Representations of finite groups
Lecture notes for Matematika++
Winter 2013/2014

1 Definitions: representations etc.

When dealing with a group, it is often convenient to represent its elements with real or complex invertible matrices in such a way that the group operation corresponds to matrix multiplication. This way we get a problem in linear algebra, which is well understood. Alternatively (but equivalently) we can think of the group elements as invertible linear transformations on some vector space V, i.e., the elements of the *general linear group* $\mathrm{GL}(V)$.

> A *representation* of a group G is a homomorphism from G to $GL(V)$ for some vector space V.

That is, a representation ϕ satisfies $\phi(st) = \phi(s)\phi(t)$ for any two elements s, t of G. It follows that $\phi(s^{-1}) = (\phi(s))^{-1}$ and $\phi(1) = \mathrm{Id}$. There is an obvious correspondence between matrices and linear transformations: if we choose a basis for V, the linear transformations correspond to matrices; on the other hand, invertible $n \times n$ matrices determine linear transformations of $\mathbb{R}^n$.

The *dimension* of the representation is the dimension of the vector space. If the mapping $G \to \mathrm{GL}(V)$ is clear from context, we sometimes talk about the space V itself as the representation of G.

Throughout this text, all groups will be finite and all vector spaces finite dimensional, over the field $\mathbb{C}$, without having to specifically mention it each time.

2 G-linear mappings, equivalent representations

A certain kind of linear mappings between vector spaces will be used often enough to deserve a name. Let $\rho : G \to GL(V)$ and $\tau : G \to GL(W)$ be two representations of a group G and $f : V \to W$ a linear map. We say that f is *G-linear* if $f \circ \rho(s) = \tau(s) \circ f$ holds for every $s \in G$. In other words, the following diagram "commutes" for every $s \in G$:

$$\begin{tikzcd} V \arrow[r, "\rho(s)"] \arrow[d, "f"'] & V \arrow[d, "f"] \\ W \arrow[r, "\tau(s)"] & W \end{tikzcd}$$

We say that two representations ρ and τ of the same group are *equivalent* if there is an isomorphism of the two vector spaces which preserves the action of the group. That is, there exists a G-linear bijection between V and W.

3 Examples of representations, irreducibility

Example 1. Note that representations do not have to be injective. As an example, the *trivial representation* assigns the identity function on V to every group element.

Example 2. Let S_n be the group of permutations of n elements, as usual. In the *alternating representation* of S_n, each permutation π is represented with $f_\pi : V \to V$ defined as $x \mapsto \operatorname{sgn}(\pi) \cdot x$. In both this and the previous example, we can take $V = \mathbb{C}$ as our vector space, so we have two 1-dimensional representations.

Example 3. An obvious n-dimensional representation of S_n is the following. Fix an n-dimensional vector space V and a basis $e_1, \ldots, e_n$ of V. The image of a permutation $\pi \in S_n$ is the mapping $f : V \to V$ that permutes the basis elements according to π (that is, $f(e_i) = e_{\pi(i)}$) and is extended linearly to the rest of V. This is the so called *permutation representation.*

Example 4. Any n-element group G is isomorphic to a subgroup of S_n. Let $g_1, \ldots, g_n$ be the elements of G. Left multiplication by an element h permutes the elements. That is, we can define a permutation π_h by $\pi_h(i) = j$ whenever $h(g_i) = g_j$. Having associated the group elements with permutations, we can use the previous paragraph to find an n-dimensional representation of G. This is called the *regular representation* of G. For example, for $G = S_n$, this yields an $n!$-dimensional representation.

Let us go back to the permutation representation $S_n \to GL(\mathbb{R}^n)$; call it ρ. The subspace $W = \{c \cdot (e_1 + \cdots + e_n); c \in \mathbb{R}\}$ has the property that for any $v \in W$, the image $\rho_\pi(v)$ is also in W, for any $\pi \in S_n$. The subspace $\{v \in \mathbb{R}^n; \sum v_i = 0\}$ (this is the orthogonal complement of W if $e_1, \ldots, e_n$ is the canonical basis) also has this property. We call such subspaces *invariant.*

If V does not have any invariant subspace, we call the representation *irreducible.*

Example 5. *TO DO: Explicit computation of the 2-dimensional representation of S_3.*

4 Making new representations out of old ones

Recall that given two vector spaces, U and W, we can define their *direct sum* $U \oplus W$ as the vector space on the pairs (u, w) with $u \in U$ and $w \in W$, with coordinate-wise addition of vectors and multiplication by scalars. If a vector space V has two subspaces, U and W, such that $U \cap W = \{0\}$ and every $v \in V$ can be written as the sum of a vector in U and a vector in W, then V is isomorphic to the direct sum of U and W. For example the real plane is (isomorphic to) the direct sum of the two lines along the coordinate axes.

Given two representations $\rho : G \to GL(V)$ and $\tau : G \to GL(W)$ of the same group, there are two basic ways to create a new representation—the direct sum and the tensor product. The direct sum $\rho \oplus \tau : G \to V \oplus W$ is defined component-wise: $(\rho \oplus \tau)(g)(v, w) = (\rho(g)(v), \tau(g)(w))$. If the dimensions of the two representations are $\dim(\rho)$ and $\dim(\tau)$, their direct sum has dimension $\dim(\rho) + \dim(\tau)$ and contains invariant subspaces isomorphic to V and W.

Let $\{v_1, \ldots, v_k\}$ and $\{w_1, \ldots, w_\ell\}$ be bases of V and W, respectively. The tensor product of V and W is the vector space over the same field as V and W, with a basis of the $k\ell$ symbols $v_i \otimes w_j$. These should be understood as purely formal symbols. The elements of the vector space are the formal linear combinations $\sum c_{ij}(v_i \otimes w_j)$. If $v = \sum a_i v_i$ and $w = \sum b_j w_j$ (where a_i, b_j are scalars), we define $v \otimes w$ to be the sum $\sum a_i b_j (v_i \otimes w_j)$. If we choose different bases of V and W and follow the above recipe, the resulting vector space is isomorphic to the one above.

It follows from the definition that the operator $\otimes$ is billinear, that is, it satisfies $(au + bv) \otimes w = a(u \otimes w) + b(v \otimes w)$ for all $u, v \in V$, $w \in W$ and scalars a, b (and symmetrically for W). In this sense, it is a generalization of the usual multiplication of, say, real numbers.

As a side remark, let us mention another way to visualize the tensor product. The space $V \otimes W$ can be also regarded as the set of $k \times \ell$ matrices. If $v = \sum \lambda_i v_i$ and $w = \mu_j w_j$, the i, j-th entry of the matrix $v \otimes w$ is equal to $\lambda_i \mu_j$. The mapping $(\rho \otimes \tau)(s)$ sends $v \otimes w$ to $\rho(s) \otimes \tau(s)$.

5 Representations as direct sums of irreducibles

Irreducible representations play a crucial role in the theory of representations. They have two properties that make them extremely convenient to use—first, there is a lot of useful theory developed around them, and second, we can reduce problems involving general representations to problems involving irreducible representations. More specifically, in this section we will prove that every representation is a direct sum of irreducible representations. In this respect, the role of irreducible representations is similar to the role of prime numbers in number theory.

Theorem 1. *Suppose that $\rho : G \to GL(V)$ is a representation and V has an invariant subspace. Then W has a complementary subspace that is also invariant.*

Since V is finite-dimensional, we get the following very important fact by induction.

Every representation can be written as a direct sum of irreducible representations.

Theorem 1 is a powerful theorem with an easy proof. We start with an easy lemma.

Lemma 1. *Let V be a vector space, W a subspace of V, and $f : V \to W$ a linear function such that $f(x) = x$ whenever $x \in W$. Then $V = W \oplus \mathrm{Ker}(f)$.*

Proof. It is easy to check that $W \cap \mathrm{Ker}(f) = \{0\}$. One of the basic theorems of linear algebra (sometimes called the first isomorphism theorem) says that, if X and Y are vector spaces (it is enough that they are groups, really) and $f : X \to Y$ a linear map, then the quotient $V/\mathrm{Ker}(f)$ is isomorphic to the image of f. It follows that the dimension of V equals the sum the the dimensions of the kernel of f and the image of f. The image of f is W, so

$$\dim W + \dim \mathrm{Ker}(f) = \dim V.$$

We have two trivially-intersecting subspaces of V, and the sum of their dimensions is the dimension of V. Every vector of V can therefore be written as a sum of a vector in W and a vector in $\mathrm{Ker}(f)$. We already mentioned that in such case, V is isomorphic to the direct sum of the two subspaces. □

Proof of Theorem 1. For each $s \in G$, let $\rho'(s)$ be the restriction of $\rho(s)$ to W. This subspace is invariant, so ρ' is another representation of G.

In Lemma 1 we described a way to obtain a subspace complementary to W, but this subspace is not necessarily closed under the linear transformations $\rho(s)$. To obtain this property, we add the assumption that f is G-linear with respect to ρ and ρ'. With this assumption, if $x \in \mathrm{Ker}(f)$, then

$$f(\rho(s)(x)) = \rho(s)(f(x)) = \rho(s)(0) = 0.$$

The only thing left to prove the theorem is to choose a suitable G-linear f. Let $f_1 : V \to W$ be a projection on the subspace W. (That is, choose a basis $\{v_1, \dots, v_k\}$ of W and extend it to a basis $\{v_1, \dots, v_k, v_{k+1}, \dots, v_n\}$ of V. Define $f_1(v_i) = v_i$ for $i \leq k$, $f_1(v_i) = 0$ for $i > k$ and extend f_1 linearly to the rest of V.) This is linear but probably not G-linear. The G-linearity is achieved by the following "averaging" trick. Define $f_2 : V \to V$ by $f_2 = \sum_{g \in G} \frac{1}{|G|}(\rho(g) \circ f_1 \circ \rho(g^{-1}))$. Check that this is indeed a G-linear projection from V to W. □

We have already seen one example of a decomposition: the decomposition of S_3 into 1- and 2-dimensional representations. Here is one more.

Example 6. Consider the representation $\rho : S_n \to GL(\mathbb{R}^2)$, given by $\rho(\pi)(x, y) = \mathrm{sgn}(\pi)(x, y)$. The line $\{(x, y); x = y\}$ is an invariant subspace. The subspace complementary under the usual inner product $\{(x, y); x = -y\}$ is also invariant. This choice is not unique. For example, the line $\{(x, y); x = -2y\}$ is another invariant complement.

6 Invariant inner product, Weyl unitarity trick

An inner product on a complex vector space is generally required to be sesquilinear. That is, linear in the first argument (if ϕ is our inner product on V, this means that $\phi(ax+by,z) = a\phi(x,z) + b\phi(y,z)$ holds for any $a,b \in \mathbb{C}$ and any $x,y,z \in V$) and conjugate-linear in the second argument (i.e., $\phi(x, ay+bz) = \overline{a}\phi(x,y) + \overline{b}\phi(x,z)$, where the overline denotes a complex conjugate).

We can always define an inner product on any (finite dimensional) vector space V: pick a basis of V, say $e_1, \dots, e_n$, define $\langle e_i, e_j \rangle_1 = \delta_{ij}$ and extend this sesquilinearly to the whole vector space. In other words, $\langle \sum c_i e_i, \sum d_j e_j \rangle_1 = \sum c_i \overline{d_i}$.

If G is a finite group, we can use this inner product to define another one as kind of an average over the elements of the group: $\langle u, v \rangle = \sum_{s \in G} \langle \rho_s(u), \rho_s(v) \rangle_1$. This new inner product has the advantage that it is invariant, i.e., $\langle \rho_s(u), \rho_s(v) \rangle = \langle u, v \rangle$ for any $s \in G$ and any $u, v \in V$.

If W is an invariant subspace, its orthogonal complement with respect to the inner product $\langle , \rangle$ is also invariant. This gives us another proof of Theorem 1.

Recall that a matrix is *unitary* if its inverse is created by transposing the matrix and replacing each entry with its complex conjugate.

If $e_1, \dots, e_n$ is a basis of V, then the matrix A of the mapping $\rho(s)$ with respect to this basis has columns $\rho_s(e_1), \dots, \rho_s(e_n)$. If this basis is moreover orthonormal with respect to the inner product $\langle , \rangle$, then $\langle \rho_s(e_i), \rho_s(e_j) \rangle = \langle e_i, e_j \rangle = \delta_{ij}$. That is, the columns form an orthonormal basis of V. This is equivalent to A being unitary. This is called the

Weyl unitarity trick: if G is "nice" (in our case finite), we can assume that the matrices in its representation are unitary.

We will use this fact later.

7 Schur's lemma

Any introductory course on representation theory would not be complete without mentioning Schur's lemma, a tool of paramout importance.

Theorem 2 (Schur's lemma)**.** *Let $\rho : G \to GL(V)$ and $\tau : G \to GL(W)$ be two irreducible representations and let $f : V \to W$ be a G-linear mapping with respect to the two representations.*

1. *Either f is a bijection (i.e., ρ and τ are equivalent), or $f = 0$.*
2. *If $V = W$ and $\rho = \tau$, then f is a constant times the identity.*

Proof. Suppose that $f \neq 0$. The kernel of f is an invariant subspace of V and the image is an invariant subspace of W. Since ρ is irreducible, $\mathrm{Ker}(f)$ is either the trivial subspace of V itself. The latter is in contradiction to $f \neq 0$. Similarly, the image of f has to be W itself. It follows that f is a bijection.

To prove the second statement, note that since we are working over the field of complex numbers, every linear function has an eigenvalue. In fact, this was the main reason for choosing complex numbers over the more familiar reals. Let λ be an eigenvalue of f and define $f_1 = f - \lambda \mathrm{Id}_V$. This function is G-linear, i.e., $f_1(\rho(s)(x)) = \rho(s)(f_1(x))$ for every $x \in V$ and $s \in G$. The eigenvector corresponding to λ belongs to the kernel of f_1, so the kernel is non-trivial and by part (i), f_1 is the zero mapping. □

Before dissecting the Schur's lemma further, let us list an interesting corollary.

Corollary 1. *Every irreducible complex representation of a finite abelian group G is one-dimensional.*

Proof. For $s \in G$, $\rho(s)$ is a G-linear mapping (this uses the assumption that G is abelian), and by the second part of Schur's lemma we have $\rho(s)(v) = c \cdot v$. If follows that $\{cv; c \in \mathbb{C}\}$ is an invariant subspace. We complete the proof by remembering that ρ is irreducible. □

We will need some more technical facts that follow from Schur's lemma.

Recall that the *trace* of a matrix A is equal to the sum of the elements on the main diagonal of A. Suppose that V is a d-dimensional vector space, and fix a basis of V. Let $f : V \to V$ be an endomorphism and A the $d \times d$ matrix that corresponds to f with respect to the chosen basis. If we choose a different basis, f corresponds to a different matrix, say B, but there exists an invertible matrix C such that $A = CBC^{-1}$. It is well known that for square matrices X and Y, we have $\mathrm{Tr}(XY) = \mathrm{Tr}(YX)$. It follow that, in our case, $\mathrm{Tr}(B) = \mathrm{Tr}(C^{-1}CB) = \mathrm{Tr}(CBC^{-1}) = \mathrm{Tr}(A)$. We can therefore define the *trace of an endomorphism* as the trace of any corresponding matrix.

The order of G will be denoted by $|G|$.

Corollary 2. *Let $\rho : G \to GL(V)$ and $\tau : G \to GL(W)$ be two irreducible representations and $h : V \to W$ any linear mapping. Define*

$$f = \frac{1}{|G|} \sum_{t \in G} (\tau(t))^{-1} \circ h \circ (\rho(t)).$$

1. *If ρ and τ are not equivalent, then $f = 0$.*
2. *If $V = W$ and $\rho = \tau$ with dimension d, then f is a constant times the identity, with the constant equal to* $\mathrm{Tr}(h)/d$.

Proof. The function f was already defined in the proof of Theorem 1 and we have seen that it is G-linear. The first part follows from the first part of Theorem 7.

As for the second part, we already know that $f = c \cdot \mathrm{Id}$. Taking the trace of both sides, we have $c = \mathrm{Tr}(f)/d$ and $\mathrm{Tr}(h) = (1/|G|) \sum \mathrm{Tr}(\rho(t)^{-1} \circ h \circ \rho(t)) = (1/|G|) \sum \mathrm{Tr}(h) = \mathrm{Tr}(h)$. □

Suppose that the linear mapping $\rho(s)$ is given by a matrix $R(s)$, with entries $r_{ij}(s)$. Similarly, let the mappings $\tau(s)$ be given by matrices $T(s)$ with entries $t_{ij}(s)$, h by a matrix H, and f by F. Writing out the formula in Corollary (2) as matrix multiplication, the (i,j)-th entry of the matrix F is given by

$$f_{ij} = \frac{1}{|G|} \sum_{\substack{s \in G \\ 1 \le p,q \le n}} t_{iq}(s^{-1}) \cdot h_{qp} \cdot r_{pj}(s).$$

Unless ρ and τ are equivalent, $f \equiv 0$ for all choices of linear h. That is, $f_{ij} = 0$ for all i, j. Fixing p, q and choosing a matrix H such that $h_{qp} = 1$, and all other entries of H are zero, we get the following:

$$0 = f_{ij} = \frac{1}{|G|} \sum_{s \in G} t_{iq}(s^{-1}) r_{pj}(s). \tag{1}$$

Similarly, if $\rho = \tau$ and $V = W$, then

$$\frac{1}{|G|} \sum_{s \in G} t_{iq}(s^{-1}) r_{pj}(s) = \begin{cases} \frac{1}{d} & \text{if } i = j \text{ and } p = q \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

A convenient way to view the matrices $R(s)$ is to consider a single matrix R, whose entries are the functions $r_{ij} : G \to \mathbb{C}$. The above formulas in particular tell us that if the matrices $\rho(s)$ for an irreducible representation ρ are unitary (and we can assume this due to the Weyl unitarity trick from section 6), then their entries are orthogonal as functions $G \to \mathbb{C}$.

These formulas will be more useful later, in Section 9.

8 Characters, basic properties

We define the *character* χ of a representation ρ as the function $G \to \mathbb{C}$ mapping $s \in G$ to $\mathrm{Tr}(\rho(s))$.

We have already argued in Section 7 that the trace of an endomorphism is well-defined.

Characters encode essential information about representations, and as such are one of the central tools in representation theory. We will begin with a few basic facts. Recall that two elements s, t of a group G are *conjugate* if there exists an element $r \in G$ such that $s = rtr^{-1}$.

Lemma 2. *Let ρ be a d-dimensional representation.*

- $\chi(\mathrm{Id}) = d$
- $\chi(s^{-1}) = \overline{\chi(s)}$
- *If s, t are conjugate in G, then $\chi(s) = \chi(t)$.*

The first part is obvious: $\rho(\mathrm{id})$ is a $d \times d$ identity matrix. The second part follows from the Weyl unitarity trick described in Section 6. The third one follows from basic properties of trace, similarly to the reasoning used to establish the fact that the trace of an endomorphism is well-defined.

Lemma 3. *Let ρ and τ be two representations of G. Then*

- $\chi(\rho \oplus \tau) = \chi(\rho) + \chi(\tau)$ *and*
- $\chi(\rho \otimes \tau) = \chi(\rho) \cdot \chi(\tau)$.

9 More on characters

Let us define an inner product of two complex-valued functions ϕ, ψ on G in the usual way:

$$\langle \phi, \psi \rangle = \frac{1}{|G|} \sum_{s \in G} \phi(s) \overline{\psi(s)}.$$

We have already seen in Section 7 that the matrix entries of unitary irreducible representations are orthogonal as functions $G \to \mathbb{C}$.

Theorem 3. *The characters of every two nonequivalent irreducile representations are orthonormal.*

Proof. Let χ and χ' be two characters of irreducible representations. Expand $\langle \chi | \chi' \rangle$ and use formulas (1) and (2) from Section 7.

□

Theorem 4. *Let $\rho : G \to GL(V)$ be a representation with character ϕ. Let $V = W_1 \oplus \cdots \oplus W_k$ be a decomposition of V into irreducible representations. If W is an irreducible representation with character χ_i, then the number of W_i equivalent to W is given by $\langle \phi, \chi \rangle$.*

Proof. If χ_i is the character of W_i, then $\phi = \chi_1 + \cdots + \chi_k$ and

$$\langle \phi, \chi \rangle = \sum_{i=1}^{k} \langle \chi_i, \chi \rangle.$$

It follows from the previous theorem that each of the terms is either 0 or 1, depending on whether W_i is equivalent to W. □

The number of W_i equivalent to W does not depend on the particular decomposition of V, because the inner product of the characters does not depend on it. The decomposition is therefore unique in a certain sense.

Suppose that two representations have the same character, and decompose them into irreducible representations. The two decompositions contain the same irreducible representations, and each of these appears with the same multiplicity in the two decompositions. We get the following important fact:

Two representations are equivalent if and only if they have the same character.

We can now derive an easy and powerful irreducibility criterion. If V is decomposed as $V = m_1 W_1 \oplus \cdots \oplus m_k W_k$, where W_i are nonequivalent irreducible representations, then $\langle \phi, \phi \rangle = \sum m_i^2$. This is equal to 1 if and only if ρ is irreducible.

10 Regular representation

Recall that in a regular representation, we choose a vector space V of dimension $|G|$, we index the basis vectors by the elements of G, and define $\rho_s(e_t) = e_{st}$. Let $W_1, \ldots, W_k$ a maximal set of nonequivalent irreducible representations of G. Denote χ_i the character of W_i, and d_i its degree.

Theorem 5. *The character r_G of the regular representation is given by $r_G(1) = |G|$ and $r_G(s) = 0$ whenever $s \neq 1$.*

Proof. The first statement is easy. For the second one, $\rho(s)$ is a permutation matrix with zeros on diagonal. If not, there is a $t \in G$ such that $\rho(s)(e_t) = e_t$. By definition, this means that $st = t$ and $s = 1$, a contradiction. □

Corollary 3. *Every irreducible representation is contained in the regular representation with multiplicity equal to its degree.*

Proof. $\langle r_g, \chi_i \rangle = d_i$. □

Corollary 4. $\sum d_i^2 = |G|$, *and whenever* $s \neq 1$, $\sum d_i \chi_i(s) = 0$.

Proof. According to the previous corollary, $r_G(s) = \sum d_i \chi_i(s)$. Substituting $s = 1$ and $s \neq 1$, we get the two statements. □

Let us recall that, using the matrices $R(s)$ corresponding to the mappings $\rho(s)$, we can define a single matrix R whose entries are functions $r_{ij} : G \to \mathbb{C}$.

Corollary 5. *The functions r_{ij} for nonequivalent unitary irreducible representations form a basis for the vector space of all functions from G to $\mathbb{C}$.*

Proof. We already know that the functions are orthogonal, and hence also linearly independent. There are $\sum d_i^2 = |G|$ of them. This is equal to the dimension of the vector space in question. □

11 Communication complexity. Does the rank of the input matrix determine it?

We will use representation theory to derive a result in communication complexity. This result was obtained by Raz and Spieker and published in their paper, *On the "log rank"–Conjecture in Communication Complexity"*.

The basic scenario in communication complexity is the following: a real-valued function f defined on $\{1,\ldots,n\}^2$ is known to both Alice and Bob. Alice is given the number x while Bob is given y, both numbers between 1 and n. Their task is to determine together the value $f(x,y)$ while having communicated as little as possible. After the players have seen the function (but before they are given the input numbers), they decide on a communication protocol. This protocol specifies in each step which one of the two players sends information to the other one, and the information sent, as a function of the information exchanged so far. The goal is that at the end of their communication, both players know the value $f(x,y)$. The (deterministic) communication complexity is the minimum number of bits that they need to exchange in the worst-case scenario.

Example 7. Consider the function f_1 that specifies the parity of $x+y$. For both players to know the answer, information has to flow in both directions, so the complexity is at least equal to 2. On the other hand, it suffices to exchange two bits - each player sends the parity of his input.

As a second example, consider the function f_2 that equals 1 if $x=y$ and 0 otherwise. It is not hard to see that n bits are necessary.

These two examples support the intuition that communication complexity should increase with increasing rank of the $n\times n$ matrix that specifies the values of f. A well-known result by Mehlhorn and Schmidt is that the complexity is at least the binary logarithm of this rank (over any field). It has been a long–standing open problem whether this bound is tight for every function when the rank is computed over $\mathbb{R}$ (this value is greater or equal to the value of rank computed over any finite field). This would reduce the question of communication complexity to the question of the rank of the corresponding matrix.

12 Raz and Spieker: "Unfortunately, no."

Raz and Spieker showed that the answer is negative by producing and example of function such that the rank of the matrix is $2^{O(n)}$ and the communication complexity is $\Theta(n\log\log n)$.

The function is described in the following way. Each of the two players is given a perfect matching in the complete bipartite graph $K_{n,n}$. The function is 1 if the union of the two matchings is a Hamiltonian cycle, and 0 otherwise. The input matrix D is therefore an $n!$-by-$n!$ matrix with entries in $\{0,1\}$, whose rows and columns are indexed by permutations of n. To compute the rank of this matrix, we will use representation theory. The proof of the second part of the result (the determination of the communication complexity) is done by an information-theoretic argument; this is outside the scope of these notes and we will refer the interested reader to the original paper.

We will prove the following.

Theorem 6. *The rank of D is $\binom{2n-2}{n-1}$.*

Let us consider the regular representation of S_n. As usual, we choose some basis of $\mathbb{R}^{n!}$ and each permutation π corresponds to the linear mapping permuting the basis vectors. This is given by a permutation matrix P_π.

The conjugacy classes of the group S_n are fully determined by the cycle structure of the permutations. That is, two permutations are conjugates if and only if they have the same cycle structure. The conjugacy class of n-cycles, denoted X_n, will be of a particular interest. It is easy to verify the following fact:

$$D = \sum_{\pi \in X_n} P_\pi.$$

This follows from the observation that $D_{\pi\sigma} = 1$ if and only if the permutation $\pi \circ \sigma^{-1}$ is an n-cycle, together with the definition of the matrices P_π—that is, P_π has a one in the row corresponding to σ and the column corresponding to τ if and only if $\pi \circ \tau = \sigma$. Remember that if we decompose the regular representation into irreducible representations, each of them appears with multiplicity equal to its dimension. If V is the regular representation and W_i the distinct irreducible representations with dimensions d_i respectively, we can write

$$V = d_1 W_1 \oplus d_2 W_2 \oplus \cdots \oplus d_k W_k.$$

Consider a new basis of V such that the first d_1 vectors form a basis of the first copy of W_1, the next d_1 vectors form a basis of the second copy of W_1, etc. After we exhaust copies of W_1, the next d_2 vectors form a basis of the first copy of W_2, and continue in a similar fashion. Recall that if A is a matrix of a linear transformation with respect to a certain basis, the matrix of the same transformation with respect to a different basis is of the form UAU^{-1}, for a suitable matrix U (that only depends on the two bases, not on the linear transformation in question). Let U be the appropriate matrix in our scenario—that is, a matrix such that $UP_\pi U^{-1}$ is the matrix corresponding to P_π, with respect to our new basis. Since W_i are invariant subspaces, the matrices $UP_\pi U^{-1}$ are block matrices, with d_1 blocks of size $d_1 \times d_1$ followed by d_2 blocks of size $d_2 \times d_2$ etc. We can also assume that we have chosen the basis so conveniently that the first d_1 blocks are identical and so forth.

Let us look more closely at the matrix UDU^{-1}. Since it is the sum of the block matrices $UP_\pi U^{-1}$, it is also a block matrix, with blocks of the same sizes as $UP_\pi U^{-1}$. It is not hard to verify that D commutes with all matrices P_π. The same is of course true for UDU^{-1} and the matrices $UP_\pi U^{-1}$. It follows that the first block of UDU^{-1} commutes with the first block of $UP_\pi U^{-1}$ (for every π), and so on for all the other blocks. Schur's lemma tells us that the first block of UDU^{-1} is a constant times the identity matrix. There are d_1 copies of this block, followed by d_2 copies of a block that is another constant times the $d_2 \times d_2$ identity matrix, and so forth. The matrix UDU^{-1} is therefore a diagonal matrix, and computing its rank is equivalent to determining the number of nonzero diagonal entries.

If K is a conjugacy class on S_n, let us define $\chi^K(i)$ to be the trace of the block corresponding to W_i (that is, i-th distinct block) of the matrix $UP_\pi U^{-1}$, where π belongs to K. Since trace is constant on a conjugacy class, this is well-defined. Let us observe that

$$\textit{The block corresponding to } W_i \textit{ in } UDU^{-1} \textit{ is a zero matrix} \iff \chi^{X_n}(i) = 0$$

where, again, X_n is the conjugacy class consisting of n-cycles.

The number of nonzero rows of UDU^{-1} is given by

$$\sum d_i \cdot d_i, \tag{3}$$

where the sum is taken over all i such that the block of $UP_\pi U^{-1}$ corresponding to W_i has nonzero trace, for $\pi \in X_n$.

Now, let us step aside for a moment. A partition of a natural number is a way to write the number as a sum of natural numbers, with their order being irrelevant. We can sort these numbers in decreasing order and draw a diagram consisting of rows of squares, corresponding to the partition in a way that is obvious from the following picture. This is a diagram corresponding to the partition $10 = 5 + 4 + 1$:

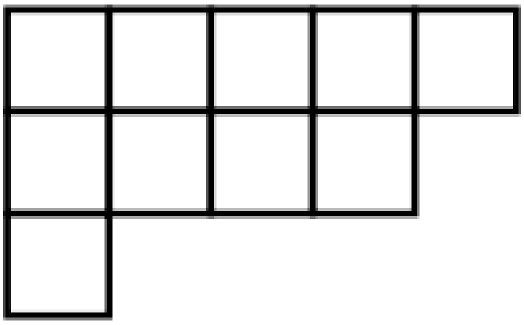

Such pictures are called *Ferrer's diagrams.* We get a *standard Young tableau* by writing numbers $1, \dots, n$ in the squares in such a way that each row (from left to right) and each column (from top to bottom) contain an increasing sequence.

These seemingly innocent cute drawings play an important role in the theory of represesentations of S_n. It can be proven (although we won't do it here) that the irreducible representations of S_n are in bijection with its conjugacy classes, i.e., with the partitions of n. Moreover, the dimension of an irreducible representation is given by the number of standard Young tableaux of the shape given by the corresponding partition—in other words, the number of ways to "legally" fill the corresponding Ferrer's diagram with numbers $1, \dots, n$.

Another useful fact (also taken on faith in this text) is that the irreducible representations such that the corresponding blocks of $UP_\pi U^{-1}$ for $\pi \in X_n$ are nonzero, correspond to Ferrer's diagrams of an upside-down "L" shape, that is, with at most one row and at most one column with more than one square. It is easy to count the number of ways to legally fill such a shape with numbers. We put the number 1 in the corner and if the diagram has j rows, we have $\binom{n-1}{j-1}$ choices of which numbers to put in the first column. As we have already mentioned, this is equal to the dimension of the irreducible representation in question. The sum in (3) is therefore equal to

$$\sum_{j=1}^{n} \binom{n-1}{j-1}^2 = \binom{2n-2}{n-1}.$$