Combinatorial Maps

Tomaz Pisanski   Brigitte Servatius

University of Ljubljana and University of Primorska
and
Worcester Polytechnic Institute

1. Maps

1.1. Flags.

Given a connected graph $G = (V, E)$, we associate to each edge e a set $e' = \{x, \theta x, \phi x, \theta\phi x\}$ of four elements called *flags*, making the eight flags on any two distinct edges all distinct. We obtain a set Φ of flags and we have $|\Phi| = 4|E|$.

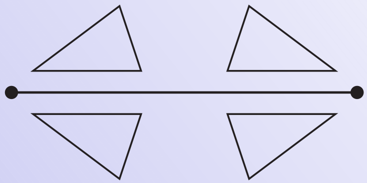

Fig 1 *Four flags for each edge*

The idea behind flags is simple. Suppose we want to embed G on some surface, say draw it on a piece of paper representing the edges by line segments (or curves) and the vertices by points, imposing the additional requirement that the edges only intersect at their endpoints. Such a drawing, if it can be produced, contains information not contained in the graph. Each edge drawn on the surface has a left and a right side, where left and right depend on the direction of edge traversal. We can think of the local picture of an edge as in Figure 1. This extra information will be encoded in the flags.

We consider ϕ and θ as permutations on Φ satisfying

1. $\phi^2 = \theta^2 = \mathrm{Id}$,
2. $\theta\phi = \phi\theta$.

For each edge we partition its flags $\{x, \theta x\} \cup \{\phi x, \theta\phi x\}$ and assign one of these pairs to each of the endpoints of the edge, so each vertex has associated to it a θ invariant pair of flags for each edge incident to it.
The local picture of a drawing of G at a vertex v consists of the set of semi-edges with endpoint v drawn in a particular clockwise order and labelled on each side by a flag. Using this local picture, we may associate to each flag x at v the opposite flag θx.

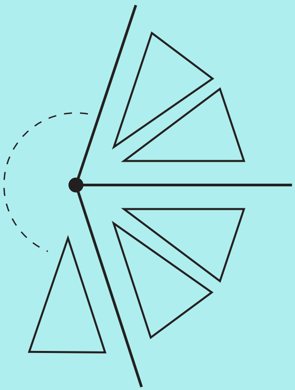

Fig 2 *vertex v with semi-edges*

This induces a permutation P on the set of flags, namely P rotates each flag x one place in the cyclic order around v, with half the flags rotating clockwise, and the other half counterclockwise. See Figure 2.

We require that the orbits of the flags x and θx under P are distinct and of the same size, i.e.

3. $P\theta = \theta P^{-1}$ and
4. $\{P^i x\} \cap \{P^i \theta x\} = \emptyset$

must hold.

The permutations $\{\theta, \phi, P\}$ generate a group A of permutations of Φ. Since G is connected, A acts transitively on Φ.

We call $M = M(\theta, \phi, P)$ a *map* if Axiom 1–4 are satisfied and if A acts transitively on Φ.

The permutations $\theta\phi$ and P generate a subgroup Θ of A. consider two flags x and y. Since A acts transitively there is an α in A such that $\alpha(x) = y$. Since ϕ and θ commute and $\theta P = P^{-1}\theta$, we can write α either as ωP, for some $\omega \in \Theta$, or α itself is an element of Θ. We conclude that for each pair of flags x and y there is an element of Θ mapping x into y or θy. This means that Θ partitions Φ into either one or two equal equivalence classes, the *orientation classes* of M.

If there are two orientation classes and the flag x is in one of them, then the flag θx must be in the other, as well as $\phi x = (\theta\phi)\theta x$. We call the map *orientable* if $\Theta = \{\theta\phi, P\}$ generates two orientation classes.

M is *non-orientable* if Θ only generates one orientation class.

Given a map $M = (\theta, \phi, P)$ on a set Φ of flags, we can construct other maps from it. For example, we could replace P by P^2, or by P^{-1}, and all axioms would be satisfied.

Let us study the permutation $P^* = P\theta\phi$. In order for P^* to define a map on Φ, we need P^* to satisfy Axiom 3, not necessarily together with θ. We have

$$P^*\phi = P\theta\phi^2 = P\theta = \theta P^{-1} = \phi(\phi\theta P^{-1}) = \phi(P^*)^{-1}.$$

This means that P^* satisfies Axiom 3 with the role of θ replaced by ϕ.
Interchanging the roles of ϕ and θ means interchanging the roles of the endpoints of an edge with that of its left and right side. We need to prove Axiom 4, namely that the orbits of P^* through x and ϕx are distinct. Assume to the contrary that $(P^*)^m x = \phi x$ and assume that x is chosen so that m is minimal. If $m = 1$, we have $P^* x = P\theta(\phi x) = \phi x$, contradicting Axiom 4 for $M = (\theta, \phi, P)$.

If $m \geq 2$, we calculate as follows:

$$
\begin{aligned}
(P^*)^m x &= \phi x \\
(P^*)^{-1}(P^*)^m x &= (P^*)^{-1}\phi x \\
(P^*)^{m-1} x &= \phi P^* x \\
(P^*)^{m-2}(P^* x) &= \phi(P^* x)
\end{aligned}
$$

So we have found a flag, namely P^*x, whose exponent is strictly smaller than m, contradicting our choice of x and m.

Given our connected graph $G = (V, E)$, together with local information of a cyclical order of edges incident to a vertex, as well as local information concerning each edge, we obtained the map $M(\theta, \phi, P)$. We can recover our graph G from it by observing that each vertex corresponds to a pair of conjugate orbits of $\theta\phi$. We call the orbits of P^* the *faces* of G. The map $M^* = (\phi, \theta, P^*)$ defines the graph G^*, with the conjugate orbits of P^* as the vertex set and the orbits of $\phi\theta$ as the edge set. We call the G^* the *geometric dual* of G with respect to the map M. Note that G and G^* have the same edge set. The faces of G are the vertices of G^*. Moreover, $(G^*)^* = G$.

If G is connected, then G^* is connected as well because the permutation group generated by $\{\theta, \phi, P\}$ is the same as that generated by $\{\phi, \theta, P\}$.
Similarly, the permutations $\phi\theta$ and P^* generate the same group Θ as do the permutations $\theta\phi$ and P, implying that both M and M^* are either both orientable or both non-orientable.

The flag graph.

We are now in a position to define the flag graph F associated with the map M. The vertex of F is the set of flags, $VF = \Phi$. We distinguish between three kinds of edges.

Type 1 Every flag x is adjacent to θx.

Type 2 Every flag x is adjacent to ϕx.

Type 3 Every flag x is adjacent to $P\theta x$.

Since by Axiom 3, $\theta P = P^{-1}\theta$, we see that $(\theta P)^2 = \theta P P^{-1}\theta = \theta^2 = \mathrm{Id}$, so θP is also an involution.
Altogether, we have four involutions, namely ϕ, θ, $\theta\phi$ and θP. With the help of these four involutions it is easy to describe, from the flag graph, the vertices, edges, faces, and the Petrie paths of the map M. The edges are the orbits generated by $\{\phi, \theta\}$. The vertices are the orbits generated by $\{\theta, P\theta\}$. The faces are the orbits generated by $\{\phi, P\theta\}$. The Petrie walks are the orbits generated by $\{\theta\phi, P\theta\}$.

Map projections

To depict a map obtained from a graph by choosing a particular clockwise order of half-edges around each vertex, we may start out by drawing each vertex, on a piece of paper, together with its half edges in the correct cyclic order with each side of a half-edge marked with a flag. Then connect corresponding half-edges by line segments, disregarding the perhaps unavoidable line crossings. If x and θx appear in clockwise order around v, but $\theta\phi x$ and ϕx appear in counterclockwise order around w, mark the edge corresponding to $\{x, \theta x, \phi x, \theta\phi x\}$ with a little cross in the middle. The resulting drawing is called a *map projection*.

It allows us in small cases to easily deduce the set of faces from the diagram: We start at one side of an edge,

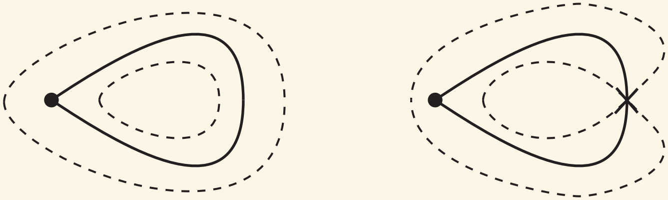

staying on the same side of the edge until we reach a vertex where we turn to follow the next edge as prescribed by the corner. However, if an edge is marked by a cross, we use the cross to switch to the other side and proceed as before. We keep walking on the same side (using all crosses as they come along) until we reach the starting point again. The edges we have travelled along are the boundary edges of the corresponding face.

Vertex Splitting

Given a map $M = (\theta, \phi, P)$ with v vertices, e edges and f faces, we call the alternating sum $v - e + f$ the *Euler characteristic* of M.
We now informally describe an operation and its inverse on maps which will allow us combinatorially describe the different types of surfaces (canonical normal forms) on which our graphs are embedded.

Let M be a map on a set Φ of flags, $|\Phi| = 4n$. We want to construct from M a map M_1 on Φ_1, $|\Phi_1| = 4(n+1)$, by choosing a vertex v of M, separating it into two vertices and inserting a new edge as described by the following local picture, see Figure 3. Nothing else of M is changed.

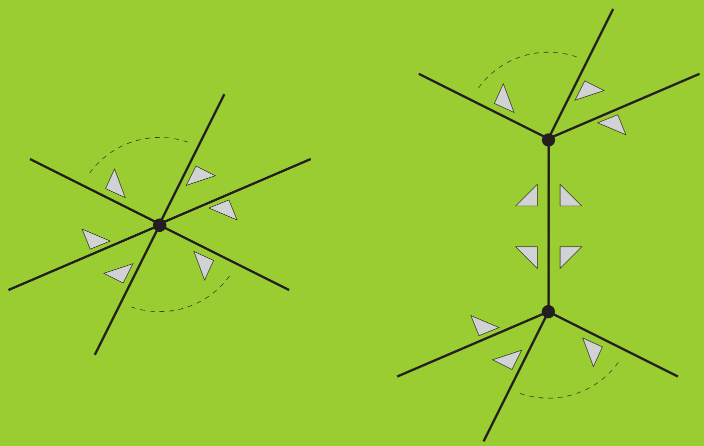

Fig 3

As indicated in Figure 3, $\Phi_1 = \Phi \cup \{z, \phi z, \theta z, \theta\phi z\}$ and θ and ϕ are extended to θ_1 and ϕ_1.
The two conjugate orbits of P corresponding to v are split into two pairs of conjugate orbits of P_1 with z, ϕz, θz, and $\theta\phi z$ inserted into exactly one of these four orbits as indicated by the drawing.

M_1 has one more vertex than M and one more edge than M, but the number of faces of M_1 equals the number of faces of M. Therefore M and M_1 have the same Euler characteristic. The inverse operation of vertex splitting is called *edge contraction*. Keeping in mind that a map has at least one edge, edge contraction is not defined for the link map.

Clearly, vertex splitting or edge contraction do not change the orientability character of the map. We define a *combinatorial surface* to be the (non-empty) class of all maps with a given Euler characteristic and a given orientability character. We say that a map is on the corresponding combinatorial surface.

Any map on a combinatorial surface can be transformed, by a finite sequence of edge contractions and dual replacements, to a map, on the same surface, which is either the link or the loop map, or a map consisting of one vertex and one face (and one or more edges.) A map with exactly one vertex and one face is called *unitary.*[a]

[a]Topologists call the 1 skeleton a unitary map *bouquet of circles*.

The permutation P of a unitary map consists of two conjugate orbits of equal length. If the pair $\{x, \phi x\}$ of flags is in the same orbit, $\{x, \phi x\}$ is called a *crosscap*. If the pair $\{x, \theta\phi x\}$ is contained in the same orbit of P, then there must be some other pair $\{y, \phi y\}$ contained in that orbit, and the orbit must be of the form $(x, R_1, y, R_2, \theta\phi x, R_3, \theta\phi y, R_4)$, where the R_i's are (possibly empty) sequences of flags.

To prove this important observation, let us assume to the contrary that the orbit of P through x is of the form $(x, R, \theta\phi x, S)$ and that with any flag y in R also $\theta\phi y$ is in R. Now for each flag y in R, consider P^*y. Since $P^*y = P\theta\phi y$ and both y and $\theta\phi y$ are in R, P^*y is a flag in $\{R, \theta\phi x\}$. Therefore one orbit of P^* contains only flags of $\{R, \theta\phi x\}$, but it does not contain x, contrary to the assumption that the map under consideration is unitary, i.e., P^* has only one vertex.

We call a quadruple $(x, y, \theta\phi x, \theta\phi y)$ a *handle* if there are integers $i < j < k$ such that $P^i x = y$, $P^j x = \theta\phi x$, $P^k x = \theta\phi y$. The handle is called *assembled* if $\{i, j, k\} = \{1, 2, 3\}$. Likewise, a crosscap $\{x, \phi x\}$ is assembled if $P(x) = \phi x$.

Let M be a unitary map with some unassembled crosscap $\{x, \phi x\}$. M can be transformed, by a vertex splitting and an edge contraction, into a map where $\{x, \phi x\}$ is replaced by an assembled crosscap denoted $\{z, \phi z\}$. The resulting map is unitary and no other crosscaps of M are disturbed. By induction, we can show that any unitary map can be transformed into a unitary map where all crosscaps are are assembled by a finite sequence of vertex splittings and edge contractions. Likewise, we can assemble all handles.

However, assembling crosscaps and handles is not enough to canonically classify unitary maps. A unitary map with one assembled handle and one assembled cross cap can be transformed, using only vertex splittings and edge contractions, into a unitary map containing three crosscaps. Such a transformation is possible on a unitary map M as long as M contains both a crosscap and a handle.

We see that we can transform a given unitary map into one containing only assembled crosscaps or only assembled handles. So we can describe canonical forms of combinatorial unitary maps as follows:
Either M is orientable and consists of g assembled handles and has Euler characteristic $2-2g$, in which case g is called the *orientable genus* of M, or M is unorientable and consists of g assembled crosscaps, hence has Euler characteristic $2-g$ and g is called the *unorientable genus* of M. For a unitary map, the genus is positive. We define the genus of the link and loop map to be zero. A combinatorial surface of genus zero is called a *sphere*. The surface of orientable genus 1 is called a *torus*, and a *double torus* if the orientable genus is two. The surfaces of unorientable genus 1 and 2 are called *projective planes* and *Klein bottles*, respectively.

Map projections of unitary maps on these surfaces are drawn in Figure 4. It is customary and useful to draw instead of the map projection, the single face of the unitary map as a polygon, where the corners of the polygon all correspond to the unique vertex of the map and are to be identified. Likewise every edge occurs twice on the boundary of the polygon and edges are to be identified pairwise, respecting the orientation of face traversal. Equivalently, one could represent the unitary map as a vertex in the interior of the polygon and edges to the boundary edges of the polygon, where these boundary edges indicate how the ends of the edge are attached to the vertex, twisted (in the unorientable case) or conforming (orientable case,) yielding our standard local picture of the map. See Figure 4.

GENFIG?

Fig 4