4.1. Gelfand-Zetlin Algebras

The chain (4.1) gives rise to a corresponding chain of group algebras,

$$\Bbbk = \Bbbk\mathcal{S}_1 \subseteq \Bbbk\mathcal{S}_2 \subseteq \cdots \subseteq \Bbbk\mathcal{S}_{n-1} \subseteq \Bbbk\mathcal{S}_n \subseteq \cdots . \tag{4.2}$$

The ***Gelfand-Zetlin*** (GZ) ***algebra*** $\mathcal{GZ}_n$ [**85**], [**86**] is defined to be the subalgebra of $\Bbbk\mathcal{S}_n$ that is generated by the centers $\mathscr{Z}_k := \mathscr{Z}(\Bbbk\mathcal{S}_k)$ for $k \le n$:

$$\boxed{\mathcal{GZ}_n \stackrel{\text{def}}{=} \Bbbk[\mathscr{Z}_1, \mathscr{Z}_2, \ldots, \mathscr{Z}_n] \subseteq \Bbbk\mathcal{S}_n .}$$

Note that all $\mathscr{Z}_k$ commute elementwise with each other: if $\alpha \in \mathscr{Z}_k$ and $\beta \in \mathscr{Z}_l$ with $k \le l$, say, then $\alpha \in \Bbbk\mathcal{S}_k \subseteq \Bbbk\mathcal{S}_l$ and $\beta \in \mathscr{Z}(\Bbbk\mathcal{S}_l)$ and hence $\alpha\beta = \beta\alpha$. Therefore, $\mathcal{GZ}_n$ is certainly commutative; in fact, the same argument will work for any chain of algebras in place of (4.2). In order to derive more interesting facts about $\mathcal{GZ}_n$, we will need to use additional properties of (4.2). For example, we will show that $\mathcal{GZ}_n$ is a *maximal* commutative subalgebra of $\Bbbk\mathcal{S}_n$ and that $\mathcal{GZ}_n$ is semisimple; see Theorem 4.4 below.

4.1.1. Centralizer Subalgebras

Our first goal is to exhibit a more economical set of generators for the algebra $\mathcal{GZ}_n$. This will be provided by the so-called ***Jucys-Murphy*** (JM) ***elements***, which will play an important role throughout this chapter. The n^{th} JM-element, denoted by X_n, is defined as the orbit sum of the transposition $(1, n) \in \mathcal{S}_n$ under the conjugation action $\mathcal{S}_{n-1} \circlearrowright \Bbbk\mathcal{S}_n$:

$$\boxed{X_n \stackrel{\text{def}}{=} \sum_{i=1}^{n-1} (i, n) \in (\Bbbk\mathcal{S}_n)^{\mathcal{S}_{n-1}} .}$$

Here, $(\Bbbk\mathcal{S}_n)^{\mathcal{S}_{n-1}}$ denotes the subalgebra of $\Bbbk\mathcal{S}_n$ consisting of all $\mathcal{S}_{n-1}$-invariants. Evidently, $(\Bbbk\mathcal{S}_n)^{\mathcal{S}_{n-1}}$ is contained in the invariant algebra $(\Bbbk\mathcal{S}_n)^{\mathcal{S}_k}$ of the conjugation action $\mathcal{S}_k \circlearrowright \Bbbk\mathcal{S}_n$ for all $k < n$, and $(\Bbbk\mathcal{S}_n)^{\mathcal{S}_k}$ can also be described as the centralizer of the subalgebra $\Bbbk\mathcal{S}_k$ in $\Bbbk\mathcal{S}_n$:

$$(\Bbbk\mathcal{S}_n)^{\mathcal{S}_k} = \{a \in \Bbbk\mathcal{S}_n \mid ab = ba \text{ for all } b \in \Bbbk\mathcal{S}_k\}.$$

By the foregoing, the JM-elements $X_{k+1}, \ldots, X_n$ all belong to $(\Bbbk\mathcal{S}_n)^{\mathcal{S}_k}$, and this algebra clearly also contains the center $\mathscr{Z}_k = (\Bbbk\mathcal{S}_k)^{\mathcal{S}_k}$ as well as the subgroup $\mathcal{S}'_{n-k} \le \mathcal{S}_n$ consisting of the permutations of $[n] = \{1, 2, \ldots, n\}$ that fix all elements of $[k]$. The following theorem is due to Olshanskiĭ [**166**].

Theorem 4.1. *The $\Bbbk$-algebra $(\Bbbk\mathcal{S}_n)^{\mathcal{S}_k}$ is generated by the center $\mathscr{Z}_k = \mathscr{Z}(\Bbbk\mathcal{S}_k)$, the subgroup $\mathcal{S}'_{n-k} \le \mathcal{S}_n$, and the JM-elements $X_{k+1}, \ldots, X_n$.*

Note that, for $m \geq k+1$,

$$X_m - (k+1,m) - \cdots - (m-1,m) = (k+1,m)X_{k+1}(k+1,m). \tag{4.3}$$

Since $(i,m) \in \mathcal{S}'_{n-k}$ for $k < i < m$, all but one of the JM-elements could be deleted from the list of generators in the theorem. However, our focus later on will be on the JM-elements rather than the other generators.

Before diving into the proof of Olshanskiĭ's Theorem, we remark that the $\mathcal{S}_k$-conjugacy class of any $s \in \mathcal{S}_n$ can be thought of in terms of "marked cycle shapes". In detail, if s is given as a product of disjoint cycles, possibly including 1-cycles, then we can represent ${}^{\mathcal{S}_k}s$ by the shape that is obtained by keeping each of $k+1,\ldots,n$ in its position in the given product while placing the symbol $*$ in all other positions. For example, the marked cycle shape

$$(*,*,*)(*,*)(*,*)(*)(12,*,15)(13,*,*)(14)$$

represents the $\mathcal{S}_{11}$-conjugacy class consisting of all permutations of $[15]$ that are obtained by filling the positions marked $*$ by the elements of $[11]$ in some order.

Proof of Theorem 4.1. We already know that $\mathcal{A} := \Bbbk[\mathscr{Z}_k, \mathcal{S}'_{n-k}, X_{k+1},\ldots,X_n] \subseteq \mathcal{B} := (\Bbbk\mathcal{S}_n)^{\mathcal{S}_k}$. In order to prove the inclusion $\mathcal{B} \subseteq \mathcal{A}$, observe that $\Bbbk\mathcal{S}_n$ is a permutation representation of $\mathcal{S}_k$. Hence a $\Bbbk$-basis of $\mathcal{B}$ is given by the $\mathcal{S}_k$-orbit sums (§3.3.1),

$$\boldsymbol{\sigma}_s := \sum_{t \in {}^{\mathcal{S}_k}s} t \qquad (s \in \mathcal{S}_n),$$

where ${}^{\mathcal{S}_k}s$ denotes the $\mathcal{S}_k$-conjugacy class of s. Our goal is to show that $\boldsymbol{\sigma}_s \in \mathcal{A}$ for all $s \in \mathcal{S}_n$.

To this end, we use a temporary notion of length[1] for elements $s \in \mathcal{S}_n$, defining $l(s)$ to be the number of points from $[n]$ that are moved by s or, equivalently, the number of symbols occurring in the disjoint cycle decomposition of s with all 1-cycles omitted. Clearly, $l(\,\cdot\,)$ is constant on conjugacy classes of $\mathcal{S}_n$. Moreover, $l(ss') \leq l(s) + l(s')$ for $s, s' \in \mathcal{S}_n$ and equality holds exactly if s and s' do not move a common point. Letting $\mathscr{F}_l \subseteq \Bbbk\mathcal{S}_n$ denote the $\Bbbk$-linear span of all $s \in \mathcal{S}_n$ with $l(s) \leq l$, we have $\mathscr{F}_0 = \Bbbk \subseteq \cdots \subseteq \mathscr{F}_{l-1} \subseteq \mathscr{F}_l \subseteq \cdots \subseteq \mathscr{F}_n = \Bbbk\mathcal{S}_n$ and $\mathscr{F}_l\mathscr{F}_{l'} \subseteq \mathscr{F}_{l+l'}$. Moreover, all subspaces $\mathscr{F}_l$ are $\mathcal{S}_k$-stable. Put

$$\mathcal{B}_l = \mathcal{B} \cap \mathscr{F}_l = \mathscr{F}_l^{\mathcal{S}_k}.$$

A basis of $\mathcal{B}_l$ is given by the orbit sums $\boldsymbol{\sigma}_s$ with $l(s) \leq l$. We will show by induction on l that $\mathcal{B}_l \subseteq \mathcal{A}$ or, equivalently, $\boldsymbol{\sigma}_s \in \mathcal{A}$ for all $s \in \mathcal{S}_n$ with $l(s) = l$.

To start, if $l = l(s) \leq 1$, then $\boldsymbol{\sigma}_s = s = 1 \in \mathcal{A}$. More generally, if $s = rt$ with $r \in \mathcal{S}_k$ and $t \in \mathcal{S}'_{n-k}$, then $\boldsymbol{\sigma}_s = \boldsymbol{\sigma}_r t$ and $\boldsymbol{\sigma}_r \in \mathscr{Z}_k$. Hence, $\boldsymbol{\sigma}_s \in \mathcal{A}$ again. Thus, we may assume that $s \notin \mathcal{S}_k \times \mathcal{S}'_{n-k}$; so the disjoint cycle decomposition of s involves a cycle of the form $(\ldots,i,m)$ with $i \leq k < m$. If $l = 2$, then $s = (i,m)$ and $\boldsymbol{\sigma}_s$

[1]This is not to be confused with another notion of "length", considered in Example 7.10.

is identical to the left-hand side of (4.3), which belongs to $\mathcal{A}$. Therefore, we may assume that $l > 2$ and that $\mathcal{B}_{l-1} \subseteq \mathcal{A}$.

Next, assume that $s = rt$ with $r, t \neq 1$ and $l(r) + l(t) = l$. Then $\boldsymbol{\sigma}_r, \boldsymbol{\sigma}_t \in \mathcal{A}$ by induction, and hence $\mathcal{A} \ni \boldsymbol{\sigma}_r \boldsymbol{\sigma}_t = \sum_{r',t''} r't''$, where r' and t'' run over the $\mathcal{S}_k$-conjugacy classes of r and t, respectively. If r' and t'' move a common point from $[n]$, then $l(r't'') < l(r') + l(t'') = l$. The sum of all these products $r't''$ is an $\mathcal{S}_k$-invariant belonging to $\mathcal{B}_{l-1} \subseteq \mathcal{A}$. Therefore, it suffices to consider the sum of the nonoverlapping products $r't''$. Each of these products has the same marked cycle shape as s, and hence it belongs to the $\mathcal{S}_k$-conjugacy class of s. By $\mathcal{S}_k$-invariance, the sum of the nonoverlapping products $r't''$ is a positive integer multiple of $\boldsymbol{\sigma}_s$. This shows that $\boldsymbol{\sigma}_r \boldsymbol{\sigma}_t \equiv z\boldsymbol{\sigma}_s \mod \mathcal{A}$ for some $z \in \mathbb{N}$, and we conclude that $\boldsymbol{\sigma}_s \in \mathcal{A}$.

It remains to treat the case where s is a cycle, say $s = (j_1, \ldots, j_{l-2}, i, m)$ with $i \leq k < m$. Write $s = rt$ with $r = (i, m)$ and $t = (j_1, \ldots, j_{l-2}, m)$. Since $\boldsymbol{\sigma}_r, \boldsymbol{\sigma}_t \in \mathcal{A}$ by induction, we once again have $\mathcal{A} \ni \boldsymbol{\sigma}_r \boldsymbol{\sigma}_t = \sum_{i',t'} (i', m)t'$ with $i' \leq k$ and t' running over the $\mathcal{S}_k$-conjugacy class of t. As above, the sum of all these products $(i', m)t'$ having length less than l belongs to $\mathcal{A}$ by induction. The products of length equal to l all have the form $(i', m)t' = (i', m)(j_1', \ldots, j_{l-2}', m) = (j_1', \ldots, j_{l-2}', i', m)$, and these products form the $\mathcal{S}_k$-conjugacy class of s. Therefore, we again have $\boldsymbol{\sigma}_r \boldsymbol{\sigma}_t \equiv z\boldsymbol{\sigma}_s \mod \mathcal{A}$ for some $z \in \mathbb{N}$, which finishes the proof. □

4.1.2. Generators of the Gelfand-Zetlin Algebra

As a consequence of Theorem 4.1, we obtain the promised generating set for the Gelfand-Zetlin algebra: $\mathcal{GZ}_n$ is generated by the JM-elements X_k with $k \leq n$. Even though $X_1 = 0$ is of course not needed as a generator, it will be convenient to keep this element in the list.

Corollary 4.2. $\mathcal{GZ}_n = \mathbb{k}[X_1, X_2, \ldots, X_n]$.

Proof. First note that

$$X_k = \sum \{\text{all transpositions of } \mathcal{S}_k\} - \sum \{\text{all transpositions of } \mathcal{S}_{k-1}\}.$$

Since the first sum belongs to $\mathscr{Z}_k$ and the second to $\mathscr{Z}_{k-1}$, it follows that $X_k \in \mathcal{GZ}_n = \mathbb{k}[\mathscr{Z}_1, \ldots, \mathscr{Z}_n]$. For the inclusion $\mathcal{GZ}_n \subseteq \mathbb{k}[X_1, \ldots, X_n]$, we proceed by induction on n. The case of $\mathcal{GZ}_1 = \mathbb{k}$ being clear, assume that $n > 1$ and that $\mathcal{GZ}_{n-1} \subseteq \mathbb{k}[X_1, \ldots, X_{n-1}]$. Since $\mathcal{GZ}_n = \mathbb{k}[\mathcal{GZ}_{n-1}, \mathscr{Z}_n]$ by definition, it suffices to show that $\mathscr{Z}_n \subseteq \mathbb{k}[\mathcal{GZ}_{n-1}, X_n]$. But

$$\mathscr{Z}_n = (\mathbb{k}\mathcal{S}_n)^{\mathcal{S}_n} \subseteq (\mathbb{k}\mathcal{S}_n)^{\mathcal{S}_{n-1}} \underset{\text{Theorem 4.1}}{=} \mathbb{k}[\mathscr{Z}_{n-1}, X_n] \subseteq \mathbb{k}[\mathcal{GZ}_{n-1}, X_n]$$

as desired. □

Exercises for Section 4.1

4.1.1 (Relations between JM-elements and Coxeter generators)**.** The transpositions $s_i = (i, i+1)$ $(i = 1, \dots, n-1)$ are called the ***Coxeter generators*** of $\mathcal{S}_n$. Show that the following relations hold for the Coxeter generators and the JM-elements: $s_i X_i + 1 = X_{i+1} s_i$ and $s_i X_j = X_j s_i$ if $j \neq i, i+1$.

4.1.2 (Product of the JM-elements)**.** Show that $X_2 X_3 \cdots X_n$ is the sum of all n-cycles in $\mathcal{S}_n$.

4.1.3 (Semisimplicity of some subalgebras of $\Bbbk\mathcal{S}_n$)**.** Show that any subalgebra of $\Bbbk\mathcal{S}_n$ that is generated by some of the JM-elements X_i $(i \leq n)$ and some subgroups of $\mathcal{S}_n$ and the centers of some subgroup algebras of $\Bbbk\mathcal{S}_n$ is semisimple. In particular, the centralizer algebras $(\Bbbk\mathcal{S}_n)^{\mathcal{S}_k}$ and $\mathcal{GZ}_n$ are semisimple. (Use Exercise 3.4.2 and the fact that all these subalgebras are stable under the standard involution of $\Bbbk\mathcal{S}_n$ and defined over $\mathbb{Q}$.)

4.2. The Branching Graph

In this section, we define the first of two graphs that will play a major role in this chapter: the branching graph $\mathbb{B}$. This graph efficiently encodes a great deal of information concerning the irreducible representations of the various symmetric groups $\mathcal{S}_n$.

4.2.1. Restricting Irreducible Representations

The developments in this section hinge on the following observation.

Multiplicity-Freeness Theorem. *For each $V \in \operatorname{Irr} \mathcal{S}_n$, the restriction $V{\downarrow}_{\mathcal{S}_{n-1}}$ is a direct sum of nonisomorphic irreducible representations of $\mathcal{S}_{n-1}$.*

Proof. Since $\Bbbk\mathcal{S}_{n-1}$ is split semisimple, we have $V{\downarrow}_{\mathcal{S}_{n-1}} \cong \bigoplus_{W \in \operatorname{Irr} \mathcal{S}_{n-1}} W^{\oplus m(W)}$ with $m(W) \in \mathbb{Z}_+$ and so $\operatorname{End}_{\mathcal{S}_{n-1}}(V) \cong \prod_W \operatorname{Mat}_{m(W)}(\Bbbk)$ (Proposition 1.33). The theorem states that $m(W) \leq 1$ for all W, which is equivalent to the assertion that the algebra $\operatorname{End}_{\mathcal{S}_{n-1}}(V)$ is commutative. Similarly, since $\Bbbk\mathcal{S}_n$ is split semisimple, we have the standard isomorphism (1.46) of $\Bbbk$-algebras:

$$
\begin{array}{ccc}
\Bbbk\mathcal{S}_n & \xrightarrow{\sim} & \prod_{V \in \operatorname{Irr} \mathcal{S}_n} \operatorname{End}_{\Bbbk}(V) \\
\cup & & \cup \\
a & \longmapsto & (a_V)
\end{array}
\tag{4.4}
$$

Under this isomorphism, the conjugation action $\mathcal{S}_n \circlearrowright \Bbbk\mathcal{S}_n$ translates into the standard $\mathcal{S}_n$-action on each component $\operatorname{End}_{\Bbbk}(V)$: $({}^s a)_V = s_V \circ a_V \circ s_V^{-1} = s.a_V$.

Therefore, the isomorphism (4.4) restricts to an isomorphism of algebras of $\mathcal{S}_{n-1}$-invariants,

$$(\Bbbk\mathcal{S}_n)^{\mathcal{S}_{n-1}} \xrightarrow{\sim} \prod_{V \in \operatorname{Irr} \mathcal{S}_n} \operatorname{End}_{\Bbbk}(V)^{\mathcal{S}_{n-1}} \underset{(3.30)}{=} \prod_{V \in \operatorname{Irr} \mathcal{S}_n} \operatorname{End}_{\mathcal{S}_{n-1}}(V).$$

By Theorem 4.1, $(\Bbbk\mathcal{S}_n)^{\mathcal{S}_{n-1}} = \Bbbk[\mathscr{Z}_{n-1}, X_n]$ is a commutative algebra. Consequently, all $\operatorname{End}_{\mathcal{S}_{n-1}}(V)$ are commutative as well, as desired. □

4.2.2. The Graph $\mathbb{B}$

Consider the following graph $\mathbb{B}$, called the ***branching graph*** of the chain (4.2). The set of vertices of $\mathbb{B}$ is defined by

$$\operatorname{vert} \mathbb{B} = \bigsqcup_{n \geq 1} \operatorname{Irr} \mathcal{S}_n .$$

For given vertices $W \in \operatorname{Irr} \mathcal{S}_{n-1}$ and $V \in \operatorname{Irr} \mathcal{S}_n$, the graph $\mathbb{B}$ has a directed edge $W \to V$ if and only if there is a nonzero map $W \to V{\downarrow}_{\mathcal{S}_{n-1}}$ in $\mathsf{Rep}\,\mathcal{S}_{n-1}$, that is, W is an irreducible constituent of $V{\downarrow}_{\mathcal{S}_{n-1}}$. Thus, the vertices of $\mathbb{B}$ are organized into levels, with $\operatorname{Irr} \mathcal{S}_n$ being the set of level-n vertices, and all arrows in $\mathbb{B}$ are directed toward the next higher level. Figure 4.1 shows the first five levels of $\mathbb{B}$ (Exercise 4.2.1).

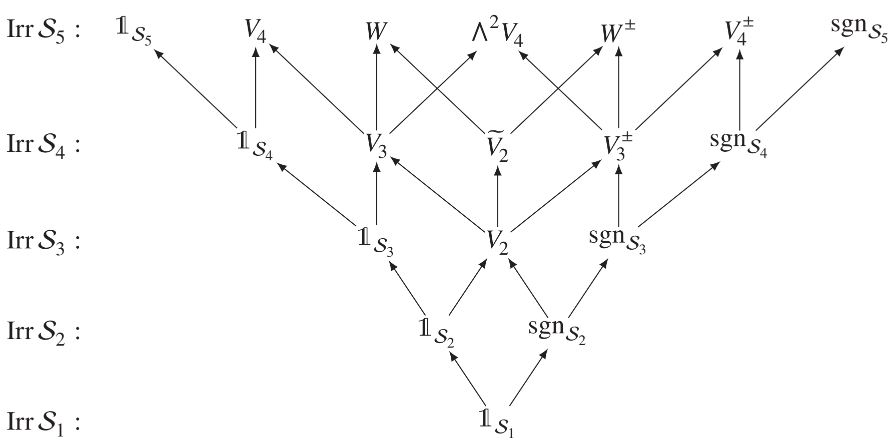

Figure 4.1. Bottom of the branching graph $\mathbb{B}$ (notation of §3.5.2)

The Multiplicity-Freeness Theorem can now be stated as follows:

$$V{\downarrow}_{\mathcal{S}_{n-1}} \cong \bigoplus_{\substack{W \to V \\ \text{in } \mathbb{B}}} W. \tag{4.5}$$

Note that the decomposition (4.5) is *canonical*: the image of W in $V{\downarrow}_{\mathcal{S}_{n-1}}$ is uniquely determined as the W-homogeneous component of $V{\downarrow}_{\mathcal{S}_{n-1}}$ and the map $W \to V{\downarrow}_{\mathcal{S}_{n-1}}$ is a monomorphism in $\mathsf{Rep}\,\mathcal{S}_{n-1}$ that is uniquely determined up to a scalar multiple by Schur's Lemma.

4.2.3. Gelfand-Zetlin Bases

Let $V \in \operatorname{Irr} \mathcal{S}_n$ be given. The following procedure yields a canonical decomposition of V into 1-dimensional subspaces. Start by decomposing $V{\downarrow}_{\mathcal{S}_{n-1}}$ into irreducible constituents as in (4.5); then, for each arrow $W \to V$ in $\mathbb{B}$, decompose $W{\downarrow}_{\mathcal{S}_{n-2}}$ into irreducibles. Proceeding in this manner all the way down to $\mathcal{S}_1 = 1$, we obtain the desired decomposition of the vector space $V{\downarrow}_{\mathcal{S}_1}$ into 1-dimensional subspaces, one for each path $T\colon \mathbb{1}_{\mathcal{S}_1} \to \cdots \to V$ in $\mathbb{B}$. The resulting decomposition of V is uniquely determined, because the various decompositions at each step are unique. Choosing $0 \neq v_T$ in the subspace of V corresponding to the path T, we obtain a basis (v_T) of V and each v_T is determined up to a scalar multiple. This basis is called "the" ***Gelfand-Zetlin*** (GZ) ***basis*** of V; of course, any rescaling of this basis would also be a GZ-basis. To summarize,

(4.6)
$$\boxed{V{\downarrow}_{\mathcal{S}_1} = \bigoplus_{\substack{T\,:\,\mathbb{1}_{\mathcal{S}_1} \to \cdots \to V \\ \text{in } \mathbb{B}}} \Bbbk\, v_T\,.}$$

It follows in particular that

(4.7)
$$\dim V = \#\{\text{paths } \mathbb{1}_{\mathcal{S}_1} \to \cdots \to V \text{ in } \mathbb{B}\}\,.$$

The reader is invited to check that there are five such paths in Figure 4.1 for $V = W$, and six paths for $V = \Lambda^2 V_4$. This does of course agree with what we know already about the dimensions of these representations (§3.5.2). For a generalization of (4.7), see Exercise 4.2.3.

Example 4.3 (GZ-basis of the standard representation V_{n-1})**.** Consider the chain $M_1 \subseteq M_2 \subseteq \cdots \subseteq M_n \subseteq \cdots$, where $M_n = \bigoplus_{i=1}^n \Bbbk b_i$ is the standard permutation representation of $\mathcal{S}_n$ (§3.2.4). Working inside $\bigcup_{n\geq 1} M_n = \bigoplus_{i\geq 1} \Bbbk b_i$, we have $V_{n-1} = \{\sum_i \lambda_i b_i \mid \sum_i \lambda_i = 0 \text{ and } \lambda_i = 0 \text{ for } i > n\}$ and so $\cdots \subseteq V_{n-2} \subseteq V_{n-1} \subseteq \cdots$. Thus, V_{n-2} provides us with an irreducible component of $V_{n-1}{\downarrow}_{\mathcal{S}_{n-1}}$. The vector

$$v_{n-1} = \sum_{i=1}^{n-1} (b_i - b_n) = \sum_{i=1}^{n-1} b_i - (n-1) b_n \in V_{n-1}$$

is a nonzero $\mathcal{S}_{n-1}$-invariant that does not belong to V_{n-2}. For dimension reasons, we conclude that $V_{n-1}{\downarrow}_{\mathcal{S}_{n-1}} = \Bbbk v_{n-1} \oplus V_{n-2} \cong \mathbb{1}_{\mathcal{S}_{n-1}} \oplus V_{n-2}$ is the decomposition of $V_{n-1}{\downarrow}_{\mathcal{S}_{n-1}}$ into irreducible constituents. Inductively we further deduce that

$V_{n-1} = \bigoplus_{j=1}^{n-1} \Bbbk v_j$ and that $(v_1, \dots, v_{n-1})$ is the GZ-basis of V_{n-1}. It is straightforward to check that the Coxeter generator $s_i = (i, i+1) \in \mathcal{S}_n$ acts on this basis as follows:

$$s_i.v_j = \begin{cases} v_j & \text{for } j \neq i-1, i, \\ \frac{1}{i} v_{i-1} + (1 - \frac{1}{i})\, v_i & \text{for } j = i-1, \\ (1 + \frac{1}{i}) v_{i-1} - \frac{1}{i}\, v_i & \text{for } j = i. \end{cases} \tag{4.8}$$

These equations determine the vectors $v_j \in V_{n-1}$ up to a common scalar factor: if (4.8) also holds with w_j in place of v_j, then $v_j \mapsto w_j$ is an $\mathcal{S}_n$-equivariant endomorphism of V_{n-1} and hence an element of $D(V_{n-1}) = \Bbbk$. We shall discuss some rescalings of the GZ-basis (v_j) in Examples 4.17 and 4.19.

4.2.4. Properties of $\mathcal{GZ}_n$

We have seen that the Gelfand-Zetlin algebra $\mathcal{GZ}_n$ is commutative and generated by the JM-elements $X_1, \dots, X_n$. Now we derive further information about $\mathcal{GZ}_n$ from the foregoing.

Theorem 4.4. (a) *$\mathcal{GZ}_n$ is the set of all $a \in \Bbbk\mathcal{S}_n$ such that the GZ-basis of each $V \in \operatorname{Irr}\mathcal{S}_n$ consists of eigenvectors for a_V.*

(b) *$\mathcal{GZ}_n$ is a maximal commutative subalgebra of $\Bbbk\mathcal{S}_n$.*

(c) *$\mathcal{GZ}_n$ is semisimple: $\mathcal{GZ}_n \cong \Bbbk^{\times d_n}$ with $d_n = \sum_{V \in \operatorname{Irr}\mathcal{S}_n} \dim V$.*

Proof. For each $V \in \operatorname{Irr}\Bbbk\mathcal{S}_n$, identify $\operatorname{End}_{\Bbbk}(V)$ with the matrix algebra $\operatorname{Mat}_{\dim V}(\Bbbk)$ via the GZ-basis of V. Then the isomorphism (4.4) identifies the group algebra $\Bbbk\mathcal{S}_n$ with the direct product of these matrix algebras. Let $\mathcal{D}$ denote the subalgebra of $\Bbbk\mathcal{S}_n$ that corresponds to the direct product of the algebras of diagonal matrices in each component. Part (a) asserts that $\mathcal{D} = \mathcal{GZ}_n$:

$$\begin{array}{ccc} \Bbbk\mathcal{S}_n & \xrightarrow[\text{via GZ-bases}]{\sim} & \prod_{V \in \operatorname{Irr}\mathcal{S}_n} \operatorname{Mat}_{\dim V}(\Bbbk) \\ \cup\!\text{I} & & \cup\!\text{I} \\ \mathcal{GZ}_n & \xrightarrow{\sim} & \prod_{V \in \operatorname{Irr}\mathcal{S}_n} \begin{pmatrix} \Bbbk & & \\ & \ddots & \\ & & \Bbbk \end{pmatrix} \end{array} \tag{4.9}$$

The isomorphism $\mathcal{GZ}_n \cong \Bbbk^{\times d_n}$ in (c) is then clear and so is the maximality assertion in (b). Indeed, the subalgebra of diagonal matrices in any matrix algebra $\operatorname{Mat}_d(\Bbbk)$ is self-centralizing: the only matrices that commute with all diagonal matrices are themselves diagonal. Therefore, $\mathcal{D}$ is a self-centralizing subalgebra of $\Bbbk\mathcal{S}_n$, and hence $\mathcal{D}$ is a maximal commutative subalgebra. In particular, in order to prove the equality $\mathcal{D} = \mathcal{GZ}_n$, it suffices to show that $\mathcal{D} \subseteq \mathcal{GZ}_n$, because we already know that $\mathcal{GZ}_n$ is commutative.

To prove the inclusion $\mathcal{D} \subseteq \mathcal{GZ}_n$, let $\boldsymbol{e}(V) \in \mathscr{Z}_n$ denote the primitive central idempotent of $\Bbbk\mathcal{S}_n$ corresponding to $V \in \operatorname{Irr}\mathcal{S}_n$ (§1.4.4). Recall that, for any $W \in \mathsf{Rep}\,\mathcal{S}_n$, the operator $\boldsymbol{e}(V)_W$ projects W onto the V-homogeneous component $W(V)$, annihilating all other homogeneous components of W. Therefore, for any path $T\colon \mathbb{1}_{\mathcal{S}_1} = W_1 \to W_2 \to \cdots \to W_n = V$ in $\mathbb{B}$, the element $\boldsymbol{e}(T) := \boldsymbol{e}(W_1)\boldsymbol{e}(W_2)\cdots\boldsymbol{e}(W_n) \in \Bbbk[\mathscr{Z}_1, \mathscr{Z}_2, \dots, \mathscr{Z}_n] = \mathcal{GZ}_n$ acts on V as the projection $\pi_T : V \twoheadrightarrow V_T = \Bbbk v_T$ in (4.6) and $\boldsymbol{e}(T)_{V'} = 0_{V'}$ for all $V \neq V' \in \operatorname{Irr}\mathcal{S}_n$. Thus, in (4.9), we have:

$$\begin{array}{ccc} \Bbbk\mathcal{S}_n & \xrightarrow{\ \sim\ } & \prod_{V\in\operatorname{Irr}\mathcal{S}_n} \operatorname{Mat}_{\dim V}(\Bbbk) \\ \cup\!\!\!\!\! \text{\tiny I} & & \cup\!\!\!\!\! \text{\tiny I} \\ \boldsymbol{e}(T) & \longmapsto & (0,\dots,\pi_T,\dots,0) \end{array}$$

This shows that the idempotents $\boldsymbol{e}(T)$ form the standard basis of the diagonal algebra $\mathcal{D}$, consisting of the diagonal matrices with one entry equal to 1 and all others 0, which proves the desired inclusion $\mathcal{D} \subseteq \mathcal{GZ}_n$. □

4.2.5. The Spectrum of $\mathcal{GZ}_n$

We now give a description of $\operatorname{Spec}\mathcal{GZ}_n = \operatorname{MaxSpec}\mathcal{GZ}_n \cong \operatorname{Hom}_{\mathsf{Alg}_\Bbbk}(\mathcal{GZ}_n, \Bbbk)$ (§1.3.2) that will play an important role in Section 4.4. For this, we elaborate on some of the properties of $\mathcal{GZ}_n$ stated in Theorem 4.4. First, the fact that the GZ-basis (v_T) of any $V \in \operatorname{Irr}\mathcal{S}_n$ consists of eigenvectors for $\mathcal{GZ}_n$ says that each v_T is a weight vector for a suitable ***weight*** $\phi_T \in \operatorname{Hom}_{\mathsf{Alg}_\Bbbk}(\mathcal{GZ}_n, \Bbbk)$:

$$a.v_T = \phi_T(a)v_T \qquad (a \in \mathcal{GZ}_n). \tag{4.10}$$

Moreover, in view of (4.7), the dimension $d_n = \dim\mathcal{GZ}_n$ is equal to the total number of paths in $\mathbb{B}$ from $\mathbb{1}_{\mathcal{S}_1}$ to some vertex $\in \operatorname{Irr}\mathcal{S}_n$. Finally, the isomorphism $\mathcal{GZ}_n \xrightarrow{\sim} \Bbbk^{\times d_n}$ in (4.9) is given by $a \mapsto (\phi_T(a))_T$. Therefore, each ϕ_T is a weight of a unique $V \in \operatorname{Irr}\mathcal{S}_n$, the endpoint of the path $T\colon \mathbb{1}_{\mathcal{S}_1} \to \cdots \to V$ in (4.6), and

$$\boxed{\operatorname{Hom}_{\mathsf{Alg}_\Bbbk}(\mathcal{GZ}_n, \Bbbk) = \left\{\phi_T \;\middle|\; \begin{array}{l} T \text{ a path } \mathbb{1}_{\mathcal{S}_1} \to \cdots \text{ in } \mathbb{B} \\ \text{with endpoint} \in \operatorname{Irr}\mathcal{S}_n \end{array}\right\}.} \tag{4.11}$$

Since the algebra $\mathcal{GZ}_n$ is generated by the JM-elements $X_1, \dots, X_n$ (Corollary 4.2), each weight ϕ_T is determined by the n-tuple $(\phi_T(X_i))_1^n \in \Bbbk^n$. Therefore, the spectrum $\operatorname{Hom}_{\mathsf{Alg}_\Bbbk}(\mathcal{GZ}_n, \Bbbk)$ of $\mathcal{GZ}_n$ is in one-to-one correspondence with the set

$$\boxed{\operatorname{Spec}(n) \stackrel{\text{def}}{=} \{(\phi(X_1), \phi(X_2), \dots, \phi(X_n)) \in \Bbbk^n \mid \phi \in \operatorname{Hom}_{\mathsf{Alg}_\Bbbk}(\mathcal{GZ}_n, \Bbbk)\}.}$$

To summarize, we have the following bijections:

$$
(4.12)\qquad
\begin{array}{ccccc}
\operatorname{Spec}(n) & \overset{\sim}{\longleftrightarrow} & \operatorname{Hom}_{\mathsf{Alg}_{\Bbbk}}(\mathcal{GZ}_n, \Bbbk) & \underset{\text{via GZ-bases}}{\overset{\sim}{\longleftrightarrow}} & \left\{ \begin{array}{c} \text{paths } \mathbb{1}_{\mathcal{S}_1} \to \cdots \text{ in } \mathbb{B} \\ \text{with endpoint} \in \operatorname{Irr} \mathcal{S}_n \end{array} \right\} \\
\Cup & & \Cup & & \Cup \\
(\phi_T(X_i))_1^n & \longleftrightarrow & \phi_T & \longleftrightarrow & T
\end{array}
$$

Exercises for Section 4.2

4.2.1 (Bottom of $\mathbb{B}$)**.** Verify the bottom of the branching graph $\mathbb{B}$ as in Figure 4.1.

4.2.2 (Dimension of the Gelfand-Zetlin algebra)**.** Note that $d_n = \dim \mathcal{GZ}_n$ is also the length of the group algebra $\Bbbk\mathcal{S}_n$. Show that the first five values of the sequence d_n are: $1, 2, 4, 10, 26$.[2]

4.2.3 (Lengths of homogeneous components)**.** Let $V \in \operatorname{Irr} \mathcal{S}_n$ and $W \in \operatorname{Irr} \mathcal{S}_k$ and assume that $k \leq n$. Show that the length of the W-homogeneous component of $V{\downarrow_{\mathcal{S}_k}}$ is equal to the number of paths $W \to \cdots \to V$ in $\mathbb{B}$.

4.2.4 (Orthogonality of GZ-bases)**.** Let $V \in \operatorname{Irr} \mathcal{S}_n$ and let $(\,.\,,\,.\,) : V \times V \to \Bbbk$ be any $\mathcal{S}_n$-invariant bilinear form; so $(s.v, s.v') = (v, v')$ for all $v, v' \in V$ and $s \in \mathcal{S}_n$. Show that the GZ-basis (v_T) of V is orthogonal: $(v_T, v_{T'}) = 0$ for $T \neq T'$. (Use the fact that representations of symmetric groups are self-dual by Lemma 3.24.)

4.2.5 (Weights of the standard representation)**.** Let V_{n-1} be the standard representation of $\mathcal{S}_n$ and let $v_j = b_1 + b_2 + \cdots + b_j - jb_{j+1}$ as in Example 4.3. Show that v_j has weight $(0, 1, \ldots, j-1, -1, j, \ldots, n-2)$.

4.3. The Young Graph

We now start afresh, working in purely combinatorial rather than representation-theoretic territory.

4.3.1. Partitions and Young Diagrams

The main player in this section is the following set of nonnegative integer sequences:

$$
\boxed{\mathscr{P} \overset{\text{def}}{=} \{(\lambda_1, \lambda_2, \ldots) \in \mathbb{Z}_+^{\mathbb{N}} \mid \lambda_1 \geq \lambda_2 \geq \cdots \text{ and } \textstyle\sum_i \lambda_i < \infty\}.}
$$

We will denote the sequence $(\lambda_1, \lambda_2, \ldots) \in \mathscr{P}$ by λ and write $|\lambda| = \sum_i \lambda_i$. Thus,

$$
\mathscr{P} = \bigsqcup_n \mathscr{P}_n \qquad \text{with} \qquad \mathscr{P}_n \overset{\text{def}}{=} \{\lambda \in \mathscr{P} \mid |\lambda| = n\}.
$$

[2]For more information on this sequence, see The On-Line Encyclopedia of Integer Sequences [**194**, Sequence A000085].

The members of $\mathscr{P}_n$ are called ***partitions*** of n. For $\lambda \in \mathscr{P}_n$, we will also use the standard notation $\lambda \vdash n$. Partitions will be visualized by ***Young diagrams***: the Young diagram of $\lambda = (\lambda_1, \lambda_2, \dots)$ consists of rows of boxes that are aligned on the left, with λ_1 boxes in the first row, λ_2 in the second, etc. The unique partition with $|\lambda| = 0$ thus has an empty Young diagram. We will generally only consider partitions with $|\lambda| \geq 1$ (except in Exercise 4.3.1) and we will typically write partitions as finite sequences, omitting a tail of 0-components. Here, for example, is the Young diagram of the partition $(7, 5, 4, 4, 2)$:

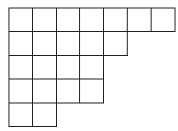

Young diagrams are also called Ferrers diagrams, particularly when represented using dots instead of boxes, and sometimes the rows are arranged in the reverse order (French notation), but we will use the above convention.

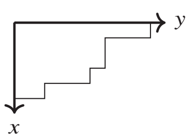

The columns and rows of a given Young diagram will be numbered left-to-right and top-to-bottom, respectively, starting with number 1. The box in the x^{th} row and y^{th} column will also be referred to as the box in position (x, y) or the (x, y)-box.

Reflecting a given partition λ across the line $y = x$ yields the so-called ***conjugate partition***; it will be denoted by λ^c. For example, $(7, 5, 4, 4, 2)^c = (5, 5, 4, 4, 2, 1, 1)$.

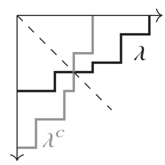

4.3.2. The graph $\mathbb{Y}$ and the Graph Isomorphism Theorem

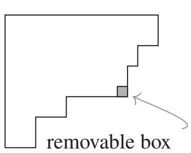

The ***Young graph*** $\mathbb{Y}$ has vertex set

$$\operatorname{vert} \mathbb{Y} = \mathscr{P},$$

with each $\lambda \in \mathscr{P}$ represented by its Young diagram. An arrow $\mu \to \lambda$ in $\mathbb{Y}$ means that the Young diagram of μ is obtained from the one of λ by removing one box, necessarily a "southeast" corner box. Note that the number of removable boxes of $\lambda = (\lambda_i)_{i \geq 1}$ is equal to the number of distinct values among the λ_i. We define a partial order $\leq$ on $\mathscr{P}$ by declaring $\mu \leq \lambda$ if the Young diagram of μ fits into the diagram of λ or, more formally, $\mu_i \leq \lambda_i$ for all i. Thus, there is an arrow $\mu \to \lambda$ in $\mathbb{Y}$ iff $\mu < \lambda$ but no $\nu \in \mathscr{P}$ satisfies $\mu < \nu < \lambda$.

The vertices of $\mathbb{Y}$ are divided into levels, with $\mathscr{P}_n$ at level n. The first five levels of $\mathbb{Y}$, as displayed in Figure 4.2, show a striking similarity to the corresponding levels of the branching graph $\mathbb{B}$ (Figure 4.1). In fact, we have the following fundamental theorem.

Graph Isomorphism Theorem. *The graphs $\mathbb{Y}$ and $\mathbb{B}$ are isomorphic.*

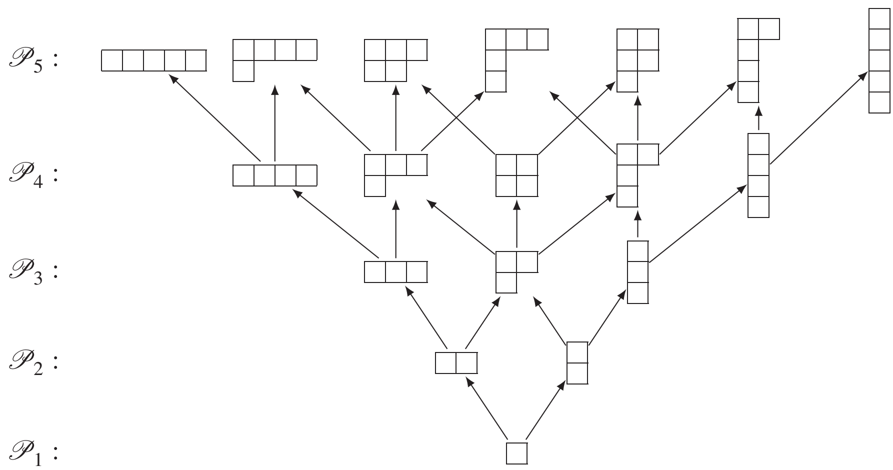

Figure 4.2. Bottom of the Young graph $\mathbb{Y}$

Explicitly, the theorem asserts the existence of a bijection $\phi\colon \operatorname{vert}\mathbb{Y} \xrightarrow{\sim} \operatorname{vert}\mathbb{B}$ such that there is an arrow $\mu \to \lambda$ in $\mathbb{Y}$ if and only if there is an arrow $\phi(\mu) \to \phi(\lambda)$ in $\mathbb{B}$. We will then write $\phi\colon \mathbb{Y} \xrightarrow{\sim} \mathbb{B}$. We may of course also speak of automorphisms $\mathbb{B} \xrightarrow{\sim} \mathbb{B}$ and likewise for $\mathbb{Y}$. In fact, it is not hard to see that conjugation, $\lambda \mapsto \lambda^c$, is the only nonidentity automorphism of $\mathbb{Y}$ (Exercise 4.3.2). Thus, there are at most two possible graph isomorphisms $\mathbb{Y} \xrightarrow{\sim} \mathbb{B}$.

The proof of the Graph Isomorphism Theorem will be given in Section 4.4. In the remainder of this section, we will discuss some consequences. Throughout, let us write the bijection on vertices given by the graph isomorphism $\mathbb{Y} \xrightarrow{\sim} \mathbb{B}$ as

$$\lambda \longmapsto V^\lambda \qquad (\lambda \in \mathscr{P}).$$

For example, we clearly must have $V^\square = \mathbb{1}_{\mathcal{S}_1}$, because $\square$ and $\mathbb{1}_{\mathcal{S}_1}$ are the sole vertices of $\mathbb{Y}$ and $\mathbb{B}$ with no incoming arrows. More generally, any isomorphism $\mathbb{Y} \xrightarrow{\sim} \mathbb{B}$ must bijectively map the n-vertex paths $\square \to \mu_2 \to \cdots \to \mu_n$ in $\mathbb{Y}$ to the corresponding paths in $\mathbb{B}$, and hence it will match the level-n vertices of $\mathbb{Y}$ with the level-n vertices of $\mathbb{B}$, giving bijections $\mathscr{P}_n \xrightarrow{\sim} \operatorname{Irr}\mathcal{S}_n$ for all n. Of course, we already know that such bijections exist: there are as many irreducible representations of $\mathcal{S}_n$ as there are conjugacy classes of $\mathcal{S}_n$ (Corollary 3.21) and the conjugacy classes in turn are in bijection with the partitions of n (§3.5.2). However, the full form of the the Graph Isomorphism Theorem gives much more detailed information and its proof will require more work.

4.3.3. Consequences of the Graph Isomorphism Theorem

The Graph Isomorphism Theorem allows us to derive information about the irreducible representations of all $\mathcal{S}_n$ from combinatorial features of $\mathbb{Y}$. Here are some examples.

The Branching Rule. Since $\mu \to \lambda$ in $\mathbb{Y}$ is equivalent to $V^\mu \to V^\lambda$ in $\mathbb{B}$, we may rewrite (4.5) as

$$V^\lambda\!\downarrow_{\mathcal{S}_{n-1}} \cong \bigoplus_{\substack{\mu\to\lambda \\ \text{in } \mathbb{Y}}} V^\mu. \tag{4.13}$$

This formula is referred to as the ***branching rule***. The number of arrows $\mu \to \lambda$ in (4.13) is equal to the number of removable boxes of λ, which in turn is equal to the number of distinct values (row lengths in the Young diagram) of λ. Therefore,

$$\text{length } V^\lambda\!\downarrow_{\mathcal{S}_{n-1}} = \#\{\text{distinct values of } \lambda\}.$$

In particular, $V^\lambda\!\downarrow_{\mathcal{S}_{n-1}}$ is irreducible if and only if the Young diagram of λ is a rectangle. The branching rule (4.13) has the following reformulation:

$$V^\mu\!\uparrow^{\mathcal{S}_n} \cong \bigoplus_{\substack{\mu\to\lambda \\ \text{in } \mathbb{Y}}} V^\lambda. \tag{4.14}$$

Indeed, (4.13) says that the multiplicity of V^μ in $V^\lambda\!\downarrow_{\mathcal{S}_{n-1}}$ is equal to 1 if there is an arrow $\mu \to \lambda$ in $\mathbb{Y}$ and equal to 0 otherwise; (4.14) makes the same statement about the multiplicity of V^λ in $V^\mu\!\uparrow^{\mathcal{S}_n}$. The equivalence of (4.13) and (4.14) thus follows from Frobenius reciprocity (Corollary 1.37): $m(V^\mu, V^\lambda\!\downarrow_{\mathcal{S}_{n-1}}) = m(V^\lambda, V^\mu\!\uparrow^{\mathcal{S}_n})$.

Dimension. Any graph isomorphism $\mathbb{Y} \xrightarrow{\sim} \mathbb{B}$ will induce bijections between the set of all paths $\square \to \cdots \to \lambda$ in $\mathbb{Y}$ and the set of all paths $\mathbb{1}_{\mathcal{S}_1} \to \cdots \to V^\lambda$ in $\mathbb{B}$. Since the size of the latter set of paths equals $\dim V^\lambda$ by (4.7), we obtain

$$\dim V^\lambda = f^\lambda \stackrel{\text{def}}{=} \#\{\text{paths } \square \to \cdots \to \lambda \text{ in } \mathbb{Y}\}. \tag{4.15}$$

The number f^λ will be determined in §4.3.5. Note that the dimension $d_n = \dim \mathcal{GZ}_n$ (Theorem 4.4) can now also be written as $d_n = \sum_{\lambda \vdash n} f^\lambda$.

4.3.4. Paths in $\mathbb{Y}$ and Standard Young Tableaux

In this subsection, we will describe the number f^λ defined in (4.15) in terms of ***standard Young tableaux***. By definition, a standard Young tableau of ***shape*** $\lambda \vdash n$, or ***λ-tableau*** for short, is obtained by filling the numbers $1, 2, \ldots, n$ into the boxes of the Young diagram of λ in such a way that the numbers increase along rows (left to right) and along columns (top to bottom). Clearly, the $(1, 1)$-box must contain the number 1 and n must occur in some removable corner box. Removing this box, we obtain the Young diagram of a partition μ with $\mu \to \lambda$. Continuing in this manner, successively removing the boxes containing the highest number, we eventually end up with the tableau $\boxed{1}$. This process is easily seen to yield a bijection

$$\{\text{paths } \square \to \cdots \to \lambda \text{ in } \mathbb{Y}\} \xleftrightarrow{\sim} \{\lambda\text{-tableaux}\}. \tag{4.16}$$

Just as we have identified partitions with their Young diagrams in the description of the Young graph $\mathbb{Y}$, we will also oftentimes not distinguish between paths in $\mathbb{Y}$ and standard Young tableaux. In particular, we may rewrite the definition of f^λ in (4.15) as follows:

$$\boxed{f^\lambda = \#\{\lambda\text{-tableaux}\}.} \tag{4.17}$$

Example 4.5. Here are all standard Young tableaux of shape $\lambda = (2, 2, 1)$ along with the corresponding paths in $\mathbb{Y}$:

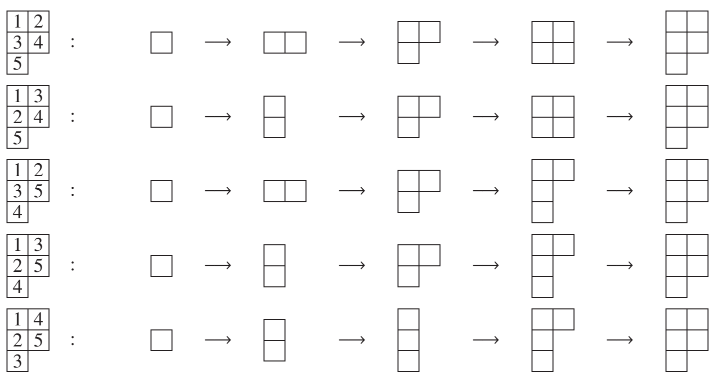

4.3.5. The Hook-Length Formula

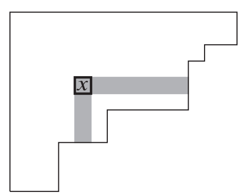

Let $\lambda \vdash n$ be given by its Young diagram. The ***hook*** at a box x of λ consists of the boxes occupying the gray area in the picture on the right. The ***hook length*** at x is defined by

$$h(x) \stackrel{\text{def}}{=} \#\{\text{boxes in the hook at } x\}.$$

Corner boxes or, equivalently, removable boxes are exactly the boxes such that $h(x) = 1$. The following formula is due to Frame, Robinson, and Thrall [**77**]; the probabilistic proof that we shall present below is due to Greene, Nijenhuis, and Wilf [**94**].

Hook-Length Formula. *For $\lambda \vdash n$, we have* $f^\lambda = \dfrac{n!}{\prod_{x \in \lambda} h(x)}$.

6	4	1
4	2	
3	1	
1		

For example, consider the partition $\lambda = (3, 2, 2, 1)$. Filling each box in the Young diagram of λ with the length of the hook at this box, we obtain the scheme on the left. Hence, $f^\lambda = \frac{8!}{6 \cdot 4^2 \cdot 3 \cdot 2 \cdot 1^3} = 70$.

Here is some brief background on probability and the description of the particular experiment that is used in the proof of the hook-length formula.

The Hook-Walk Experiment. Suppose that a certain experiment has a finite set Ω of possible outcomes; the set Ω is called the ***sample space*** and subsets $E \subseteq \Omega$ are referred to as ***events***. Assume further that, for each $\omega \in \Omega$, there is a probability value $p(\omega) \in \mathbb{R}_{\geq 0}$ such that $\sum_{\omega\in\Omega} p(\omega) = 1$. Then the probability of a given event $E \subseteq \Omega$ is defined by

$$P(E) \stackrel{\text{def}}{=} \sum_{\omega\in E} p(\omega)\,.$$

For example, if all outcomes $\omega \in \Omega$ have the same probability, then $P(E) = \frac{|E|}{|\Omega|}$.

The particular experiment that we will consider is, in probability parlance, a memoryless random walk or Markov chain: each step in the walk depends only on the current position and not on the sequence of steps that preceded it. Specifically, consider a partition $\lambda \vdash n$, identified with its Young diagram. To start with, choose a box, $x = x_0$, among the n boxes of λ with uniform probability $\frac{1}{n}$. If x is a corner box, then stop; otherwise, choose a different box, $x_1 \neq x$, in the hook at x with uniform probability $q(x) := \frac{1}{h(x)-1}$. If x_1 is a corner box, then stop; otherwise, choose a different box, $x_2 \neq x_1$, in the hook at x_1 with uniform probability $q(x_1)$, etc. Each step moves either down or right, and the walk will terminate at some corner box $x_t = c \in \lambda$ after finitely many steps:

$$\omega\colon\ x = x_0 \to x_1 \to \cdots \to x_{t-1} \to x_t = c\,.$$

We will refer to ω as a ***hook walk*** in λ. Our sample set Ω consists of all such hook walks; this is clearly a finite set. The probability of the hook walk ω is given by

$$(4.18)\qquad p(\omega) = \frac{1}{n}q(\omega) \quad \text{with} \quad q(\omega) := q(x_0)q(x_1)\cdots q(x_{t-1}).$$

It is not hard to see that $\sum_{\omega\in\Omega} p(\omega) = 1$. For each corner box c, consider the event $E_c = \{\text{all hook walks in } \lambda \text{ that end at } c\}$. These events form a partition of the sample set Ω. Therefore,

$$(4.19)\qquad 1 = \sum_{\omega\in\Omega} p(\omega) = \sum_c P(E_c) \underset{(4.18)}{=} \sum_c \frac{1}{n} \sum_{\omega\in E_c} q(\omega)\,,$$

where c runs over the corner boxes of λ.

Proof of the Hook-Length Formula. Our goal is to prove the formula $f^\lambda = \frac{n!}{H(\lambda)}$, where we have put

$$H(\lambda) := \prod_{x\in\lambda} h(x)\,.$$

We proceed by induction on n. The formula is trivially true for $n = 1$. To deal with $n > 1$, we use the following recursion, which is evident from the definition of f^λ

as the number of paths $\square \to \cdots \to \lambda$ in $\mathbb{Y}$:

$$f^\lambda = \sum_{\mu \to \lambda} f^\mu .$$

By induction, we know that $f^\mu = \frac{(n-1)!}{H(\mu)}$ for all μ in the above sum. Thus, we need to show that $\frac{n!}{H(\lambda)} = \sum_{\mu\to\lambda} \frac{(n-1)!}{H(\mu)}$ or, equivalently, $1 = \sum_{\mu\to\lambda} \frac{1}{n} \frac{H(\lambda)}{H(\mu)}$. Recall that $\mu \to \lambda$ means that μ arises from λ by removing one corner box, say c. Denoting the resulting μ by $\lambda \setminus c$, our goal is to show that

$$1 = \sum_c \frac{1}{n} \frac{H(\lambda)}{H(\lambda \setminus c)} ,$$

where c runs over the corner boxes of λ. Comparison with (4.19) shows that it suffices to prove the following equality, for each corner box c:

$$\sum_{\omega \in E_c} q(\omega) = \frac{H(\lambda)}{H(\lambda \setminus c)} . \tag{4.20}$$

First, let us consider the right-hand side of (4.20). Note that $\lambda \setminus c$ has the same hooks as λ, except that the hook at c, of length 1, is missing and the hooks at all boxes in the gray region marked $\mathbf{B} = \mathbf{B}(c)$ on the right have lengths shorter by 1 than the corresponding hooks of λ. Therefore, the right-hand side of (4.20) can be written as follows:

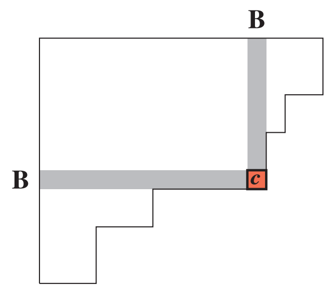

$$\frac{H(\lambda)}{H(\lambda \setminus c)} = \prod_{b \in \mathbf{B}} \frac{h(b)}{h(b) - 1} .$$

Using the notation $q(b) = \frac{1}{h(b)-1}$ from the hook-walk experiment, we can write the product on the right as $\prod_{b\in\mathbf{B}}(1 + q(b)) = \sum_{S\subseteq\mathbf{B}} \prod_{b\in S} q(b)$. Hence,

$$\frac{H(\lambda)}{H(\lambda \setminus c)} = \sum_{S \subseteq \mathbf{B}} \prod_{b \in S} q(b) . \tag{4.21}$$

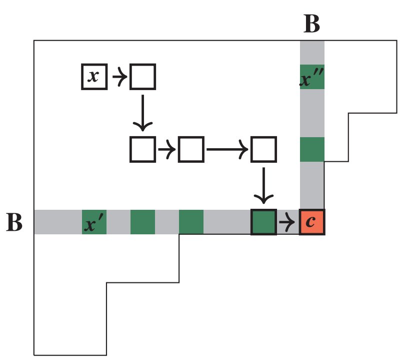

Now for the left-hand side of (4.20). For each hook walk $\omega \in E_c$, let $S_\omega \subseteq \mathbf{B}$ denote the set of boxes that arise as the horizontal and vertical projections of the boxes of ω into $\mathbf{B}$; see the green boxes on the left. Note that, while S_ω generally does not determine the entire walk ω, the starting point, x, is certainly determined, and if $x \in \mathbf{B}$, then ω is determined. We claim that, for each subset $S \subseteq \mathbf{B}$,

$$\sum_{\substack{\omega\in E_c\\ S_\omega=S}} q(\omega) = \prod_{b\in S} q(b)\,. \tag{4.22}$$

This will give the following expression for the left-hand side of (4.20):

$$\sum_{\omega\in E_c} q(\omega) = \sum_{S\subseteq \mathbf{B}} \sum_{\substack{\omega\in E_c\\ S_\omega=S}} q(\omega) = \sum_{S\subseteq \mathbf{B}} \prod_{b\in S} q(b)\,.$$

By (4.21) this is identical to the right-hand side of (4.20), thereby proving (4.20).

We still need to justify the claimed equality (4.22). For this, we argue by induction on $|S|$. The only hook walk $\omega \in E_c$ with $S_\omega = \varnothing$ is the walk starting and ending at c without ever moving; so the claim is trivially true for $|S| = 0$. The claim is also clear if the starting point, x, of ω belongs to $\mathbf{B}$, because then the sum on the left has only one term, which is equal to the right-hand side of (4.22). So assume that $x \notin \mathbf{B} \cup \{c\}$. Then there are two kinds of possible hook walks $\omega \in E_c$ with $S_\omega = S$: those that start with a move to the right and those that start with a move down. Letting η denote the remainder of the walk and letting $x', x'' \in \mathbf{B}$ be the vertical and horizontal projections of x into $\mathbf{B}$, we have $S_\eta = S \setminus \{x'\}$ in the former case and $S_\eta = S \setminus \{x''\}$ in the latter. Therefore,

$$\begin{aligned}
\sum_{\substack{\omega\in E_c\\ S_\omega=S}} q(\omega) &= q(x)\Bigg(\sum_{\substack{\eta\in E_c\\ S_\eta=S\setminus\{x'\}}} q(\eta) + \sum_{\substack{\eta\in E_c\\ S_\eta=S\setminus\{x''\}}} q(\eta)\Bigg)\\
&\underset{\text{by induction}}{=} q(x)\Bigg(\prod_{b\in S\setminus\{x'\}} q(b) + \prod_{b\in S\setminus\{x''\}} q(b)\Bigg)\\
&= q(x)\Bigg(\frac{1}{q(x')} + \frac{1}{q(x'')}\Bigg)\prod_{b\in S} q(b)\,.
\end{aligned}$$

To complete the proof of (4.22), it remains to observe that $q(x)\big(\frac{1}{q(x')} + \frac{1}{q(x'')}\big) = 1$ or, equivalently, $h(x) + 1 = h(x') + h(x'')$, which is indeed the case:

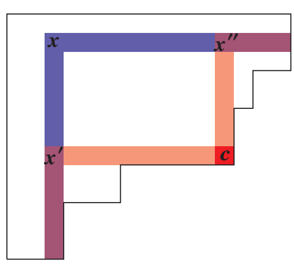

This finishes the proof of the hook-length formula. □

Exercises for Section 4.3

4.3.1 (Up and down operators)**.** Let $\mathbb{Z}\mathscr{P} = \bigoplus_n \mathbb{Z}\mathscr{P}_n$ denote the $\mathbb{Z}$-module of all formal $\mathbb{Z}$-linear combinations of partitions λ. Here, $\mathbb{Z}\mathscr{P}_n = 0$ for $n < 0$, because $\mathscr{P}_n$ is empty in this case, and $\mathbb{Z}\mathscr{P}_0 \cong \mathbb{Z}$, because $\mathscr{P}_0$ contains only $(0, 0, \dots)$, with Young diagram $\varnothing$. Thus, we have added $\varnothing$ as a root vertex to $\mathbb{Y}$ and a unique arrow $\varnothing \to \square$. Consider the operators $U, D \in \operatorname{End}_{\mathbb{Z}}(\mathbb{Z}\mathscr{P})$ that are defined by $U(\lambda) = \sum_{\lambda\to\mu} \mu$ and $D(\lambda) = \sum_{\mu\to\lambda} \mu$. Show that these operators satisfy the Weyl algebra relation $DU = UD + 1$.

4.3.2 (Automorphisms of $\mathbb{Y}$)**.** (a) Show that each $\lambda \in \mathscr{P}_n$ $(n > 2)$ is determined by the set $S(\lambda) := \{\mu \in \mathscr{P}_{n-1} \mid \mu \to \lambda \text{ in } \mathbb{Y}\}$.

(b) Conclude by induction on n that the graph $\mathbb{Y}$ has only two automorphisms: the identity and conjugation.

4.3.3 (Rectangle partitions)**.** Show:

(a) $f^{(n,n)} = \frac{1}{n+1}\binom{2n}{n}$, the n^{th} Catalan number.

(b) $f^{\lambda} > rc$ for $\lambda = (c^r) := (\underbrace{c, \dots, c}_{r})$ with $c, r \geq 2$ and $rc \geq 8$.

4.3.4 (Dimensions of the irreducible representations of $\mathcal{S}_6$)**.** Extend Figure 4.2 up to layer $\mathscr{P}_6$ and find the dimensions of all irreducible representations of $\mathcal{S}_6$ (assuming the Graph Isomorphism Theorem).

4.3.5 (Hook partitions and exterior powers of the standard representation)**.** Let $\lambda \mapsto V^{\lambda}$ be the bijection on vertices that is given by a graph isomorphism $\mathbb{Y} \xrightarrow{\sim} \mathbb{B}$ as in §4.3.2. Assume that the isomorphism has been chosen so that $\square\square \mapsto \mathbb{1}_{\mathcal{S}_2}$. Show that $V^{(n-k,1^k)} = \wedge^k V_{n-1}$ holds for all n and $k = 0, \dots, n-1$, where $(n-k, 1^k)$ is the "hook partition" $(n-k, \underbrace{1, \dots, 1}_{k})$.

4.3.6 (Irreducible representations of dimension $< n$)**.** Let $n \geq 7$. Assuming the Graph Isomorphism Theorem, show that $\mathbb{1}$, sgn, the standard representation V_{n-1}, and its sign twist $V_{n-1}^{\pm}$ are the only irreducible representations of dimension $< n$ of $\mathcal{S}_n$. (Use Exercise 4.3.3(b).)

4.4. Proof of the Graph Isomorphism Theorem

The main goal of this section is to provide the proof of the Graph Isomorphism Theorem. This will be accomplished in Corollary 4.13 after some technical tools have been deployed earlier in this section. In short, the strategy is to set up, for all n, a bijection

$$\Gamma_n : \quad \left\{ \begin{matrix} \text{paths } \square \to \cdots \text{ in } \mathbb{Y} \\ \text{with endpoint} \in \mathscr{P}_n \end{matrix} \right\} \xleftrightarrow{\sim} \left\{ \begin{matrix} \text{paths } \mathbb{1}_{\mathcal{S}_1} \to \cdots \text{ in } \mathbb{B} \\ \text{with endpoint} \in \operatorname{Irr} \mathcal{S}_n \end{matrix} \right\}$$