Math 821, Spring 2013, Lecture 1

Karen Yeats
(Scribe: Iain Crump)

January 15, 2013

Generating Functions

Definition. A *combinatorial class* $\mathcal{C}$ is a countable set with a size function $|\cdot| : \mathcal{C} \to \mathbb{Z}_{\geq 0}$ with the property that $\mathcal{C}_n = \{a \in \mathcal{C} : |a| = n\}$ is finite. We will write $c_n = |\mathcal{C}_n|$.

Definition. Let $\mathcal{C}$ be a combinatorial class. The *(ordinary) generating function* is $c(x) = \sum_{c \in \mathcal{C}} x^{|c|}$.

Proposition 1. *Let $\mathcal{C}$ be a combinatorial class. Then, $c(x) = \sum_{n=0}^{\infty} c_n x^n$.*

Proof. $c(x) = \sum_{c \in \mathcal{C}} x^{|c|} = \sum_{n=0}^{\infty} \sum_{c \in \mathcal{C}_n} x^n = \sum_{n=0}^{\infty} c_n x^n$ $\square$

Example. Consider binary rooted trees with distinct left and right children at each vertex. Lets say the size of a tree is the number of vertices. Call this class $\mathcal{T}$.

The generating function of this class is

$$T(x) = \sum_{t \in \mathcal{T}} x^{|t|} = 1 + x + 2x^2 + \cdots .$$

To build a generating function for this class, note that every non-empty tree has the form shown in Figure 1. Then,

$$\begin{aligned} T(x) &= \sum_{t \in \mathcal{T}} x^{|t|} \\ &= 1 + \sum_{t \in \mathcal{T} - \epsilon} x^{1+|L|+|R|} \\ &= 1 + x \left(\sum_{L \in \mathcal{T}} x^{|L|} \right) \left(\sum_{R \in \mathcal{T}} x^{|R|} \right) \\ &= 1 + xT(x)^2. \end{aligned}$$

Strictly speaking, what we are using here a bijection between $\mathcal{T} - \epsilon$ and $\mathcal{T} \times \mathcal{T}$. As $T(x) = 1 + T(x)^2$, we may solve using the quadratic formula, $T(x) = \frac{1 \pm \sqrt{1-4x}}{2x}$. Checking the positive root, $\frac{1+1+O(x)}{2x} = \frac{1}{x} + O(1)$. The $\frac{1}{x}$

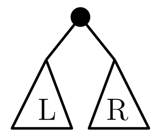

Figure 1

term immediately indicates that this is incorrect. Hence, we must take the negative root; $\frac{1-1+O(x)}{2x} = O(1)$. We choose our sign based on which has the correct behaviour near zero.

Definition. If $A(x)$ is a formal power series $A(x) = \sum a_n x^n$, then $[x^n]A(x) = a_n$.

Proposition 2. *$[x^n](1+x)^r = \binom{r}{n} = \frac{r(r-1)(r-2)\cdots(r-n+1)}{n!}$ for all real r and non-negative n.*

Example (Continued). From before, $[x^n]T(x) = [x^n]\left(\frac{1-\sqrt{1-4x}}{2x}\right)$. Then,

$$\begin{aligned}
[x^n](1-4x)^{1/2} &= \binom{1/2}{n}(-4)^n \\
&= (-4)^n \cdot \frac{\frac{1}{2}\left(\frac{1}{2}-1\right)\cdots\left(\frac{1}{2}-n+1\right)}{n!} \\
&= (-1)^n 2^n \cdot \frac{1(1-2)\cdots(1-2n+2)}{n!} \\
&= (-1)2^n \cdot \frac{(1)(3)\cdots(2n-3)}{n!} \cdot \frac{(2)(4)\cdots(2n-4)(2n-2)}{2^{n-1}(n-1)!} \\
&= (-1)\cdot\frac{(2n-2)!}{n!(n-1)!} \\
&= -\frac{2}{n}\binom{2n-2}{n-1}.
\end{aligned}$$

It follows that

$$[x^n]T(x) = -\frac{1}{2}[x^{n+1}](1-4x)^{1/2} = \frac{1}{n+1}\binom{2n}{n}.$$

This is the n^{th} Catalan number.

Definition. Let $\mathcal{A}$ and $\mathcal{B}$ be combinatorial classes. We say that $\mathcal{A}$ and $\mathcal{B}$ are *combinatorially equivalent* if $a_n = b_n$ for all $n \geq 0$.

Example. Consider paths beginning at $(0, 0)$, ending at $(n, 0)$, never strictly below the x-axis, and built of $(1, 1)-$ and $(1, -1)-$steps. We call these *Dyck paths*. The size of a Dyck path is the number of steps taken. An example of such a path is shown in Figure 2. Let $\mathcal{D}$ denote this combinatorial class.

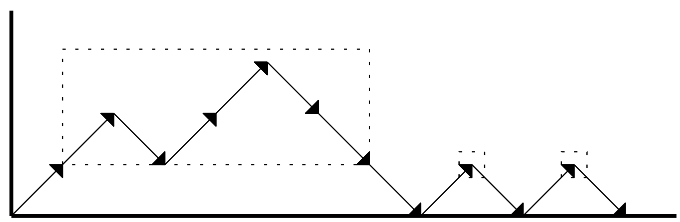

Figure 2

A Dyck path is either empty or it is built of some number of blocks of the form shown in Figure 3. These blocks are highlighted in the previous figure. Then,

$$\begin{aligned} D(x) &= 1 + x^2 D(x) + x^4 D(x)^2 + \cdots \\ &= \sum_{k=0}^{\infty} x^{2k} D(x)^k \\ &= \frac{1}{1 - x^2 D(x)}. \end{aligned}$$

So, $D(x) - x^2 D(x)^2 = 1$, or $1 - D(x) + x^2 D(x)^2 = 0$.

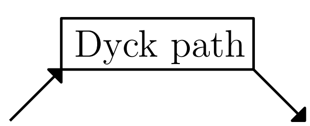

Figure 3

To compare with our previous example, recall $1 - T(x) + xT(x)^2 = 0$ for binary rooted trees. Hence, these are the same up to how we define size. That is, $[x^{2k}]D(x) = [x^k]T(x)$.

Another decomposition uses the first return to the x-axis, and writes a non-empty Dyck path as in Figure 4. Then, $D(x) = 1 + x^2 D(x)^2$, the same equation.

Lemma 3. *Let A and B be sets and $f : A \to B$, $g : B \to A$. If $f \circ g = id_B$ and $g \circ f = id_A$, then f and g are both bijections.*

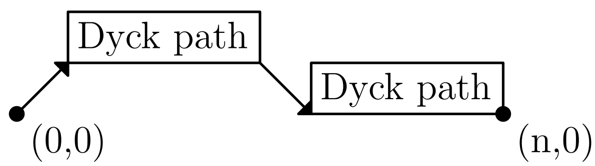

Figure 4

Proof. By symmetry it suffices to show f is a bijection. Suppose $f(a_1) = f(a_2)$ for $a_1, a_2 \in A$. Then, $a_1 = g(f(a_1)) = g(f(a_2)) = a_2$, and hence f is one-to-one. Now, take $b \in B$. As $f(g(b)) = b$, f is onto. Hence, f is a bijection. $\square$

Example. We can use Lemma 3 to give a bijection between $\mathcal{D}_{2n}$ and $\mathcal{T}_n$ for all $n \geq 0$.

[**The following was inserted after class by Dr. Yeats.**]

In class today we tried to write down a bijection between binary rooted trees and Dyck paths in an explicit nonrecursive way, but we accidentally wrote down a different bijection, one between rooted (not necessarily binary) trees counted by edges and Dyck paths.

Let's use the same basic idea but actually get rooted trees.

Let $\mathcal{D}$ be the class of Dyck paths counted by length. Let $\mathcal{T}$ be the class of binary rooted trees (with distinct left and right children at each vertex) counted by number of vertices. Then from the generating functions we know $|\mathcal{D}_{2n}| = |\mathcal{T}_n|$.

Let

$$g : \mathcal{T}_n \to \mathcal{D}_{2n}$$

be defined as follows. If $n = 0$ then $g(\epsilon) = \epsilon$. Assume $n > 0$. Take $t \in \mathcal{T}_n$. Since left and right are defined at each vertex, t has a canonical embedding in the plane. Fatten t by a sufficiently small δ and remove it from the plane. Now we want to traverse the boundary of this hole in a counterclockwise manner. Begin the traversal at the top above the root.

In this traversal we meet each vertex at most 3 times. We need some vocabulary for these meetings. For a non-root vertex v with two nonempty children, it will be met *on the left* as we approach from its parent, *on the bottom* between its two children, and *on the right* as we return to its parent. If the left child is empty, then the meeting on the left and on the bottom are joined. Likewise if the right child is empty, and if the vertex is a leaf then all the meetings are jointed into a single meeting.

Now build a path as follows: each time we meet a vertex on the left (including our initial beginning at the root and whether or not this is joined with another meeting) take a $(1, 1)$-step. Each time we meet a vertex on the bottom (whether or not this is joined with another meeting) take a $(1, -1)$-step. When a left and bottom meeting are joint put in the $(1, 1)$-step first and then the $(1, -1)$-step. In particular, when we meet a leaf we take a $(1, 1)$-step immediately followed by a $(1, -1)$ step.

This path is $g(t)$. It is a Dyck path because we meet each vertex on the left before we can meet it on the bottom; it is length $2n$ since each vertex contributes 2 steps.

Let

$$f : \mathcal{D}_{2n} \to \mathcal{T}_n$$

be defined as follows. If $n = 0$ then $f(\epsilon) = \epsilon$. Assume $n > 0$. View the Dyck path as a sequence w of steps, from left to right. Build a tree as follows. Read w from left to right. On the first $(1, 1)$-step build the root and call it the current vertex. Each time we meet a $(1, 1)$-step

- build a new vertex w
- if the previous step was $(1, 1)$ then make w the right child for the current vertex
- if the previous step was $(1, -1)$ then make w the left child for the current vertex
- call w the current vertex.

If we meet a $(1, -1)$-step, then

- if the current vertex is not marked, then mark this vertex and this vertex remains the current vertex.
- if the current vertex is marked, then find the first ancestor v of the current vertex which has no left child. Mark v and set the current vertex to be v.

Now lets check this builds a binary rooted tree.

After each $(1, 1)$-step the current vertex is a leaf, so if we are on a $(1, 1)$-step and the previous step was a $(1, 1)$-step then the current vertex does not have a right child, and so the construction in this case is valid. After a $(1, -1)$-steps (provided that step was valid), the current vertex has no left child, so if we are on a $(1,1)$-step after a $(1,-1)$-step then the current vertex does not have a left child so the construction in this case is valid.

For a vertex v, let $m(v)$ be the number of unmarked ancestors of v with no left child, including v itself.

After a given step, let c be the new current vertex and d the previous one. Then if the step was a $(1,1)$ step we have $m(c) = m(d) + 1$ and if the step was a $(1,-1)$ step we have $m(c) = m(d) - 1$. Thus by the Dyck path property at every $(1,-1)$-step there is at least one unmarked ancestor of the current vertex with no left child (including the current vertex itself), and so the construction is valid in this case also.

A nice picture probably makes this much clearer:

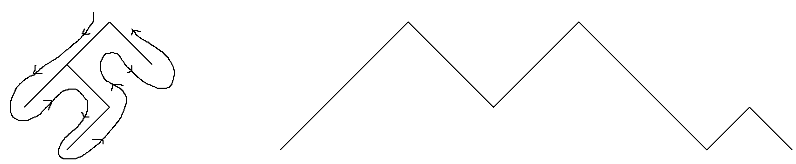

Try going through both maps, step by step, on this example to really see it.

Now we need to check $f \circ g = \mathrm{id}$ and $g \circ f = \mathrm{id}$. The key idea in both directions is that building a vertex in the construction for f corresponds to meeting the vertex on the left in the traversal for g, while marking a vertex in the construction for f corresponds to meeting the vertex on the bottom in the traversal for g. Since all of this is something of a digression anyway, I'll leave the details to you.

[This ends the newly inserted material.]

Another construction; define a function $h : \mathcal{D}_{2n} \to \mathcal{T}_n$ by $h(\epsilon) = \epsilon$, and for non-empty Dyck paths as in Figure 5.

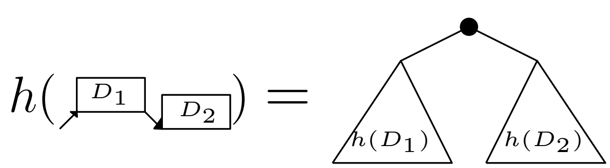

Figure 5

Similarly, define a function $j : \mathcal{T}_n \to \mathcal{D}_{2n}$ by $j(\epsilon) = \epsilon$, and for non-empty trees as in Figure 6.

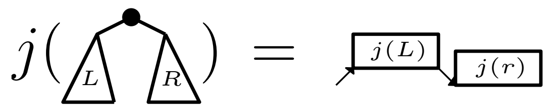

Figure 6

Some Combinatorial Constructions

Definition. Let Φ be a map which takes any collection of m combinatorial classes $\mathcal{B}^{(1)}, \mathcal{B}^{(2)}, ..., \mathcal{B}^{(m)}$ to a combinatorial class $\mathcal{A} = \Phi(\mathcal{B}^{(1)}, \mathcal{B}^{(2)}, ..., \mathcal{B}^{(m)})$. We say that Φ is *admissible* if the sequence $(a_n)_{n=0}^{\infty}$ depends only on the sequence $(b_n^{(1)})_{n=0}^{\infty}, ..., (b_n^{(m)})_{n=0}^{\infty}$.

So, this is a map on isomorphism classes under combinatorial isomorphism, and further there is (at least in principle) a map of generating functions $B^{(1)}(x), B^{(2)}(x), ..., B^{(m)}(x) \mapsto A(x)$.

Definition. Let $\mathcal{B}$ and $\mathcal{C}$ be combinatorial classes, and define $\mathcal{A} = \mathcal{B} \times \mathcal{C}$ to be the combinatorial class with underlying set $\mathcal{B} \times \mathcal{C} = \{(b, c) : b \in \mathcal{B}, c \in \mathcal{C}\}$ and size function $|(b, c)|_{\mathcal{A}} = |b|_{\mathcal{B}} + |c|_{\mathcal{C}}$.

Proposition 4. *Let $\mathcal{B}$ and $\mathcal{C}$ be combinatorial classes and let $A = \mathcal{B} \times \mathcal{C}$. Then, $A(x) = B(x)C(x)$.*

Proof. First,

$$\begin{aligned}\mathcal{A}_n &= \{(b, c) : b \in \mathcal{B}, c \in \mathcal{C}, |b| + |c| = n\} \\ &= \bigcup_{k=0}^{n} \{(b, c) : b \in \mathcal{B}_k, c \in \mathcal{C}_{n-k}\} \\ &= \bigcup_{k=0}^{n} \mathcal{B}_k \times \mathcal{C}_{n-k}.\end{aligned}$$

This union is disjoint, as the different terms have different sizes in the first coordinate. So, $a_n = \sum_{k=0}^{n} b_k c_{n-k}$ and

$$A(x) = \sum_{n=0}^{\infty} \left(\sum_{k=0}^{n} b_k c_{n-k} \right) x^n = B(x)C(x).$$

$\square$

Note that the function $\times$ is admissible.

Notation. Use $\mathcal{E}$ for a combinatorial class with one element of size zero.
Use $\mathcal{Z}$ for a combinatorial class with one element of size one.
Use subscripts to distinguish between copies of $\mathcal{E}$ and $\mathcal{Z}$.

Note that $\mathcal{E} \times \mathcal{A}$ is combinatorially isomorphic to $\mathcal{A}$.

Definition. Let $\mathcal{B}$ and $\mathcal{C}$ be combinatorial classes, and define $\mathcal{A} = \mathcal{B} + \mathcal{C}$ to be the combinatorial class with underlying set $\mathcal{B} + \mathcal{C} = (\mathcal{E}_1 \times \mathcal{B}) \cup (\mathcal{E}_2 \times \mathcal{C})$ and size function

$$|a|_{\mathcal{A}} = \begin{cases} |b|_{\mathcal{B}}, & \text{if } a = (\mathcal{E}_1, b) \\ |c|_{\mathcal{C}}, & \text{if } a = (\mathcal{E}_2, c) \end{cases}.$$

Note that the function $+$ is admissible.

Remark 1. The purpose of the union is to force disjointness. If we had chosen $\mathcal{A} = \mathcal{B} \cup \mathcal{C}$, then

$$\begin{aligned} a_n = |\mathcal{A}_n| &= |\mathcal{B}_n| + |\mathcal{C}_n| - |\mathcal{B}_n \cap \mathcal{C}_n| \\ &= b_n + c_n - |\mathcal{B}_n \cap \mathcal{C}_n|, \end{aligned}$$

which is not a function of b_i and c_i terms. Hence, this is definition is not admissible.

Proposition 5. *Let $\mathcal{B}$ and $\mathcal{C}$ be combinatorial classes and let $\mathcal{A} = \mathcal{B} + \mathcal{C}$. Then, $A(x) = B(x) + C(x)$.*

Proof. As $(\mathcal{E}_1 \times \mathcal{B}) \cap (\mathcal{E}_2 \times \mathcal{C}) = \emptyset$,

$$|\mathcal{A}_n| = |(\mathcal{E}_1 \times \mathcal{B})_n| + |(\mathcal{E}_2 \times \mathcal{C})_n| = |\mathcal{B}_n| + |\mathcal{C}_n|.$$

Then, $a_n = b_n + c_n$ and hence $A(x) = B(x) + C(x)$. $\square$

References: Flajolet and Sedgewick, *Analytic combinatorics*, Cambridge (2009). I1-I2.