Discrete Mathematics lecture notes 10-2

November 19, 2013

28. Introduction to the symmetric group

In this (much belated; apologies) installment, we'll introduce the main topic of study for the remainder of the course. In fact, we've already met this character, but now we'll be delving deeper into its structure.

Definition 1. Recall that Σ_n is the set of self-bijections of $[n]$ (or more generally, Σ_X is the set of bijections $X \xrightarrow{\sim} X$). Σ_n has a *multiplication* or *product* $\Sigma_n \times \Sigma_n \to \Sigma_n$, written $\circ$ or simply by concatenation[1], given by composition of functions.

Lemma 2. *The product on Σ_n is associative, unital, and has inverses, but not commutative if $n > 2$:*

- *($\circ$ is associative) for all $\sigma, \tau, \chi \in \Sigma_n$, $\sigma \circ (\tau \circ \chi) = (\sigma \circ \tau) \circ \chi$,*
- *($\circ$ is unital) there is $\iota \in \Sigma_n$ such that for all $\sigma \in \Sigma_n$, $\sigma \circ \iota = \iota \circ \sigma = \sigma$,*
- *($\circ$ has inverses) for all $\sigma \in \Sigma_n$, there is $\xi \in \Sigma_n$ such that $\sigma \circ \xi = \xi \circ \sigma = \iota$,*
- *($\circ$ is not commutative) if $n > 2$, there exist $\sigma, \tau \in \Sigma_n$ such that $\sigma \circ \tau \neq \tau \circ \sigma$.*

Proof. Exercise. □

Exercise. Show the following:

- The unit of the multiplication is unique.
- For each $\sigma \in \Sigma_n$, the inverse of Σ is unique.
- If $\sigma \circ \tau = \sigma \circ \tau'$, then $\tau = \tau'$, and similarly if $\tau \circ \sigma = \tau' \circ \sigma$.

Notation 3. The unit of the multiplication is of course the identity map $\mathrm{id} = \mathrm{id}_{[n]}$, and the inverse of σ is commonly written σ^{-1}. Thus σ^{-1} is defined by the property $\sigma \circ \sigma^{-1}(i) = \sigma^{-1} \circ \sigma(i) = i$ for all $i \in [n]$.

Σ_n is called the symmetric *group* (of degree n, or on n letters), and now we're in the position to identify exactly what the italicized word means.

Definition 4. A *group* is a set G together with a *multiplication* (or *product*) $- \cdot - : G \times G \to G$. The multiplication is required to be associative, unital, and have inverse, but it is not require to be commutative. If the multiplication is commutative (so that $g \cdot h = h \cdot g$ for all $g, h \in G$), the group is called *abelian*.

In general, whenever the word "group" is used in a mathematical context, the implication is that there is some multiplicative structure satisfying the above properties, and that that structure is of at least incidental interest to us.

We'll need a good way of representing elements of Σ_n: There are two primary methods we'll use. First, as any $\sigma \in \Sigma_n$ is first and foremost a map $[n] \to [n]$, we could use the row expression of σ to indicate directly where σ sends each element of $[n]$. Thus

$$\sigma = \begin{bmatrix} 1 & 2 & 3 & 4 & 5 \\ 4 & 2 & 1 & 5 & 3 \end{bmatrix} = \begin{bmatrix} 4 & 2 & 1 & 5 & 3 \end{bmatrix}$$

is shorthand for $\sigma(1) = 4$, $\sigma(2) = 2$, $\sigma(3) = 1$, $\sigma(4) = 5$, $\sigma(5) = 3$. In the first array, we list all the elements of $[n]$ in order, and then the image of each element under σ beneath it. The second array has the same data, but the first row (which never changes) is omitted.

[1] That is, given elements $\sigma, \tau \in \Sigma_n$, we will write the product as either $\sigma \circ \tau$ or $\sigma\tau$, depending on our mood.

A variant of this is to list both the domain and range of σ as rows, both in order, but connect each element on the top row to its image under σ on the bottom. The same permutation is then written

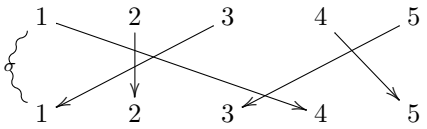

The advantage of this variant is that, while it is not so compact, it is easier to compute the product of two permutations. For instance, if $\tau \in \Sigma_5$ has row expression $\tau = \begin{bmatrix} 3 & 4 & 5 & 2 & 1 \end{bmatrix}$, we can compute $\tau \circ \sigma$ by simply placing the arrow presentation of σ on top of that of τ,[2] then tracing each path out:

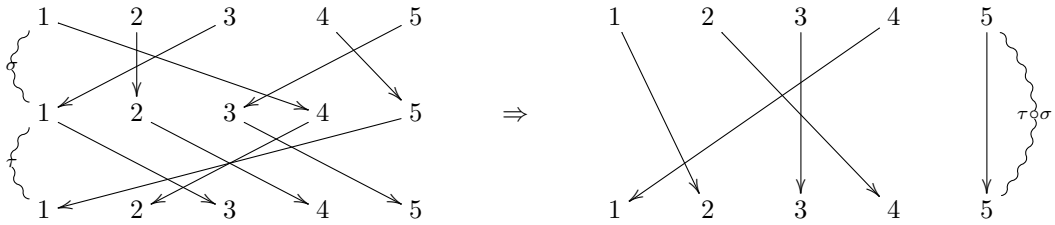

Of course, doing calculations with this presentation isn't really an option.[3] We need a method of writing permutations that is both compact and lends itself easily to computing the product.

Fix some $i \in \Sigma_n$, and consider the following sequence of elements of $[n]$:

$$1, \quad \sigma(1), \quad \sigma^2(1) := (\sigma \circ \sigma)(1), \quad \sigma^3(1), \quad \dots$$

By the pigeonhole principal, there must be some minimal i such that $\sigma^i(1) = 1$.[4] We call the ordered list $(1, \sigma(1), \sigma^2(1), \dots, \sigma^{i-1}(1))$ a *cycle*, and understand it to mean the permutation of $[n]$ sending 1 to $\sigma(1)$, $\sigma(1)$ to $\sigma^2(1)$,..., and $\sigma^{i-1}(1)$ to 1. Of course, there may be elements of $[n]$ that do not appear in our list; these must be accounted for in their own cycles.

In the above examples, we have $\sigma = (1\ 4\ 5\ 3)(2)$ and $\tau = (1\ 3\ 5)(2\ 4)$. This means that if we feed in a number $i \in [n]$, we read *from right to left*, looking for the cycle that contains i, then send i to the following term within that cycle. If it happens that i occurs at the end of the cycle, we send it to the term that begins the cycle. Thus $\tau(1) = 3$ and $\tau(5) = 1$.

We adopt the convention that if a cycle ever consists of a single element, we will drop it and understand it to be fixed by the permutation. Thus we could also write $\sigma = (1\ 4\ 5\ 3)$, with the implicit understanding that $\sigma(2) = 2$.

To define composition, concatenate: $\tau \circ \sigma = (1\ 3\ 5)(2\ 4)(1\ 4\ 5\ 3)$. Note that this will in general yield an expression in which a given element appears in more than one cycle; the cycles are not *disjoint*. To correct this, we can reexpress the composition in terms of disjoint cycles by feeding an element (say 1) into the rightmost cycle (getting 4), feeding that into the next rightmost cycle (getting 2), feeding *that* into the next rightmost cycle (getting 4 again, because 4 does not appear in $(1\ 3\ 5)$), and continuing until all cycles are exhausted. The term you end with is where i is sent; place that number next to i in the composite cycle.

In our example, we have $\tau \circ \sigma = (1\ 3\ 5)(2\ 4)(1\ 4\ 5\ 3) = (1\ 2\ 4)$.

Proposition 5. *Any $\sigma \in \Sigma_n$ can be written as a concatenation of disjoint cycles. This expression is unique up to the ordering of the cycles, and cyclic shifts within the cycle.*

Proof. Exercise. The uniqueness claims reflect the fact that that *disjoint* cycles will always commute (so $\tau = (1\ 3\ 5)(2\ 4) = (2\ 4)(1\ 3\ 5)$), and up to cyclic shifting each cycle is the same no matter where we choose to start (so $(1\ 4\ 5\ 3) = (4\ 5\ 3\ 1) = (5\ 3\ 1\ 4) = (3\ 1\ 4\ 5)$). □

[2]Note that our convention that we read composition of maps right to left is what determines which sequence of arrows goes on top. In other words, since $\tau \circ \sigma$ means "first do σ, then do τ to the result," we have to place the presentation of σ above that of τ in order for this tracing procedure to accurately reflect the composition.

[3]Try to prove that the composition in Σ_{10} is associative by recording generic elements in this manner.

[4]Why?