M3P8 LECTURE NOTES 7A: MODULES OVER EUCLIDEAN DOMAINS

1. Smith normal form

Given a presentation of a module M over a ring R, we have seen that we can try to simplify the presentation of M by performing elementary row and column operations on the presentation matrix of M. This is quite difficult to do in a systematic way in general- one needs assumptions on the ring R.

The simplest case is when R is a field; then elementary row and column operations suffice to put any matrix into a form whose entries are only 1 and 0, and whose only nonzero entries are a consecutive set of entries on the diagonal. Of course, doing this requires division, so we can't expect to do this for rings that are not fields.

The next simplest case is when R is a principal ideal domain. Here we make the following definition:

Definition 1.1. Let R be a principal ideal domain and A an n by m matrix with entries in R. We say A is in *Smith Normal Form* if the only nonzero entries of A are diagonal entries $a_{i,i}$, and for each i, $a_{i,i}$ divides $a_{i+1,i+1}$.

When R is a PID, it is possible to put any matrix A into Smith normal form by elementary row and column operations. In general this is somewhat hard to prove, and we will not do so here. Instead, we will focus on the somewhat easier case of a Euclidean domain R. The existence of a Euclidean norm makes it more straightforward to describe an algorithm for putting A into Smith normal form, although the algorithm we describe will be far from optimal.

The algorithm is based on the following lemma:

Lemma 1.2. *Let A be a matrix with at least one nonzero entry. Then, after elementary row and column operations, we can place A in a form where the upper left entry a_{11} is the only nonzero entry in the first row or column, and a_{11} divides every entry of A.*

Proof. We describe an algorithm to produce such a matrix from A via row and column operations. This algorithm proceeds in two phases:

In the first phase we begin by exchanging rows and columns so that the upper left entry a_{11} is the nonzero entry of A with smallest norm. Then, adding multiples of the first row and column to each other row and column, we can arrange that every entry of the first row and column has norm strictly smaller than a_{11}. If all the entries in the first row and column are now zero, we move to the next step; otherwise we again exchange rows and columns

so that a_{11} is the nonzero entry of A with smallest norm, and return to the start of phase one. Each time we do this the norm of a_{11} gets smaller, so this eventually terminates; when it does a_{11} is the only nonzero entry in the first row or column. At this point we proceed to the second phase.

In the second phase we suppose we have a nonzero a_{ij} element of A that is not divisible by a_{11}. (If this does not happen, we are done!) Adding the ith column (which contains a_{ij} to the first column does not change a_{11} (since the first row is zero except for a_{11}.) We then repeat the entire first phase. Notice that when we do this the norm of a_{11} has decreased; in particular the final value of a_{11} must divide the element a_{ij} we began this process with. We continue to find nonzero elements a_{ij} not divisible by a_{11}, adding their column to the first column, and repeating the first phase. Since the second phase only terminates when every entry of A is divisible by a_{11} and the first column and row are zero except for a_{11}, the result follows. □

To construct the Smith normal form of A we now proceed inductively: applying the lemma, we arrive in a situation where the first row and column of A are nonzero except for the entry a_{11}, and a_{11} divides every element of the submatrix B obtained from A by deleting the first row and column. By the inductive hypothesis we can put B in Smith normal form, so we are done.

2. The classification

Let R be a Euclidean domain, and let M be a finitely generated R-module. We can then find a presentation matrix A for M. Moreover, since changing a presentation matrix by elementary row and column operations gives rise to different presentations for the same module M, the results of the previous section show that M has a presentation by a matrix A in Smith Normal Form. Let $a_1, \dots, a_n$ be the nonzero diagonal entries of A, *excluding those that are units*, and let r be the number of zeroes in A. Then M is generated by $n + r$ elements $x_1, \dots, x_n, y_1, \dots, y_r$ subject to the relations $a_i x_i = 0$ for all i, with no other relations. In particular M is isomorphic to the module

$$R^r \oplus r/\langle a_1 \rangle \oplus R/\langle a_2 \rangle \oplus \cdots \oplus R/\langle a_n \rangle.$$

(Note that if we had included units among the a_i, this isomorphism would still hold, but some of the summands would be the zero module.)

We have thus shown:

Theorem 2.1. *Let M be a finitely generated module over a Euclidean domain R. Then there exists an integer r, and nonunit elements $a_1, \dots, a_n$ of R, such that a_i divides a_{i+1} for all i, and an isomorphism:*

$$M \cong R^r \oplus r/\langle a_1 \rangle \oplus R/\langle a_2 \rangle \oplus \cdots \oplus R/\langle a_n \rangle.$$

This result is called the *classification of finitely generated modules over a Euclidean Domain*. It holds as well for principal ideal domains, but we will not prove this here. A natural question to ask is how unique this

classification is. Since the integer r in the above theorem is equal to the rank of M, it does not depend on the choice of presentation of M. The elements $a_1, \ldots, a_n$ are called the *invariant factors* of M. They are only well-defined up to multiplication by units, but note that for any positive integer m, and any prime p of R, we have:

- $p^{m-1}R/p^mR$ is a one-dimensional $R/\langle p\rangle$-vector space,
- if $p^m|a$, then $p^{m-1}(R/\langle a\rangle)/p^m(R/\langle a\rangle)$ is a one-dimensonal $R/\langle p\rangle$-vector space, and
- if p^m does not divide a, then $p^{m-1}R/\langle a\rangle = p^mR/\langle a\rangle$, so $(p^{m-1}R/\langle a\rangle)/(p^mR/\langle a\rangle) = 0$.

Thus the $R/\langle p\rangle$-vector space $p^{m-1}M/p^mM$ has dimension equal to r plus the number of i such that p^m divides a_i. In particular, the number of i such that p^m divides a_i is independent of the choice of presentation and the reduction process, so the factorizations of the a_i are uniquely determined up to multiplication by units.

There is another formulation of the classification that works "one prime at a time". Let $p_1, \ldots, p_t$ be the primes of R dividing a_r. For each j, write $a_j = u_j p_1^{k_{1j}} \ldots p_t^{k_{tj}}$. Then the Chinese remainder theorem gives isomorphisms:

$$R/\langle a_j\rangle \cong \bigoplus_i R/\langle p_i^{k_{ij}}\rangle,$$

$$M \cong R^s \oplus \bigoplus_i \bigoplus_j R/\langle p_i^{k_{ij}}\rangle.$$

The powers $p_i^{k_{ij}}$ appearing here are called the *elementary divisors* of M; they give an alternative to the invariant factors as a way of describing the torsion of M. The submodule $\bigoplus_j R/\langle p_i^{k_{ij}}\rangle$ is called the *p_i-primary part* of M; it is the largest submodule of M that is annihilated by a power of p_i.

Since abelian groups are simply $\mathbb{Z}$-modules, and $\mathbb{Z}$ is Euclidean, we immediately obtain the classification of finitely generated abelian groups:

Theorem 2.2. *Let A be a finitely generated abelian group. Then there is a unique integer r, and unique integers $a_1, \ldots, a_t > 1$, with $a_1|a_2|\ldots|a_t$, such that there exists an isomorphism:*

$$A \cong \mathbb{Z}^r \oplus \mathbb{Z}/a_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/a_t\mathbb{Z}.$$

3. Linear transformations on finite dimensional vector spaces and $K[T]$-modules

A perhaps more surprising application of the theory of modules over Euclidean domains is to linear algebra. Let K be a field, let V be a finite dimensional K-vector space, and let $L : V \to V$ be a K-linear map. We will study the map L by relating it to a certain $K[T]$-module M_L and applying the classification. As a consequence, we'll obtain an easy proof of the Jordan Canonical Form when K is algebraically closed and even see how to generalize it to non-algebraically closed K!

The key point is that a linear map $L : V \to V$ lets us give V the structure of a $K[T]$-module. The key idea is as follows: let $P(T) \in K[T]$ be the polynomial $a_0 + a_1T + \cdots + a_dT^d$. Then we define a K-linear map $P(L) : V \to V$ by $P(L) = a_0 + a_1L + a_2L^2 + \cdots + a_dL^d$.

Define a $K[T]$-module M_L as follows: the underlying set of M_L is simply the vector space V. Addition on M_L is given by addition in V. For any $v \in M_L$ and $P(T)$ in $K[T]$, we define $P(T) \cdot v = P(L)(v)$. This makes M_L into a $K[T]$-module.

We can find a presentation for M_L as follows: Let $v_1, \ldots, v_m$ be a basis for V, and A the matrix of L with respect to the v_i, so that $Lv_i = \sum_{j=1}^m a_{ji}v_j$. Then the v_i generate M_L as a K-vector space, so also as a $K[T]$-module. Moreover for each i we have the relation $Tv_i = \sum_{j=1}^m a_{ji}v_j$.

I claim that these relations, for $i = 1, \ldots, m$, generate all of the relations on the generators v_i. To see this, let f be the map from $K[T]^m$ to M_L that takes $(P_1(X), \ldots, P_m(X))$ to $\sum_{i=1}^r P_i(X)v_i$, and let N be its kernel. Let N' be the subset of K generated, as a $K[T]$-module, by the relations $Tv_i = \sum_{j=1}^m a_{ji}v_j$ (that is, by the elements $r_i = (-a_{i1}, -a_{i2}, \ldots, -a_{i,i-1}, T - a_{ii}, -a_{i,i+1}, \cdots - a_{im})$ of $K[T]^m$.) We need to show that $N = N'$. Since $N' \subseteq N$, and the quotient $K[T]^m/N$ is isomorphic to M_L (and thus m-dimensional as a K-vector space) it suffices to show that the dimension of $K[T]^m/N'$ as a K-vector space is less than or equal to m; in particular it suffices to show that this quotient is spanned, *as a K-vector space*, by the elements e_i of $K[T]^n$ that are zero except in position i and 1 in position i.

Let $r = \sum_i P_i(T)e_i$ be an arbitrary element of $K[T]^m$, and let d be the maximum among the degrees of the P_i. If $d > 0$, and P_i has degree d we can subtract a monomial times r_i from r to obtain a new element $\sum_i P_i'(T)e_i$ where the degree of P_i' is strictly less than d, and where for $i \neq j$, $P_i(t)$ differs from $P_j'(t)$ by a monomial of degree $d - 1$. Doing this for all i such that P_i has degree d gives us a new element r' that differs from r be an element of N', and where the maximum degree among its entries is at most $d - 1$. Repeating as necessary we find an element r'' that is congruent to r modulo N', and all of whose entries lie in K. Thus r'' is a K-linear combination of the e_i, so the quotient $K[T]^m/N'$ is spanned over K by the e_i as claimed.

The upshot is that M_L has the presentation matrix whose columns are the r_i; this matrix is the matrix $T\,\mathrm{Id}_m - A$. (This should be familiar from linear algebra: it is the matrix whose determinant is the characteristic polynomial of L. This is not a coincidence!)

The module M_L is finite-dimensional as a K-vector space, so is certainly finitely generated as a $K[T]$-module. Since $K[T]$ is infinite-dimensional over K, the classification shows that M_L has rank zero as a $K[T]$-module. Thus the classification of finitely generated $K[T]$-modules shows there are polynomials $P_1(T), \ldots, P_t(T)$ such that $P_1(T)|P_2(T)|\ldots|P_t(T)$, and M_L is isomorphic to the $K[T]$-module:

$$K[T]/\langle P_1(T)\rangle \oplus K[T]/\langle P_2(T)\rangle \oplus \cdots \oplus K[T]/\langle P_t(T)\rangle.$$

We can (and do) take the P_i to be *monic*; they are then uniquely determined by L. For each i, let d_i denote the degree of $P_i(T)$.

The polynomials $P_i(T)$ contain information about L. To make this more precise we need to recall some facts from linear algebra: Note that the set $I_L := \{P(T) \in K(T) : P(L) = 0\}$ is a nonzero ideal of $K[T]$. The unique monic generator $Q(T)$ of I_L is called the *minimal polynomial* of L. We have $Q(L) = 0$, and if $P(T)$ is any polynomial with $P(L) = 0$, then $Q(T)$ divides $P(T)$. (Note, however, that unlike minimal polynomials of elements of field extensions, the minimal polynomial of a linear map: $V \to V$ need *not* be irreducible.)

Note that I_L is the annihilator of M_L; that is, the ideal of polynomials $P(T)$ such that $P(T)M_L = 0$. It is clear from the isomorphism:

$$M_L \cong K[T]/\langle P_1(T)\rangle \oplus K[T]/\langle P_2(T)\rangle \oplus \cdots \oplus K[T]/\langle P_t(T)\rangle$$

that this annihilator is generated by $P_t(T)$, so $P_t(T)$ is the minimal polynomial of L.

Since M_L is naturally isomorphic to V as K-vector spaces, we can find a K-basis for V by finding a K-basis for M_L. Even better, since our isomorphism of M_L with V carries "multiplication by T" to L, the matrix of L with respect to the basis for V is *the same* as the matrix of "multiplication by T" on M_L with respect to the corresponding basis of M_L. By choosing a nice basis of M_L, we can thus find a basis for V that gives a nice matrix for L!

For instance, the $K[T]$-module $K[T]/\langle P_i(T)\rangle$ has as a basis $1, T, T^2, \ldots, T^{d_i-1}$. With respect to this basis, "multiplication by T" has as its matrix:

$$\begin{pmatrix} 0 & 0 & 0 & \ldots & 0 & -a_{0i} \\ 1 & 0 & 0 & \ldots & 0 & -a_{1i} \\ 0 & 1 & 0 & \ldots & 0 & -a_{2i} \\ 0 & 0 & 1 & \ldots & 0 & -a_{3i} \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & \ldots & 1 & a_{d_i-1,i} \end{pmatrix}$$

where $P_i(T) = a_{0i} + a_{1i}T + \cdots + a_{d_i-1,i}T^{d_i-1} + T^{d_i}$. (Recall that this matrix is called the *companion matrix* for $P_i(T)$).

If we let M_i denote the compation matrix for $P_i(T)$, then (with respect to the above chosen basis for each summand $K[T]/\langle P_i(T)\rangle$ of M_L) the matrix of "multiplication by T" on M_L is "block diagonal", with blocks $M_1, M_2, \ldots, M_t$. That is, we have:

Theorem 3.1. *There exists a basis for V such that the map $L : V \to V$ has matrix:*

$$\begin{pmatrix} M_1 & 0 & 0 & \dots & 0 \\ 0 & M_2 & 0 & \dots & 0 \\ 0 & 0 & M_3 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & M_t \end{pmatrix}$$

where M_i is the companion matrix of the ith invariant factor for M_L.

This matrix is called a *rational canonical form* for L; it exists for *any* L, over *any* field K, and is uniquely determined by L. In particular, given any n by n matrix A with cofficients in K, there is a unique matrix in rational canonical form that is conjugate to A by an element of $\mathrm{G}L_n(K)$.

Now let us assume K is algebraically closed; in particular every irreducible monic polynomial in $K[T]$ is of the form $T-\lambda$ for some $\lambda \in K$. By factoring each $P_i(T)$ into irreducibles, we find that there is a finite collection of elements $\lambda_1, \dots, \lambda_r$ of K, and for each λ_i a nondecreasing sequence $n_{i1}, \dots, n_{it}$ of nonnegative integers such that

$$P_j(T) = \prod_i (T - \lambda_i)^{n_{ij}}.$$

We thus obtain an isomorphism:

$$M_L \cong \prod_i \prod_j K[T]/\langle (T - \lambda_i)^{n_{ij}} \rangle.$$

The module $K[T]/\langle (T-\lambda)^n \rangle$ has as a basis $(T-\lambda)^{n-1}, (T-\lambda)^{n-2}, \dots, (T-\lambda), 1$. The matrix of "multiplication by T" on $K[T]/\langle (T-\lambda)^n \rangle$ with respect to this basis is the "Jordan block of size n and eigenvalue λ"; that is, the matrix:

$$\begin{pmatrix} \lambda & 1 & 0 & \dots & 0 \\ 0 & \lambda & 1 & \dots & 0 \\ 0 & 0 & \lambda & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ 0 & 0 & 0 & \dots & \lambda \end{pmatrix}.$$

Conclude that when K is algebraically closed and L is any linear map $V \to V$, there exists a basis of V such that the matrix of L is in Jordan canonical form (that is, is block diagonal with each block a Jordan block.