ra \in N \cap Ra_t$ and by assumption must be 0. Then recall that a is an element annihilated by p^t but not p^{t-1}, thus $r = pr'$ for some $r' \in R$, since we can write $p^{t-1} pr' a = r'(p^t a) = 0$. Thus, $n = (n + ra) - ra = pb - pr'a = p(b - r'a)$. Thus, $p(b - r'a) \in N$. Let $c = b - r'a$, then we have $pc \in N$. Now consider $N' = N + Rc$. Since $b \notin N + Ra$, $c = b - r'a \notin N$. Therefore, $N + Rc \not\subset N$, but $N + Rc \subsetneq N'$, and we have constructed a superset of N: $N \subsetneq N'$. However, we assumed that N was a maxmial submodule satisfying linear disjointness. Therefore, N'

must not be linearly disjoint: $N' \cap Ra \neq \{0\}$. Let us exhibit such an element in terms of both of the sets: Take $z \in N, r_1, r_2 \in R$ such that $z + r_1 c \in N'$, $r_2 a \in Ra$, and additionally $z + r_1 c = r_2 a \neq 0$. Then $z + r_1(b - r'a) = r_2 a$, or $r_1 b = -z + (r_1 r' + r_2)a \in N + Ra$. Therefore, $r_1 c = r_1 b - (r_1 r')a \in N + Ra$. Since we showed before that $pc \in N$ and $N \cap Ra = \{0\}$, we must have that $pc \notin Ra$ or $pc = 0$. If $pc \notin Ra$, then $pc \neq z + r_1 c \in Ra$. Thus, $z \neq (p - r_1)c$, so $(p - r_1)c \notin N$. If r_1 is a power of p, we get a contradiction since $pc \in N$. Thus we must have $p \nmid r_1$. Let $x, y \in R$ be elements such that $px + r_1 y = 1$ (Bézout's Lemma holds in PID: This is clear from the fact that each ideal is generated by one element, so a linear combination of relatively prime elements must live in $(1) = R$, the ideal generating the whole space). Multiplying by c, we get $pcx + r_1 cy = c$. Recall $pc \in N$, thus, $pcx \in N$ by closure property of submodules. Since we also have $r_1 c \in N + Ra$, $r_1 cy \in N + Ra$ because of closure of N and commutativity of R: $r_1 c = \tilde{n} + \tilde{r}a \implies r_1 cy = y\tilde{n} + (\tilde{r}y)a \in N + Ra$. Therefore $c \in N + Ra$. But then $b - r'a \in N + Ra$, so we have $b \in N + Ra$, contradiction. In the case $pc = 0$, we have $r_1 y = 1$ which implies that r_1 is a unit, and therefore that since $r_1 b \in N + Ra$, $y r_1 b = b \in y(N + Ra) \implies b \in N + Ra$, a contradiction.

Therefore, $N + Ra_t = M_p$ (is not a subset) and by assumption are linearly disjoint. Thus by definition of direct sum, we have

$$M_p \cong N_t \oplus Ra_t$$

as desired.

□

Now what we're going to do is repeatedly apply the previous theorem to break off chunks Ra_t one at a time. We have:

$$\begin{aligned} M_p &\cong N_1 \oplus Ra_1 \\ M_p &\cong (N_2 \oplus Ra_2) \oplus Ra_1 \\ M_p &\cong ((N_3 \oplus Ra_3) \oplus Ra_2) \oplus Ra_1 \\ &\cdots \\ M_p &\cong \bigoplus_{j=1}^{l} Ra_j \end{aligned} \tag{1}$$

We can do this since M_p is finitely generated (we showed before that all submodules will also be p-primary and finitely generated). Thus l must be finite. Now, assuming that a_j is annihilated by p^{t_j}, we can write

$$M_p \cong \bigoplus_{j=1}^{l} R/p^{t_j}$$

2.5 Concluding the Proof

If we now plug the cyclic decomposition into the torsion and primary decompositions, we recover the Structure Theorem for Modules over PIDs:

$$\begin{aligned} M &\cong M/M^{\text{tor}} \oplus M^{\text{tor}} \\ M &\cong R^b \oplus \left[M^{\text{tor}}\right] \\ M &\cong R^b \oplus \bigoplus_{i=1}^{k} M[p_i^{\infty}] \\ M &\cong R^b \oplus \bigoplus_{i=1}^{k} \left[\bigoplus_{j=1}^{l} R/p_i^{t(i)_j}\right] \end{aligned} \tag{2}$$

3 Smith Normal Form: An Algorithmic Approach to Proving Classification

Now that we know finitely generated R-modules over PIDs decompose in the way given by the Structure Theorem, we might want to know how to actually find such a decomposition. In this section, we prove the Structure Theorem via a different approach which gives such an algorithmic decomposition for free.

First, we need some new terminology. We want a way to talk about equivalent modules in linear algebraic terms:

Definition 3.1. Relation matrix.
Let R be a principal ideal domain and let M be a finitely generated R-module. If $\{m_1, \cdots, m_n\}$ is a set of generators of M, then we have a surjective R-module homomorphism $\varphi : R^n \to M$ by sending $(r_1, \cdots, r_n) \to \sum_{i=1}^{n} r_i m_i$. Let $K = \ker(\varphi)$. Then by the first Isomorphism Theorem, we have

$$M \cong R^n/K$$

Note that if $(r_1, \cdots, r_n) \in K$, we have a relation on the generators given by

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

K is finitely generated since it is a submodule of R^n, which is finitely generated since we are operating over PIDs. Let $k_1, \cdots, k_m \in K$ be a generating set for K. Then, for each k_j, we have

$$\sum_{i=1}^{n} \alpha_{j,i} m_i = 0$$

where $k_j = (\alpha_{j,1}, \alpha_{j,2}, \cdots, \alpha_{j,n})$. Define the **relation matrix** to be

$$A = \begin{bmatrix} \alpha_{1,1} & \alpha_{1,2} & \cdots & \alpha_{1,n} \\ \alpha_{2,1} & \alpha_{2,2} & \cdots & \alpha_{2,n} \\ \vdots & & \ddots & \vdots \\ \alpha_{m,1} & \alpha_{m,2} & \cdots & \alpha_{m,n} \end{bmatrix}$$

However, these generator matrices are not unique for module M and relation submodule K. It turns out that we can argue that multiplying on the left and right by elementary row and column operations P, Q yields

$$B = PAQ$$

where B is also a relation matrix for module M and a corresponding K, for a different set of generators of M and K. We can think of P as a "change of basis" matrix for the generators of K, and Q as a "change of basis" matrix for the generators of M. Note that these are not actually bases, unless the module is free.

Now, how can we relate relation matrices with the Structure Theorem?

Theorem 3.2. *Decomposition via Matrix.*
Suppose A is a relation matrix for R-module M. If there are invertible P, Q for which $B = PAQ$ is a diagonal matrix with diagonal elements $a_1, a_2, \cdots, a_n$, then

$$M \cong R/(a_1) \oplus R/(a_2) \oplus \cdots \oplus R/(a_n)$$

Proof. PAQ is a relation matrix for generating set $(m_1, \cdots, m_n)$. If $\varphi : R^n \to M$ is the natural homomorphism, then K is kernel of φ and $M \cong R^n/K$ as before. But K is also the kernel of the homomorphism

$$R^n \to R/(a_1) \oplus \cdots \oplus R/(a_n)$$

achieved by sending

$$(r_1, \cdots, r_n) \to (r_1 + (a_1), \cdots, r_n + (a_n))$$

Thus

$$M \cong R^n/K \cong R/(a_1) \oplus R/(a_2) \oplus \cdots \oplus R/(a_n)$$

□

We can also think of this proof in terms of exact sequences.

Proof. Essentially, we have

$$R^m \to^B R^n \to M \to 0$$

And we note that the image of B $\mathrm{Im}(B)$ is the kernel of the natural homomorphism between R^n and M, since B is just another representation of A (recall that A is not unique). Then we see that

$$M \cong R^n/\mathrm{Ker}(\varphi) \cong R^n/\mathrm{Im}(B)$$

Since B is a diagonal matrix, the image of B is just a direct sum of the images of each column: $Ra_1 \oplus Ra_2 \oplus \cdots Ra_n$. But each Ra_i by definition is (a_i). Thus

$$M \cong R^n/\mathrm{Im}(B) \cong R^n/\left(R(a_1) \oplus \cdots \oplus R^n/(a_n)\right) \cong R/(a_1) \oplus \cdots \oplus R/(a_n)$$

Note that we allow a_i to be 0 (this corresponds to matrices where the diagonal elements are not all non-zero). Then, we see the free terms of the decomposition emerge: Quotienting by (0) just gives $R/(0) = R$, where is a free term of the direct sum. □

Now all we have to do is prove that such a matrix decomposition for generator matrices always exists over R a PID. It turns out that this is true and is called the Smith Normal Form.

Theorem 3.3. *Smith Normal Form.*
Suppose we consider a generator-relation matrix A, an $m \times n$ matrix over R, a PID. Then, we can write

$$A = PDQ$$

where P is $m \times m$, D is $m \times n$ and diagonal, and Q is $n \times n$. Futhermore, P, Q are constructed of row operation matrices (i.e. swap two rows or columns, multiply a row or column by an element of R, or permutation matrix). D is special because

$$D = \begin{bmatrix} d_{11} & 0 & 0 & & \cdots & & 0 \\ 0 & d_{22} & 0 & & \cdots & & 0 \\ 0 & 0 & \ddots & & & & 0 \\ \vdots & & & d_{rr} & & & \vdots \\ & & & & 0 & & \\ & & & & & \ddots & \\ 0 & & & \cdots & & & 0 \end{bmatrix}$$

where $d_{11} | d_{22} | d_{33} | \cdots | d_{bb}$ where d_{ii} are unique up to multiplication by units and are called the ***invariant factors****, also known as* ***elementary divisors****.*

We see how to construct the SNF of a relation matrix via the following algorithm.

3.1 The SNF Algorithm

The algorithm takes advantage of the fact that Bézout's Lemma holds in PIDs, as we showed before, and also of the fact that PIDs are UFDs: Thus, the Euclidean division algorithm will work in a PID.

Essentially, our approach is as follows: Take the smallest entry and using row operations, move it to the top left corner. Apply the Euclidean Algorithm over R to make the first column all 0, except for the top element. Repeat for the first row. If any divisors remain in the body (after first row and column), swap that element into the top left and repeat. Repeat until all elements in the body are larger than the top left element. Now repeat the whole process with the subblock. We will be able to finish since our R-module is Noetherian. Then we can think of D as a linear map between R modules, where one is isomorphic to R^m and the other to R^n. An easier approach computationally may be to just diagonalize the matrix first, and then fix each term row and column by row and column so that the divisibility property holds.

For ruther details, check the Wikipedia page on Smith Normal Form.

3.2 Decomposing with the SNF Algorithm

In particular, this algorithm, applied to the torsion M^{tor} portion of M, allows us to write

$$M^{\text{tor}} \cong \bigoplus_{a=1}^{r} R/(d_{aa})$$

where the d_{aa} are invariant factors. Note that we can recover the original statement of the Structure Theorem by decomposing each d_{aa} into powers of prime ideals. Smith Normal Form can be applied to compute homology of a chain complex in algebraic topology, when the chain modules are finitely generated. The moral of this story is that any operations you can do over integers, you can do in PIDs, including this particular matrix decomposition.