Irreducible Representations of S_n

Report for PMATH745

Aleksa Vujičić

December 14, 2020

1 Introduction

The aim of this report is to describe all of the irreducible representations of the symmetric group S_n. Each conjugacy class of S_n has an associated partition λ, and we will use this partition to construct an irreducible representation V_λ. Because of this correspondence between conjugacy classes and irreducible representations, we know that these V_λ form the complete set of irreducible representations of S_n. Furthermore, we shall state Frobenius's formula - a way to compute the character of V_λ. We then use this to find a nice algorithm for computing the dimension of V_λ.

This report assumes the reader is familiar with representation theory (such as that which is presented in PMATH445/745). We will follow closely the material laid out in 'Representation Theory: A First Course' by Fulton and Harris [1, §4]. Unless otherwise noted, the theorems and proofs in this report are taken from this book.

2 Partitions of n and Young Diagrams

Recall that for a finite group G, the number of irreducible representations of G is equal to the number of conjugacy classes of G. In the case where $G = S_n$, we know exactly what the conjugacy classes are. They are determined precisely by *cycle type* of an element of S_n.

Definition 2.1. *Let $\sigma \in S_n$. We can always write σ as a product of k disjoint cycles. Then the* ***cycle type*** *of σ is a k-tuple which contains the lengths of these cycles, ordered from highest to lowest (including any trailing ones).*

For example, if $\sigma = (13)(246) \in S_6$, then its cycle type is $(3, 2, 1)$. Note that while not explicitly written, we include the trivial cycle (5). The reason for this is so that the sum of the elements of a cycle type is always n, which will be useful later when we study partitions.

For now we have the following theorem that allows us to easily find the conjugacy classes of S_n.

Proposition 2.2. *Let $\sigma, \tau \in S_n$. Then $\sigma \in C(\tau)$ if and only if σ and τ have the same cycle type.*

Proof. Let us consider the case where σ is a single cycle, and we will write $\sigma = (a_1 a_2 \dots a_k)$ where $a_1, \dots a_n \in \{1, \dots, n\} =: N$. Let $g \in S_n$ and consider how $g\sigma g^{-1}$ acts on the set N. Take any $b \in N$ and suppose that $b = g(a_i)$ for some i. It then follows that

$$(g\sigma g^{-1})(b) = (g\sigma)(a_i) = g(a_{i+1})$$

For convenience, if $i = k$, we define $i + 1 = 1$. Now consider if b is not of the form $g(a_i)$ for any i. Then $\sigma(g^{-1}(b)) = g^{-1}(b)$. And so in this case we have

$$(g\sigma g^{-1})(b) = g(\sigma g^{-1})(b) = gg^{-1}(b) = b$$

Thus it follows that $g\sigma g^{-1}$ is precisely the cycle $(g(a_1)g(a_2)\dots g(a_k))$. This argument generalises if σ is a product of disjoint cycles.

So now, if σ and τ are in the same conjugacy class, then $\tau = g\sigma g^{-1}$ for some g and thus they have the same cycle type. Conversely, if σ and τ have the same cycle type, we can let g be the permutation that swaps the entries of σ for the entries of τ. It would then follow that $\tau = g\sigma g^{-1}$, and thus they would belong to the same conjugacy class. $\square$

This gives us a nice characterisation of the conjugacy classes of S_n. This will be useful, as we will use these conjugacy classes to construct the irreducible representations directly. This is somewhat unusual, as although we know that the number of conjugacy classes is the same as the number of irreducible representations, we rarely get a direct correspondence between the two.

To this end, let us first look at the defining feature of a conjugacy class of S_n - it's cycle type. As hinted at previously, this cycle type is really just a *partition* of n elements.

Definition 2.3. *A* ***partition*** λ *of* $n \in \mathbb{N}$ *is a sequence of integers* $\lambda = (\lambda_1, \dots, \lambda_p)$ *such that*

$$\sum_{i=1}^{p} \lambda_i = n \quad \text{and} \quad p \geq \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p \geq 1$$

We also have a way of ordering partitions lexicographically:

Definition 2.4. *Let λ and μ be partitions of n. We write $\lambda > \mu$ if there is an i so that $\lambda_i > \mu_i$ and $\lambda_j = \mu_j$ for all $j < i$.*

As we shall later see, every partition λ has a distinct irreducible representation V_λ. And since the number of partitions is equal to the number of irreducible representations, this means that these V_λ form the complete set of irreducible representations of S_n.

First however, we need a way of visualising any particular partition λ. One common method is by way of *Young diagrams*. In these diagrams, we simply draw boxes in a grid, where the number of boxes in row i is just λ_i. So for instance, if $n = 10$ and $\lambda = (5, 2, 2, 1)$, then the corresponding Young diagram would be

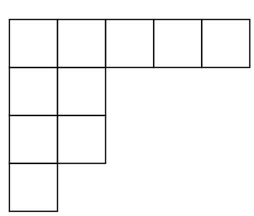

Notice that this diagram reveals a natural way of finding a *conjugate* partition. Namely we can simply interchange the rows and columns. So for this particular example, the conjugate Young diagram would be

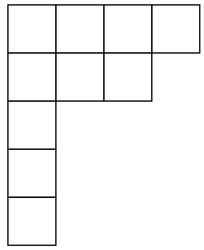

and so the conjugate partition would be $(4, 3, 1, 1, 1)$. Algebraically, if $\lambda = (\lambda_1, \ldots, \lambda_p)$, then the conjugate partition $\lambda' = (\lambda'_1, \ldots, \lambda'_q)$, where λ'_i is the number of λ_j which are greater than i (i.e. $\lambda'_i = |\{j : \lambda_j \geq i\}|$).

Now we want to consider each box as an element of S_n. There is an obvious canonical way of doing so, where we simply label each box from 1 to n starting at the top, going from left to right. So in the previous example this would look like

$$\begin{array}{|c|c|c|c|c|}\hline 1 & 2 & 3 & 4 & 5 \\ \hline 6 & 7 \\ \cline{1-2} 8 & 9 \\ \cline{1-2} 10 \\ \cline{1-1}\end{array}$$

Such a labelling is called a *tableau*. We can of course create different tableaux by applying a permutation. For instance if T is the tableau given above, and $\sigma = (237)$, then the tableau $T' = \sigma T$ would be

$$\begin{array}{|c|c|c|c|c|}\hline 1 & 3 & 7 & 4 & 5 \\ \hline 6 & 2 \\ \cline{1-2} 8 & 9 \\ \cline{1-2} 10 \\ \cline{1-1}\end{array}$$

If no tableau is given, we shall assume that we are working with the canonical tableau. Note that there is nothing particularly special about the canonical tableau - we could choose it to be any tableau we wish. It is for convenience that we fix a canonical tableau. However, we will occasionally need to explicitly refer to others.

3 The Irreducible Representation from λ

Recall that for any finite group G, the group algebra $\mathbb{C}[G]$ contains all irreducible representations of G. As alluded to previously, each partition λ has its own irreducible representation V_λ. This representation will be constructed as a subspace of $\mathbb{C}[S_n]$.

First however, we will define two subgroups P_λ and Q_λ of S_n based on the Young diagrams discussed in the previous section.

Definition 3.1. *Let λ be a partition of n. Let T be a tableau for the Young diagram of λ. We define the subgroups $P_{\lambda,T}$ and $Q_{\lambda,T}$ as follows*

$$\begin{aligned} P_{\lambda,T} &:= \{\sigma \in S_n : \sigma \text{ preserves all rows of the Young tableau } T \text{ of } \lambda\} \\ Q_{\lambda,T} &:= \{\sigma \in S_n : \sigma \text{ preserves all columns of the Young tableau } T \text{ of } \lambda\} \end{aligned}$$

Keeping to the previous example of $\lambda = (5, 2, 2, 1)$ with T being the canonical tableau, we would have that any element of $P_{\lambda,T}$ must fix the sets

$$\{1, 2, 3, 4, 5\}, \{6, 7\}, \{8, 9\}, \{10\}$$

whereas any element of $Q_{\lambda,T}$ would fix the sets

$$\{1, 6, 8, 10\}, \{2, 7, 9\}, \{3\}, \{4\}, \{5\}$$

Note that these subgroups are always disjoint, as any element which fixes both the rows and columns must necessarily be the identity.

Now we turn our to attention to the group algebra $\mathbb{C}[S_n]$. We can use the subgroups $P_{\lambda,T}$ and $Q_{\lambda,T}$ to define the following elements of the group algebra

$$a_{\lambda,T} := \sum_{p \in P_{\lambda,T}} p \quad \text{and} \quad b_{\lambda,T} := \sum_{q \in Q_{\lambda,T}} \operatorname{sgn}(q)q$$

and from these we define the *Young symmetriser* to be

$$c_{\lambda,T} := a_{\lambda,T} \cdot b_{\lambda,T}$$

For the sake of simplicity, we will sometimes drop some or all of the subscripts So when we refer to the canonical tableau, we will often write c_λ or even just c if the partition λ is unambiguous. This also applies to the elements a, and b, and the subgroups P and Q as well.

The elements a, b, and c, have an interesting absorption property, which we state in the following lemma.

Lemma 3.2. *Let λ be a partition of S_n, and define a, b, c as above. Then we have the following:*

(i) For any $p \in P$, we have $p \cdot a = a \cdot p = a$.

(ii) For any $q \in Q$, we have $\operatorname{sgn}(q)q \cdot b = b \cdot \operatorname{sgn}(q)q = b$.

(iii) For any $p \in P$, and $q \in Q$, we have $p \cdot c \cdot \operatorname{sgn}(q)q = c$. Moreover, c is the only such element of $\mathbb{C}[S_n]$ with this property (up to scalar multiplication).

Proof. The proofs of (i) and (ii) follow immediately from the fact that P and Q are subgroups. For the third assertion, take some element $s \in \mathbb{C}[S_n]$ which satisfies $\operatorname{sgn}(q)psq = s$ for every $p \in P, q \in Q$. We can write s as

$$s = \sum_{g \in G} \alpha_g g$$

In order to verify that $s = c$, we would have to show that $\alpha_{pq} = \mathrm{sgn}(q)$ and $\alpha_g = 0$ for $g \notin PQ$.

Now the condition on s gives us that $\alpha_{pgq} = \mathrm{sgn}(q)\alpha_g$ for all g, p, q. In particular, when g is the identity e, we have $\alpha_{pq} = \mathrm{sgn}(q)\alpha_e$. So all that remains is to show that $\alpha_g = 0$ whenever $g \notin PQ$.

So take any $g \notin PQ$. Suppose we could find a transposition t so that $t \in P$ and $g^{-1}tg \in Q$. Then if we set $p = t$ and $q = g^{-1}tg$, we would then have that $g = pgq$, and so

$$\alpha_g = \alpha_{pgq} = \mathrm{sgn}(q)\alpha_g = -\alpha_g$$

and thus $\alpha_g = 0$, and so we would be done. We will show that if no such transposition exists, then it would necessarily imply that $g \in PQ$.

Let T be the canonical tabulation of λ, and let $T' = gT$. Let $Q' = gQg^{-1}$ which is precisely those elements which fix the columns of T' (so $Q' = Q_{\lambda,T'}$). Now the existence of t would imply that t transposes two integers which are in the same row in T, and in the same column in T'. Conversely, if we assume no such t exists, then that means that any two integers belonging to the same row of T must belong to different columns of T'.

Now, for each integer i in the first row of T, let m_i be the column at which that integer i appears in T'. As we have assumed that no transposition t exists, these m_i are all distinct. So we can find a $p_1 \in P$ that moves the integer i to position m_i, and choose $q'_1 \in Q'$ so that each i is moved to the first row. Wit this choice we would we have that p_1T and q'_1T' agree on the first row.

We can then continue this process so that we find $p \in P$ and $q' \in Q'$ so that $pT = q'T'$. But this means that $pT = (q' \cdot g)T$ and so $p = q' \cdot g$. Hence $g = p \cdot q$ where $q = g^{-1} \cdot (q')^{-1} \cdot g$, so we are done. □

The third part of this lemma gives us a sort of idempotent property of c. This property is as follows.

Lemma 3.3. *Let λ be a partition. Then for any $x \in \mathbb{C}[S_n]$, we have that cxc is a scalar multiple of c. In particular, there is some $r = r_\lambda \in \mathbb{C}$ so that $c^2 = rc$.*

Proof. For any $p \in P$, $q \in Q$ we have that $p \cdot (c \cdot x \cdot c) \cdot q = (p \cdot c) \cdot x \cdot (c \cdot q) = c \cdot x \cdot c$. However by the previous lemma, this means that $c \cdot x \cdot c$ is some scalar multiple of c. □

We also have the following independence property.

Lemma 3.4. *Let λ and μ be partitions so that $\lambda > \mu$. Then for every $x \in \mathbb{C}[S_n]$ we have that $a_\lambda \cdot x \cdot b_\mu = 0$, and in particular, $c_\lambda \cdot c_\mu = 0$.*

Proof. Assume without loss of generality that $x = g$ for some $g \in S_n$. Now let $T' = gT$, where T is the canonical tableau. It is clear then that $g \cdot b_\mu \cdot g^{-1} = b_{\mu,T'}$. So it is sufficient to show that $a_{\lambda,T} \cdot b_{\mu,T'} = 0$.

Now since we have $\lambda > \mu$, we claim that there must be two integers which are in the same row of T, and the same column of T' (note that this is true for any T and T' we choose). To see this, consider the first row of these tableaux. If $\lambda_1 > \mu_1$, then notice that T has strictly more columns than T'. By the pigeonhole principle, two of the integer in the first row of T

must share a column in T'. In the other case we have $\lambda_1 = \mu_1$, but we can continue this argument down the rows until we eventually have $\lambda_i > \mu_i$.

With this established, let t be a transposition of the integers mentioned. Then by Lemma 3.2, we have that $a_{\lambda,T} \cdot t = a_{\lambda,T}$ and $t \cdot b_{\mu,T'} = -b_{\mu,T'}$. Since $t^2 = 1$, we have that

$$a_{\lambda,T} \cdot b_{\mu,T'} = a_{\lambda,T} \cdot t \cdot t \cdot b_{\mu,T'} = -a_{\lambda,T} \cdot b_{\mu,T'}$$

and thus it follows that $a_{\lambda,T} \cdot b_{\mu,T'} = 0$ as required. □

With this we can finally construct our irreducible representation.

Definition 3.5. *Given a partition λ, its corresponding irreducible representation V_λ is given by*

$$V_\lambda := \mathbb{C}[S_n]c_\lambda$$

We can show that these are indeed irreducible. Furthermore, we can also show that they are distinct for each partition.

Theorem 3.6. *Let λ be a partition of n. Then we have that V_λ is an irreducible subrepresentation of $\mathbb{C}[S_n]$. Moreover, if $\mu \neq \lambda$ is also a partition, then V_λ and V_μ are not isomorphic.*

Proof. To show that V_λ is irreducible, let us suppose it has a subrepresentation $W \subseteq V_\lambda$. Recalling that $\mathbb{C}[S_n]$ is left Artinian and semiprimitive, we know that we can write $W = \mathbb{C}[S_n]e$ for some idempotent element $e \in \mathbb{C}[S_n]$. By Lemma 3.3, we have that $c_\lambda V_\lambda = \mathbb{C}c_\lambda$, so that $c_\lambda W \subseteq \mathbb{C}c_\lambda$. So we have that $c_\lambda W$ is either $\mathbb{C}c_\lambda$ or 0. In the former case where $c_\lambda W = \mathbb{C}c_\lambda$, we have that

$$V_\lambda = \mathbb{C}[S_n]c_\lambda = \mathbb{C}[S_n]c_\lambda W \subseteq W$$

and so $W = V_\lambda$. On the other hand if $c_\lambda W = 0$, then it follows that $WW \subseteq \mathbb{C}[S_n]c_\lambda W = 0$. But since $W = \mathbb{C}[S_n]e$, it follows then that $e = e^2 = 0$, and hence $W = 0$.

For the second half of the proof, let us assume without loss of generality that $\lambda > \mu$. Now we have that $c_\lambda V_\lambda = \mathbb{C}c_\lambda$ by Lemma 3.3. But by Lemma 3.4, we also have $c\lambda V_\mu = c_\lambda \mathbb{C}[S_n]c_\mu = 0$. So V_λ and V_μ are not isomorphic as $\mathbb{C}[S_n]$ modules (and are thus not isomorphic as representations). □

At this point we have constructed a unique irreducible representation of S_n for each partition of n. As previously mentioned, there is a partition for each conjugacy class of S_n, and so we know this collection of irreducible representations is complete.

We also have the following fact which allow us to relate the scalar r_λ to the dimension of V_λ.

Proposition 3.7. *Let λ be a partition, with $c^2 = rc$. Then we have that $r = n!/\dim V_\lambda$.*

Proof. Let $f : \mathbb{C}[S_n] \to \mathbb{C}[S_n]$ be the homomorphism defined by $f(x) = xc$. Note that if $v \in V_\lambda$, then as $V_\lambda = \mathbb{C}[S_n]c$, we have that $f(v) = rv$. Furthermore, since $Im(f) \subseteq V_\lambda$, it then follows that $\operatorname{Tr} f = r \dim V_\lambda$.

However, we can easily compute the coefficient of g in $g \cdot c$ which is always 1. So it follows that we also have that $\operatorname{Tr} f = |S_n| = n!$. Hence $r = n!/\dim V_\lambda$. □

4 Frobenius's Formula

What can we say about these representations V_λ? It turns out that we have a nice way to compute the character χ_λ known as Frobenius's formula. In this section we shall define and state this formula, though for the sake of brevity we shall not prove why it holds. We also use this formula to find the dimension of V_λ, via the hook length formula.

Frobenius's formula is defined in terms of polynomials of independent variables. If we take a partition λ, we can let $x_1, \ldots, x_k$ be independent variables such that k is the number of rows in the Young diagram of λ. For $1 \leq j \leq n$, we can define the power sums and discriminant by

$$P_j(x) := x_1^j + \ldots + x_k^j$$
$$\Delta(x) := \prod_{i<j}(x_i - x_j)$$

The Frobenius formula takes some product of these formulae and extracts the coefficient of a particular term. With this in mind, we define the following notation. If $p(x) = p(x_1, \ldots, x_k)$ is a polynomial of $x_1, \ldots, x_k$, and $(l_1, \ldots, l_k)$ a k-tuple of integers, then define $[p(x)]_{(l_1,\ldots,l_k)}$ to be the coefficient of the $x_1^{l_1} \cdot \cdots \cdot x_k^{l_k}$ term.

Now, recall that any character function is a class function, so we can define it on the conjugacy classes of S_n. However, we will need a slightly different way of describing these classes (though still based on cycle type). Let $i = (i_1, i_2, \ldots, i_n)$ be a sequence of integers so that $\sum_j j i_j = n$, and let C_i be the conjugacy class of S_n whose cycle types consist of i_1 1-cycles, i_2 2-cycles and so on. For example, if $n = 10$, we could take $i = (1, 3, 1) = (1, 3, 1, 0, \ldots, 0)$. Noting that $1 \cdot 1 + 2 \cdot 3 + 3 \cdot 1 = 10$, this means that C_i is the conjugacy class which has one 1-cycle, three 2-cycles, and one 3-cycle.

We can now state **Frobenius's formula**.

Theorem 4.1 (Frobenius's Formula)**.** *Given a partition λ, take any conjugacy class C_i as defined above, and define the sequence*

$$l_i = \lambda_i + k - i$$

We then have that the following formula holds.

$$\chi_\lambda(C_i) = \left[\Delta(x) \prod_j P_j(x)^{i_j}\right]_{(l_1,\ldots,l_k)}$$

As previously mentioned, we shall not prove this formula here. One may find the full proof of this formula in Fulton and Harris [1, §4.3].

We will give an example to illustrate this process. Let $n = 6$ and $\lambda = (4, 2)$. Let $i = (2, 2) = (2, 2, 0, 0, 0, 0)$ and C_i the associated conjugacy class (i.e. the conjugacy class of (12)(34)). In this case we will have two independent variables x_1, x_2 with $P_j(x) = x_1^j + x_2^j$ and $\Delta(x) = (x_1 - x_2)$. We also have the sequence $l = (l_1, l_2) = (5, 2)$. So using the formula we get

$$\begin{aligned}
\chi_\lambda(C_i) &= \left[\Delta(x)\prod_j P_j(x)^{i_j}\right]_{(l_1,\ldots,l_k)} \\
&= \left[(x_1-x_2)P_1(x)^2P_2(x)^2\right]_{(5,2)} \\
&= \left[(x_1-x_2)(x_1+x_2)^2(x_1^2+x_2^2)^2\right]_{(5,2)} \\
&= \left[x_1^7+x_1^6x_2+x_1^5x_2^2+x_1^4x_2^3-x_1^3x_2^4-x_1^2x_2^5-x_1x_2^6-x_2^7\right]_{(5,2)} \\
&= 1
\end{aligned}$$

Now, we know that for any representation, we have that evaluation of the character on the identity will always be the dimension of the representation. In our case, the identity element is made of n 1-cycles, so if we let $i=(n)$, then C_i will be the corresponding conjugacy class. In other words we have that

$$\dim V_\lambda = \chi_\lambda(C_{(n)}) = [\Delta(x)P_1(x)^n]_{(l_1,\ldots,l_k)}$$

One can rearrange this formula in the following form

$$\dim V_\lambda = \frac{n!}{l_1!\cdots l_k!}\prod_{i<j}(l_i-l_j) \tag{1}$$

We omit the details of this rearrangement here, which one may also find in Fulton and Harris [1, p. 49-50]. So in our example of $\lambda=(4,2)$ where $l=(5,2)$, we have that

$$\dim V_\lambda = \frac{6!}{5!2!}(5-2) = \frac{6\cdot 3}{2} = 9$$

The formula in (1) also has an alternative form, in terms of a more accessible *hook length.* Given a partition λ, we can assign a **hook length** to each box, which is the number of boxes directly to its right and directly below it, counting itself once. Using an earlier example of $\lambda=(5,2,2,1)$, its hook lengths are as follows:

8	6	3	2	1
4	2			
3	1			
1				

We can then compute $\dim V_\lambda$ as follows.

Proposition 4.2. [1] *Given a partition λ, we have that*

$$\dim V_\lambda = \frac{n!}{\prod \text{Hook lengths}}$$

[1]This only appears as an exercise in Fulton and Harris, and the proof is my own.

Proof. We prove this using equation (1).

First however, let $h_{i,j}$ denote the hook number of the box in row i and column j. We claim that

$$\prod_j h_{1,j} = \frac{l_1!}{\prod_j (l_1 - l_j)}$$

The proof of this will generalise to any row i.

Now let $r_{i,j}$ be the number of boxes strictly to the right of the box in position (i, j), and similarly let $b_{i,j}$ be the number of boxes strictly below it, so that $h_{i,j} = r_{i,j} + b_{i,j} + 1$. Now, notice that $r_{1,j} = \lambda_j - 1$ and $b_{1,j} = k - j - 1$. It then follows that

$$h_{1,j} = r_{i,j} + b_{i,j} + 1 = (\lambda_j - 1) + (k - j - 1) + 1 = \lambda_j + k - j - 1 = l_j$$

So the first column of hook numbers is precisely the sequence l, and in particular we have $h_{1,1} = l_1$.

Now, suppose we have a 2x2 square of boxes, then we have the following relation between them:

$$\begin{aligned}
h_{i+1,j+1} &= r_{i+1,j+1} + b_{i+1,j+1} + 1 \\
&= (r_{i,j+1} + 1) + (b_{i+1,j} + 1) + 1 \\
&= (h_{i,j+1} - b_{i,j+1}) + (h_{i+1,j} - r_{i+1,j}) + 1 \\
&= (h_{i,j+1} - b_{i,j} - 1) + (h_{i+1,j} - r_{i,j} - 1) + 1 \\
&= h_{i,j+1} + h_{i+1,j} - (b_{i,j} + r_{i,j} + 1) \\
&= h_{i,j+1} + h_{i+1,j} - h_{i,j}
\end{aligned}$$

Or in other words, the sum of their opposite corners are equal. By induction, we can extend this to larger rectangles (though we require that all corner boxes exist in the diagram).

With this we can prove the following claim. If we let $N = h_{1,1}$ and define the sets

$$\begin{aligned}
R_1 &:= \{h_{1,j} : 1 \leq j \leq \lambda_1\} \\
C_1 &:= \{N - h_{i,1} : 1 \leq i \leq k\}
\end{aligned}$$

then our claim is that

$$R_1 \cup C_1 = \{0, 1, \ldots, N\}$$

and that moreover this union is disjoint. Indeed, we have that $|R_1| = \lambda_1$ and $|C_1| = k$. And since $N = h_{1,1} = \lambda_1 + k - 1$, all we need to do is verify that the first two sets are disjoint (since every element of R_1 and C_1 is less than N).

So suppose that $x \in R_1 \cap C_1$. Then there are integers i, j so that $x = h_{1,j}$ and $N - x = h_{i,1}$. Now if we draw a rectangle with corners $(1, 1), (1, j), (i, 1), (i, j)$, then we know that the sum of the opposite corners of this rectangle are equal. So we would have $h_{i,j} = h_{i,1} + h_{1,j} - h1, 1 = N - x + x - N = 0$. This is clearly impossible, so the box (i, j) does not appear in the Young diagram. So this means that $r_{i,1} + 1 < j$ and $b_{1,j} + 1 < i$. Now, since $b_{i,1} = k - i$, it follows that

$$h_{i,1} = b_{i,1} + r_{i,1} + 1 < k + j - i$$

Similarly we get

$$h_{1,j} = b_{i,1} + r_{i,1} + 1 < \lambda_1 + i - j$$

And so we have that

$$N = h_{i,1} + h_{1,j} < (k + j - i) + \lambda_1 + i - j = k + \lambda_1 = N - 1$$

which is clearly a contradiction. So no such x exists, and the sets R_1 and C_1 are disjoint.

To complete the proof of our original claim, we have that

$$\begin{aligned} l_1! = N! &= \prod_{x \in R_1} x \prod_{y \in C_1} y \\ &= \prod_j h_{1,j} \prod_j (N - h_{i,j}) \\ &= \prod_j h_{1,j} \prod_j (l_1 - l_j) \end{aligned}$$

and thus our claim is proven. More generally, one can use this argument to show that for any i, we have

$$\prod_j h_{i,j} = \frac{l_i!}{\prod_{j>i}(l_i - l_j)}$$

and taking the product over all i, we obtain the hook length formula as desired. □

Recalling the example of $\lambda = (5, 2, 2, 1)$ given before, we can use this proposition to compute the dimension of V_λ as

$$\dim V_\lambda = \frac{n!}{8 \cdot 6 \cdot 4 \cdot 3 \cdot 3 \cdot 2 \cdot 2} = 525$$

We conclude this report by summarising the main topics we have covered. For the symmetric group S_n, we constructed the complete set of irreducible representations V_λ. These representations correspond to different partitions λ of n, which in turn correspond to the conjugacy classes of S_n. We have also stated a (somewhat involved) formula to calculate the character function χ_λ of these representations. We then use this to find the dimension of V_λ, which as we have seen can be computed relatively easily.

References

[1] Fulton, W., and Harris, J. *Representation Theory: A First Course.* Graduate texts in mathematics. Springer, 1991.