Chapter 2. Basics of Representations and Characters

A. Definitions and examples.

We start with the notion of a *group*: a set G with an associative multiplication $s, t \to st$, an identity id, and inverses s^{-1}. A *representation* ρ of G assigns an invertible matrix $\rho(s)$ to each $s \in G$ in such a way that the matrix assigned to the product of two elements is the product of the matrices assigned to each element: $\rho(st) = \rho(s)\rho(t)$. This implies that $\rho(id) = I$, $\rho(s^{-1}) = \rho(s)^{-1}$. The matrices we work with are all invertible and are considered over the real or complex numbers. We thus regard ρ as a homomorphism from G to $GL(V)$ — the linear maps on a vector space V. The dimension of V is denoted d_ρ and called the *dimension of* ρ.

If W is a subspace of V stable under G (i.e., $\rho(s)W \subset W$ for all $s \in G$), then ρ restricted to W gives a *subrepresentation.* Of course the zero subspace and the subspace $W = V$ are trivial subrepresentations. If the representation ρ admits no non-trivial subrepresentation, then ρ is called *irreducible.* Before going on, let us consider an example.

Example. S_n the permutation group on n letters.

This is the group S_n of $1-1$ mappings from a finite set into itself; we will use the notation $\left[\begin{smallmatrix}1 & 2 & \cdots & n\\ \pi(1) & \pi(2) & \cdots & \pi(n)\end{smallmatrix}\right]$. Here are three different representations. There are others.

(a) The *trivial representation* is 1-dimensional. It assigns each permutation to the identity map $\rho(\pi)x = x$.

(b) The *alternating representation* is also 1-dimensional. To define it, recall the sign of a permutation π is $+1$ if π can be written as a product or an even number of transpositions $\pi = \overbrace{(ab)(cd)\dots(ef)}^{\textbf{even \# of factors}}$. The sign of π is -1 if π can be written as an odd number of transpositions. Elementary books on group theory show that $\text{sgn}(\pi)$ is well defined and that $\text{sgn}(\pi_1\pi_2) = \text{sgn}(\pi_1)\text{sgn}(\pi_2)$. It follows that $x \to \text{sgn}(\pi)\cdot x$ is a 1-dimensional representation.

(c) The *permutation representation* is an n-dimensional representation. To define it, consider the standard basis $e_1, \dots, e_n$ of $\mathbb{R}^n$. It is only necessary to define the linear map $\rho(\pi)$ on the basis vectors. Define $\rho(\pi)e_j = e_{\pi(j)}$. The matrix of a linear map L is defined by $L(e_j) = \Sigma L_{ij}e_i$. With this convention, $\rho(\pi)_{ij}$ is zero or one. It is one if and only if $\pi(j) = i$, so $\rho(\pi)_{ij} = \delta_{i\pi(j)}$. I will write permutations right to left. Thus $\pi_2\pi_1$ means first perform π_1 and then perform π_2.

We will also be using cycle notation for permutations, $(a_1a_2\dots a_k)$ means $a_1 \to a_2$, $a_2 \to a_3 \dots a_k \to a_1$. Thus $(1\ 2)(2\ 3) = (1\ 2\ 3)$ (and *not* $(1\ 3\ 2)$).

Under the permutation representation this last equation transforms into

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}.$$

Observe that the permutation representation has subspaces that are sent into themselves under the action of the group: the 1-dimensional space spanned by $e_1 + \cdots + e_n$, and its complement $W = \{\underset{\sim}{x} \in \mathbb{R}^n : \ \Sigma x_i = 0\}$ both have this property. A representation ρ is *irreducible* if there is no non-trivial subspace $W \subset V$ with $\rho(s)W \subset W$ for all $s \in G$. Irreducible representations are the basic building blocks of any representation, in the sense that any representation can be decomposed into irreducible representations (Theorem 2 below). It turns out (Exercise 2.6 in Serre or "a useful fact" in 7-A below) that the restriction of the permutation representation to W is an irreducible $n-1$-dimensional representation. For S_3, there are only three irreducible representations; the trivial, alternating, and 2-dimensional representation (Corollary 2 of Proposition 5 below).

EXPLICIT COMPUTATION OF THE 2-DIMENSIONAL REPRESENTATION OF S_3

Let $W = \{x \in \mathbb{R}^3 : \ x_1 + x_2 + x_3 = 0\}$. Let $w_1 = e_1 - e_2$, $w_2 = e_2 - e_3$. Clearly $w_i \in W$. They form a basis for W, for if $v = xe_1 + ye_2 + ze_3 \in W$, then $v = xe_1 + ye_2 + (-x-y)e_3 = x(e_1 - e_2) + (x+y)(e_2 - e_3)$. In this case, it is easy to argue that the restriction of the permutation representation to W is irreducible. Let (x, y, z) be nonzero in W (suppose, say $x \neq 0$) and let W_1 be the span of this vector. We want to show that W_1 is not a subrepresentation. Suppose it were. Then, we would have $(1, y', z')$ and so $(y', 1, z')$ and so $(1 - y', y' - 1, 0)$ in W_1. If $y' \neq 1$, then $e_1 - e_2$ and so $e_2 - e_3$ and $e_1 - e_2$ are in W_1. So $W_1 = W$. If $y' = 1$, then $(1, 1, -2) \in W_1$. Permuting the last two coordinates and subtracting shows $e_2 - e_3$ and so $e_1 - e_2$ are in W_1, so $W_1 = W$.

Next consider the action of π on this basis

π	$\rho(\pi)w_1$	$\rho(\pi)w_2$	$\rho(\pi)$
id	w_1	w_2	$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$
(1 2)	$-w_1$	$w_1 + w_2$	$\begin{pmatrix} -1 & 1 \\ 0 & 1 \end{pmatrix}$
(2 3)	$w_1 + w_2$	$-w_2$	$\begin{pmatrix} 1 & 0 \\ 1 & -1 \end{pmatrix}$
(1 3)	$-w_2$	$-w_1$	$\begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}$
(1 2 3)	w_2	$-(w_1 + w_2)$	$\begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}$
(1 3 2)	$-(w_1 + w_2)$	w_1	$\begin{pmatrix} -1 & 1 \\ -1 & 0 \end{pmatrix}$

CONVOLUTIONS AND FOURIER TRANSFORMS

Throughout we will use the notion of convolution and the Fourier transform. Suppose P and Q are probabilities on a finite group G. Thus $P(s) \geq 0, \Sigma_s P(s) = 1$. By the *convolution* $P * Q$ we mean the probability $P * Q(s) = \Sigma_t P(st^{-1})Q(t)$: "first pick t from Q, then independently pick u from P and form the product ut." Note that in general $P * Q \neq Q * P$. Let the order of G be denoted $|G|$. The *uniform* distribution on G is $U(s) = 1/|G|$ for all $s \in G$. Observe that $U * U = U$ but this does not characterize U-the uniform distribution on any subgroup satisfies this as well. However, $U * P = U$ for any P and this characterizes U.

Let P be a probability on G. The *Fourier transform* of P at the representation ρ is the matrix

$$\hat{P}(\rho) = \Sigma_s P(s)\rho(s).$$

The same definitions works for any function P. In Proposition 11, we will show that as ρ ranges over irreducible representations, the matrices $\hat{P}(\rho)$ determine P.

EXERCISE 1. Let ρ be any representation. Show $\widehat{P * Q}(\rho) = \hat{P}(\rho)\hat{Q}(\rho)$.

EXERCISE 2. Consider the following probability (random transpositions) on S_3

$$P(\text{id}) = p,\ P(12) = P(13) = P(23) = (1-p)/3.$$

Compute $\hat{T}(\rho)$ for the three irreducible representations of S_3. (You'll learn something.)

B. THE BASIC THEOREMS.

This section follows Serre quite closely. In particular, the theorems are numbered to match Serre.

Theorem 1. *Let $\rho : G \to GL(V)$ be a linear representation of G in V and let W be a subspace of V stable under G. Then there is a complement W^0 (so $V = W + W^0$, $W \cap W^0 = 0$) stable under G.*

Proof. Let $<,>_1$ be a scalar product on V. Define a new inner product by $< u, v >= \Sigma_s < \rho(s)u,\ \rho(s)v >_1$. Then $<,>$ is invariant: $< \rho(s)u,\ \rho(s)v > = < u, v >$. The orthogonal complement of W in V serves as W^0. □

Remark 1. We will say that the representation V splits into the *direct sum* of W and W^0 and write $V = W \oplus W^0$. The importance of this decomposition cannot be overemphasized. It means we can study the action of G on V by separately studying the action of G on W and W^0.

Remark 2. We have already seen a simple example: the decomposition of the permutation representation of S_n. Here is a second example. Let S_n act on $\mathbb{R}^2$ by $\rho(\pi)(x, y) = \text{sgn}(\pi)(x, y)$. The subspace $W = \{(x, y) : x = y\}$ is invariant. Its complement, under the usual inner product, is $W^0 = \{(x, y) : x = -y\}$ is also invariant. Here, the complement is not unique. For example, $W^{00} = \{(x, y) : 2x = -y\}$ is also an invariant complement.

Remark 3. The proof of Theorem 1 uses the "averaging trick;" it is the standard way to make a function of several variables invariant. The second most widely used approach, defining $< u, v >_2 = \max_g < \rho(g)u,\ \rho(g)v >_1$, doesn't work here since $<,>_2$ is not still an inner product.

Remark 4. The invariance of the scalar product $<,>$ means that if e_i is chosen as an orthonormal basis with respect to $<,>$, then $< \rho(s)e_i,\ \rho(s)e_j > = \delta_{ij}$. It follows that the matrices $\rho(s)$ are unitary. Thus, if ever we need to, we may assume our representations are unitary.

Remark 5. Theorem 1 is true for compact groups. It can fail for noncompact groups. For example, take $G = \mathbb{R}$ under addition. Take V as the set of linear polynomials $ax + b$. Define $\rho(t)f(x) = f(x + t)$. The constants form a non-trivial subspace with no invariant complement. Theorem 1 can also fail over a finite field.

Return to the setting of Theorem 1 by induction we get:

Theorem 2. *Every representation is a direct sum of irreducible representations.*

There are two ways of taking two representations (ρ, V) and (η, W) of the same group and making a new representation. The *direct sum* constructs the vector space $V \oplus W$ consisting of all pairs (v, w), $v \in V$, $w \in W$. The direct sum representation $\rho \oplus \eta(s)(v, w) = (\rho(s)v, \eta(s)w)$. This has dimension $d_\rho + d_\eta$ and clearly contains invariant subspaces equivalent to V and W.

The *tensor product* constructs a new vector space $V \otimes W$ of dimension $d_\rho d_\eta$ which can be defined as the set of formal linear combinations $v \otimes w$ subject to the rules $(av_1 + bv_2) \otimes w = a(v_1 \otimes w) + b(v_2 \otimes w)$ (and symmetrically). If $v_1, \ldots, v_a$ and $w_1, \ldots, w_b$ are a basis for V and W, then $v_i \otimes w_j$ is a basis for

$V \otimes W$. Alternatively, $V \otimes W$ can be regarded as the set of a by b matrices were $v \otimes w$ has ij entry $\lambda_i \mu_j$ if $v = \Sigma \lambda_i v_i$, $w = \Sigma \mu_j w_j$. The representation operates as $\rho \otimes \eta(s)(v \otimes w) = \rho(s)v \otimes \eta(s)w$.

The explicit decomposition of tensor products into direct sums is a booming business. New irreducible representations can be constructed from known ones by tensoring and decomposing.

The notion of the *character* of a representation is extraordinarily useful. If ρ is a representation, define $\chi_\rho(s) = \text{Tr } \rho(s)$. This doesn't depend on the basis chosen for V because the trace is basis free.

PROPOSITION 1. *If χ is the character of a representation ρ of degree d then*

$$(1)\ \chi(\text{id}) = \text{d}; \quad (2)\ \chi(\text{s}^{-1}) = \chi(\text{s})^*; \quad (3)\ \chi(\text{tst}^{-1}) = \chi(\text{s}).$$

Proof. (1) $\rho(\text{id}) = \text{id}$. (2) First $\rho(s^a) = I$ for a large enough. It follows that the eigenvalues λ_i of $\rho(s)$ are roots of unity. Then, with * complex conjugation,

$$\chi(s)^* = \text{Tr } \rho(s)^* = \Sigma \lambda_i^* = \Sigma 1/\lambda_i = \text{Tr } \rho(s)^{-1} = \text{Tr } \rho(s^{-1}) = \chi(s^{-1}).$$

(3) $\text{Tr}(AB) = \text{Tr}(BA)$. □

PROPOSITION 2. *Let $\rho_1 : G \to GL(V_1)$ and $\rho_2 : G \to GL(V_2)$ be representations with characters χ_1 and χ_2. Then* (1) the character of $\rho_1 \oplus \rho_2$ is $\chi_1 + \chi_2$ and (2) the character of $\rho_1 \otimes \rho_2$ is $\chi_1 \cdot \chi_2$.

Proof. (1) Choose a basis so the matrix of $\rho_1 \oplus \rho_2$ is given as $\begin{pmatrix} \rho_1 & 0 \\ 0 & \rho_2 \end{pmatrix}$. (2) The matrix of the linear map $\rho_1(s) \otimes \rho_2(s)$ is the tensor product of the matrices $\rho_1(s)$ and $\rho_2(s)$. This has diagonal entries $\rho_1^{i_1 i_1}(s)\rho_2^{j_2 j_2}(s)$. □

Consider two representations ρ based on V and τ based on W. They are called *equivalent* if there is a 1-1 linear map f from V onto W such that $\tau_s \circ f = f \circ \rho_s$. For example, consider the following two representations of the symmetric group: ρ, the 1-dimensional trivial representation (so $V = \mathbb{R}$ and $\rho(\pi)x = x$) and τ, the restriction of the n-dimensional permutation representation to the subspace W spanned by the vector $e_1 + \cdots + e_n$. Here $\tau(\pi)x(e_1 + \cdots + e_n) = x(e_1 + \cdots + e_n)$. The isomorphism can be taken as $f(x) = x(e_1 + \cdots + e_n)$.

The following "lemma" is one of the most used elementary tools.

SCHUR'S LEMMA

Let $\rho^1 : G \to GL(V_1)$ and $\rho^2 : G \to GL(V_2)$ be two irreducible representations of G, and let f be a linear map of V_1 into V_2 such that

$$\rho_s^2 \circ f = f \circ \rho_s^1 \text{ for all } s \in G.$$

Then

(1) If ρ^1 and ρ^2 are not equivalent, we have $f = 0$.

(2) If $V_1 = V_2$ and $\rho^1 = \rho^2$, f is a constant times the identity.

Proof. Observe that the kernel and image of f are both invariant subspaces. For the kernel, if $f(v) = 0$, then $f\rho_s^1(v) = \rho_s^2 f(v) = 0$, so $\rho_s^1(v)$ is in the kernel. For the image, if $w = f(v)$, then $\rho_s^2(w) = f\rho_s^1(v)$ is in the image too. By irreducibility, both kernel and image are trivial or the whole space. To prove (1) suppose $f \neq 0$. Then Ker $= 0$, image $= V_2$ and f is an isomorphism. To prove (2) suppose $f \neq 0$ (if $f = 0$ the result is true). Then f has a non-zero eigenvalue λ. The map $f^1 = f - \lambda I$ satisfies $\rho_s^2 f^1 = f^1 \rho_s^1$ and has a non-trivial kernel, so $f^1 \equiv 0$. □

EXERCISE 3. Recall that the uniform distribution is defined by $U(s) = 1/|G|$, where $|G|$ is the order of the group G. Then at the trivial representation $\hat{U}(\rho) = 1$ and at any non-trivial irreducible representation $\hat{U}(\rho) = 0$.

There are a number of useful ways of rewriting Schur's lemma. Let $|G|$ be the order of G.

COROLLARY 1. *Let h be any linear map of V_1 into V_2. Let*

$$h^0 = \frac{1}{|G|}\Sigma(\rho_t^2)^{-1} h \rho_t^1.$$

Then
(1) *If ρ^1 and ρ^2 are not equivalent, $h^0 = 0$.*
(2) *If $V_1 = V_2$ and $\rho^1 = \rho^2$, then h^0 is a constant times the identity, the constant being $Tr\ h/d_\rho$.*

Proof. For any s, $\rho_{s^{-1}}^2 h^0 \rho_s^1 = \frac{1}{|G|}\Sigma \rho_{s^{-1}t^{-1}}^2 h \rho_{ts}^1 = \frac{1}{|G|}\Sigma(\rho_{ts}^2)^{-1} h \rho_{ts}^1 = h^0$. If ρ^1 and ρ^2 are not isomorphic then $h^0 = 0$ by part (1) of Schur's lemma. If $V_1 = V_2$, $\rho_1 = \rho_2 = \rho$, then by part (2), $h^0 = cI$. Take the trace of both sides and solve for c. □

The object of the next rewriting of Schur's lemma is to show that the matrix entries of the irreducible representations form an orthogonal basis for all functions on the group G. For compact groups, this sometimes is called the Peter-Weyl theorem.

Suppose ρ^1 and ρ^2 are given in matrix form

$$\rho_t^1 = (r_{i_1 j_1}(t)), \quad \rho_t^2 = (r_{i_2 j_2}(t)).$$

The linear maps h and h^0 are defined by matrices $x_{i_2 i_1}$ and $x_{i_2 i_1}^0$. We have

$$x_{i_2 i_1}^0 = \frac{1}{|G|}\sum_{t j_1 j_2} r_{i_2 j_2}(t^{-1}) x_{j_2 j_1} r_{j_1 i_1}(t).$$

In case (1), $h^0 \equiv 0$ for *all* choices of h. This can only happen if the coefficients of $x_{j_2 j_1}$ are all zero. This gives

Corollary 2. *In case* (1)

$$\frac{1}{|G|}\sum_{t\in G} r_{i_2j_2}(t^{-1})r_{j_1i_1}(t)=0 \text{ for all } i_1,\ i_2,\ j_1,\ j_2.$$

Corollary 3. *In case* (2)

$$\frac{1}{|G|}\sum_{t\in G} r_{i_2j_2}(t^{-1})r_{j_1i_1}(t)=\begin{cases}\frac{1}{d_\rho} & \text{if } i_1=i_2 \text{ and } j_1=j_2\\ 0 & \text{otherwise.}\end{cases}$$

Proof. In case (2), $h^0=\lambda I$, or $x^0_{i_2i_1}=\lambda\delta_{i_2i_1}$, with $\lambda=\frac{1}{d\rho}\Sigma\delta_{j_2j_1}x_{j_2j_1}$. This gives

$$\frac{1}{|G|}\sum_{tj_1j_2} r_{i_2j_2}(t^{-1})x_{j_2j_1}r_{j_1i_1}(t)=\frac{\delta_{i_1i_2}}{d_\rho}\sum_{j_1j_2}\delta_{j_1j_2}x_{j_1j_2}.$$

Since h is arbitrary, we get to equate coefficients of $x_{j_2j_1}$. □

Orthogonality Relations for Characters.

Corollaries 2 and 3 above assume a neat form if the representations involved are unitary, so that $r(s)^*=r(s^{-1})$ where * indicates conjugate transpose. Remark 4 to Theorem 1 implies this can always be assumed without loss of generality. Introduce the usual inner product on functions

$$(\phi|\psi)=\frac{1}{|G|}\Sigma\phi(t)\psi(t)^*.$$

With this inner product, Corollaries 2 and 3 say that the matrix entries of the unitary irreducible representations are orthogonal as functions from G into C.

Theorem 3. *The characters of irreducible representations are orthonormal.*

Proof. Let ρ be irreducible with character χ and given in matrix form by $\rho_t=r_{ij}(t)$. So $\chi(t)=\Sigma r_{ii}(t)$, $(\chi|\chi)=\Sigma_{i,j}(r_{ii}|r_{jj})$. From Corollary 3 above $(r_{ii}|r_{jj})=\frac{1}{d_\rho}\delta_{ij}$. If χ, χ' are characters of non-equivalent representations, then in obvious notation

$$(\chi|\chi')=\sum_{ij}(r_{ii}|r'_{jj}).$$

Corollary 2 shows each term $(r_{ii}|r'_{jj})=0$. □

Theorem 4. *Let ρ, V be a representation of G with character ϕ. Suppose V decomposes into a direct sum of irreducible representations:*

$$V=W_1\oplus\cdots\oplus W_k.$$

Then, if W is an irreducible representation with character χ, the number of W_i equivalent to W equals $(\phi|\chi)$.

Proof. Let χ_i be the character of W_i. By Proposition 2, $\phi = \chi_1 + \cdots + \chi_k$, and $(\chi_i|\chi)$ is 0 or 1 as W_i is not, or is, equivalent to W. □

COROLLARY 1. *The number of W_i isomorphic to W does not depend on the decomposition (e.g., the basis chosen).*

Proof. $(\phi|\chi)$ does not depend on the decomposition. □

COROLLARY 2. *Two representations with the same character are equivalent.*

Proof. They each contain the same irreducible representations the same number of times. □

We often write $V = m_1 W_1 \oplus \cdots \oplus m_n W_n$ to denote that V contains W_i m_i times. Observe that $(\phi|\phi) = \Sigma m_i^2$. This sum equals 1 if and only if ϕ is the character of an irreducible representation.

Theorem 5. *If ϕ is the character of a representation then $(\phi|\phi)$ is a positive integer and equals 1 if and only if the representation is irreducible.*

EXERCISE 4. Do exercises 2.5 and 2.6 in Serre. Use 2.6 to prove that the $n-1$-dimensional part of the n-dimensional permutation representation is irreducible. (Another proof follows from "A useful fact" in Chapter 7-A.)

C. DECOMPOSITION OF THE REGULAR REPRESENTATION AND FOURIER INVERSION.

Let the irreducible characters be labelled χ_i. Suppose their degrees are d_i. The *regular representation* is based on a vector space with basis $\{e_s\}$, $s \in G$. Define $\rho_s(e_t) = e_{st}$. Observe that the underlying vector space can be identified with the set of all functions on G.

PROPOSITION 5. *The character r_G of the regular representation is given by*

$$r_G(1) = |G|$$
$$r_G(s) = 0, \quad s \neq 1.$$

Proof. $\rho_1(e_s) = e_s$ so Tr $\rho_1 = |G|$. For $s \neq 1$, $\rho_s e_t = e_{st} \neq e_t$ so all diagonal entries of the matrix for ρ_s are zero. □

COROLLARY 1. *Every irreducible representation W_i is contained in the regular representation with multiplicity equal to its degree.*

Proof. The number in question is

$$(r_G|\chi_i) = \frac{1}{|G|} \sum_{s \in G} r_G(s) \chi_i^*(s) = \chi_i^*(1) = d_i. \qquad \square$$

Remark. Thus, in particular, there are only finitely many irreducible representations.

COROLLARY 2.
(a) *The degrees d_i satisfy $\Sigma d_i^2 = |G|$.*
(b) *If $s \in G$ is different from 1, $\Sigma d_i \chi_i(s) = 0$.*

Proof. By Corollary 1, $r_G(s) = \Sigma d_i \chi_i(s)$. For (a) take $s = 1$, for (b) take any other s. □

In light of remark 4 to Theorem 1, we may always choose a basis so the matrices $r_{ij}(s)$ are unitary.

COROLLARY 3. *The matrix entries of the unitary irreducible representations form an orthogonal basis for the set of all functions on G.*

Proof. We already know the matrix entries are all orthogonal as functions. There are $\Sigma d_i^2 = |G|$ of them, and this is the dimension of the vector space of all functions. □

In practice it is useful to have an explicit formula expressing a function in this basis. The following two results will be in constant use.

PROPOSITION.
(a) *Fourier Inversion Theorem. Let f be a function on G, then*

$$f(s) = \frac{1}{|G|} \Sigma d_i \ \mathrm{Tr}(\rho_i(s^{-1}) \hat{f}(\rho_i)).$$

(b) *Plancherel Formula. Let f and h be functions on G, then*

$$\Sigma f(s^{-1}) h(s) = \frac{1}{|G|} \Sigma d_i \ \mathrm{Tr}(\hat{f}(\rho_i) \hat{h}(\rho_i)).$$

Proof. Part (a). Both sides are linear in f so it is sufficient to check the formula for $f(s) = \delta_{st}$. Then $\hat{f}(\rho_i) = \rho_i(t)$, and the right side equals

$$\frac{1}{|G|} \Sigma d_i \chi_i(s^{-1} t).$$

The result follows from Corollary 2.

Part (b). Both sides are linear in f; taking $f(s) = \delta_{st}$, we must show

$$h(t^{-1}) = \frac{1}{|G|} \Sigma d_i \ \mathrm{Tr}(\rho_i(t) \hat{h}(\rho_i)).$$

This was proved in part (a). □

Remark 1. The inversion theorem shows that the transforms of f at the irreducible representations determine f. It reduces to the well known discrete Fourier inversion theorem when $G = Z_n$.

Remark 2. The right hand side of the inversion theorem gives an explicit recipe for expressing a function f as a linear combination of the basis functions of Corollary 3. The right hand side being precisely the required linear combination as can be seen by expanding out the trace.

Remark 3. The Plancherel Formula says, as usual, that the inner product of two functions equals the "inner product" of their transforms. For real functions and unitary representations it can be rewritten as $\Sigma f(s)h(s) = \frac{1}{|G|}\Sigma d_i \operatorname{Tr}(\hat{h}(\rho_i)\hat{f}(\rho_i)^*)$. The theorem is surprisingly useful.

EXERCISE 5. The following problem comes up in investigating the distribution of how close two randomly chosen group elements are. Let P be a probability on G. Define $\overline{P}(s) = P(s^{-1})$. Show that $U = P * \overline{P}$ if and only if P is uniform.

EXERCISE 6. Let H be the eight element group of quarternions $\{\pm 1, \pm i, \pm j, \pm k\}$ with $i^2 = j^2 = k^2 = -1$ and multiplication given by 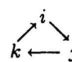

so $ij = k$, $ji = -k$, etc. How many irreducible representations are there? What are their degrees? Give an explicit construction of all of them. Show that if P is a probability on H such that $P * P = U$, then $P = U$. Hint: See Diaconis and Shahshahani (1986b).

D. NUMBER OF IRREDUCIBLE REPRESENTATIONS.

Conjugacy is a useful equivalence relation on groups: s and t are called *conjugate* if $usu^{-1} = t$ for some u. This is an equivalence relation and splits the group into conjugacy classes. In an Abelian group, each class has only one element. In non-Abelian groups, the definition lumps together sizable numbers of elements. For matrix groups, the classification of matrices up to conjugacy is the problem of "canonical forms." For the permutation group, S_n, there is one conjugacy class for each partition of n: thus the identity forms a class (always), the transpositions $\{(ij)\}$ form a class, the 3 cycles $\{(ijk)\}$, products of 2–2 cycles $\{(ij)(k\ell)\}$, and so on. The reason is the following formula for computing the conjugate: if η, written in cycle notation is $(a \dots b)(c \dots d) \dots (e \dots f)$, then $\pi\eta\pi^{-1} = (\pi(a) \dots \pi(b))(\pi(c) \dots \pi(d)) \dots (\pi(e) \dots \pi(f))$. It follows that two permutations with the same cycle lengths are conjugate, so there is one conjugacy class for each partition of n.

A function f on G that is constant on conjugacy classes is called a *class function.*

PROPOSITION 6. *Let f be a class function on G. Let $\rho : G \to GL(V)$ be an irreducible representation of G. Then $\hat{f}(\rho) = \lambda I$ with*

$$\lambda = \frac{1}{d_\rho}\Sigma f(t)\chi_\rho(t) = \frac{|G|}{d_\rho}(f|\chi_\rho^*).$$

Proof. $\rho_s \hat{f}(\rho)\rho_s^{-1} = \Sigma f(t)\rho(s)\rho(t)\rho(s^{-1}) = \Sigma f(t)\rho(\text{sts}^{-1}) = \hat{f}(\rho)$. So, by part 2 of Schur's lemma $\hat{f}(\rho) = \lambda I$. Take traces of both sides and solve for λ. □

Remark. Sometimes in random walk problems, the probability used is constant on conjugacy classes. An example is the walk generated by random transpositions: this puts mass $1/n$ on the class of {id} and $2/n^2$ on {(id)}. Proposition 6 says that the Fourier transform $\hat{f}(\rho)$ is a constant times the identity. So $\hat{P}^{*k}(\rho) = \lambda^k I$ and there is every possibility of a careful analysis of the rate of convergence. See Chapter 3-D.

EXERCISE 7. Show that the convolution of two class functions is again a class function. Show that f is a class function if and only if $f * h = h * f$ for all functions h.

Theorem 6. *The characters of the irreducible representations: $\chi_1, \ldots, \chi_h$ form an orthonormal basis for the class functions.*

Proof. Proposition 1 shows that characters are class functions and Theorem 3 shows that they are orthonormal. It remains to show there are enough. Suppose $(f|\chi_i^*) = 0$, for f a class function. Then Proposition 6 gives $\hat{f}(\rho) = 0$ for every irreducible ρ and the inversion theorem gives $f = 0$. □

Theorem 7. *The number of irreducible representations equals the number of conjugacy classes.*

Proof. Theorem 6 gives the number h of irreducible representations as the dimension of the space of class functions. Clearly, a class function can be defined to have an arbitrary value on each conjugacy class, so the dimension of the class function equals the number of classes. □

Theorem 8. *The following properties are equivalent*
(1) *G is Abelian.*
(2) *All irreducible representations of G have degree 1.*

Proof. We have $\Sigma d_\rho^2 = |G|$. If G is Abelian, then there are $|G|$ conjugacy classes, and so G terms in the sum, each of which must be 1. If all $d_\rho = 1$, then there must be $|G|$ conjugacy classes, so for each $s, t, sts^{-1} = t$, or G is Abelian. □

Example. The irreducible representations of Z_n — the integers mod n.

This is an Abelian group, so all irreducible representations have degree 1. Any ρ is determined by the image of 1: $\rho(k) = \rho(1)^k$, and $\rho(1)^n = 1$, so $\rho(1)$ must be an n^{th} root of unity. There are n such: $e^{2\pi ij/n}$. Each gives an irreducible representation: $\rho_j(k) = e^{2\pi ijk/n}$ (any 1-dimensional representation is irreducible). They are in-equivalent, since the characters are all distinct (not allowed) or $\rho^1(k) = \rho^2(k)$. The Fourier transform is the well known discrete Fourier transform and the inversion theorem translates to the familiar result: If f is a function on Z_n, and $\hat{f}(j) = \Sigma_k f(k)e^{2\pi ijk/n}$, then $f(k) = \frac{1}{n}\Sigma_j \hat{f}(j)e^{-2\pi ijk/n}$.

E. Product of Groups.

If G_1 and G_2 are groups, their *product* is the set of pairs (g_1, g_2) with multiplication defined coordinate-wise. The following considerations show that the representation theory of the product is determined by the representation theory of each factor.

Let $\rho^1 : G_1 \to GL(V_1)$ and $\rho^2 : G_2 \to GL(V_2)$ be representations. Define $\rho^1 \otimes \rho^2 : G_1 \times G_2 \to GL(V_1 \otimes V_2)$ by

$$\rho^1 \otimes \rho^2_{(s,t)}(v_1 \otimes v_2) = \rho^1_s(v_1) \otimes \rho^2_t(v_2).$$

This is a representation with character $\chi_1(s) \cdot \chi_2(t)$.

Theorem 9.
(1) *If ρ^1 and ρ^2 are irreducible, then $\rho^1 \otimes \rho^2$ is irreducible.*
(2) *Each irreducible representation of $G_1 \times G_2$ is equivalent to a representation $\rho^1 \otimes \rho^2$ where ρ^i is an irreducible representation of G_i.*

Proof.
(1) $(\chi_1|\chi_1) = (\chi_2|\chi_2) = 1$, but the norm of the character of $\rho_1 \otimes \rho_2$ is $\frac{1}{|G_1||G_2|}\Sigma\chi_1(s)\chi_2(t)\chi_1(s)^*\chi_2(t)^* = (\chi_1|\chi_1) \cdot (\chi_2|\chi_2) = 1$. So Theorem 5 gives irreducibility.
(2) The characters of the product representation are of the form $\chi_1 \cdot \chi_2$. It is enough to show these form a basis for the class functions on $G_1 \times G_2$. Since they are all characters of irreducible representations, they are orthonormal, so it must be proved that they are it all of the possible characters. If $f(s,t)$ is a class function orthogonal to all $\chi_1(s)\chi_2(t)$, then

$$\Sigma f(s,t)\chi_1(s)^*\chi_2(t)^* = 0.$$

Then for each t, $\Sigma f(s,t)\chi_1(s)^* = 0$, so $f(s,t) = 0$ for each t. □

Exercise 8. Compute all the irreducible representations of Z_2^k, explicitly.

We now leave Serre to get to applications, omitting the very important topic of induced representations. The most relevant material is Section 3.3, Chapter 7, and Sections 8.1, 8.2. A bit of it is developed here in Chapter 3-F.