M7210 Lecture 23 Monday October 15, 2012

Abelian Groups III

Today, we are going to set up the machinery we will use to prove:

Theorem. *If S is subgroup of $\mathbb{Z}^n$, then we can choose a new basis $\{b_1, \dots, b_n\}$ of $\mathbb{Z}^n$ and a new generating set $\{t_1, \dots, t_\ell\}$ ($\ell \le n$) for S such that for each $i = 1, \dots, \ell$, t_i is an integer multiple of b_i, i.e., $t_i = d_i b_i$ for some $d_i \in \mathbb{Z}$.*

This theorem ought to have a distinguished name—maybe, "Fundamental Theorem of Finitely-Generated Abelian Groups." But there are numerous ways of stating it and its immediate consequences, and maybe this stands in the way of a standard name. The theorem is closely related to "Smith Normal Form," but I will not work with this concept in the present lecture.[1]

As I said at the end of the last lecture, the theorem has remarkable consequences. First, it means that every subgroup of a free abelian group is free, for the t_i form a basis for S. (A non-tivial $\mathbb{Z}$-linear relation among the t_i would produce a non-trivial relation among the b_i.) Second, the theorem implies that every finitely generated abelian group is a direct sum of cyclic subgroups. For suppose G is any finitely-generated abelian group. Then there is a surjection $\mathbb{Z}^n \twoheadrightarrow G$. Let S be the kernel of this map, and choose $\{b_1, \dots, b_n\}$ and $\{t_1, \dots, t_\ell\}$ as in the theorem. Then

$$S = \bigoplus_{i=1}^{\ell} \mathbb{Z} t_i,$$

and it follows (as we shall show in more detail later) that

$$G \cong \mathbb{Z}^n / S \cong \mathbb{Z}/d_1\mathbb{Z} \oplus \mathbb{Z}/d_2\mathbb{Z} \oplus \cdots \mathbb{Z}/d_\ell\mathbb{Z} \oplus \mathbb{Z}^{n-\ell}.$$

Observe that some of the d_i may be equal to 1, in which case the corresponding summand is $\{0\}$. Finally, let me point out that the proof of the theorem is constructive, allowing us to compute the basis $\{b_1, \dots, b_n\}$ and the d_i from any list of generators for S. The procedure that I describe below allows some choices to the user, but it can be turned into a computer program. Finding fast implementations is an area of current research.

Using matrices to describe morphisms between free R-modules.

Let R be a ring with multiplicative identity. We have shown that R^n is the free R-module on the standard basis $\{e_1, \dots, e_n\} \subset R^n$, where

$$e_{ij} = (e_i)_j = \pi_j(e_i) = \begin{cases} 1_R, & \text{if } i = j; \\ 0, & \text{otherwise.} \end{cases}$$

We are going to introduce a way of displaying the data required to describe elements and morphisms of free R-modules relative to a basis. It's much like the notation introduced earlier for vector spaces,

[1] You might look up "Smith Normal Form" on the web. You can even read Smith's original paper:

`http://www.jstor.org/stable/108738`.

but we switch the roles of rows and columns. This seems to be done mainly to remain consistent with the usual matrix multiplication conventions when R is not commutative.

Let $B = \{b_1, \dots, b_m\}$ be a basis for R^m. Each element $x \in R^m$ has a unique expression as $r_1 b_1 + \cdots + r_m b_m$, with $r_i \in R$. The row vector with the r_i as entries is denoted

$$\big(\, x\,;\; B\,\big) := (r_1 \;\cdots\; r_m).$$

Suppose R^n has basis $C = \{c_1, \dots, c_n\}$ and $\phi : R^m \to R^n$. Let $\big(\,\phi\,;\; B\,C\big)$ denote the matrix whose m rows are the $\big(\,\phi(b_1)\,;\; C\,\big), \dots, \big(\,\phi(b_m)\,;\; C\,\big)$:

$$\big(\,\phi\,;\; B\,C\big) := \begin{pmatrix} \big(\,\phi(b_1)\,;\; C\,\big) \\ \vdots \\ \big(\,\phi(b_m)\,;\; C\,\big) \end{pmatrix}.$$

In other words, the entries of $\big(\,\phi\,;\; B\,C\big)$ are—in the i^{th} row and j^{th} column—the elements $s_{ij} \in R$ defined by

$$\phi(b_i) = s_{i1}c_1 + \ldots + s_{in}c_n.$$

This results in conventions that are compatible with non-commutative rings. Indeed, if $x = r_1 b_1 + \cdots + r_m b_m$, then

$$\begin{aligned}
\phi(x) &= \phi(r_1 b_1 + \cdots + r_m b_m) \\
&= r_1 \phi(b_1) + \cdots + r_m \phi(b_m) \\
&= r_1(s_{11}c_1 + \cdots s_{1n}c_n) + \cdots + r_m(s_{m1}c_1 + \cdots s_{1n}c_n) \\
&= (r_1 s_{11} + r_2 s_{21} + \cdots + r_m s_{m1})c_1 + \cdots + (r_m s_{m1} + \cdots r_m s_{mn})c_n.
\end{aligned}$$

This is consistent with matrix multiplication:

$$\begin{aligned}
\big(\, x\,;\; B\,\big)\big(\,\phi\,;\; B\,C\,\big) &= (r_1 \;\cdots\; r_m) \begin{pmatrix} s_{11} & \cdots & s_{1n} \\ \vdots & \vdots & \vdots \\ s_{m1} & \cdots & s_{mn} \end{pmatrix} \\
&= (r_1 s_{11} + r_2 s_{21} + \cdots + r_m s_{m1} \quad \cdots \quad r_m s_{m1} + \cdots r_m s_{mn}) \\
&= \big(\,\phi(x)\,;\; C\,\big).
\end{aligned}$$

Exercise. Suppose R is not commutative. What would go wrong in attempting to represent morphisms of free R-modules by matrix multiplication if we chose to represent elements by columns and the application of a morphism by matrix multiplication with the matrix on the left.

Suppose $\phi : R^m \to R^n$ is as above, R^p has basis D and $\psi : R^n \to R^p$. Then we have the following formulae:

$$\big(\, x\,;\; B\,\big)\big(\,\phi\,;\; B\,C\,\big)\big(\,\psi\,;\; C\,D\,\big) = \big(\,\psi(\phi(x))\,;\; D\,\big);$$
$$\big(\,\phi\,;\; B\,C\,\big)\big(\,\psi\,;\; C\,D\,\big) = \big(\,\psi\phi\,;\; B\,D\,\big).$$

Just as previously described for vector spaces, if B and C are two bases for R^n, the matrix $\big(\,\mathrm{id}_{R^n}\,;\; B\,C\,\big)$ effects a change of base:

$$\big(\, x\,;\; B\,\big)\big(\,\mathrm{id}_{R^n}\,;\; B\,C\,\big) = \big(\, x\,;\; C\,\big).$$

The i^{th} row of $\big(\,\mathrm{id}_{R^n}\,;\; B\,C\,\big)$ is the n-tuple of coefficients required to write b_i as an R-linear combination of the $c_1, \dots, c_n$. This is an $n \times n$ matrix with right inverse $\big(\,\mathrm{id}_{R^n}\,;\; C\,B\,\big)$. In case R is commutative, these matrices are inverses. In the non-commutative case, there are complications: a matrix may have a right (resp., left) inverse that is not a left (resp., right) inverse; see Jacobson, *Basic Algebra I*, page 97, Exercise 2. Even worse, in the non-commutative case, it is possible for $R^m \cong R^n$ with $m \neq n$; see Jacobson, page 171.

Diagonalizing integer matrices

Suppose $A = \{a_{ij}\}$ is an $m \times n$ integer matrix.

How to add a multiple of one row (column) of A to another row (column). Let $T^n_{ij}(k)$, $i \neq j$, $i, j \in \{1, \ldots, n\}$, be the $n \times n$ matrix with k in the $(i,j)^{th}$ position and 1s on the diagonal. The inverse of $T^n_{ij}(k)$ is $T^n_{ij}(-k)$. Note that $T^m_{ij}(k)A$ has the same rows as A, except for the i^{th} row, which is the i^{th} row of A plus k times the j^{th} row of A. Similarly, $AT^n_{ij}(k)$ has the same columns as A, except for the j^{th} column, which is the j^{th} column of A plus k times the i^{th} column of A.

How to switch two rows (columns) of A. Let P^n_{ij} be the $n \times n$ identity matrix with rows i and j switched. The same matrix arises by switching columns i and j. P^n_{ij} is its own inverse. Notice that $P^m_{ij}A$ is the same as A, but with rows i and j switched. AP^n_{ij} is the same as A, but with columns i and j switched.

How to multiply a row (column) by -1. Let U^n_i be the diagonal matrix in which the ii^{th} entry is -1 and all others are 1. U^n_i is its own inverse. $U^m_i A$ is the same as A, but with its i^{th} row multiplied by -1. AU^n_j is the same as A, but with its j^{th} column multiplied by -1.

Column Step. By repeatedly multiplying A on the left by various matrices of the form $T^m_{ij}(k)$ and P^m_{ij}, we can bring it to the form

$$A' = \begin{pmatrix} a'_{11} & a'_{12} & \cdots & a'_{1n} \\ 0 & a'_{22} & \cdots & a'_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ 0 & a'_{m2} & \cdots & a'_{mn} \end{pmatrix},$$

where the first column has zero in every place but the first and a'_{11} is the greatest common divisor of the entries in the first column of A. The reason that we can do this is essentially that the GCD of any finite set of integers can be expressed as a $\mathbb{Z}$-linear combination of them. I will not specify the precise steps used to select the matrices to multiply by. One might, for example, subtract (or add) a row with the smallest (in absolute value) initial entry from the other rows, and do this over and over until only one row with a non-zero initial entry remained. In examples that are done by hand, it may be more convenient or faster to use some other method. The specific numbers themselves may suggest a strategy.

Row Step. By repeatedly multiplying A on the right by various matrices of the form $T^n_{ij}(k)$ and P^n_{ij}, we can bring it to the form

$$A'' = \begin{pmatrix} a''_{11} & 0 & \cdots & 0 \\ a''_{21} & a''_{22} & \cdots & a''_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ a''_{m1} & d_{m2} & \cdots & d_{mn} \end{pmatrix},$$

where the first row has zero in every place but the first and a''_{11} is the greatest common divisor of the entries in the first row of A.

We can perform these steps without multiplying rows or columns by -1, but it may be useful or convenient to make some entries positive, and this can clearly be done by using U^n_i.

Barring the possibility that some $a_{11}^{(s)}$ is zero, $|a_{11}^{(s+1)}| < |a_{11}^{(s)}|$. Thus, by repeatedly performing row steps and column steps, we may bring A to the form

$$B = \begin{pmatrix} b_{11} & 0 & \cdots & 0 \\ 0 & b_{22} & \cdots & b_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ 0 & b_{m2} & \cdots & b_{mn} \end{pmatrix},$$

where b_{11} is the greatest common divisor of all the elements in the first column and first row of A. If some $a_{11}^{(s)}$ is zero, we may either make it non-zero by exchanging t

We an then apply the same process to the submatrix $\begin{pmatrix} b_{22} & \cdots & b_{2n} \\ \vdots & \vdots & \vdots \\ b_{m2} & \cdots & b_{mn} \end{pmatrix}$. We can do this "in place" by multiplying B on the right and left by $T_{ij}^n(k)$ and P_{ij}^n with i and j never equal to 1. We get a matrix of the form:

$$B' = \begin{pmatrix} b_{11} & 0 & 0 & \cdots & 0 \\ 0 & c_{22} & 0 & \cdots & 0 \\ 0 & 0 & c_{33} & \cdots & c_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & c_{m3} & \cdots & c_{mn} \end{pmatrix}.$$

Continuing, we get a diagonal matrix D and a relationship:

$$PAQ = D,$$

where P is $m \times m$, Q is $n \times n$ and each is a product of matrices of the form T_{ij} and P_{ij} and therefore is invertible.

Closing remarks. We have not required any conditions on the diagonal entries in PAQ. If the entries are such that $d_i | d_{i+1}$ for $i = 1, \ldots, \ell$, $\ell \leq m$, and $d_i = 0$ for $i > \ell$, then we say that the matrix is in Smith Normal Form. We can assure that PAQ winds up in this form by a slightly more complex algorithm than the one we described, but we do not need this detail to complete our analysis of finitely-generated abelian groups.