MPRI course 2.10 Algorithmic aspects of combinatorics

Exam 1 — Friday 29 November 2024 — duration: 2,5 hours

Written by J. Bouttier and G. Chapuy

Please write your name on every sheet of paper you return to us. You can answer questions in either English or French.

Exercise 1 Peaks in m-ary trees

For $m \geq 2$, an m-ary tree is a rooted plane tree having only vertices of arity zero (*leaves*) or arity m (*nodes*). The *size* of a tree is its number of nodes. A node is a *peak* if its leftmost child is a leaf.

1. Explain why it is natural to use the word *peak*.

In the Lukaciewicz path associated to the tree, this corresponds to an up-step (of height $m-1$) followed by a down-step, hence, a peak.

2. Let $T(z,u)$ be the generating function of m-ary trees, where z marks the size, and u the number of peaks. Write an equation for $T(z,u)$. Check that for $u = 1$ you recover the usual equation for m-ary trees.

We start from the equation $T = 1 + zT^m$ of m-ary tree, that we modify to take into account the case where the root is a peak. This case corresponds to a contribution of $z \times 1 \times T^{m-1}$ and should receive an additional weight u. We get $T(z,u) = 1 + zuT^{m-1} + z(T^m - T^{m-1})$.

3. Express the coefficient of z^n in $T(z,u)$ as a coefficient extraction using a tool seen in class.

We use the Lagrange inversion formula. For this we need the function to have no constant term so we introduce $S = T - 1 = z((u-1)(1+S)^{m-1} + (1+S)^m) = z\phi(S)$ with $\phi(y) = (1+y)^{m-1}(u+y)$. By the FIL we have, for $n \geq 1$, $[z^n]T = [z^n]S = \frac{1}{n}[y^{n-1}](1+y)^{nm-n}(u+y)^n$

4. Extract the coefficient of u^k in the previous formula, and finish the computation, to obtain an explicit formula for the number of m-ary trees of size n having k peaks. What formula do you recognize for $m = 2$?

As in class we use the fact that $[u^k][z^n]\cdot = [z^n][u^k]\cdot$. We have $[u^k][z^n]T = [u^k][z^n]S = \frac{1}{n}[y^{n-1}](1+y)^{nm-n}[u^k](u+y))^n = \frac{1}{n}\binom{n}{k}[y^{n-1}]y^{n-k}(1+y)^{nm-n} = \frac{1}{n}\binom{n}{k}[y^{k-1}](1+y)^{nm-n} = \frac{1}{n}\binom{n}{k}\binom{nm-n}{k-1}$. For $m = 2$ we recognize the Narayana number seen in class, that counts Dyck paths by the number of peaks (which makes sense, given question 1 above).

5. (bonus, will be marked only if you did a substantial amount of other exercises) Give a bijective proof of this result.

One easily adapts the proof given in the corrigé of the TD for $m = 2$.

Exercise 2 Generalized Temperley bijection

(Adapted from the paper *Trees and perfect matchings* by R. Kenyon, J. Propp and D. Wilson.)

Let G be a connected plane graph, which is unoriented for now. We denote by V, E, F its sets of vertices, faces and faces, respectively. Let G^* denote the dual of G.

1. Given a spanning tree T of G, recall how its dual tree T^* is defined. (By the way, why does G always admit a spanning tree?)

Generally speaking, the dual of a spanning subgraph S of G (defined by a subset of E) is by definition the spanning subgraph of G^* consisting of the duals of the edges *not* in S. The dual of a connected graph is acyclic, and vice versa. In particular, the dual T^* of a spanning tree T is also a tree. Any connected graph G admits at least one spanning tree, as a spanning tree is a connected spanning subgraph having a minimal number of edges.

2. What is the relation between the numbers of vertices and edges of a tree?

A tree has one more vertex than edge.

3. By applying the result of the previous question to both T and T^*, deduce Euler's relation

$$|V| - |E| + |F| = 2.$$

As T is a spanning tree of G, its number of edges is $|V| - 1$. Similarly, as T^* is a spanning tree of G^*, its number of edges is $|F| - 1$. Now, each edge of G is either in T, (x)or its dual is in T^*. Thus, the numbers of edges of T and T^* sum up to $|E|$, giving Euler's relation.

Now, let G' denote the plane graph obtained by superimposing G and G^*. Since the edges of G and their duals cross each other, we create a new (tetravalent) vertex at their intersection. Observe that the vertex set of G' is in bijection with $V \cup E \cup F$.

4. Illustrate this construction by drawing the G' associated with the following G:

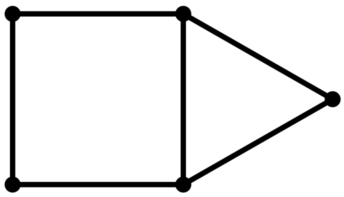

On the left of the figure below is the graph G', where the dashed edges connect to the vertex dual to the outer face of G, which we place "at infinity":

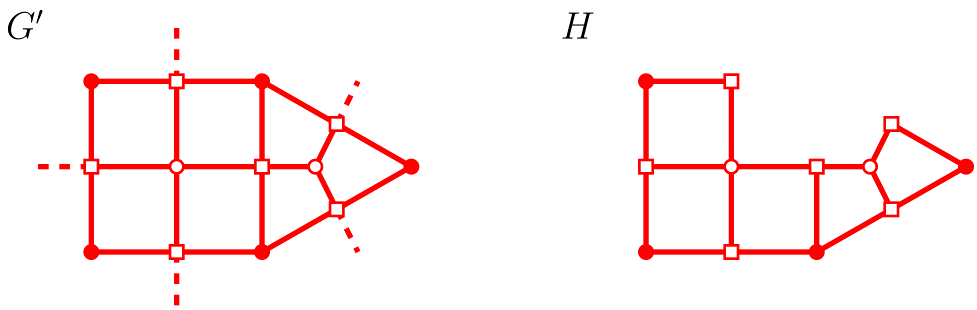

5. Returning to the general case, explain why G' does not admit any dimer configuration.

Observe that G' is a bipartite: viewing its vertex set as $V \cup E \cup F$, every edge of G' connects an element of E to an element of $V \cup F$. But, by Euler's relation, the cardinality of these two sets differ by two. Thus, there cannot be any dimer configuration, since each dimer covers exactly one element of both sets.

Let v (resp. f) be a vertex (resp. face) of G, such that v and f are incident to each other in G. Let v^* be the vertex of G^* corresponding to f. Let H be the plane graph obtained from G' by removing v, v^* and all their incident edges in G'.

6. Update the illustration made at question 4 by drawing H, when f is the outer face of the displayed graph, and v the vertex at the top-right corner of the square.

See the right of the figure above.

Let us denote by $\mathcal{T}$ the set of spanning trees of G, and by $\mathcal{D}$ the set of dimer configurations on H. Our purpose is to show that there is a bijection between $\mathcal{T}$ and $\mathcal{D}$. Let us consider a spanning tree T of G, and let T^* be its dual tree. We orient all the edges of T (resp. T^*) towards v (resp. v^*). To each edge e of T or T^*, we associate an edge e' of G' in this way:

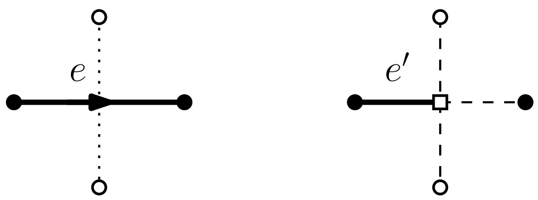

Namely, e' is obtained by "shortening" e by a half: it remains incident to the origin of e, but its other endpoint is now the intersection of e with its dual edge. Let D denote the set of all edges e' obtained in this way.

7. Illustrate this construction by drawing your favorite spanning tree of the graph of question 4, and then by drawing the corresponding dimer configuration.

Here is a spanning tree T (shown with the edge orientations) of the G above, and the corresponding dimer configuration D on G'.

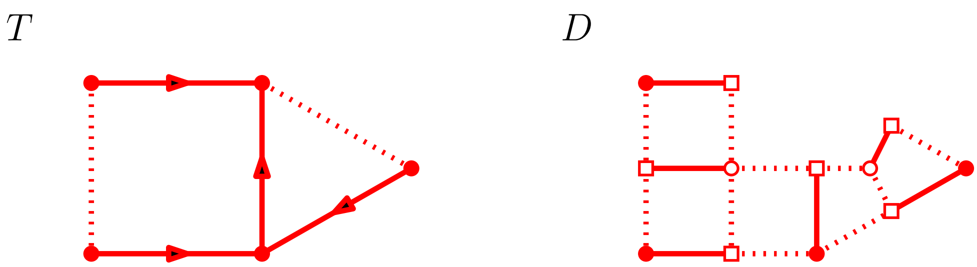

8. Show that, in general, D is a dimer configuration on H. Thus, $T \mapsto D$ is a mapping from $\mathcal{T}$ to $\mathcal{D}$. Show that it is injective.

By the way the edges of T and T' are oriented, every vertex of G or of G^* other than v and v^* has exactly one outgoing edge e. Thus, by the construction, it is incident to exactly one edge e' in D (and v and v^* are incident to no such edge). Every edge of G is either in T, (x)or its dual is in T^*. Thus, again by the construction, its corresponding vertex in G' is incident to exactly one edge in D. Hence, D is a dimer configuration on H. The mapping is injective since, by reversing the construction rule, we recover T from D.

9. To prove that the mapping is surjective, let us now consider a dimer configuration D on H. By applying the above construction backwards on each edge e' in D, we get a collection of oriented edges belonging either to G or to G^*. It is a priori not obvious that these form spanning trees: show that it is nevertheless the case. (For acyclicity, reason by contradiction, applying Euler's relation to an appropriate graph.)

As said in the question, we start from a dimer configuration D on H, perform the construction backwards to get a collection of oriented edges belonging either to G or to G^*. They form two spanning subgraphs of respectively G and G^*, which we denote by T and T^* respectively (anticipating on the fact that they are trees). Since D is a dimer configuration, every vertex of G (resp. G^*) other than v (resp. v^*) has exactly one outgoing edge in T (resp. T^*). It thus suffices to show that T is acyclic, since this will imply that it is connected by the previous remark, and that T^* is a tree by duality.

Suppose by contradiction that T contains a cycle C, which is necessarily oriented. Without loss of generality, we may assume f to be the outer face of G (perform an inversion of the plane otherwise). By the Jordan curve theorem, C divides the plane into two regions, one of them bounded which we call the interior region. Let us denote by $\tilde{V}, \tilde{E}, \tilde{F}$ the sets of vertices, edges and faces of G contained in the closure of interior region. Note that v does not belong to $\tilde{V}$, since v is assumed incident to the outer face f in G, and it cannot belong to C as it has no outgoing edge. The contradiction is then the following:

- on the one hand, the elements of $\tilde{V} \cup \tilde{E} \cup \tilde{F}$ correspond to vertices of H which are necessarily paired together in the dimer configuration D, and thus $|\tilde{V}| + |\tilde{E}| + |\tilde{F}|$ is necessarily even,
- on the other hand, $\tilde{V}$ and $\tilde{E}$ form the vertex and edge sets of a plane graph $\tilde{G}$, which is a subgraph of G. Its set of faces is $\tilde{F} \cup \{\tilde{f}\}$, where $\tilde{f}$ is the outer face (i.e. the exterior region delimited by C). By Euler's relation we have $|\tilde{V}| - |\tilde{E}| + |\tilde{F}| + 1 = 2$ which, taken modulo 2, contradicts the previous item.

Hence T is acyclic, as wanted.

10. How does does the above argument fail when the vertex v and the face f of G chosen above to construct H are not incident to each other? Find a counterexample showing that the mapping $T \mapsto D$ is no longer surjective.

When v and f are not incident to each other, it is possible to have dimer configurations corresponding to spanning subgraphs of G containing cycles which “separate” v from f. Here is an example of such situation:

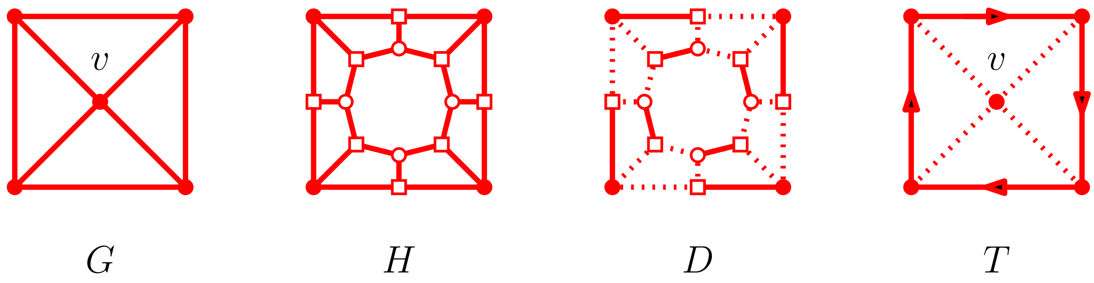

Namely, in the graph G on the left, we choose v to be the central vertex and f to be the outer face. We then display the corresponding H, and a certain dimer configuration D on H. By applying the construction rule backwards, we obtain an oriented subgraph T of G containing a cycle separating v from f. For completeness, let us make the following general remarks:

- T still has the property that every vertex other than v has exactly one outgoing edge,
- there may exist several separating cycles nested into one another,
- while we do not display it on the figure, there is also a dual subgraph T^* which has as many cycles as T,
- the cycles in T and T^* can be oriented in two possible ways, and this information is needed to reconstruct the dimer configuration. If we just keep the graph T and forget about the orientations, then it corresponds to 4^c dimer configurations, where c is the number of cycles of c (two orientations per cycle of T and two orientations per dual cycle of T^*).
- such spanning subgraphs are sometimes called "cycle-rooted spanning forest" (CRSF). They can be enumerated by a variant of the matrix-tree theorem involving "monodromies". We leave this as food for thought to the reader.

11. Let us return to the situation where v and f are incident, so that the mapping $T \mapsto D$ is a bijection between $\mathcal{T}$ and $\mathcal{D}$. To each edge e' of H, we assign a weight $w_{e'}$ with the constraint that $w_{e'} = 1$ for those edges which are incident to a vertex belonging to G^* (the weight remains arbitrary for those incident to a vertex belonging to G). Write the dimer partition function

$$Z_{H,w} = \sum_{D \in \mathcal{D}} \prod_{e' \in D} w_{e'}$$

as a determinant of a matrix which you will explicit, using the oriented version of the matrix-tree theorem.

Probably this is best explained on our running example. On the left of figure below, we display the graph H together with the edge weights $w_1, \ldots, w_9$ which are arbitrary. When no weight is displayed, the corresponding edge has weight 1 (these are precisely the edges incident to vertices of G^*, here shown as circles).

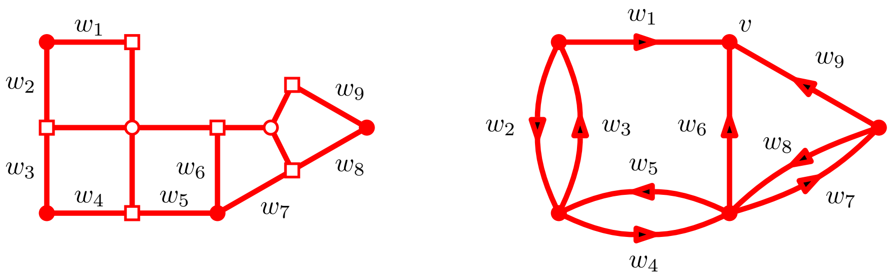

By the bijection between dimer configurations on H and trees on G, computing the dimer partition function $Z_{H,w}$ amounts to counting spanning trees of H, with a weight depending on the orientation of their edges. This may be done by "converting" G into an oriented graph, as displayed on the right of the figure. We basically duplicate each edge of G into a pair of edges with opposite orientations, and different weights. Edges of G incident to v need not be duplicated, they are just oriented towards v. The dimer partition function is equal to the generating function of spanning trees on this graph oriented towards v. In turn, this generating function can be computed by the weighted oriented version of the matrix tree theorem. On our running example, the Laplacian matrix reads

$$\begin{pmatrix} w_1+w_2 & -w_2 & 0 & 0 \\ -w_3 & w_3+w_4 & -w_4 & 0 \\ 0 & -w_5 & w_5+w_6+w_7 & -w_7 \\ 0 & 0 & -w_8 & w_8+w_9 \end{pmatrix}.$$

Its determinant is equal to $w_1w_3w_5w_8 + w_1w_3w_6w_8 + w_1w_4w_6w_8 + w_2w_4w_6w_8 + w_1w_3w_5w_9 + w_1w_3w_6w_9 + w_1w_4w_6w_9 + w_2w_4w_6w_9 + w_1w_3w_7w_9 + w_1w_4w_7w_9 + w_2w_4w_7w_9$ (up to possible typos). We leave it to the reader to check that these eleven terms indeed correspond to the weights for the eleven possible dimer configurations of the weighted graph H.

Exercise 3 Spanning trees and rational classes

We consider the "fan graph" F_n, which is obtained from a path of n vertices $(v_1, \ldots, v_n)$ by adding a vertex v_0 linked to all other vertices, as on the figure. We let a_n be the number of spanning trees of this graph.

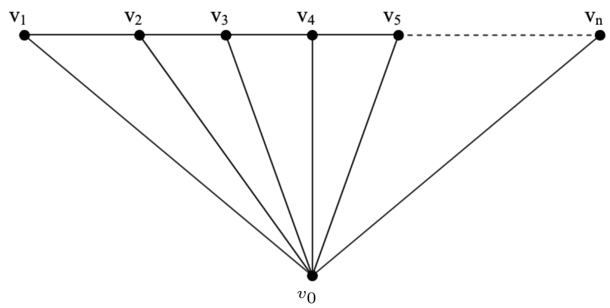

1. Show that $a_n = a_{n-1} + \sum_{k=1}^{n} a_{n-k}$ for $n \geq 2$, with $a_0 = a_1 = 1$.

Look at vertex v_n in a spanning tree of F_n. If it is not connected to v_0, then it is connected v_{n-1} and what remains is a spanning tree of F_{n-1}, this gives the first term. Otherwise, consider the maximal k such that it is connected to a chain of horizontal edges of the form

$v_n, v_{n-1}, \ldots, v_{n+1-k}$ (we have $1 \le k \le n$). What remains is a spanning tree of F_{n-k}, hence the relation.

2. Convert this equation into an equation for the generating function $A(z) := \sum_{n\ge 0} a_n z^n$.

We multiply the recurrence relation by z^n and we sum over $n \ge 2$. We obtain $A(z) - 1 = z(A(z) - 1) + \sum_{k\ge 1, p\ge 0} a_p z_{k+p} = z(A(z) - 1) + \frac{z}{1-z} A(z)$

3. What is the nature of the series $A(z)$? What is the form of its coefficients? (we do not ask you to explicitly calculate all the constants).

The series $A(z)$ is rational! We have $A(z) = (1 - z)/(1 - z - \frac{1}{1-z}) = \frac{(1-z)^2}{1-3z+z^2}$. We have $a_n = c_1 \alpha_1^n + c_2 \alpha_2^n$ where $\alpha_{1,2}$ are the (distinct!) roots of the denominator.

4. The Fibonacci sequence is defined by the recurrence $F_{n+2} = F_{n+1} + F_n$ for $n \ge 0$, with $F(0) = 0$ and $F(1) = 1$. Show that $a_n = F_{2n}$ for $n \ge 1$. [Hint: introduce a combinatorial sequence that will play the role of the Fibonacci numbers of odd index.]

We call b_n the number of spanning trees of F_n that contain the edge (v_0, v_n). We then have $a_n = b_n + a_{n-1}$, because the spanning trees of F_n that do not contain the edge (v_0, v_n) are in bijection with the spanning trees of F_{n-1}. Moreover, the spanning trees counted by b_n are of two types. Either they also contain the edge (v_n, v_{n-1}): these cases are in bijection with the spanning trees of F_{n-1} containing the rightmost edge, counted by b_{n-1}. Either they do not contain this edge, and by removing it we just have a spanning tree of F_{n-1}, counted by a_{n-1}. We therefore have $b_n = b_{n-1} + a_{n-1}$. We deduce by induction that $a_n = F_{2n}$ and $b_n = F_{2n-1}$ (be careful to check the initial conditions).