

Math 551 – Algebra – Fall 2002

Richard Lyons
Rutgers University
New Brunswick, New Jersey, USA

B. Abelian Groups and Modules over Principal Ideal Domains

1. Abelian Groups.

1a. Direct Products

The nicest way a group G can decompose is as the direct product of a number of other groups (which are then isomorphic to subgroups of G). We give a criterion here for G to be isomorphic to $G_1 \times \cdots \times G_r$ for certain subgroups G_1, \dots, G_r of G .

Definition. Let G_1, \dots, G_r be groups. The (external) direct product $G_1 \dot{\times} \cdots \dot{\times} G_r$ is the group whose underlying set is the Cartesian product $G_1 \times \cdots \times G_r$, with the operation

$$(g_1, \dots, g_r)(h_1, \dots, h_r) = (g_1 h_1, \dots, g_r h_r),$$

the product in the i -th place being formed in the group G_i .

A similar definition can be made for an infinite family of groups $\{G_i\}_{i \in I}$: the external direct product

$$\prod_{i \in I} G_i$$

is the group whose underlying set is the Cartesian product of the G_i . The elements of this set are functions f on I such that $f(i) \in G_i$ for each i ; the multiplication is pointwise, i.e. $(ff')(i) = f(i)f'(i)$.

(External) direct products come with a family of projection homomorphisms

$$\pi_j : G_1 \dot{\times} \cdots \dot{\times} G_r \rightarrow G_j, \quad j = 1, \dots, r$$

or more generally

$$\pi_j : \prod_{i \in I} G_i \rightarrow G_j, \quad j \in I,$$

and satisfy the universal property that for any group G and family of homomorphisms $\gamma_j : G \rightarrow G_j$, $j \in I$, there is a unique homomorphism

$$\gamma : G \rightarrow \prod_{i \in I} G_i \text{ such that } \pi_j \gamma = \gamma_j \quad \forall j \in I.$$

Namely, for any $g \in G$, $\gamma(g)$ has $\gamma_j(g)$ for its j -th coordinate.

There is another *internal* notion of direct product, at least for the case of direct products of finitely many groups:

Definition. Let G be a group and G_1, \dots, G_r be subgroups of G . We write

$$G = G_1 \times \cdots \times G_r$$

and say that G is the direct product of its subgroups G_1, \dots, G_r if and only if the mapping

$$\phi : G_1 \times \cdots \times G_r \rightarrow G, \quad \phi(g_1, \dots, g_r) = g_1 \cdots g_r$$

is an isomorphism.

It is not always the case that ϕ is even a homomorphism; the requirements that it be a homomorphism, and that it be injective and surjective, each carry weight. When $G = G_1 \times \cdots \times G_r$, we are free to think of G interchangeably with the external direct product.

Theorem. Let G_1, \dots, G_r be subgroups of a group G . Then (a), (b) and (c) are equivalent.

- (a) $G = G_1 \times \cdots \times G_r$.
- (b) (1) $G_i \triangleleft G$ for all i ;
 (2) $G = \langle G_1, \dots, G_r \rangle$;
 (3) For each $i = 1, \dots, r$, if we put $G^i = \prod_{j \neq i} G_j$ (in order), then $G^i \cap G_i = 1$.
- (c) (1) For any $i \neq j$, and any $g_i \in G_i$ and $g_j \in G_j$, $g_i g_j = g_j g_i$;
 (2) $G = G_1 \cdots G_r$;
 (3) $G^i \cap G_i = 1$ for all i (with G^i defined as in (b)).

Proof. (a) implies (b): Let $\Gamma = G_1 \times \cdots \times G_r$, and let Γ_i be the subgroup consisting of all elements of Γ all of whose coordinates are 1, except possibly the i -th coordinate. Clearly $\Gamma_i \leq \Gamma$ and ϕ maps Γ_i isomorphically onto G_i . Now if (a) holds, then $\Gamma \cong_\phi G$, and so to prove the various statements of (b) it suffices to check the corresponding statements for Γ and the Γ_i . First, Γ_i is the intersection of the kernels of the projections π_j , $j \neq i$, so $\Gamma_i \triangleleft \Gamma$. A typical element $(g_1, \dots, g_r) \in \Gamma$ is the product of the elements $(g_1, 1, \dots, 1)(1, g_2, 1, \dots, 1) \cdots$ so $\Gamma = \Gamma_1 \cdots \Gamma_r$. Under ϕ , G^i corresponds to the kernel of the projection π_i , and $\Gamma_i \cap \ker \pi_i = 1$.

(b) implies (c): For $g_i \in G_i$, $g_j \in G_j$ and $i \neq j$, $[g_i, g_j] = g_i g_j (g_j^{-1}) = g_i (g_j g_j^{-1}) \in G_i \cap G_j \leq G_i \cap G^i = 1$. This implies (1). Then any word in elements of G_1, \dots, G_r can be shuffled to a word $g_1 \cdots g_r$, so (b2) implies (c2). Trivially (b3) implies (c3).

(c) implies (a): The three properties imply in turn that ϕ is a homomorphism, ϕ is surjective, and ϕ is injective. Namely,

$$\begin{aligned} \phi[(g_1, \dots, g_r)(h_1, \dots, h_r)] &= \phi(g_1 h_1, \dots, g_r h_r) = g_1 h_1 \cdots g_r h_r = g_1 \cdots g_r \cdot h_1 \cdots h_r; \\ &= \phi(g_1, \dots, g_r) \phi(h_1, \dots, h_r) \end{aligned}$$

$\phi(\Gamma)$ is a subgroup containing G_1, \dots, G_r so equals G ;

and if $g = (g_1, \dots, g_r) \in \ker \phi$, then $g_1 \cdots g_r = 1$, and for each i if we solve for g_i we get $g_i \in G_i \cap G^i = 1$, so $\ker \phi = 1$. \square

The case of two factors is important:

$$G = H \times K \iff H \triangleleft G, K \triangleleft G, H \cap K = 1 \text{ and } G = \langle H, K \rangle.$$

Of course $H \times K$ has the normal subgroup $H \times 1 \cong H$, and the quotient $H \times K / H \times 1 \cong K$ (by the first isomorphism theorem applied to the projection onto K). We routinely identify these subgroups and quotients with H and K . What distinguishes the direct product among all “extensions of H by K ” are the facts that (1) H has a complement (a subgroup K such that $HK = G$ and $H \cap K = 1$), and (2) that complement commutes elementwise with H .

Not every subgroup of a direct product is a direct product of subgroups. However, the following is fundamental.

Proposition. *Suppose that $H_1 \triangleleft G_1$ and $H_2 \triangleleft G_2$. Then $H_1 \times H_2 \triangleleft G_1 \times G_2$, and*

$$(G_1 \times G_2) / (H_1 \times H_2) \cong G_1 / H_1 \times G_2 / H_2.$$

Proof. Define $\phi : G_1 \times G_2 \rightarrow G_1 / H_1 \times G_2 / H_2$ by $\phi(g_1, g_2) = (g_1 H_1, g_2 H_2)$. It is easily checked that this is a surjective homomorphism whose kernel is $H_1 \times H_2$. Now apply the first isomorphism theorem.

Finite direct products have another universal property, for which they are also called “sums”, particularly in the context of abelian groups. Namely, given groups G_1, \dots, G_n , there are canonical injections

$$\iota_j : G_j \rightarrow G_1 \times \cdots \times G_n,$$

with $\iota_j(g)$ being the n -tuple whose j -th coordinate is g and all of whose other coordinates are 1. The images of the various ι_j are “supported” in different coordinates and so commute elementwise with one another. The universal property is that for any group H , and any n -tuple of homomorphisms $\phi_i : G_i \rightarrow H$ such that

$$[\phi_i(g_i), \phi_j(g_j)] = 1 \text{ for all } i \neq j, \text{ all } g_i \in G_i \text{ and all } g_j \in G_j,$$

there is a unique homomorphism

$$\Phi : G_1 \times \cdots \times G_n \rightarrow H, \quad (g_1, \dots, g_n) \mapsto \phi_1(g_1)\phi_2(g_2)\cdots\phi_n(g_n)$$

such that $\phi_i = \Phi \circ \iota_i$ for each $i = 1, \dots, n$. The unique mapping Φ is sometimes notated $\phi_1 \times \cdots \times \phi_n$, so that by definition

$$\phi_1 \times \cdots \times \phi_n(g_1, \dots, g_n) = \phi_1(g_1)\phi_2(g_2)\cdots\phi_n(g_n).$$

For example the mapping constructed in the proof of the preceding proposition is $\pi_{H_1} \times \pi_{H_2}$. The commutator condition $[\phi_i(g_i), \phi_j(g_j)] = 1$ is needed to prove (actually equivalent to the statement) that $\phi_1 \times \cdots \times \phi_n$ is a homomorphism.

1b. Direct Sums and Free Abelian Groups

For small abelian groups there are decisive structure theorems, giving essentially unique decompositions as direct products of cyclic groups. By “decisive” is meant that the theorems are strong enough in many cases to enable one to check whether a given statement about (small) abelian groups is true, to check whether two abelian groups arising in different contexts are isomorphic, etc. The word “small” here could be interpreted in two senses: either finitely generated, or of finite exponent. We shall prove the decomposition theorem for the first of these.

It is customary to use additive notation for abelian groups, and we shall do so. Nevertheless, the direct product of a family $\{G_i\}_{i \in I}$ of abelian groups will still be denoted

$$\prod_{i \in I} G_i.$$

However, for infinite index sets I , it is the direct sum (or coproduct) which best suits the theory of abelian groups.

Definition. Let $\{G_i\}_{i \in I}$ be a family of abelian groups. The direct sum

$$\prod_{i \in I} G_i \text{ (or } \bigoplus_{i \in I} G_i$$

is the subgroup of the direct product $\prod_{i \in I} G_i$ consisting of all elements which have only finitely many non-identity coordinates.

Then there are canonical (injective) homomorphisms $\iota_j : G_j \rightarrow \prod_{i \in I} G_i$, one for each $j \in I$, such that $\iota_j(g)$ is the element whose j -th coordinate is g and all of whose other coordinates are the identity (in the appropriate group). These mappings have the universal property that for any abelian group H and any family of homomorphisms $\{\phi_i\}_{i \in I}$ with $\phi_i : G_i \rightarrow H$, there is a unique homomorphism

$$\Phi = \bigoplus_{i \in I} \phi_i : \bigoplus_{i \in I} G_i \rightarrow H \text{ such that } \phi_i = \Phi \circ \iota_i \text{ for all } i \in I;$$

namely $\bigoplus_{i \in I} \phi_i$ maps $(g_i)_{i \in I}$ to $\sum_{i \in I} \phi_i(g_i)$, this (possibly infinite) sum being well-defined since $g_i = 0$ for all but finitely many i , and hence $\phi_i(g_i) = 0$ for all but finitely many i .

This direct sum therefore plays the same role in the theory of abelian groups that the free product plays in the theory of (all) groups. We have really defined an “external” direct sum, and then we say that an abelian group G is the “internal” direct sum of subgroups G_i if and only if the mapping $\bigoplus_{i \in I} : G_i \rightarrow G$ induced by the inclusion mappings (i.e., $\bigoplus_{i \in I} : (g_i)_{i \in I} \mapsto \sum_{i \in I} g_i$) is an isomorphism. Generally the distinction between internal and external is blurred, since it can always be bridged by replacing groups by isomorphic copies.

Definition. A free abelian group is the direct sum

$$\bigoplus_{i \in I} \mathbf{Z}$$

of copies of \mathbf{Z} (indexed by some set I).

If we let e_i be the element of this direct sum which is 1 in the i -th coordinate and 0 in every other coordinate, then every element of $\bigoplus_{i \in I} \mathbf{Z}$ is uniquely expressible as a sum $\sum_{i \in I} n_i e_i$, where the n_i are all integers, all but finitely many (“almost all”) of them being 0. The set $\{e_i\}$ is called a basis, by analogy with the theory of vector spaces.

A free abelian group has the expected (?) universal property, which as usual characterizes it: there is a (set) mapping $\iota : I \rightarrow \bigoplus_{i \in I} \mathbf{Z}$, taking i to the element e_i which is 1 in the i -th coordinate and 0 in every other coordinate; moreover for any abelian group H and any (set) mapping $\psi : I \rightarrow H$ there exists a unique homomorphism $\Psi : \bigoplus_{i \in I} \mathbf{Z} \rightarrow H$ such that $\Psi \circ \iota = \psi$. Namely,

$$\Psi\left(\sum_i n_i e_i\right) = \sum_i n_i \psi(e_i).$$

This has the following consequence, completely analogous to the corresponding result for arbitrary groups and free groups.

Theorem 0. Every abelian group is a quotient of a free abelian group. More precisely, if G is an abelian group and $S \subseteq G$ is a subset such that $\langle S \rangle = G$, then there is a surjective homomorphism

$$\phi : F \rightarrow G,$$

where F is free abelian on S ; thus $G \cong F / \ker \phi$.

Pursuing the analogy with vector spaces a bit further, we may define a subset $\{f_i\}_{i \in I} \subseteq H$ of an arbitrary abelian group H to be linearly independent if and only if whenever $\sum_{i \in I} n_i f_i = 0$, the n_i being integers almost all zero, it follows that $n_i = 0$ for all i . Then a basis is just a linearly independent generating set. It is easily seen that a basis can be equivalently defined as a subset $\{f_i\}$ such that every element of the group is uniquely expressible as a sum $\sum_i n_i f_i$, $n_i \in \mathbf{Z}$, almost all n_i being 0.

Proposition. An abelian group is free abelian on some set if and only if it has a basis. In that case, the cardinality of a basis is uniquely determined.

The cardinality of a basis is called the “rank” of a free abelian group.

Proof. We saw above that a free abelian group has a basis $\{e_i\}$. Conversely, if G has the basis $\{f_i\}_{i \in I}$, then G , together with the mapping $i \mapsto f_i$, has the above universal property since every element of G may be uniquely expressed $\sum_i n_i f_i$, so that the analogue of Ψ above is well-defined.

Next suppose that $F = \bigoplus_{i \in I} \langle e_i \rangle$ is free abelian with basis $\{e_i\}_{i \in I}$. Thus each $e_i \cong \mathbf{Z}$. Pick a prime $p \in \mathbf{Z}$ and set $pF = \{pg \mid g \in F\}$. It is easily checked that $H \leq F$ and indeed $pF = \bigoplus_{i \in I} \langle pe_i \rangle$. Consequently

$$F/pF \cong \bigoplus_{i \in I} \langle e_i \rangle / \langle pe_i \rangle \cong \bigoplus_{i \in I} \mathbf{Z}/p\mathbf{Z},$$

the direct sum of $|I|$ copies of $\mathbf{Z}/p\mathbf{Z}$. But this direct sum may be considered a vector space over $\mathbf{Z}/p\mathbf{Z}$, and $|I|$ is its dimension. so is uniquely determined. Thus the rank of F is $\dim_{\mathbf{Z}/p\mathbf{Z}}(F/pF)$, an expression independent of the basis.

(In the finite-dimensional case it is familiar that the dimension is uniquely determined. In the infinite-dimensional case it is also true, and just as easy to prove given the Schröder-Bernstein theorem in set theory. For vector spaces over $\mathbf{Z}/p\mathbf{Z}$, one can also easily argue that if V is an infinite-dimensional vector space, then $|V| = \dim V$. However, this last approach doesn't work for PID's in general, see next section.)

An important property of free abelian groups is the following splitting property (free abelian groups are “projective” abelian groups):

Theorem 1. *Suppose that G is an abelian group and H is a subgroup such that G/H is free. Then $G = H \oplus K$ for some subgroup $K \leq G$.*

Proof. We prove the equivalent statement: if $\phi : G \rightarrow F$ is a surjective homomorphism of abelian groups with F free abelian, then $G = \ker \phi \oplus K$ for some subgroup $K \leq G$. (This implies the desired statement when applied to the projection $\pi_H : G \rightarrow G/H$.)

Let $\{f_i\}_{i \in I}$ be a basis of F . Choose for each $i \in I$ an element $g_i \in G$ such that $\phi(g_i) = f_i$. Let $K = \langle g_i \mid i \in I \rangle$. Then $G = H + K$: given $g \in G$ we write $\phi(g) = \sum n_i f_i$, set $g' = \sum n_i g_i$ and observe that $\phi(g') = \phi(g)$. Thus $\phi(g - g') = 0$, and $g = (g - g') + g'$ with $g - g' \in H$ and $g' \in K$. Also $H \cap K = 0$: if $g \in H \cap K$, then $g = \sum n_i g_i$ for some integers n_i , and applying ϕ gives $0 = \sum n_i f_i$, so $n_i = 0$ for all i whence $g = 0$.

1c. Finitely Generated Free Abelian Groups

The main theorems on finitely generated abelian groups rest on (and are equivalent to) theorems on finitely generated *free* abelian groups.

Theorem 2. *Any subgroup H of a finitely generated free abelian group G is a free abelian group. Moreover the rank of H is at most the rank of G .*

Proof. Let $\{e_1, \dots, e_n\}$ be a basis of G . The proof is by induction on n . Let $G_0 = \mathbf{Z}e_1 + \dots + \mathbf{Z}e_{n-1}$ and $H_0 = H \cap G_0$. Then by induction H_0 is free abelian of rank $r \leq n - 1$. Moreover $H/H_0 \cong H + G_0/G_0 \leq G/G_0 \cong \mathbf{Z}e_n \cong \mathbf{Z}$. Hence $H/H_0 \cong \mathbf{Z}$ or 0 . In the first case there exists $K \leq H$ such that $H = H_0 \oplus K$, so H is free abelian of rank $r + 1 \leq n$. In the second case the desired conclusion is obvious.

Corollary. *Any subgroup of a finitely generated abelian group is finitely generated.*

Proof. If $H \leq G$ with G finitely generated, we may write $G = F/R$ where F is finitely generated free abelian. Then $H = F_0/R$ for some $R \leq F_0 \leq F$ by the third isomorphism theorem, and the theorem implies that F_0 is finitely generated. *A fortiori*, H is finitely generated.

The main theorem is formulated as a theorem about the relationship between appropriate bases of a free abelian group and a subgroup. It will be applied later in the context of Theorem 0 to the inclusion $\ker \phi \rightarrow F$ (note that Theorem 2 tells us that $\ker \phi$ is itself finitely generated and free, given that F is). It will yield structural information about the quotient $F/\ker \phi$, which as Theorem 0 shows is an arbitrary finitely generated abelian group.

The theorem has an existence and a uniqueness statement.

Fundamental theorem on finitely generated abelian groups (I: free abelian group version). *Let F be a finitely generated free abelian group, of rank s , and E a subgroup of F . Then E is a finitely generated free abelian group, of rank $r \leq s$. Moreover there exist bases e_1, \dots, e_r of E and f_1, \dots, f_s of F and uniquely determined positive integers m_1, \dots, m_r such that*

- a) $e_i = m_i f_i, i = 1, \dots, r$, and
- b) $m_1 | m_2 | \dots | m_r$.

However, the bases $\{e_i\}$ and $\{f_i\}$ are not uniquely determined. The integers

$$m_1, m_2, \dots, m_r$$

of the theorem are sometimes called the “invariant factors” of the inclusion mapping $E \rightarrow F$.

Before proceeding with the proof of this theorem, we make two digressions.

1d. Integer matrices and equivalence

The first digression is to interpret the Fundamental Theorem (version Ia) in matrix language. The reader should be able to supply proofs for all statements in this section. Like finite-dimensional vector spaces and linear transformations, free abelian groups of finite rank and homomorphisms between them have concrete realizations as column vectors and matrices; setting up these realizations for abelian groups is essentially identical to setting them up for vector spaces, except that the entries of the matrices and column vectors now come from \mathbf{Z} instead of from a field.

The group \mathbf{Z}^n of all integer $n \times 1$ column vectors is clearly a free abelian group, and one of its bases is $\{e_1^n, \dots, e_n^n\}$, where e_i^n is the $n \times 1$ column vector whose i -th entry is 1 and all of whose other entries are 0.

Moreover, given any rank n free abelian group G and basis $B = \{e_1, \dots, e_n\}$, and given any $g \in G$, we write $g = \sum_{i=1}^n m_i e_i$ and define $[g]^B = [m_1 \ \dots \ m_n]^T$. The mapping

$$G \rightarrow \mathbf{Z}^n, \quad g \mapsto [g]^B,$$

is then an isomorphism.

Now, given a homomorphism $\phi : G \rightarrow H$ and bases $B = \{e_1, \dots, e_n\}$ and $B' = \{f_1, \dots, f_m\}$ of G and H , we write $\phi(e_j) = \sum_{i=1}^m c_{ij} f_i$, ($c_{ij} \in \mathbf{Z}$) for each $j = 1, \dots, n$, and define

$$[\phi]_B^{B'} = [c_{ij}]_{1 \leq i \leq m, 1 \leq j \leq n}.$$

The mapping

$$\text{Hom}(G, H) \rightarrow \mathbf{Z}^{m \times n}, \quad \phi \mapsto [\phi]_B^{B'}, \quad 1A$$

is then an isomorphism of abelian groups, and

$$[\phi(g)]^{B'} = [\phi]_B^{B'} [g]^B \text{ for all } g \in G,$$

with ordinary matrix multiplication on the right. Indeed, given ϕ , $[\phi]_B^{B'}$ is the unique $m \times n$ matrix for which this equation is true for all $g \in G$. This last remark (or direct computation) quickly implies in turn that for any $\phi, \psi \in \text{Hom}(G, H)$ and $\chi \in \text{Hom}(H, K)$, and for any bases B, B', B'' of G, H, K , respectively, we have

$$[\phi + \psi]_B^{B'} = [\phi]_B^{B'} + [\psi]_B^{B'} \text{ and } [\chi \circ \phi]_B^{B''} = [\chi]_B^{B''} [\phi]_B^{B'},$$

again with matrix multiplication on the right.

If B and B' are two bases of the same free abelian group H , then

$$[1_H]_B^{B'}$$

is the “change of basis” matrix whose columns give the B -expansion of the elements of B' . In particular

$$[1_H]_B^B = I,$$

the identity matrix. Moreover different choices C, C' of bases of H and K , respectively, give the related matrix

$$[\phi]_{C'}^C = [id_K]_{C'}^{B'} [\phi]_{B'}^B [id_H]_B^C = P [\phi]_{B'}^B Q, \quad 1B$$

where $P = [id_K]_{C'}^{B'}$ and $Q = [id_H]_B^C$.

If H and K are free abelian groups of ranks m and n , respectively, and $\phi : H \rightarrow K$ and $\psi : K \rightarrow H$ are homomorphisms, then with respect to bases B and B' of H and K respectively we have

$$[\phi]_{B'}^B [\psi]_B^{B'} = [\phi \circ \psi]_{B'}^{B'} \text{ and } [\psi]_B^{B'} [\phi]_{B'}^B = [\psi \circ \phi]_B^{B'}.$$

Because of the isomorphisms $\phi \mapsto [\phi]_{B'}^B$ and $\psi \mapsto [\psi]_B^{B'}$ of $\text{Hom}(H, K)$ and $\text{Hom}(K, H)$, respectively, with $\mathbf{Z}^{m \times n}$ and $\mathbf{Z}^{n \times m}$, we conclude that

$$\phi \text{ is invertible with inverse } \psi \text{ if and only if } [\phi]_{B'}^B \text{ is invertible with inverse } [\psi]_B^{B'}.$$

In particular,

$$[id_H]_B^{B'} = ([id_H]_{B'}^B)^{-1}.$$

for any two bases B, B' of H .

This permits us to establish bijections (not canonical) between the set of all bases of H and the set of all invertible matrices in $\mathbf{Z}^{m \times m}$. Namely, fix a basis B of H ; the correspondence

$$B' \mapsto [id_H]_B^{B'}$$

is a bijection. So is the correspondence $B' \mapsto [id_H]_{B'}^B$.

As a consequence we can determine, given $\phi : H \rightarrow K$, all the matrices representing ϕ with respect to all bases of H and K . We first define an equivalence relation \sim (called “equivalence”) on $\mathbf{Z}^{m \times n}$ by:

$$A \sim A' \iff A' = PAQ \text{ for some invertible } P \in \mathbf{Z}^{m \times m} \text{ and } Q \in \mathbf{Z}^{n \times n}.$$

Fix bases B and C of H and K respectively. Let $\phi : H \rightarrow K$ be a homomorphism and $A = [\phi]_C^B$. The matrices P and Q allowable above are precisely the matrices $[id_H]_B^{B'}$ and $[id_H]_{C'}^C$, as B' and C' range over the sets of bases of H and K , respectively. Because of (1B), we conclude:

Proposition. *Let H and K be free abelian groups of finite ranks m and n , respectively. Let $\phi : H \rightarrow K$ be a homomorphism, and let $A = [\phi]_C^B$ for some bases B and C of H and K , respectively. Then for any $m \times n$ matrix A' over \mathbf{Z} , $A \sim A'$ if and only if $A' = [\phi]_{C'}^{B'}$ for some bases B' and C' of H and K , respectively.*

The following theorem therefore largely follows from the main theorem of the last section.

Fundamental theorem on finitely generated abelian groups (Ib: \mathbf{Z} -matrix version). *Let A be an integer $m \times n$ matrix. Then there exists a uniquely determined integer $r \geq 0$, $r \leq \min(m, n)$, and uniquely determined positive integers m_1, \dots, m_r such that*

- a) $A \sim \text{diag}(m_1, m_2, \dots, m_r, 0, \dots, 0)$; and
- b) $m_1 \mid m_2 \mid \dots \mid m_r$.

Here $\text{diag}(m_1, m_2, \dots, m_r, 0, \dots, 0)$ is the $m \times n$ matrix with all entries 0 except the first r entries down the main diagonal, which are m_1, \dots, m_r , in that order.

Proof. Let H and K be free abelian groups of ranks m and n , respectively. By the isomorphism (1A), there is a homomorphism $\phi : H \rightarrow K$ and bases B and C of H and K , respectively, such that $[\phi]_C^B = A$.

Set $G = \phi(H) \leq K$ and $r = \text{rank}(G)$. By Theorem 2, G is free, and hence by Theorem 1, $H = H_1 \oplus \ker \phi$ for some subgroup H_1 . Then $\psi = \phi|_{H_1}$ is an isomorphism $H_1 \cong G$.

By version Ia, there exists $r \geq 0$ and uniquely determined positive integers m_1, \dots, m_r such that there are bases $C' = \{e_1, \dots, e_n\}$ of K and $\{f_1, \dots, f_r\}$ of G with $f_i = m_i e_i$

for each $i = 1, \dots, r$. The elements $\psi^{-1}f_1, \dots, \psi^{-1}f_r$, together with an arbitrarily chosen basis of $\ker \phi$ form a basis B' of H . We then have

$$[\phi]_{C'}^{B'} = \text{diag}(m_1, \dots, m_r, 0, \dots, 0), \quad 1C$$

establishing existence. Conversely, if A is equivalent to such a matrix, then bases B' and C' exist such that (1C) holds. Writing $C' = \{e_1, \dots, e_n\}$ and $B' = \{f_1, \dots, f_r, \dots\}$ we have that $\phi(f_1), \dots, \phi(f_r)$ form a basis of G , and $\phi(f_i) = m_i e_i$ for each $i = 1, \dots, r$. So the uniqueness follows from version Ia.

In using this theorem, the following observation is useful:

Proposition. *Let P be a square matrix with entries in \mathbf{Z} . Then P is invertible in $\mathbf{Z}^{m \times m}$ if and only if $\det(P)$ is a unit in \mathbf{Z} , i.e., $\det(P) = \pm 1$.*

On the one hand if P is invertible then $\det(P) \det(P^{-1}) = \det I = 1$, with both $\det(P)$ and $\det(P^{-1})$ in \mathbf{Z} . On the other hand if $\det(P)$ is a unit in \mathbf{Z} then the familiar formula

$$P^{-1} = \det(P)^{-1} \text{adj}(P),$$

$\text{adj}(P)$ being the transpose of the matrix of cofactors, shows that $P^{-1} \in \mathbf{Z}^{m \times m}$.

1e. Principal Ideal Domains

The second digression is to observe that the proofs we are giving generalize without change to modules over PID's (instead of just abelian groups). We might as well formalize this.

Definition. *A principal ideal domain (PID) is a commutative ring R with unit 1, such that R is an integral domain ($xy = 0$ implies either $x = 0$ or $y = 0$) and such that every ideal in R is principal, i.e. consists of all the R -multiples of a single element of R .*

If we consider R to be an R -group (an abelian group with operators R , acting by left multiplication) then the principal ideal condition is just the condition that every R -subgroup (i.e. ideal) is “ R -cyclic”—i.e., is generated as an R -group by a single element.

Examples of principal ideal domains include the following two essential examples: \mathbf{Z} , and $k[X]$, the polynomial ring in one variable over a field. In both cases a stronger statement is actually true: there is a division algorithm. In the case of \mathbf{Z} , for every $d \neq 0$ and every $n \in \mathbf{Z}$, there exist $q, r \in \mathbf{Z}$ such that $n = qd + r$ and $|r| < |d|$. In the case of $k[X]$, for every $d, n \in k[X]$ with $d \neq 0$, there exist $q, r \in k[X]$ such that $n = qd + r$ and $\deg r < \deg d$. Thus these are “Euclidean domains”.

Definition. *A Euclidean domain (ED) is an integral domain R possessing a function $\phi : R - \{0\} \rightarrow \mathbf{Z}^+$, the set of nonnegative integers, such that*

- a) $\phi(ab) \geq \phi(a)$ for all $a, b \in R - \{0\}$;
- b) for every $d, n \in R$ with $d \neq 0$ there exist $q, r \in R$ such that $n = qd + r$, and either $r = 0$ or $\phi(r) < \phi(d)$.

Proposition. *Every ED is a PID.*

Proof. Given an ideal $I \subseteq R$, we must find a generator for it. If $I = 0$, then 0 is a generator. Otherwise choose $x \in I$ such that $\phi(x) \leq \phi(y)$ for all $y \in I$, $y \neq 0$. Then for any $u \in I$, write $u = qx + r$ with $\phi(r) < \phi(x)$ or $r = 0$. But then $r = u - qx \in I$ so $r = 0$ and $u = qx \in Rx$. Therefore $I \subseteq Rx$, and the reverse inclusion is obvious.

In any integral domain R we may define $a \mid b$, for $a, b \in R$, to mean: $b = qa$ for some $q \in R$. Equivalently:

$$a \mid b \iff Rb \subseteq Ra.$$

We may also define a unit to be an element u such that $Ru = R$, or equivalently such that $uv = 1$ for some $v \in R$. Two elements $a, b \in R$ are said to be associates if and only if $a = ub$ for some unit u .

Exercise. *The relation*

$$a \sim b \iff a \text{ and } b \text{ are associates}$$

is an equivalence relation.

Then the relation \mid is transitive, and also is “antisymmetric”, at least relative to: if $a \mid b$ and $b \mid a$, then $a = qb$ and $b = q'a$ for some q, q' , so $a = qq'a$ and $a(1 - qq') = 0$. If $a = 0$, then $b = 0$. If $a \neq 0$, then $qq' = 1$ so q is a unit and a and b are associates.

We may also define a gcd of two nonzero elements $a, b \in R$ to be an element d such that

a) $d \mid a$ and $d \mid b$;

b) for any $d' \in R$ such that $d' \mid a$ and $d' \mid b$, we have $d' \mid d$.

In any integral domain, GCD's are unique up to associates (if they exist). For if d and d' are two gcd's of the same a, b , then by definition $d \mid d'$ and $d' \mid d$, so d and d' are associates.

Proposition. *Let R be a PID and let $a, b \in R$. Then there exists a gcd d of a and b in R . Moreover, $Rd = Ra + Rb$, so that $d = ma + nb$ for some $m, n \in R$.*

Proof. In a PID, the set of ideals is in one-to-one correspondence with the set of associate-classes of elements of R , under $Rx \leftrightarrow [x]$, where $[x]$ denotes the set of associates of x . Moreover $x \mid y \iff Rx \supseteq Ry$. Moreover, by definition a gcd is simply a “greatest lower bound” with respect to the divisibility relation. Hence we must prove that in the set of ideals, $Ra + Rb$ is the least ideal containing Ra and Rb . But this is obvious.

The maximum condition holds in any PID:

Theorem. *In a PID, the maximum condition holds on the set of all ideals.*

Proof. Let R be a PID and let $I_1 \subseteq I_2 \subseteq I_3 \subseteq \dots$ be an infinite ascending chain of ideals of R . Set

$$I = \bigcup_{n=1}^{\infty} I_n.$$

Then I is an ideal, and so $I = Rx$ for some $x \in R$ as R is a PID. But then $x \in I_i$ for some i , and therefore $I \subseteq I_i$, so $I = I_i$. QED

Remark. Virtually the same proof shows that for any commutative ring R , if every ideal of R is finitely generated then R satisfies the maximum condition on ideals. The converse to this statement is also true:

Theorem–Definition. If R is a commutative ring, then the following statements about R are equivalent:

- a) Every ideal of R is finitely generated.
- b) The set of ideals of R satisfies the maximum condition.
- c) R is a noetherian ring.

(Part c) is just the definition of “noetherian ring”.)

To show that b) implies a), take an ideal I . If I were not finitely generated then we could find an infinite sequence of elements x_1, \dots, x_n, \dots of I such that if I_i is the ideal generated by x_1, \dots, x_i , then

$$I_1 < I_2 < \dots < I_n \dots < I,$$

contradicting the maximum condition.

The maximum condition has two equivalent formulations: a) no infinite ascending chain $I_1 < I_2 < \dots$ of ideals exists; b) every nonempty set of ideals has a maximal element. Indeed an infinite ascending chain has no maximal element, so b) implies a). If some nonempty set of ideals had no maximal element, then an infinite ascending chain could be obtained, starting with any ideal in the set, then any ideal proving that the first wasn't maximal in the set, then any ideal proving that *that* one wasn't maximal, etc.

Now we can generalize the theorem of the previous section, with the identical proof. First of all we make the following definition.

Definition. Let R be a ring with 1. A (left) R -module is an abelian group M equipped with a “scalar multiplication”

$$R \times M \rightarrow M, \quad \text{written } (r, m) \mapsto rm,$$

such that the following conditions hold for all $r, s \in R$ and all $m, n \in M$:

- a) $r(m + n) = rm + rn$;
- b) $(r + s)m = rm + sm$;
- c) $r(sm) = (rs)m$;
- d) $1m = m$.

The usual trivial consequences follow, as with vector spaces: $0m = 0$, $r0 = 0$ and $r(-m) = -(rm) = (-r)m$ for all $r \in R$ and $m \in M$.

Definition. Let M and N be R -modules. An R -homomorphism from M to N , or a homomorphism of R -modules from M to N , is a homomorphism $\phi : M \rightarrow N$ of abelian groups such that $\phi(rm) = r\phi(m)$ for all $r \in R$, $m \in M$.

Thus an R -module is in particular an abelian group with operators R (satisfying the extra conditions b), c) and d) above), and an R -homomorphism is precisely a homomorphism of R -groups.

Definition. If M is an R -module, then an R -submodule of M is an additive subgroup N of M such that $rN \subseteq N$ for all $r \in R$.

An R -submodule is thus an R -subgroup, in the terminology of groups with operators.

Thus

The Noether isomorphism theorems hold for R -modules.

and the Jordan-Hölder theorem does as well.

The terms “homomorphism”, “subgroup” have to be replaced by “ R -homomorphism” and “ R -submodule”. The parallelogram law reads

$$(N + P)/N \cong P/(N \cap P)$$

for any R -submodules N and P of an R -module M . (The analogue of the hypothesis $H \leq N_G(K)$ is automatically true since the group operation on M is commutative.)

Given an R -module M and a subset $S \subseteq M$, the R -submodule of M generated by S is defined to be the intersection of all submodules of M containing S , or equivalently the set of all elements of M representable as “ R -linear combinations” $\sum_i r_i s_i$ with all $r_i \in R$, all s_i in S , and almost all $r_i = 0$.

An R -module M is “cyclic” if and only if it is generated by a single element. We claim that M is cyclic if and only if $M \cong {}_R R/A$ for some ideal A of R . If M is generated by m , then

$${}_R R \rightarrow M, r \mapsto rm$$

is a surjective homomorphism and so $M \cong R/A$ where $A = \{r \in R \mid rm = 0\}$ is the “annihilator” of m . Conversely, Conversely, ${}_R R$ is generated as an R -module by 1, hence is cyclic, and so any quotient of it is generated by a single element as well.

The notions of direct sum and product go over as well to R -modules: the direct product and sum of R -modules are both R -modules, with scalar multiplication being defined coordinate by coordinate, e.g., in the R -module $M \oplus N$, $r(m, n) = (rm, rn)$ by definition. The usual universal properties hold for direct sums and products. In particular, if N and P are submodules of the R -module M , then $M = N \oplus P$ if and only if $N + P = M$ and $N \cap P = 0$.

The ring R is itself a (left) R -module; its submodules are precisely its (left) ideals. The notation for this module is ${}_R R$.

Definition. Let R be a ring. A free (left) R -module (on a set I) is the direct sum of copies of the R -module ${}_R R$ (indexed by I).

We may also define a basis of an R -module M to be a subset $B \subseteq M$ such that each $m \in M$ has a unique expression

$$m = \sum_{b \in B} r_b b$$

with almost all coefficients $r_b = 0$.

As with free abelian groups, M is a free (left) R -module on I if and only if it has a (left) basis indexed by I , and these conditions are equivalent to the existence of a mapping $\iota : I \rightarrow M$ satisfying the usual universal mapping property.

Namely, if $M = \bigoplus_{i \in I} R$, then ι takes $i \in I$ to $\iota(i) =$ the element of M which is 0 in all coordinates but the i -th coordinate, and which is 1 in that coordinate. This mapping has the usual universal property: for any R -module N and any (set) map $\psi : I \rightarrow N$ there is a unique homomorphism $\Psi : M \rightarrow N$ (of R -modules) such that $\Psi \circ \iota = \psi$.

As with abelian groups any two free R -modules on the same set I are isomorphic. The commutativity of R allows us to prove as well that the cardinality of I is determined by the isomorphism type of the corresponding free module.

Proposition. *Let R be a PID, or indeed any commutative ring with 1. Let M be an R -module which is free on a set I and also free on a set J . Then $|I| = |J|$.*

Proof. Choose any maximal ideal* A of R . Define AM to be the submodule of M generated by all products am , $a \in A$, $m \in M$. Set $M_I = \bigoplus_{i \in I} R$, so that $M \cong M_I$. Under such an isomorphism, AM corresponds to AM_1 , so $M/AM \cong M_I/AM_I$. However, it is clear that $AM_I = \bigoplus_{i \in I} A$, and so $M_I/AM_I \cong \bigoplus_{i \in I} R/A$.

Now M/AM is an R -module and $ax = 0$ for all $a \in A$ and $x \in M/AM$. We may therefore try (and succeed) making M/AM into an R/A -module by defining

$$(r + A)x = rx$$

for all $r \in R$. The fact that $Ax = 0$ makes this a good definition, and the module axioms may be checked without incident.

In the same way we may make M_1/AM_1 an R/A -module, and our R -isomorphism between M/AM and M_1/AM_1 is also an R/A -isomorphism, because of the way the R/A -module structure has been defined on these modules. Therefore, as R/A -modules, we have

$$M/AM \cong M_1/AM_1 \cong \bigoplus_{i \in I} R/A$$

But since A is a maximal ideal, R/A is a field, and the three objects above are vector spaces over R/A . Thus $|I| = \dim_{R/A}(M/AM)$, which equals $|J|$ by a similar argument.

□

Moreover, the universal property as usual implies:

* That is, $A < R$ and there exist no ideals B of R such that $A < B < R$. The existence of maximal ideals in a (commutative) ring follows quickly from Zorn's Lemma, q.v. In the case of PID's, we know that the set of all ideals satisfies the maximum condition and so maximal ideals of course exist.

Theorem 0^R. *Let R be any ring. Then every (left) R -module is a quotient of a free (left) R -module. More precisely, if M is an R -module and $S \subseteq G$ is a subset such that $\langle S \rangle = M$, then there is a surjective homomorphism*

$$\phi : F \rightarrow M,$$

where F is a free R -module on S ; thus $M \cong F/\ker \phi$.

Now we can repeat the development of the previous section.

Theorem 1^R. *Suppose that G is an R -module and H is a submodule such that G/H is free. Then $G = H \oplus K$ for some R -submodule $K \leq G$.*

Proof. Choose a basis B for G/H and choose an arbitrary preimage of each element of B , thereby forming a subset $C \subset G$. Let K be the R -submodule generated by C . The freeness of G/H implies that $\pi_H|_K$ is an isomorphism between K and G/H . This in turn means that $H \cap K = 0$ and $H + K = G$, as required.

Theorem 2^R. *Let R be a PID. Any submodule H of a finitely generated free R -module G is a free R -module. Moreover the rank of H is at most the rank of G .*

Proof. Let $\{e_1, \dots, e_n\}$ be a basis of G . The proof is by induction on n . Let $G_0 = Re_1 + \dots + Re_{n-1}$ and $H_0 = H \cap G_0$. Then by induction H_0 is a free R -module of rank $r \leq n-1$. Moreover $H/H_0 \cong H + G_0/G_0 \leq G/G_0 \cong Re_n \cong R$. Hence $H/H_0 \cong R$ or 0 . In the first case there exists $K \leq H$ such that $H = H_0 \oplus K$, so H is a free R -module of rank $r+1 \leq n$. In the second case the desired conclusion is obvious.

Corollary. *Let R be a PID. Then any submodule of a finitely generated R -module is finitely generated.*

Proof. If $H \leq G$ with G finitely generated, we may write $G = F/R$ where F is finitely generated free. Then $H = F_0/R$ for some $R \leq F_0 \leq F$ by the third isomorphism theorem, and the theorem implies that F_0 is finitely generated. *A fortiori*, H is finitely generated.

The main theorem is formulated as a theorem about the relationship between appropriate bases of a free module over a PID R , and a submodule. It will be applied in the next section in the context of Theorem 0 to the inclusion $\ker \phi \rightarrow F$ (note that Theorem 2 tells us that $\ker \phi$ is itself finitely generated and free, given that F is). It will yield structural information about the quotient $F/\ker \phi$, which as Theorem 0 shows is an arbitrary finitely generated R -module.

The theorem has an existence and a uniqueness statement. We shall prove the existence now; the uniqueness assertion will be argued after we formulate and prove the main theorem (see Section 1g) for arbitrary finitely generated modules over the PID R .

Fundamental theorem on finitely generated modules over a PID (Ia: free module version). *Let R be a PID. Let M be a finitely generated free R -module, of rank s ,*

and N a submodule of M . Then N is a finitely generated free R -module, of rank $r \leq s$. Moreover there exist bases h_1, \dots, h_r of N and e_1, \dots, e_s of M , and nonzero elements $m_1, \dots, m_r \in R$, the m_i 's being uniquely determined up to associates, such that

- a) $f_i = m_i e_i$, $i = 1, \dots, r$, and
- b) $m_1 | m_2 | \dots | m_r$.

Again, the bases $\{e_i\}$ and $\{h_i\}$ are not uniquely determined.

Proof of existence. First some observations about (R -module) homomorphisms $M \rightarrow {}_R R$. These form an R -module $\text{Hom}_R(M, {}_R R)$ under $(\phi + \psi)(m) = \phi(m) + \psi(m)$ and $(r\phi)(m) = r(\phi(m))$ for all $m \in M$ and $r \in R$. Checking that $\phi + \psi$ is a homomorphism requires the commutativity of addition in the image R ; checking that $r\phi$ is a homomorphism (in particular, that it preserves scalar multiplication) requires the commutativity of multiplication in R ; for any $r, s \in R$ and $m \in M$, and $\phi \in \text{Hom}_R(M, {}_R R)$,

$$(r\phi)(sm) = r(\phi(sm)) = r(s\phi(m)) = (rs)\phi(m) = (sr)\phi(m) = s(r\phi(m)) = s[(r\phi)(m)].$$

The identity element of $\text{Hom}_R(M, {}_R R)$ is the 0 mapping $0(m) = 0$ for all $m \in M$; the inverse of ϕ is $-\phi$, defined by $(-\phi)(m) = -(\phi(m))$ for all $m \in M$. Furthermore, every choice of an (ordered) basis $B = \{e_1, \dots, e_n\}$ of M gives rise to coordinate mappings $\phi_i : M \rightarrow R$ (depending on B), namely with $\phi_i(\sum_j n_j g_j) = n_i$; each such coordinate mapping lies in $\text{Hom}_R(M, {}_R R)$.

Now to the proof of the theorem. By Theorem 2, N is free abelian of rank $r \leq n$.

We consider the set of all images

$$\Phi = \{\phi(N) \mid \phi \in \text{Hom}(M, {}_R R).\}$$

Each $\phi(N)$ is an ideal of R , and so the maximum condition implies that we may select and fix $\phi \in \text{Hom}(M, {}_R R)$ such that $\phi(N)$ is maximal in Φ , i.e., whenever $\phi' \in \text{Hom}(M, {}_R R)$ and $\phi'(N) \geq \phi(N)$, then $\phi'(N) = \phi(N)$. As an ideal of R ,

$$\phi(N) = Rm_1$$

for some $m_1 \in R$.

If $m_1 = 0$, then $\phi(N) = 0$ for all $\phi \in \text{Hom}(M, {}_R R)$. But for any basis of M , each basis-coordinate function is a homomorphism, so annihilates N . Therefore $N = 0$, and the theorem holds with $r = 0$. So we may assume that

$$m_1 \neq 0.$$

Next let us show that ϕ is surjective, i.e.,

$$\phi(M) = R.$$

Of course $\phi(M) = Ra$ for some a , since $\phi(M)$ is an ideal of R . Since R is an integral domain, every element of Ra has the form ra for a unique $r \in R$. Therefore we may write

$$\phi(m) = \psi(m)a$$

where ψ is a well-defined function $M \rightarrow R$. The additivity of ϕ , and the fact that R is an integral domain so obeys the cancellation law, implies that ψ is additive, and similarly, ψ is an R -homomorphism. From the above equation, $\phi(N) = a\psi(N) \subseteq \psi(N)$. The maximality of $\phi(N)$ implies therefore that $\psi(N) = a\psi(N)$. Therefore a is a unit, and so $\phi(M) = Ra = R$, as claimed.

Now we can choose and fix $h_1 \in N$, as any element of N such that $\phi(h_1) = m_1$. The homomorphism $\phi|_N : N \rightarrow Rm_1$ maps the submodule Rh_1 of N onto $Rm_1 = \phi(N)$, and so

$$N = Rh_1 \oplus \ker(\phi|_N) = Rh_1 \oplus (K \cap N),$$

where we have put

$$K = \ker(\phi).$$

Likewise $\phi : M \rightarrow R$ is surjective, and if we choose any $f_1 \in M$ such that $\phi(f_1) = 1$, we similarly get

$$M = Rf_1 \oplus K$$

The next objective is to show that f_1 may be replaced by $e_1 \in M$ satisfying $h_1 = m_1e_1$. The submodule K of M is free, and we choose a basis f_2, \dots, f_n of it. Then f_1, \dots, f_n form a basis of M , and so $h_1 = r_1f_1 + \dots + r_nf_n$ for some $r_i \in R$. Applying ϕ we get $m_1 = r_1 + 0 = r_1$. Thus $h_1 = m_1f_1 + r_2f_2 + \dots + r_nf_n$.

We show that $m_1 | r_i$ for all $i \geq 2$. Let ϕ_i be the i -th coordinate function on M with respect to f_1, \dots, f_n . Then $\phi_2(h_1) = r_2$. Let d be a gcd of m_1 and r_2 and write $d = am_1 + br_2$, $a, b \in R$. Then $(a\phi + b\phi_2)(h_1) = d$. The image of $a\phi + b\phi_2$ thus contains Rd and hence Rm_1 . The maximality of Rm_1 then implies that $Rd = Rm_1$, so m_1 divides r_2 . Similarly it divides r_i for all $i \geq 2$, and we write $r_i = m_1s_i$.

Set $e_1 = f_1 + s_2f_2 + \dots + s_nf_n$. Then $h_1 = m_1f_1 + m_1s_2f_2 + \dots = m_1e_1$. Moreover, $\phi(e_1) = 1$ since $\phi(f_i) = 0$ for all $i \geq 2$. Therefore

$$M = Re_1 \oplus K \text{ and } h_1 = m_1e_1.$$

Clearly K has rank $n - 1$. If $K = 0$ (i.e., $n = 1$), then there is nothing more to prove. If $K \neq 0$, then by induction on n , applied to $K \cap N \subseteq K$, there are bases e_2, \dots, e_n of K and h_2, \dots, h_r of $K \cap N$ and positive integers $m_2, \dots, m_r \in R$ such that $h_j = m_je_j$, $2 \leq j \leq r$, and $m_2 \mid \dots \mid m_r$. Then e_1, e_2, \dots, e_n form a basis of M , and the h_i form a basis of N , so to complete the proof of existence it remains only to check that

$$m_1 \mid m_2.$$

Let ϕ_i now be the coordinate functions on M with respect to e_1, \dots, e_n . Then as $h_1 + h_2 = m_1e_1 + m_2e_2$, we have $\phi_1(h_1 + h_2) = m_1$ and $\phi_2(h_1 + h_2) = m_2$. Thus for suitable $a, b \in R$,

$(a\phi_1 + b\phi_2)(h_1 + h_2) = d$, a gcd of m_1 and m_2 . As before, $Rm_1 \subseteq Rd \subseteq \text{im}(a\phi_1 + b\phi_2)$, so the maximality of Rm_1 implies that $Rm_1 = Rd$, so $m_1|m_2$. This completes the proof of existence in Theorem Ia.

We have only proved existence here; the uniqueness assertion will be proved later.

The association classes $[m_1], [m_2], \dots, [m_r]$, or more sloppily the elements $m_1, \dots, m_r \in R$, are sometimes called the “invariant factors” of the inclusion mapping $E \rightarrow F$.

Exercise. *State and prove the analogue of version Ib for matrices with entries in PID.*

1f. Finitely Generated Abelian Groups and Modules over PID's

In the next two sections we apply the previous theorem to determine the structure of finitely generated modules over a PID. We also complete the (uniqueness) proof of the previous theorem!

In any abelian group G , the set

$$T(G) = \{g \in G \mid g \text{ has finite order}\}$$

is a subgroup of G , since $mg = nh = 0$ implies $mn(g - h) = 0$. It is called the torsion subgroup of G . The group G is called torsion-free if and only if $T(G) = 0$.

Likewise if M is a module over a PID R , we define

$$T(M) = \{m \in M \mid \text{for some } 0 \neq r \in R, rm = 0\}.$$

In a similar way we see that $T(M)$ is a submodule of M . (This actually only requires R to be an integral domain.) The module M is called torsion-free if and only if $T(M) = 0$.

Exercise. $T(M)$ is “fully invariant”, that is, for every homomorphism $\phi : M \rightarrow M$ of R -modules, we have $\phi(T(M)) \leq T(M)$.

Proposition. *Let R be a PID (or any integral domain). Let M be an R -module. Then $T(M)$ is a submodule of M , and $M/T(M)$ is torsion-free.*

Proof. If $rm = sn = 0$ with $r \neq 0 \neq s$, then $rs(m \pm n) = 0$ with $rs \neq 0$. Hence $T(M)$ is a submodule of M . If $x + T(M) \in T(M/T(M))$, then $rx \in T(M)$ for some $r \neq 0$. Therefore $srx = 0$ for some $s \neq 0$, and as $sr \neq 0$, this gives $x \in T(M)$. This proves that $T(M/T(M)) = 0$.

Fundamental theorem on finitely generated modules over a PID (II: invariant factor version). *Let R be a PID. Let M be finitely generated module over a PID R . Then $M = T(M) \oplus N$ for some submodule $N \leq M$. Moreover N is a free R -module, and there exist $m_1, \dots, m_r \in R$ such that no m_i is a unit, $m_1 \mid \dots \mid m_r$ and*

$$T(M) \cong Z_{m_1} \oplus \dots \oplus Z_{m_r}.$$

The rank of N , and the association classes $[m_1], \dots, [m_r]$ are uniquely determined (by M).

This theorem implies, and indeed is equivalent to the following three results:

Theorem 3. *Finitely generated torsion-free modules over a PID are free modules.*

Theorem 4. *If M is a finitely generated module over a PID, then $M \cong T(M) \oplus M/T(M)$.*

Theorem 5. *If M is a finitely generated torsion module over a PID (“torsion” means $M = T(M)$), then $M \cong R/m_1R \oplus \cdots \oplus R/m_rR$ for some $r \geq 0$ and some nonunits $m_1, \dots, m_r \in R$ such that $m_1 \mid \cdots \mid m_r$. The association classes $[m_1], \dots, [m_r]$ are uniquely determined by these conditions.*

Exercise. *Demonstrate the equivalence just asserted.*

Proof The “existence” statements are an application of version I of the theorem. Namely, by Theorem 0 there is a finitely generated free module F and a submodule K such that $M \cong F/K$. By Theorem 1, K is free, and by the main theorem there are bases $\{e_i\}_{1 \leq i \leq s}$ and $\{f_i\}_{1 \leq i \leq r}$ of F and K respectively such that $r \leq s$ and $f_i = m_i e_i$ for each $i \leq r$, where $m_i \in R$ and $m_1 \mid \cdots \mid m_r$. Letting $F_i = \langle e_i \rangle$, and setting $m_i = 0$ for $r < i \leq s$ we have

$$M \cong F/K = (\oplus_{i=1}^s F_i) / (\oplus_{i=1}^s m_i F_i) \cong \oplus_{i=1}^s F_i / m_i F_i \cong (\oplus_{i=1}^r R/m_i R) \bigoplus (\oplus_{i=r+1}^s R).$$

Obviously the first summand corresponds to $T(M)$, and the second summand is free of rank $s - r$. Furthermore, for those m_i which are units, $R/m_i R = 0$, so we may delete them from the final expression. This proves everything but uniqueness.

Now $s - r$ is the rank of $M/T(M)$, so is uniquely determined by M . It remains to show that the m_i in the statement of the theorem (i.e., the nonunits) are uniquely determined up to associates.

Notice that our argument shows more, namely that the m_i coming from a submodule $N \leq M$ in version I are the same as the m_i coming from the module M/N in version II, except that enough units are added to bring the rank of N up to $r = s - (s - r)$, the rank of M minus the torsion-free rank of M/N . Consequently knowing the rank of M and the invariant factors of M/N determines the m_i in version I. Therefore

uniqueness in version II implies uniqueness in version I.

What remains for us to do, therefore, is prove uniqueness in version II.

1g. Primary Decomposition; PID's are UFD's

A further decomposition can be made, using the Chinese Remainder Theorem. For the case $R = \mathbf{Z}$ this theorem states that \mathbf{Z}_n is isomorphic to the direct product (or sum!) of the groups $\mathbf{Z}_{p_i^{e_i}}$, where $n = \prod_i p_i^{e_i}$ is the factorization of n into powers of *distinct* primes. (In a homework exercise you showed that the direct product of two finite cyclic groups of relatively prime orders is a cyclic group, and this implies the C.R.T. for the integers.) By applying the C.R.T. to each of the summands \mathbf{Z}_{m_i} in the fundamental theorem on finitely generated abelian groups, we get a decomposition of $T(G)$ as the direct sum of cyclic groups of prime power order. Moreover, a uniqueness statement can be obtained here as well. But before doing this for PID's, we need the C.R.T. for them. Indeed we need to prove the fundamental theorem of arithmetic for them!

Definition. Let R be an integral domain, and let $p \in R$. Then p is prime if and only if $p \neq 0$, p is not a unit, and whenever $a, b \in R$ are such that $p \mid ab$, then either $p \mid a$ or $p \mid b$.

Definition. Let R be an integral domain, and let $p \in R$. Then p is irreducible if and only if $p \neq 0$, p is not a unit, and in any factorization $p = ab$, $a, b \in R$, either a or b is a unit.

These two notions, identical if $R = \mathbf{Z}$, are conceptually different and are not identical in general. For example, in $R = \mathbf{Z}[\sqrt{-5}]$, the ring of all integer combinations $a + b\sqrt{-5}$, if we put $z = 1 + 2\sqrt{-5}$, then

$$3 \cdot 7 = z \cdot \bar{z}$$

the bar denoting complex conjugation. Thus $3 \mid z \cdot \bar{z}$, although it is easily checked that 3 does not divide z or \bar{z} . Hence 3 is not prime. However, 3 is irreducible. Indeed the norm mapping $N : R \rightarrow \mathbf{Z}$, $N(x) = x\bar{x}$, is multiplicative, and $N(3) = 9$. A factorization $3 = xy$ would give $N(x)N(y) = 9$. However, using $N(a + b\sqrt{-5}) = a^2 + 5b^2$, we see that no element of R has norm 3, so either $N(x) = 1$ or $N(y) = 1$, which in turn implies that x or y is ± 1 , a unit. Hence “irreducible” does not imply “prime” in general. (The problem is that the F.T. of Arithmetic does not extend from \mathbf{Z} to R .)

Lemma. In any integral domain, prime implies irreducible.

Proof. Suppose that $p = ab$ and p is prime. Then $p \mid ab$, so p divides one of the factors, which we may as well assume is a . Thus $a = pc$ for some c . Now $p = pcb$, so $1 = cb$ as we are in an integral domain. Therefore b is a unit. \square

Lemma. In a PID, irreducible implies prime.

Proof. The key is that GCD's of two elements exist in a PID (Section 1d), and are linear combinations of the two elements. Suppose that R is a PID and $p \in R$ is irreducible, and $p \mid ab$. Suppose that p does not divide a . Let $d = \gcd(p, a)$, well-defined up to associates. If $[d] = [p]$, then $p \mid a$, contradiction. But by irreducibility the only other divisors of p are units, so $[d] = [1]$. Hence there are $m, n \in R$ such that $mp + na = 1$. Then $b = bmp + nab$, and as $p \mid ab$ we get $p \mid b$. \square

Definition. A unique factorization domain (UFD) is an integral domain R such that for every $x \in R$ such that $x \neq 0$, there exists a unit u , an integer $n \geq 0$, primes p_1, \dots, p_n , no two of which are associates, and positive integers e_1, \dots, e_n such that

$$x = up_1^{e_1} \cdots p_n^{e_n};$$

moreover, this decomposition is unique, except for “trivial changes” of the following sorts: rearrange the order of terms, and replace some p_i by an associated element (changing u in the process as well).

Theorem. Every PID is a UFD.

Proof. Let R be a PID. The existence of a factorization uses only the property that the ideals of R satisfy the maximum condition: Suppose by way of contradiction that some

element $x \neq 0$ of R has no factorization as a product of a unit and irreducibles, and among all counterexamples choose one such that the ideal Rx is maximal. Now x is itself not irreducible, otherwise it would have the factorization $x = x$. So $x = yz$ for some nonunits y and z . Clearly $x \in Ry$, so $Rx \subseteq Ry$. Likewise $Rx \subseteq Rz$. If both these inclusions are proper, then y and z would have factorizations, which when put together would give a factorization of x , contradiction. Therefore $Rx = Ry$, say. But then x and y are associates and so z is a unit, contradiction.

The uniqueness uses only the fact that every irreducible is prime. Suppose that

$$x = u \prod_i p_i^{e_i} = v \prod_j q_j^{f_j}$$

with the p_i distinct irreducibles, u a unit, the e_i positive integers, and similar conditions on the right side. Then $p_1 \mid x$. Since p_1 is prime and divides the right side, $p_1 \mid q_j$ for some j . But q_j is irreducible and so $[p_1] = [q_j]$. We may change notation so that $[p_1] = [q_1]$, so that $p_1 = wq_1$, w a unit. We may now cancel one p_1 from the left, cancel one q_1 from the right and replace v by vw^{-1} , and then complete the proof by induction on $\sum_i e_i$. □

Corollary. *In a PID, the gcd of two elements a and b is the product of all the “common” factors in the prime factorizations of a and b .*

Here “common” means “up to associates”; recall that gcd's are only well-defined up to associates, anyway.

Chinese Remainder Theorem. *Let R be a commutative ring with 1. Let I_1, \dots, I_n be ideals of R such that for each $i \neq j$, $I_i + I_j = R$. Define $I_1 \cdots I_n$ to be the ideal of R generated by all products $r_1 \cdots r_n$ with $r_i \in I_i$ for each i . Then there is an isomorphism of R -modules*

$$\frac{R}{I_1 \cdots I_n} \cong \frac{R}{I_1} \oplus \cdots \oplus \frac{R}{I_n}.$$

The hypothesis $I_i + I_j = R$ is absolutely essential here!

Proof. The R -module homomorphism

$$\phi : R \rightarrow \frac{R}{I_1} \oplus \cdots \oplus \frac{R}{I_n}, \quad r \mapsto (r + I_1, \dots, r + I_n)$$

has kernel $I_1 \cap \cdots \cap I_n$. We finish the proof by showing

- a) $I_1 \cap \cdots \cap I_n = I_1 \cdots I_n$;
- b) ϕ is onto.

Indeed,

$$R = R \cdot R \cdots R = (I_1 + I_2)(I_1 + I_3) \cdots (I_1 + I_n) \subseteq I_1 + (I_2 \cdots I_n).$$

In particular every coset of I_1 in R contains an element of $I_2 \cdots I_n$, which shows that the image of ϕ contains the first direct factor R/I_1 . Similarly it contains the others, so ϕ is

onto. To prove a), it is obvious that the product of ideals lies in their intersection, by definition of ideal. The displayed equation has an analogue $R = I_i + I_1 \cdots \hat{I}_i \cdots I_n$ for each i (\hat{I}_i means to omit this term). Multiplying these n equations we find that

$$R = \hat{I}_1 I_2 \cdots I_n + I_1 \hat{I}_2 \cdots I_n + \cdots + I_1 I_2 \cdots \hat{I}_n.$$

Putting $J = \cap_{i=1}^n I_i$ we see that the product of J with any one of these summands lies in $I_1 I_2 \cdots I_n$. Hence

$$J = RJ \subseteq I_1 I_2 \cdots I_n.$$

□

Corollary. *Let R be a PID and let $x = up_1^{e_1} \cdots p_n^{e_n}$ be the primary decomposition of $x \in R$. Then*

$$R/Rx \cong R/Rp_1^{e_1} \oplus \cdots \oplus R/Rp_n^{e_n}.$$

It is vital here that the p_i be non-associated with one another.

Proof. Since the p_i are non-associated with one another, $\gcd(p_i^{e_i}, p_j^{e_j}) = 1$ for $i \neq j$ by the last corollary. So $Rp_i^{e_i} + Rp_j^{e_j} = R$. Now apply the previous result with $I_i = Rp_i^{e_i}$. □

Now we can formulate the third and last version of our main theorem on finitely generated modules over a PID.

Fundamental theorem on finitely generated modules over a PID (III: elementary divisor version). *Let R be a PID. Let M be finitely generated module over a PID R . Then $M = T(M) \oplus N$ for some submodule $N \leq M$. Moreover N is a free R -module, and there exist prime powers $p_1^{e_1}, \dots, p_s^{e_s}$ in R such that*

$$T(M) \cong Z_{p_1^{e_1}} \oplus \cdots \oplus Z_{p_s^{e_s}}.$$

The rank of N is uniquely determined by M , as are the association classes $[p_1^{e_1}], \dots, [p_s^{e_s}]$ (except for the order in which they appear).

To prove the existence of this decomposition, use version II, factor each m_i as the product of powers of distinct primes, and apply the Chinese Remainder Theorem.

Conversely, the set of resulting prime powers (counting multiplicities!) determines the sequence m_1, \dots, m_r . That is, if $m_1 \mid \cdots \mid m_r$, we can recover m_1, \dots, m_r from the list of the prime powers appearing in the decompositions of m_1, \dots, m_r . Namely, for each prime which appears, the largest power of that prime divides some m_i and hence divides m_r . Hence m_r must be the product, over all (distinct) primes appearing, of the largest powers of those primes which occur. Removing these from consideration m_{r-1} must be the product, over all primes still remaining, of their largest powers still remaining, etc.

Thus

Uniqueness in version III implies uniqueness in version II.

Therefore it remains only to show uniqueness in version III.

Exercise. Show that the decomposition of $T(G)$ into “primary” parts does not depend on the finite generation of G . More precisely, let G be a module over the PID R , and assume that $G = T(G)$. For each association class $[p]$ of primes in R define

$$G_p = \{g \in G \mid p^n g = 0 \text{ for some positive integer } n\}.$$

Show that

- a) Each G_p is an R -submodule of G .
- b) $G = \coprod G_p$, with one summand for each association class of primes in R .

1h. Uniqueness

We finally complete the proof of the fundamental theorem by proving uniqueness in version III. We use the invariants $M[x]$ and xM of a torsion module M over a PID R .

Definition. Let M be a module over the commutative ring R . For each $x \in R$ define $M[x] = \{m \in M \mid xm = 0\}$ and $xM = \{xm \mid m \in M\}$.

It is obvious that each of these is a submodule of M for any $x \in R$. Furthermore, we may consider $M[x]$ to be an R/xR -module by defining $(r + xR)m = rm$ for each $m \in M[x]$. Moreover the following lemmas are easy to prove.

Lemma. Suppose that M and N are R -modules and $M = M[x]$, $N = N[x]$. Then a mapping $\phi : M \rightarrow N$ is an R -homomorphism if and only if it is an R/xR -homomorphism.

Lemma. If R is a commutative ring and M and M_i are R -modules such that $M = \coprod M_i$, then $M[x] = \coprod M_i[x]$ and $xM = \coprod xM_i$ for each $x \in R$.

Lemma. Let R be a PID, and $M = R/p^n R$, where p^n is a prime power in R , $n \geq 0$. Then

- a) If $q \in R$ is a prime not associated with p , then $M = qM$ and $M[q] = 0$.
- b) $M[p] \cong R/pR$ (both as R -module and as R/pR -module).
- c) If $r \in \mathbf{Z}^+$, then

$$p^r M[p] \cong \begin{cases} R/pR & \text{if } r < n \\ 0 & \text{otherwise.} \end{cases}$$

- d) If p' is a prime associated with p , then $M[p] = M[p']$ and $p^r M = (p')^r M$ for any r .

Proof. Under the canonical projection $\phi : R \rightarrow R/p^n R = M$, we have $\phi(qR) = qM$. But $\gcd(q, p^n) = 1$ so $qR + p^n R = R$. Hence $M = \phi(R) = \phi(qR) = qM$. Similarly, if $m \in M[q]$, then $m = r + p^n R$ for some $r \in R$, and $qr \in p^n R$. Therefore $p^n \mid qr$, so $p^n \mid r$ by unique factorization, so $m = 0$. This proves a).

Next, the preimage of $M[p]$ is the ideal $\{r \in R \mid pr \in p^n R\}$, i.e., the ideal $p^{n-1}R$. So $M[p] = p^{n-1}R/p^n R$. Define $\psi : R \rightarrow M[p]$ by $\psi(r) = p^{n-1}r + p^n R$. This is a surjective

module homomorphism, and $\psi(r) = 0$ if and only if $p^{n-1}r \in p^n R$, i.e., $p|r$ (using unique factorization). This proves b).

Next, it is clear that $p^n M = 0$. If $m < n$ then $p^m M = p^m R/p^n R$. But $\alpha : R \rightarrow p^m R/p^n R$ defined by $\alpha(r) = p^m r + p^n R$ is a surjective homomorphism, with kernel $p^{n-m} R$ by unique factorization. So $p^m M \cong R/p^{n-m} R$ and c) follows from b). The proof of d) is left to the reader. \square

Now suppose that

$$M \cong_{\phi} \left(\prod_i \left(\prod_j R/p_i^{n_{ij}} R \right) \right) \oplus \prod_{j=1}^m R R, \quad (1D)$$

where the p_i are pairwise non-associated primes, and the n_{ij} are positive integers. We must show that the isomorphism type of M determines the rank m , as well as the primes p_i (up to association and the order in which they occur) and the n_{ij} (up to order). This we do as follows. First, $\phi(T(M))$ is the first direct sum, so $M/T(M) \cong \prod_{j=1}^m R R$. Therefore

$$m = \text{rank}(M/T(M)).$$

Second, for any fixed i , and any integer a , we can compute $p_i^a T(M)[p_i]$, using the above lemmas. By the second lemma it is isomorphic to the direct sum of $p_i^a N[p_i]$ as N ranges over the direct summands in (1D). But the third lemma implies that these terms are trivial a) for p_k not associated to p_i ; and b) for those $N = R/p_i^{n_{ij}}$ terms for which $n_{ij} \leq a$. Moreover for the terms $N = R/p_i^{n_{ij}}$ terms for which $n_{ij} > a$, we get $N[p_i] \cong R/p_i R$. Therefore if we let $c_{i,a}$ be the number of n_{ij} for which $n_{ij} > a$, we find that $p_i^a T(M)[p_i]$ is isomorphic to the direct sum of $c_{i,a}$ copies of $R/p_i R$. By the first lemma they are isomorphic as $R/p_i R$ -modules. But $R/p_i R$ is a field:

Lemma. *In a PID, if p is a prime, then Rp is maximal.*

(Proof: if $Rp < Rx$, then x is a proper divisor of p so is a unit.)

Therefore

$$c_{i,a} = \dim_{R/p_i R}(p_i^a T(M)[p_i]).$$

Furthermore, the right side does not change if we replace p_i by an associate, by the second lemma (d).

Hence $c_{i,a}$, the is determined by the isomorphism type of M . (If $\phi : M \rightarrow N$ is an isomorphism, then ϕ carries $T(M)$ to $T(N)$, so induces an isomorphism between the vector spaces, whence they have the same dimension.

Finally, for any given i and n , the number of terms p_i^n appearing in (1D) is $c(i, a-1) - c(i, a)$ so is also determined by the isomorphism type of M .

2. Linear transformations; canonical forms.

2a. The module of a linear transformation

We shall be concerned here with a linear transformation $T : V \rightarrow V$, where V is a finite-dimensional vector space over a field F . For such a setup, the powers T^n ($n \geq 0$) i.e., the iterated composites of T with itself, are defined, and they too are linear transformations from V to V . Considering T^0 to be the identity transformation id_V , we can then make V into a module over the polynomial ring $F[X]$ in one indeterminate X . Namely given a polynomial $f \in F[X]$, say $f(X) = a_n X^n + \cdots + a_1 X + a_0$, and for any $v \in V$, define

$$fv = a_n T^n v + \cdots + a_1 T v + a_0 v.$$

It is straightforward to check that the module axioms are satisfied. We shall call this module V_T .

Because F is a field, the ring $F[X]$ is a Euclidean domain (relative to the degree function), so is a PID. Moreover:

Lemma. *In this situation, V is a finitely generated torsion module for $F[X]$.*

Proof. A basis of V (as vector space) is finite by assumption and certainly generates V as an $F[X]$ -module. So V is finitely generated. To see that it is a torsion module, either a) use the fundamental theorem to deduce that otherwise, it would have a free rank 1 submodule $V_0 \cong F[X]$, leading to the contradiction $\dim V \geq \dim V_0 = \infty$; or b), set $n = \dim V$ and argue that given any $v \in V$, the vectors $v, Tv, \dots, T^n v$ form a linearly independent set, and then a dependence relation takes the form $pv = 0$ for some nonzero polynomial $p \in F[X]$.

The fundamental theorem will therefore give a lot of information. In order to decode it, we first consider exactly what information is encoded in the module structure of V . This is the same as asking what it means for two such modules to be isomorphic.

Lemma. *Let $T : V \rightarrow V$ and $U : W \rightarrow W$ be linear transformations, where V and W are finite-dimensional vector spaces over the same field F . Then the following conditions are equivalent:*

- a) $V_T \cong W_U$ as $F[X]$ -modules;
- b) There exists an isomorphism $\phi : V \rightarrow W$ of vector spaces such that the diagram

$$\begin{array}{ccc} T : V & \rightarrow & V \\ & \downarrow & \downarrow \\ U : W & \rightarrow & W \end{array}$$

commutes;

- c) There exist ordered bases B of V and C of W such that the matrices $[T]_B^B$ and $[U]_C^C$ coincide.

Proof. If a) holds, then an isomorphism $\phi : V_T \rightarrow W_U$ is in particular a bijective linear transformation and also satisfies $\phi(Xv) = X\phi(v)$ for all $v \in V$. This means $\phi(Tv) = U\phi(v)$

so b) holds. If b) holds, then c) holds for an arbitrary basis B of V and its image $C = \phi(B)$, which is a basis of W . If c) holds, and $B = \{b_1, \dots, b_n\}$ and $C = \{c_1, \dots, c_n\}$, then the mapping $\phi(\sum \alpha_i b_i) = \sum \alpha_i c_i$ satisfies b). Finally if b) holds, then $\phi(Tv) = U\phi(v)$ for all $v \in V$, whence $\phi(T^n v) = U^n \phi(v)$ follows by an inductive argument, and then since ϕ , T and U are all linear transformations, $\phi(fv) = f\phi(v)$ for all $f \in F[X]$, and ϕ itself is the required isomorphism.

Two square matrices A, A' of the same size with coefficients in F are called similar (over F) if and only if there is an invertible D of the same size with coefficients in F such that

$$DAD^{-1} = A' \text{ or equivalently—as } D \text{ is invertible—} DA = A'D.$$

Given an ordered basis B of V , the invertible matrices of size $\dim V$ are just the matrices $D = [id_V]_B^{B'}$, as B' ranges over the set of all ordered bases of V , and so if $A = [T]_B^B$, then $DAD^{-1} = [T]_{B'}^{B'}$. Therefore

Lemma. *The conditions of the above lemma are also equivalent to:*

- c') *For some bases B of V and C of W , the matrices $[T]_B^B$ and $[U]_C^C$ are similar;*
 c'') *For any bases B of V and C of W , the matrices $[T]_B^B$ and $[U]_C^C$ are similar; The isomorphism type of V_T determines and is determined by the equivalence class of $[T]_B^B$ under similarity.*

We can therefore interpret the application of the fundamental theorem as a statement about similarity classes of matrices.

A submodule $W \subseteq V_T$ is simply a subspace such that $T(W) \subseteq W$, since this condition immediately implies that $fW \subseteq W$ for all $f \in F[X]$. Such a submodule W exists and has dimension m if and only if T , with respect to some basis, has matrix in upper-triangular “block” form

$$\begin{bmatrix} X & Y \\ 0 & Z \end{bmatrix}$$

where X is $m \times m$ and Z is $(n - m) \times (n - m)$. For if T is similar to such a matrix, then that matrix equals $[T]_B^B$ for some ordered basis B , and the zeroes in the lower left imply that the first m elements of B span a subspace W such that $T(W) \subseteq W$. Conversely, if W exists, then an ordered basis of W followed by further vectors constitutes an ordered basis B of V for which $[T]_B^B$ is in upper-triangular block form.

Similarly, to say that V_T has a direct sum decomposition $V_T = V_1 \oplus V_2$ as $F[X]$ -modules, where V_1 and V_2 have dimension m and $n - m$, respectively, is equivalent to saying that for some B , $[T]_B^B$ is in block-diagonal form

$$\begin{bmatrix} X_1 & 0 \\ 0 & X_2 \end{bmatrix}.$$

If the decomposition exists, one can take B to consist of a basis of V_1 followed by a basis of V_2 ; and conversely, if the matrix is in this form, V_1 can be recovered as the span of the first m elements in the ordered basis B , and V_2 as the span of the last $n - m$ elements.

Let $R = F[X]$. What does it mean to say that $V_T \cong R/fR$ as R -modules, where $0 \neq f \in R$? For simplicity let us assume that f is monic, i.e. has leading coefficient 1. (There is no loss of generality in this assumption since f has a (unique) associate f^* which is monic, and $fR = f^*R$.) Write $f(X) = X^n + a_{n-1}X^{n-1} + \cdots + a_1X + a_0$. Then $V_T \cong R/fR$ if and only if for some basis B ,

$$[T]_B^B = \begin{bmatrix} 0 & 0 & 0 & 0 & \cdots & 0 & -a_0 \\ 1 & 0 & 0 & 0 & \cdots & 0 & -a_1 \\ 0 & 1 & 0 & 0 & \cdots & 0 & -a_2 \\ 0 & 0 & 1 & 0 & \cdots & 0 & -a_3 \\ & & & & \cdots & & \\ 0 & 0 & 0 & 0 & \cdots & 1 & -a_{n-1} \end{bmatrix} .b$$

The matrix on the right is called the **companion** matrix of the monic polynomial f , written $C(f)$. Thus $C(f)$ is a certain $n \times n$ matrix, where $n = \deg f$.

To see this, notice that in $\overline{R} = R/fR$, every coset \overline{g} has a unique representative g which is a polynomial of degree strictly less than $n = \deg f$, by the division algorithm. Let us write $\overline{h} = h + fR$. Then the images $\overline{1}, \overline{X}, \overline{X^2}, \dots, \overline{X^{n-1}}$ form a basis of $\overline{R} = R/fR$, corresponding to a basis B of V . Multiplication by X transforms this to the basis $\overline{X}, \overline{X^2}, \dots, \overline{X^n}$. Therefore the matrix of T with respect to B has as its columns the coordinates of $\overline{X}, \overline{X^2}, \dots, \overline{X^n}$ with respect to the basis $\overline{1}, \overline{X}, \overline{X^2}, \dots, \overline{X^{n-1}}$. Since $\overline{f} = 0$ we have $f(\overline{X}) = 0$, so $\overline{X^n} = -a_0 - a_1\overline{X} - \cdots$, and we obtain the companion matrix.

2b. Rational canonical form

With this interpretation, we immediately get “rational canonical form” theorems. For example, from version II we get:

Theorem. *Let V be a finite-dimensional vector space and $T : V \rightarrow V$ a linear transformation. Then there exists a unique $r \geq 0$ and unique monic polynomials f_1, \dots, f_r such that $f_1 \mid f_2 \mid \cdots \mid f_r$ and $[T]_B^B$ is the block diagonal matrix*

$$\begin{bmatrix} C(f_1) & 0 & \cdots & 0 \\ 0 & C(f_2) & \cdots & 0 \\ & & \cdots & \\ 0 & 0 & \cdots & C(f_r) \end{bmatrix} .$$

and its matrix version:

Theorem. *Let F be a field and A a square matrix over F . Then there exist unique r and monic polynomials $f_1, \dots, f_r \in F[X]$ such that $f_1 \mid f_2 \mid \cdots \mid f_r$ and A is similar to the block-diagonal matrix whose blocks are the companion matrices $C(f_1), \dots, C(f_r)$.*

This is obtained by taking $V = F^n$ to be the space of column vectors, and $T : V \rightarrow V$ to be defined by $T(v) = Av$, and applying the previous theorem.

The polynomials f_1, \dots, f_r are called the invariant factors of T or A , as the case may be. Similarly, there is an “elementary divisor” version, corresponding to version III of the fundamental theorem, giving a canonical form for the similarity class of any square matrix over F as a block-diagonal matrix with the blocks being companion matrices of powers of monic irreducible polynomials (unique up to the order of blocks); these prime powers are the “elementary divisors” of the transformation or matrix.

Theorem. *Let A be a square matrix over a field F . Then there exist irreducible polynomials p_1, \dots, p_r and positive integers n_1, \dots, n_r such that A is similar to the block-diagonal matrix whose blocks are $C(p_1^{n_1}), \dots, C(p_r^{n_r})$. The prime powers $p_1^{n_1}, \dots, p_r^{n_r}$ are uniquely determined up to order.*

Example. *The invariant factor version implies that every 2×2 real matrix is similar (over \mathbf{R}) to exactly one of the following:*

$$\begin{bmatrix} \alpha & 0 \\ 0 & \alpha \end{bmatrix}, \alpha \in \mathbf{R} \text{ or } \begin{bmatrix} 0 & \gamma \\ 1 & \beta \end{bmatrix}, \beta, \gamma \in \mathbf{R}$$

the first having invariant factors $X - \alpha, X - \alpha$ and the second having the single invariant factor $X^2 - \beta X - \gamma$. The elementary divisor version implies that every real matrix is similar (over \mathbf{R}) to exactly one of the following:

$$\begin{bmatrix} \alpha & 0 \\ 0 & \beta \end{bmatrix}, \alpha, \beta \in \mathbf{R} \text{ or } \begin{bmatrix} 0 & \gamma \\ 1 & \beta \end{bmatrix}, \beta, \gamma \in \mathbf{R}, \beta^2 + 4\gamma \leq 0.$$

This is because the monic irreducible polynomials over \mathbf{R} are the polynomials $X - \alpha$, $\alpha \in \mathbf{R}$ and the polynomials $X^2 + \beta X + \gamma$ with $\beta^2 < 4\gamma$. (Notice that the companion matrix of this last polynomials has entries $-\beta$ and $-\gamma$ in the last column.) Now the elementary divisors of a 2×2 matrix either are a) two irreducible polynomials of degree 1; b) an irreducible polynomial of degree 2; or c) the square of an irreducible polynomial of degree 1. It is the third possibility which is covered above by the case $\beta^2 + 4\gamma = 0$.

The rational forms for 2×2 matrices over \mathbf{C} are similar, except that the quadratic irreducible possibility does not exist.

Those of the above forms with nonzero determinant therefore are a set of representatives for the conjugacy classes of the group $GL_2(\mathbf{R})$.

Caution: the term “rational canonical form” is not universally interpreted this way.

2c. Jordan canonical form

Definition. *The field F is algebraically closed if and only for every nonconstant polynomial $f \in F[X]$ there exists a zero, i.e., an element $z \in F$ such that $f(z) = 0$.*

This immediately implies that $X - z$ divides f . (Write $f(X) = q(X)(X - z) + y$, with $y \in F$. Then $0 = f(z) = y$.) Consequently over an algebraically closed field, the monic

irreducible elements of $F[X]$ are just the linear polynomials $X - \alpha$, $\alpha \in F$. Conversely, if these are all the irreducibles, then the fact that $F[X]$ is a UFD (uniqueness doesn't matter) implies that every polynomial has a linear factor, hence a zero, so F is algebraically closed.

The most famous example of an algebraically closed field is \mathbf{C} . But in fact every field can be embedded in an algebraically closed field (proof next semester).

To analyze linear transformations and matrices over an algebraically closed field, there is a more useful canonical form than rational canonical form. It uses the elementary divisor version of the fundamental theorem, and "Jordan blocks" instead of companion matrices. Every monic prime power in $F[X]$, in this case, is of the form $f(X) = (X - \alpha)^n$. Instead of the basis $\overline{1}, \overline{X}, \dots, \overline{X^{n-1}}$ of R/fR which we used for companion matrices, we can use the basis $\overline{1}, \overline{X - \alpha}, \overline{(X - \alpha)^2}, \dots, \overline{(X - \alpha)^{n-1}}$. We have $\overline{X(X - \alpha)^i} = (\alpha + \overline{X - \alpha})(\overline{X - \alpha})^i = \alpha(\overline{X - \alpha})^i + (\overline{X - \alpha})^{i+1}$. As $\overline{(X - \alpha)^n} = 0$ we get the lower triangular matrix

$$J(\alpha, n) = \begin{bmatrix} \alpha & 0 & 0 & \cdots & 0 & 0 \\ 1 & \alpha & 0 & \cdots & 0 & 0 \\ 0 & 1 & \alpha & \cdots & 0 & 0 \\ & & & \cdots & & \\ 0 & 0 & 0 & \cdots & 1 & \alpha \end{bmatrix},$$

all diagonal entries being α , all sub-diagonal entries being 1, and all other entries being 0. Version III of the fundamental theorem then implies (we state only the matrix version):

A matrix will be said to be in **Jordan canonical form** (over F) if and only if it is in block-diagonal form

$$\begin{bmatrix} J(\alpha_1, n_1) & 0 & \cdots & 0 \\ 0 & J(\alpha_2, n_2) & \cdots & 0 \\ & & \cdots & \\ 0 & 0 & \cdots & J(\alpha_r, n_r) \end{bmatrix}$$

for some $\alpha_i \in F$ and some positive integers n_1, \dots, n_r .

Version III of the fundamental theorem then implies:

Theorem. *Let F be an algebraically closed field. Then any square matrix over F is similar to a matrix in Jordan canonical form. This canonical form is unique except for the order in which the diagonal blocks appear.*

For the above canonical forms to be useful it is important that they can be calculated, or at least partial information about them can be calculated. The important invariants μ_A and χ_A (minimal and characteristic polynomials) will be discussed below. For Jordan canonical form, there is a nice characterization of the number of blocks $J(\alpha, n)$ for any $\alpha \in F$ and any positive integer n , although it requires fairly detailed information about the matrix in question. We define the nullspace N of an $n \times n$ matrix A to be the space of all vectors v such that $Av = 0$, and the nullity $\nu(A)$ to be $\dim N(A)$. If $A = [T]_B^B$ (or even $[T]_{B'}^B$!) then $\nu(A) = \dim \ker(T)$. Therefore if A and A' are similar, then $\nu(A) = \nu(A')$.

Exercise. If A is in block diagonal form,

$$A = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix},$$

then $\nu(A) = \nu(A_1) + \nu(A_2)$.

Proposition. Let F be an algebraically closed field and A a square matrix (or $T : V \rightarrow V$ a linear transformation). Let $\alpha \in F$ and let n be a positive integer. Then the number of Jordan blocks $J(\alpha, m)$ such that $m > n$ in the Jordan canonical form of A (or T) is

$$\nu((A - \alpha I)^{n+1}) - \nu((A - \alpha I)^n), \text{ or } \dim \ker((T - \alpha)^{n+1}) - \dim \ker((T - \alpha)^n).$$

Here we write α for the linear transformation $v \mapsto \alpha v$ of V .

Once we know how many blocks $J(\alpha, m)$ there are satisfying $m > n$, and we know this for all n , subtraction gives us the number of $J(\alpha, n)$ blocks.

Proof. If A and A' are similar, then so are $A - \alpha I$ and $A' - \alpha I$ and hence so are $(A - \alpha I)^n$ and $(A' - \alpha I)^n$ for every n . Hence the nullities of these two matrices coincide. The upshot of this remark is that in proving our result, we are free to replace A by any matrix similar to it, and so we may assume that A is already in Jordan canonical form. Using the exercise above, we reduce to the case of a single Jordan block $A = J(\alpha, n)$. But then $A - \beta$ is invertible for all $\beta \neq \alpha$, so the nullities $\nu((A - \beta)^n)$ are all 0 and our formula predicts no Jordan blocks $J(\beta, m)$, which is correct. For α itself, if $A = J(\alpha, m)$, then $A - \alpha I = J(0, m)$, and as one takes the powers of $J(0, m)$, the 1's slide down one diagonal at a time, until $J(0, m)^m = 0$. Therefore $\nu(J(0, m)^i) = i$ for $0 \leq i \leq m$, and $\nu(J(0, m)^i) = m$ for $i > m$. Hence

$$\nu(J(0, m)^{n+1}) - \nu(J(0, m)^n) = \begin{cases} 1 & \text{if } 0 \leq n < m \\ 0 & \text{if } n \geq m \end{cases},$$

so our formula counts blocks as claimed.

2d. Minimal polynomial

Two of the most important of the invariants of a linear transformation or square matrix (over a field F) are its minimal polynomial and characteristic polynomial, both elements of $F[X]$.

The minimal polynomial is defined as follows. Let V be a finite dimensional vector space over the field F . The space $\mathcal{A} = \text{Hom}_{\mathcal{F}}(\mathcal{V}, \mathcal{V})$, in addition to being a vector space over F , carries a (noncommutative) ring structure, with multiplication given by composition of functions. Such a ring/vector space \mathcal{A} is called an F -algebra if scalar multiplication and multiplication obey the axiom:

$$\alpha(xy) = (\alpha x)y = x(\alpha y)$$

for all $\alpha \in F$ and all $x, y \in \mathcal{A}$ (it is assumed also that the addition in the ring and the addition in the vector space are the same operation). This axiom does hold in the present case. Another algebra, isomorphic in fact to this one, is the algebra of $n \times n$ matrices over F (where $n = \dim V$). Another one is $F[X]$ itself.

Given $T \in \text{Hom}_F(V, V)$, consider the mapping

$$\phi_T : F[X] \rightarrow \text{Hom}_F(V, V), \quad \phi_T(f) = f(T).$$

This is an F -algebra homomorphism, i.e., preserves both the vector space and ring structures. The kernel of this ring homomorphism, like that of any ring homomorphism, is an ideal, so as $F[X]$ is a PID, $\ker \phi_T = F[X]\mu_T$ for a unique monic polynomial μ_T , called by definition the **minimal polynomial of T** . An equivalent definition would be: $\mu_T \in F[X]$ is the minimal polynomial of T if and only if $\mu_T(T) = 0$ and μ_T is the unique monic polynomial of least degree satisfying this equation. A further property is that if $f \in F[X]$ is any polynomial such that $f(T) = 0$, then $\mu_T \mid f$ (in $F[X]$). Similarly we can define the minimal polynomial of a square matrix A , or of any element x of an F -algebra (in the last case, there is the possibility that $\mu_x = 0$, i.e., $f(x) \neq 0$ for every nonzero polynomial $f \in F[X]$).

The minimal polynomial of a linear transformation is of course equal to the minimal polynomial of any matrix representing it. This implies that similar matrices have the same minimal polynomial (a fact which is not too hard to prove directly anyway).

The relationship between minimal polynomial and the invariants of the last sections is:

Proposition. *Let A be a square matrix over a field F . The minimal polynomial of A is the invariant factor of A of largest degree.*

The identical statement is also valid for linear transformations: $V \rightarrow V$, V a finite-dimensional vector space over F .

Proof. We give the proof for a linear transformation $T : V \rightarrow V$. Let f_1, \dots, f_r be the invariant factors of T . Thus by definition each f_i is monic and

$$V_T \cong R/f_1R \oplus \cdots \oplus R/f_rR \text{ and } f_1 \mid f_2 \mid \cdots \mid f_r.$$

Also by definition, μ_T is the unique monic polynomial of least degree such that $\mu_T V_T = 0$, this equation being equivalent to $\mu_T(T) = 0$. Since each f_i divides f_r , we have $f_r R \subseteq f_i R$ for each i and so $f_r V_T = 0$. Therefore $\mu_T \mid f_r$. Conversely $\mu_T(R/f_rR) = 0$ and so $\mu_T \cdot 1 \equiv 0 \pmod{f_r}$, that is, $f_r \mid \mu_T$. Therefore $f_r = \mu_T$.

Thus the minimal polynomial of T is analogous to the exponent of a finite abelian group.

The following observations are easy to check:

Proposition.

a) *If A is in block diagonal form,*

$$A = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix},$$

then $\mu_A = \text{lcm}(\mu_{A_1}, \mu_{A_2})$.

b) $\mu_C(f) = f$ if f is monic. (The previous proposition has actually verified this.)

c) $\mu_{J(\alpha, n)} = (X - \alpha)^n$. (Again, this follows from the previous proposition, or more messily from a matrix calculation.)

Consequently, for a matrix A , if the Jordan canonical form of A has Jordan blocks $J(\alpha_1, n_1), \dots, J(\alpha_r, n_r)$, then

$$\mu_A = \text{lcm}((X - \alpha_1)_{n_1}^n, \dots, (X - \alpha_r)_{n_r}^n) = \prod_{\alpha} (X - \alpha)^{n_{\alpha}}$$

where the product is over the distinct diagonal entries α of the Jordan form of A , and for each α , n_{α} is the size of the largest Jordan block appearing for α .

Thus for example, a complex matrix with minimal polynomial $(X^2 - 1)(X^3 - 1) = (X - 1)^2(X + 1)(X - \omega)(X - \omega^2)$ has Jordan blocks $J(-1, 1)$, $J(\omega, 1)$, $J(\omega^2, 1)$, at least one of each; at least one Jordan block $J(1, 2)$, and possibly some $J(1, 1)$'s.

This also gives the following characterization of diagonalizability:

Proposition. *A square matrix A over an algebraically closed field is diagonalizable, i.e., similar to a diagonal matrix, if and only if its minimal polynomial is square-free, i.e., has no multiple roots (i.e. no repeated linear factors).*

The same can be said for a linear transformation $T : V \rightarrow V$; we define diagonalizability of such a transformation to mean that there exists a basis v_1, \dots, v_n of V and scalars $\alpha_i \in F$ such that $T(v_i) = \alpha_i v_i$ for each $i = 1, \dots, n$.

2e. Characteristic polynomial; eigenvalues

We assume familiarity with the rudimentary theory of determinants; every square matrix A over a commutative ring R has a determinant $\det A \in R$, which can be obtained by row or column expansion. Moreover,

a) $\det(AB) = \det(A) \det(B)$,

b) $\det(I) = 1$, and

c) $A \text{adj } A = \text{adj } AA = \det A I$, where $\text{adj } A$ is the transpose of the matrix of cofactors: $(\text{adj } A)_{ij} = (-1)^{i+j} \det M_{ji}$, where M_{ji} is obtained from A by crossing out the i -th row and j -th column.

d) The matrix A is invertible if and only if $\det A \neq 0$. (In one direction compute $\det(AA^{-1})$, using a) and b); in the other direction use c).)

e) $\det \begin{bmatrix} A & 0 \\ B & C \end{bmatrix} = \det A \det C$.

Using determinants we get another useful invariant of a square matrix or linear transformation (maybe the most useful), the **characteristic** polynomial.

Definition. Let A be an $n \times n$ matrix over a commutative ring F . The characteristic polynomial of A is defined to be

$$\chi_A(X) = \det(XI - A).$$

The determinant on the right is evaluated in the polynomial ring $F[X]$. Some authors define it as $\det(A - XI)$. This is an inconsequential difference; what matters is just the association class of χ_A , and that one sticks consistently to one definition or the other.

Salient properties of characteristic polynomials:

- a) If A and A' are similar then $\chi_A = \chi_{A'}$.
- b) If A is $n \times n$, then $\chi_A(X) = X^n - (\text{Tr } A)X^{n-1} + \cdots + (-1)^n \det A$. Here $\text{Tr } A$, the trace of A , is defined as $\sum_{i=1}^n A_{ii}$, the sum of the entries of A on the main diagonal.
- c) A is invertible if and only if $\chi_A(0) \neq 0$. (This is immediate since $\chi_A(0) = \pm \det A$.)
- d) $\mu_A \mid \chi_A$. That is, $\chi_A(A) = 0$. This is the **Cayley-Hamilton Theorem**.

Proof. (of Cayley-Hamilton) (proof number 1,066) Let $V = F^n$ be the space of column vectors, with entries in F ; let v_1, \dots, v_n be the standard basis of V (i.e., the columns of the identity matrix). Make V an $F[X]$ -module in the usual way: for an arbitrary $f(X) = a_n X^n + \cdots + a_1 X + a_0 \in F[X]$, define $fv = f(A)v = a_n A^n v + \cdots + a_1 A v + a_0 v$.

Given any matrix $C = [c_{ij}(X)]$ over $F[X]$ (we shall presently take $C = XI - A$), there is an F -homomorphism $\psi_C : V \rightarrow V$ defined by $\psi_C(v_j) = \sum_i c_{ij}(A)v_i$, $j = 1, \dots, n$. (To define homomorphisms out of free modules, precisely what is required is to specify the images of the basis elements.) The following facts follow directly from the definition:

- 1) If $C = fI$ for some polynomial $f \in F[X]$ then ψ_C is left multiplication by $f(A)$. ($\psi_C(v_j) = f(A)v_j$.)
- 2) $\psi_{CD} = \psi_C \circ \psi_D$ for any $n \times n$ matrices C and D over $F[X]$.
- 3) If $C = XI - A$, then $\psi_C = 0$. (Namely, $\psi_C(v_j) = \sum_i c_{ij}(A)v_i = Av_j - (\sum_i a_{ij}v_i) = 0$.)

Now let $D = \text{adj}(XI - A)$, another $n \times n$ matrix over $F[X]$. Then $CD = \det(XI - A)I = \chi_A I$, all being matrices over $F[X]$. By a), ψ_{CD} is left multiplication by $\chi_A(A)$. But by b) and c), $\psi_{CD} = \psi_C \psi_D = 0$. Therefore $\chi_A(A) = 0$. QED

Exercise. If A is an $n \times n$ matrix over F and $A^m = 0$ for some m , then $A^n = 0$.

Compare the following properties of χ_A with those of μ_A in the previous section.

Proposition. If

- a) If A is in block diagonal form,

$$A = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix},$$

then $\chi_A = \chi_{A_1} \chi_{A_2}$.

b) $\chi_{C(f)} = f$ if f is monic.

c) $\chi_{J(\alpha, n)} = (X - \alpha)^n$.

Part a) follows from property (e) of determinants, above. For b) and c), you can calculate directly, or be sneaky: we know that these are the minimal polynomials, and since they have the right degree and divide the characteristic polynomials, and every polynomial in sight is monic, they must be the characteristic polynomials. Consequently:

Proposition. χ_A is

- a) the product of all the invariant factors of A ;
- b) the product $\prod_{\alpha} (X - \alpha)^{m_{\alpha}}$, the product being taken over all distinct diagonal entries of the Jordan form of A , with m_{α} being the number of times α appears there (in the case that F is algebraically closed).

Corollary. Let T be a linear transformation $V \rightarrow V$. Then $\mu_T = \chi_T$ if and only if V_T is “cyclic”, i.e., generated as $F[X]$ -module by a single element.

Exercise. Let $T : V \rightarrow V$ and let W be a T -invariant subspace of V . If there is no $w \in W$ such that the images w, Tw, T^2w, \dots generate W , then there is no $v \in V$ such that the images v, Tv, T^2v, \dots generate V . (Hint. Show that the negations of these assertions are that V_T and $W_{T|W}$ are cyclic R -modules in the sense just defined.)

Corollary. Any irreducible factor of χ_A also divides μ_A .

Proof. $\chi_A = f_1 \cdots f_r$ with $f_1 \mid f_2 \mid \cdots \mid f_r = \mu_A$.

Corollary. If A (or T) has at most two invariant factors, then μ_A and χ_A determine the similarity class of A (or the isomorphism class of V_T).

Namely, the invariant factors are μ_A and χ_A/μ_A (the latter existing only if $\mu_A \neq \chi_A$).

QED

Exercise. Show that χ_A and μ_A determine the similarity class of A if A is 3×3 , but not in general if A is 4×4 or larger.

Exercise. If A is a complex 4×4 matrix and $\mu_A = (X^2 + 1)(X - i)$, then what are the possibilities for χ_A ? Give an example to show that each possibility occurs.

The equation $\chi_A(X) = 0$ is the **characteristic equation** and its roots are the **eigenvalues of A** . When each eigenvalue α is counted with its multiplicity as a root of the characteristic equation (i.e., the number of linear factors $(X - \alpha)$ of χ_A) there are n eigenvalues in all (in an algebraically closed field containing F). The eigenvalues of T are defined similarly.

Proposition. α is an eigenvalue of A (or T) if and only if there is a nonzero column vector v such that $Av = \alpha v$ (or a nonzero $v \in V$ such that $T(v) = \alpha v$).

Proof. $\chi_A(\alpha) = 0$ if and only if $\det(\alpha I - A) = 0$ if and only if $\alpha I - A$ is not invertible if and only if there exists v such that $(\alpha I - A)v = 0$. \square

Such vectors are called eigenvectors; for a fixed α , the eigenvectors together with 0 form a subspace (“eigenspace”).

Exercise. If F is algebraically closed, then the multiplicity of an eigenvalue α is at least the dimension of the eigenspace corresponding to α . The multiplicity of α is exactly equal to the dimension of the **generalized eigenspace**

$$V_\alpha = \cup_{n=1}^{\infty} \ker(T - \alpha)^n.$$

A square matrix is nilpotent if and only if its only eigenvalue is 0.

As an example, if A is a 5×5 complex matrix which is nilpotent ($A^n = 0$ for some n), then $\chi_A = X^5$. The largest Jordan block has size m , where $\mu_A = X^m$. For each of the values $m = 5, 4$ and 1 there is only one similarity class possible for A : $J(0, 5)$; $\text{block}(J(0, 4), J(0, 1))$, and 0 , respectively. For each of the values $m = 3, 2$ there are two similarity classes possible: $\text{block}(J(0, 3), J(0, 2))$ and $\text{block}(J(0, 3), J(0, 1), J(0, 1))$ for $m = 3$, and for $m = 2$ there are either one or two blocks of size 2, and the rest are of size 1.

Exercise. The number of isomorphism classes of abelian groups of order 1024 is the same as the number of similarity classes of 10×10 complex matrices. Please do not calculate this number!

Exercise. If A is a nilpotent matrix, then $\text{Tr } A^n = 0$ for all $n \geq 0$. Show that the converse is false. (Look around: it’s true over some fields!)

Exercise. If $T : V \rightarrow V$ and W is a T -invariant subspace of V , then $\mu_{T|W} \mid \mu_T$ and $\chi_{T|W} \mid \chi_T$.

Exercise. If A and B are diagonalizable matrices and $AB = BA$, then they are “simultaneously” diagonalizable, that is, there exists D such that both DAD^{-1} and DBD^{-1} are diagonal.

2f. Bilinear forms

An important function of matrices is that they are the concrete versions not only of linear transformations, but also, in the square case, of bilinear forms. This leads to a rather different equivalence relation on the set of all $n \times n$ matrices over a field F . We briefly indicate the flavor here.

Definition. Let V be a finite-dimensional vector space over a field F . A bilinear form on V is a function $B : \mathcal{V} \times \mathcal{V} \rightarrow \mathcal{F}$ such that

$$\begin{aligned} B(v, \alpha_1 w_1 + \alpha_2 w_2) &= \alpha_1 B(v, w_1) + \alpha_2 B(v, w_2) \text{ and} \\ B(\alpha_1 v_1 + \alpha_2 v_2, w) &= \alpha_1 B(v_1, w) + \alpha_2 B(v_2, w) \end{aligned}$$

for all $v, v_1, v_2, w, w_1, w_2 \in V$; that is, for each $v \in V$ the functions $B(\cdot, v)$ and $B(v, \cdot)$ are linear functions $V \rightarrow F$.

Given a bilinear form B , if we choose a basis $C = \{v_1, \dots, v_n\}$ of V , we obtain the matrix of B with respect to the basis, namely

$$[B]_C = [B(v_i, v_j)]_{i,j=1}^n.$$

This matrix completely determines B , since $B(\sum_i \alpha_i v_i, \sum_j \beta_j v_j) = \sum_{i,j} \alpha_i \beta_j B(v_i, v_j)$ by the bilinearity of B . Furthermore, starting with any matrix A , one can define

$$B\left(\sum_i \alpha_i v_i, \sum_j \beta_j v_j\right) = \sum_{i,j} \alpha_i \beta_j A_{ij}$$

and obtain a bilinear form with matrix A . Thus as with linear transformations, the choice of a basis of V determines a bijective correspondence between the set of square matrices of size $\dim V$ and the set of bilinear forms on V .

However, if we want “equivalence” of matrices then to mean “same bilinear form, different basis”, we do not come out with similarity but rather the following equivalence relation on matrices:

$$A \sim A' \iff \exists \text{ an invertible } D \text{ such that } A' = D^T A D,$$

the “T” indicating transpose.

For if $w_j = \sum_i d_{ij} v_i$, $j = 1, \dots, n$, then

$$(w_i, w_j) = \left(\sum_k d_{ki} v_k, \sum_\ell d_{\ell j} v_\ell\right) = \sum_{k,\ell} d_{ki} (v_k, v_\ell) d_{\ell j}.$$

The “canonical form” problem for bilinear forms — i.e., the problem of classifying them — is thus completely different from the classification of linear transformations we have developed in the previous sections.

There is some connection, which can be formulated in terms of the dual space V^* .

Definition. If V is a finite-dimensional vector space over F , then $V^* = \text{Hom}_F(V, F)$.

The elements of $\text{Hom}_F(V, F)$ are called linear functionals on V . When V is infinite-dimensional, V^* can get out of hand and it is normally defined to be a subspace of $\text{Hom}_F(V, F)$, e.g. consisting in the case of Banach spaces, for example, of all bounded linear functionals.

For finite dimensions, however, the fact that direct sums are coproducts means that

$$\mathrm{Hom}_F(V_1 \oplus V_2, F) \cong_{\phi} \mathrm{Hom}_F(V_1, F) \times \mathrm{Hom}_F(V_2, F)$$

under the mapping $\phi(f) = (f|_{V_1}, f|_{V_2})$. Moreover, if V is 1-dimensional, and $0 \neq v_1 \in V$, then $\mathrm{Hom}_F(V, F) \cong F$, via $f \mapsto f(v_1)$. Note however this isomorphism is not “natural” but depends on the choice of v_1 . The displayed isomorphism is “natural”, however.

Choosing a basis v_1, \dots, v_n of V , in general, we therefore get an isomorphism

$$V^* = \mathrm{Hom}(V, F) = \mathrm{Hom}(Fv_1 \oplus \dots \oplus Fv_n, F) \cong \mathrm{Hom}(Fv_1, F) \times \dots \times \mathrm{Hom}(Fv_n, F) \cong F \oplus \dots \oplus F \cong V. \blacksquare$$

However the isomorphism we get here depends on the choice of basis. Following the progress of v_i from right to left, we see that under our isomorphism v_i corresponds to the functional v_i^* defined by

$$v_i^*(v_j) = \delta_{ij} \quad (= 1 \text{ or } 0 \text{ according as } i = j \text{ or } i \neq j).$$

The basis v_1^*, \dots, v_n^* of V^* is called the dual basis to v_1, \dots, v_n . We have

$$v = \sum_j v_j^*(v) v_j$$

for each $v \in V$; this is obvious for $v = v_i$ and then extends to all v by linearity.

Now a bilinear form B on V determines a linear transformation $\beta_B : V \rightarrow V^*$ by $\beta_B(v) = B(\cdot, v)$, that is,

$$\beta_B(v)(w) = B(v, w)$$

for all $v, w \in V$. Conversely a linear transformation $\beta_B : V \rightarrow V^*$ determines a bilinear form B by the same equation. Then for a basis $C = \{v_1, \dots, v_n\}$ and its dual basis $C^* = \{v_1^*, \dots, v_n^*\}$, we have $B(v_i, v_j) = \beta_B(v_i)(v_j)$, so

$$\beta_B(v_i) = \sum_j \beta_B(v_i)(v_j) v_j^* = \sum_j B(v_i, v_j) v_j^*.$$

Thus

$$[B]_C = [\beta_B]_{C^*}.$$

Exercise. A bilinear form B is called *nondegenerate (on the left)* if and only if the only $v \in V$ such that $B(v, w) = 0$ for all $w \in V$ is $v = 0$. Show that B is nonsingular if and only if β_B is injective (hence an isomorphism in the finite-dimensional case).

2g. Orthogonal complements

For simplicity let us write (v, w) for $B(v, w)$. In order for orthogonality to be an easy concept, we shall assume from now on that our bilinear forms B satisfy the condition:

$$\text{if } (v, w) = 0, \text{ then } (w, v) = 0. \quad 2A$$

This condition is satisfied, for example, if either

- a) B is **symmetric**: $(v, w) = (w, v)$ for all $v, w \in V$, or
- b) B is **alternating**: $(v, v) = 0$ for all $v \in V$; this implies that $(v, w) + (w, v) = (v + w, v + w) - (v, v) - (w, w) = 0$ so $(v, w) = -(w, v)$ for all $v, w \in V$, i.e., B is antisymmetric. (Conversely, antisymmetry implies the alternating condition as long as $1 \neq -1$, i.e., as long as the underlying field is not \mathbf{Z}_2 and is not any other field of “characteristic 2”.)

Under the hypothesis (2A), we can then define $v \perp w$, for $v, w \in V$, to mean $(v, w) = 0$, and have that \perp is a symmetric relation. Then for any subset (usually a subspace) $W \subseteq V$, we define

$$W^\perp = \{v \in V \mid v \perp w \text{ for all } w \in W\}.$$

The bilinearity of B implies immediately that W^\perp is a subspace of V . Notice the following properties:

- a) If $W_1 \subseteq W_2$, then $W_2^\perp \subseteq W_1^\perp$;
- b) $W \subseteq W^{\perp\perp}$.

We shall write $V = W_1 \perp W_2$ if $V = W_1 \oplus W_2$ and in addition, $W_1 \subseteq W_2^\perp$.

For the moment assume that B is nondegenerate. This means that the mapping $\beta_B : V \rightarrow V^*$ described in the previous section is an isomorphism of vector spaces. For any subspace $W \subseteq V$, the restriction mapping gives a homomorphism

$$\phi : V^* \rightarrow W^*,$$

which is surjective; choosing any complement X to W in V (i.e. writing $V = W \oplus X$) we may extend any linear functional $W \rightarrow V$ to a linear functional on V by prescribing that X go to 0. Let γ be the composite

$$\phi \circ \beta_B : V \rightarrow V^* \rightarrow W^*,$$

which is still surjective. The kernel of $\phi \circ \beta_B$ is precisely the set of those $v \in V$ such that $(v, w) = 0$ for all $w \in W$, i.e., W^\perp . Therefore $V/W^\perp \cong W^*$, and we have proved:

Proposition. *If B is a nondegenerate symmetric or alternating bilinear form on the finite-dimensional vector space V , then for any subspace $W \subseteq V$, we have*

$$\dim V = \dim W + \dim W^\perp.$$

Corollary. *If B is a nondegenerate symmetric or alternating bilinear form on the finite-dimensional vector space V , then for any subspace $W \subseteq V$, we have*

$$W = W^{\perp\perp}.$$

The proposition gives the equality of the dimensions of these subspaces, and we already know that $W \subseteq W^{\perp\perp}$.

Corollary. *Suppose that B is a nondegenerate symmetric or alternating bilinear form on the f.d.v.s. V . If $B|_W$ is nondegenerate, then $V = W \perp W^\perp$, and W^\perp is also nondegenerate.*

The first statement holds since the nondegeneracy of $B|_W$ means exactly that $W \cap W^\perp = 0$. By the corollary, $W \perp \cap W^{\perp\perp}$, proving the last statement.

Examples.

- Ex. A.** *The usual Euclidean inner product E on \mathbf{R}^n is a bilinear form; it is symmetric and nondegenerate. It is also positive definite ($(v, v) > 0$ for all $v \neq 0$), which is a stronger condition than nondegeneracy. There exist orthonormal bases C (i.e., $[E]_C = I$). Any positive definite symmetric bilinear form B on \mathbf{R}^n is equivalent to E , i.e., there exists an orthonormal basis for B . Namely, choose any $v \in V$, $v \neq 0$. Then $B(v, v) > 0$ and we let $c = B(v, v)^{-1/2}$ and $w_1 = cv$. Then $B(w_1, w_1) = 1$, and in particular B is nondegenerate on $\mathbf{R}w_1$. Let $V' = \mathbf{R}w_1^\perp$. Then $V = \mathbf{R}w_1 \perp V'$ by the proposition, and our assertion follows by induction.*
- Ex. B.** Let us consider an arbitrary bilinear symmetric form B on $V = \mathbf{R}^n$, dropping the positive definiteness and nondegeneracy conditions. First choose any vector-space complement W to V^\perp and write

$$V = V^\perp \oplus W = V^\perp \perp W;$$

observe also that W is nondegenerate, because W^\perp is orthogonal to both V^\perp and W , hence $W^\perp \subseteq V^\perp$, whence $W^\perp \cap W \subseteq V^\perp \cap W = 0$. This means that to obtain our form B , we first may study the nondegenerate case and then attach an arbitrary number of basis vectors orthogonal to everything.

In the nondegenerate case, we must have $(v, v) \neq 0$ for some v . For otherwise,

$$2(v, w) = (v + w, v + w) - (v, v) - (w, w) = 0$$

for all $v, w \in V$, by bilinearity and symmetry. Choosing any v with $(v, v) \neq 0$ and setting $w_1 = |(v, v)|^{-1/2}v$, we have $(w_1, w_1) = \pm 1$. Then $V' = \mathbf{R}w_1$ is nondegenerate, so $V = \mathbf{R}w_1 \perp V'$ and we may continue. In this way we have proved the existence part of Sylvester's Theorem:

Theorem. *Let B be a symmetric bilinear form on the real finite-dimensional vector space V . Then V has an orthogonal basis consisting of vectors v for which $(v, v) = 0$ or ± 1 . Moreover, the number of basis vectors for which (v, v) takes on each of these three values is uniquely determined by B , i.e., independent of the choice of basis.*

Proof. It remains to prove the uniqueness statement. Choose a basis C consisting of mutually orthogonal vectors $c_1, \dots, c_m, d_1, \dots, d_n$ and e_1, \dots, e_p , with $(c_i, c_i) = 1$, $(d_i, d_i) = 0$ and $(e_i, e_i) = -1$ for all i . Let $V_C^{+,0}$ be the span of the c 's and d 's; then $(v, v) \geq 0$ for all

$v \in C^{+,0}$. Likewise if V_C^- is the span of the e 's, then $(v, v) < 0$ for all $0 \neq v \in V_C^-$. We can make the same construction for any basis as in the theorem. We must prove that the dimensions m, n, p are unique; but they are determined by the four dimensions $V_C^{\pm,0}$ and V_C^\mp , so it is enough to prove the equality of these. As $V_C^{\pm,0}$ and V_C^\mp are complementary, it is enough to show the uniqueness of $\dim V_C^{\pm,0}$. Let D be another such basis. Then $V_C^{+,0} \cap V_D^- = 0$, since its nonzero vectors v satisfy $(v, v) \geq 0$ and $(v, v) < 0$ simultaneously. Therefore $\dim V_C^{+,0} \leq \dim V - \dim V_D^- = \dim V_D^{+,0}$. By symmetry, we have equality. QED

Ex. C. Suppose that B is a nondegenerate alternating form (here F can be any field). Let $v \in V, v \neq 0$. Since B is nondegenerate, $(v, v') \neq 0$ for some $v' \in V$. Replacing v' by a scalar multiple we may arrange that $(v, v') = 1$. Such a pair v, v' is called a hyperbolic pair, and satisfies

$$(v, v) = (v', v') = 0, \quad (v, v') = 1, \quad (v', v) = -1.$$

Obviously $v' \notin Fv$. Let $W = Fv + Fv'$. Then $\dim W = 2$ (W is a “hyperbolic plane”), and the matrix of the form B with respect to $\{v, v'\}$ is

$$\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}.$$

In particular $B|_W$ is nondegenerate, so $B|_{W^\perp}$ is too by a corollary in the previous section. By induction we have proved:

Theorem. *Let B be a nondegenerate alternating bilinear form on a f.d.v.s. V . Then V is the orthogonal sum of hyperbolic planes for B . In particular, $\dim V$ is even.*

Corollary. *Any two nondegenerate alternating bilinear forms on a f.d.v.s. V are equivalent.*

2h. Sesquilinear and Hermitian Forms.

The results of this section are specific to vector spaces over \mathbf{C} . For them, since $z\bar{z} \geq 0$ for all $z \in \mathbf{C}$, it is more favorable to consider “hermitian” forms. These are an example of sesquilinear forms (“ $\frac{11}{2}$ -linear”).

A mapping $\phi : V \rightarrow W$ of complex vector spaces is called conjugate-linear if and only if $\phi(v + v') = \phi(v) + \phi(v')$ (ϕ is “bi-additive”) and $\phi(\alpha v) = \bar{\alpha}\phi(v)$ for all $v, v' \in v$ and $\alpha \in \mathbf{C}$. Here $\bar{\alpha}$ is the complex conjugate of α .

Exercise. For a conjugate-linear mapping $\phi : V \rightarrow W$, $\ker \phi$ and $\text{im } \phi$ are subspaces and $\dim V = \dim \ker \phi + \dim \text{im } \phi$.

Definition. A hermitian form on a complex vector space V is a form $B : V \times V \rightarrow V$ which is linear in the first argument (for each fixed value of the second) and conjugate linear in the second (for each fixed value of the first), and which is hermitian-symmetric:

$$(w, v) = \overline{(v, w)}$$

for all $v, w \in V$.

Of these three conditions, either of the first two follows from the others. Without the hermitian-symmetry, the form still could be called sesquilinear.

For such a form, $(v, w) = 0 \iff (w, v) = 0$, so our general results on orthogonality from the previous section still hold. In particular, the definition of nondegeneracy is the same, and if V is nondegenerate, then $V = W \perp W^\perp$ for any nondegenerate subspace $W \subseteq V$.

We have $(v, v) = \overline{(v, v)} \in \mathbf{R}$ by the hermitian-symmetry, for all $v \in V$. We again call the form positive-definite if and only if

$$(v, v) > 0 \text{ for all } 0 \neq v \in V.$$

If a basis $C = \{v_1, \dots, v_n\}$ of V is given, then a hermitian form on V is determined by its values (v_i, v_j) :

$$\left(\sum_i \alpha_i v_i, \sum_j \beta_j v_j \right) = \sum_{i,j} \alpha_i \bar{\beta}_j (v_i, v_j).$$

So we can again speak of the matrix $[B]_C$ of B with respect to C .

Proposition. The matrix of a hermitian form is hermitian-symmetric.

Proof. By definition A is hermitian-symmetric if and only if $A^T = \bar{A}$, the matrix obtained by replacing each entry by its complex conjugate. In our situation this is just $(v_i, v_j) = \overline{(v_j, v_i)}$. QED

Exercise. If V is a f.d.v.s over \mathbf{C} with a positive-definite hermitian form, then V has an orthonormal basis, i.e., its matrix with respect to some basis is I . Any two positive-definite hermitian forms on V are equivalent.

A hermitian form B gives rise to a conjugate-linear mapping $\beta_B : V \rightarrow V^*$, by $\beta_B(v)(w) = (w, v)$. For each fixed v , this expression is linear in w , so ϕ is defined. The axioms imply that β_B is conjugate-linear. The form is nondegenerate if and only if this mapping is injective (equivalently, surjective; see the exercise above).

Notice that if V_0 is a real vector space with a symmetric form, with matrix A (which is then symmetric) relative to a basis $C = \{v_1, \dots, v_n\}$, then we may form a complex vector space V with the same basis, containing V_0 as an \mathbf{R} -subspace, and extend the given form on V_0 to a hermitian form on V , defined with respect to C by the same matrix (which is hermitian-symmetric, being real symmetric).

2i. The Spectral Theorem

We have earlier derived an algebraic necessary and sufficient condition for a matrix, or a linear transformation $V \rightarrow V$, to be diagonalizable: its minimal polynomial is square-free. When V comes equipped with a bilinear form (i.e., some geometry), there may be geometric sufficient conditions for diagonalizability. We give an important one here: the spectral theorem.

Suppose then that $T : V \rightarrow V$ is a linear transformation, and V is a complex vector space with a nondegenerate hermitian form B . The mapping β_B is then a bijection: $V \rightarrow V^*$. We use β_B to define the “adjoint” T^* of T .

Lemma. *There exists a unique linear transformation $T^* : V \rightarrow V$ such that*

$$(Tv, w) = (v, T^*w)$$

for all $v, w \in V$.

Proof. Our mapping $\beta_B(v)(w) = (w, v)$ is a bijection $\beta_B : V \rightarrow V^*$. We therefore want $\beta_B(T^*w)(v) = (Tv, w)$. For each $w \in V$, the mapping $\gamma_w : v \mapsto (Tv, w)$ is linear in v , so lies in V^* . Thus we want $\beta_B(T^*(w)) = \gamma_w$ for each w . As β_B is bijective, there is a unique function T^* satisfying the desired equation. Then $(v, T^*(\alpha w)) = (Tv, \alpha w) = \overline{\alpha}(Tv, w) = \overline{\alpha}(v, T^*(w)) = (v, \alpha T^*(w))$ for all $v, w \in V$. The uniqueness of T^* shows $T^*(\alpha w) = \alpha T^*(w)$. The additivity of T^* follows similarly. \square

It is easy to check that $(T_1 + T_2)^* = T_1^* + T_2^*$, $(\alpha T)^* = \overline{\alpha}T^*$ and $(T^*)^* = T$. E.g., $(v, T^{**}w) = (T^*v, w) = \overline{(w, T^*v)} = (Tw, v) = (v, Tw)$ for all $v, w \in V$, so $T^{**}w = Tw$ for all w .

The following elementary fact is important:

Proposition. $\ker T = (\text{im } T^*)^\perp$ and $(\ker T)^\perp = \text{im } T^*$.

Proof. Since $\dim V$ is finite and the form is nondegenerate, the second statement follows from the first by taking orthogonal complements. For any $v \in V$, the following statements

are equivalent: $v \in \ker T$; $Tv = 0$; $(Tv, w) = 0$ for all $w \in W$; $(v, T^*w) = 0$ for all $w \in W$; $v \in (\operatorname{im} T^*)^\perp$.

Definition. In the above situation, T is called

- a) self-adjoint if and only if $T^* = T$;
- b) unitary if and only if $T^* = T^{-1}$;
- c) normal if and only if $TT^* = T^*T$.

Thus self-adjoint and unitary transformations are normal.

Exercise. With respect to an orthonormal basis C of V , the matrices of T and T^* satisfy

$$[T^*]_C^C = \overline{[T]_C^C}.$$

Moreover, if $A = [T]_C^C$, then T is self-adjoint, unitary or normal according as $A^T = \overline{A}$, $A^T \overline{A} = I$, or $A^T \overline{A} = \overline{A} A^T$.

Spectral Theorem. Let V be a finite-dimensional complex vector space with a nondegenerate hermitian bilinear form. Let $T : V \rightarrow V$ be a linear transformation. Then T is normal if and only if V has an orthonormal basis consisting of eigenvectors of T . Moreover, if T is self-adjoint, all its eigenvalues are real; if T is unitary, then all its eigenvalues are on the unit circle.

We prove this as the culmination of several simple remarks.

1. If W is a subspace of V and $T(W) \subseteq W$, then $T^*(W^\perp) \subseteq W^\perp$.

Proof. For any $u \in W^\perp$ and $v \in W$, $(v, T^*(u)) = (Tv, u) = 0$.

2. If W is a subspace of V , and both W and W^\perp are T -invariant, then W is T^* -invariant and $(T|_W)^* = T^*|_W$.

Proof. By 1, W is T^* -invariant. For any $u, v \in W$, $(u, (T|_W)^*v) = ((T|_W)u, v) = (Tu, v) = (u, T^*v)$. Hence by the uniqueness of $(T|_W)^*$, $T^*v = (T|_W)^*v$.

3. If T is normal and W is a subspace of V such that both W and W^\perp are T -invariant, then $T|_W$ is normal.

Proof. Use 2 and the definition of normality.

4. If T is normal then so is $T - \lambda id_V$ for any $\lambda \in \mathbf{C}$.

Proof. $(T - \lambda id_V)^* = T^* - \overline{\lambda} id_V$ commutes with T and hence with $T - \lambda id_V$.

5. If $TU = UT$ then T leaves $\ker U$ and $\operatorname{im} U$ invariant.

Proof. $T(Uv) = U(Tv)$ for all $v \in V$. Let v range over $\ker U$ to get $T(\ker U) \subseteq \ker U$. Let v range over V to get $T(\operatorname{im} U) \subseteq \operatorname{im} U$.

Now we prove the “hard” direction of the spectral theorem. Suppose that T is normal. Using the fundamental theorem of algebra, take an eigenvalue λ of T . By the proposition,

$$V = \ker(T - \lambda id_V) \perp \operatorname{im}((T - \lambda id_V)^*) = \ker(T - \lambda id_V) \perp \operatorname{im}(T^* - \overline{\lambda} id_V).$$

Since T is normal, it commutes with $T^* - \bar{\lambda}id_V$, so both the displayed subspaces are T -invariant by 5. By 3, the restriction of T to each is normal. By choice the first is nontrivial, so by induction on $\dim V$ the second has an orthonormal basis of eigenvectors of T . Obviously $\ker(T - \lambda id_V)$ does too, and assembling these two bases we obtain the desired orthonormal basis of V .

If T is self-adjoint, then for any eigenvector v and corresponding eigenvalue λ ,

$$\lambda(v, v) = (Tv, v) = (v, Tv) = \overline{(Tv, v)}$$

is real. Also $(v, v) = \overline{(v, v)}$ is real, so λ is as well. Likewise if T is unitary, then T is invertible by definition, and the equation $Tv = \lambda v$ implies that $T^{-1}v = \lambda^{-1}v$. Now $\lambda(v, v) = (Tv, v) = (v, T^{-1}v) = \bar{\lambda}^{-1}(v, v)$, so $\lambda = \bar{\lambda}^{-1}$ is on the unit circle.

Conversely, if there exists an orthonormal basis v_1, \dots, v_n of eigenvectors of T , say $Tv_i = \lambda_i v_i$, then we put $Uv_i = \bar{\lambda}_i v_i$ for each i and extend U to a linear transformation on V ; then

$$(Tv_i, v_j) = \delta_{ij} \lambda_i = \delta_{ij} \lambda_j = (v_i, Uv_j)$$

for all i and j , and so $U = T^*$. But obviously $TU = UT$ since the matrices of both T and U with respect to C are diagonal. Hence T is normal. \square

In matrix terms, if we fix an orthonormal basis C and set $A = [T]_C^C$, then T is self-adjoint (unitary, normal) iff $A^T = \bar{A}$ ($A^T = \bar{A}^{-1}$, $A^T \bar{A} = \bar{A} A^T$). The “transition” matrix from C to another basis is easily seen to be unitary if and only if the other basis is orthonormal. Thus in matrix terms we have:

Matrix Spectral Theorem. *Let A be a square complex matrix. Then A is normal if and only if there is a unitary matrix U such that UAU^{-1} is diagonal. Moreover, if A is self-adjoint (resp. unitary), then the eigenvalues of A are real (resp. on the unit circle).*

Exercise. *Adapt the theorem and proof to the case of a real vector space with a positive definite symmetric bilinear (real-valued) form, and a self-adjoint linear transformation T . (Hint. The matrix of T being symmetric, all eigenvalues of T are known a priori to be real, by the theorem already proved.)*