CAMBRIDGE

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

1
Introduction

In this chapter we introduce some basic definitions for graphs, maps, and polyhedra. We present here the basic notions. Further definitions will be introduced later when needed. The reader can consult the following books for more detailed information: [Grü67], [Cox73], [Mun71], [Cro97].

1.1 Graphs

A *graph* G consists of a set V of *vertices* and a set E of *edges* such that each edge is assigned two vertices at its ends. Two vertices are *adjacent* if there is an edge between them. The *degree of a vertex* $v \in V$ is the number of edges to which it is incident. A graph is said to be *simple* if no two edges have identical end-vertices, i.e. if it has no loops and multiple edges. In the special case of simple graphs, automorphisms are permutations of the vertices preserving adjacencies. For non-simple graphs (for example, when 2-gons occur) an *automorphism* of a graph is a permutation of the vertices and a permutation of the edges, preserving incidence between vertices and edges. By $Aut(G)$ is denoted the group of automorphisms of the graph G; a synonym is *symmetry group*.

For $U \subseteq V$, let $E_U \subseteq E$ be the set of edges of a graph $G = (V, E)$ having end-vertices in U. Then the graph $G_U = (U, E_U)$ is called the *induced subgraph* (by U) of G.

A graph G is said to be *connected* if, for any two of its vertices u, v, there is a path in G joining u and v. Given an integer $k \geq 2$, a graph is said to be *k-connected*, if it is connected and, after removal of any set of $k - 1$ vertices, it remains connected.

Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs. Their *Cartesian product* $G_1 \times G_2$ is the graph $G = (V_1 \times V_2, E)$ with vertex-set:

$$V_1 \times V_2 = \{(v_1, v_2) : v_1 \in V_1 \text{ and } v_2 \in V_2\}$$

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

and whose edges are the pairs $((u_1, u_2), (v_1, v_2))$, where $u_1, v_1 \in V_1$ and $u_2, v_2 \in V_2$, such that either $(u_1, v_1) \in E_1$, or $(u_2, v_2) \in E_2$.

A subset E' of edges of a graph is called a *matching* if no two edges of E' have a common end-vertex. A *perfect matching* is a matching such that every vertex belongs to exactly one edge of the matching.

The following graphs will be frequently used:

- The *complete graph* K_n is the graph on n vertices $v_1, \dots, v_n$ with v_i adjacent to v_j for all $i \neq j$.
- The *path* $P_n = P_{v_1, v_2, \dots, v_n}$ is the graph with n vertices $v_1, \dots, v_n$ and $n - 1$ edges (v_i, v_{i+1}) for $1 \leq i \leq n - 1$.
- The *circuit* $C_n = C_{v_1, v_2, \dots, v_n}$ (or *n-gon*) is the path $P_{v_1, v_2, \dots, v_n}$ with additional edge (v_1, v_n).

A *plane graph* is a connected graph, together with an embedding on the plane such that every edge corresponds to a curve and no two curves intersect, except at their end points. A graph is *planar* if it admits at least one such embedding. It is known that any planar graph admits a plane embedding with the edges being straight lines (see [Wa36, Fa48, Tut63]). A *face* of a plane graph is a part of the plane delimited by a circuit of edges. A plane graph defines a partition of the plane into faces. If a is a vertex, edge, or face and b is a edge, face, or vertex, then a is said to be *incident* to b if a is included in b or b is included in a. Two vertices, respectively, faces are called *adjacent* if they share an edge. We will call *gonality* or *covalence* of a face the number of its vertices. A face is *exterior* if it is non-bounded. Bounded faces are called *interior*. Any finite plane graph has exactly one exterior face. An infinite plane graph can have any number, from zero to infinity, of exterior faces. A planar 3-connected graph admits exactly one plane embedding on the sphere, i.e. the set of faces is determined by the edge-set.

The *v-vector* $v(G) = (\dots, v_i, \dots)$ of a graph G enumerates the numbers v_i of vertices of degree i. A plane graph is *k-valent* if $v_i = 0$ for $i \neq k$. The *p-vector* $p(G) = (\dots, p_i, \dots)$ of a plane graph G counts the numbers p_i of faces of gonality i. For a connected plane graph G, denote its *plane dual graph* by G^* and define it on the set of faces of G with two faces being adjacent if they share an edge. Clearly, $v(G^*) = p(G)$ and $p(G^*) = v(G)$.

1.2 Topological notions

We present in this section the topological notions for the surfaces which will be used. Topology is concerned with continuous structures and invariants under continuous

© Cambridge University Press www.cambridge.org

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

deformations. Since we are working with vertices, edges, and faces, the classical definitions will be adapted to our context.

No proofs are given but we hope to compensate for this by giving some geometrical examples. More thorough explanations are available in basic algebraic topology textbooks, for example, [Hat01] and [God71].

1.2.1 Maps

A *map M* is a family of vertices, edges, and faces such that every edge is contained in at least one and at most two faces. An edge, contained in exactly one face, is called a *boundary edge*; all such edges form the *boundary*. A map is called *closed* if it has no boundary. A map is called *finite* if it has a finite number of vertices, edges, and faces. See below plane graphs related to $Prism_5$ (see Section 1.5) with same vertex- and edge-sets but different face-sets; their boundary edges are boldfaced:

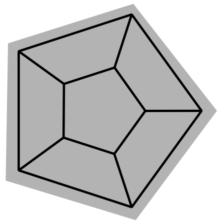

A closed map, cell-complex of a polyhedron. It is a $5R_0$, $4R_2$ plane graph (see Chapter 9)

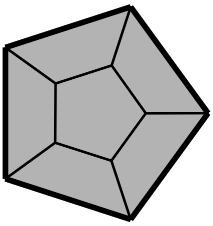

A map with boundary edges. It is a ({4, 5}, 3)-polycycle (see Chapter 7)

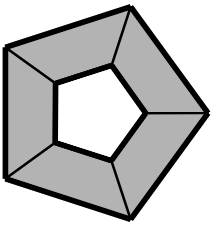

A map with boundary edges; not simply connected. It is a $(4, 3)_{gen}$-polycycle (see Section 4.5)

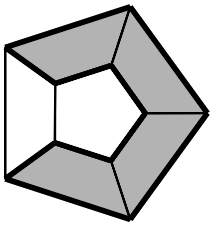

Not a map because two edges are not contained in a face. It is not considered

If M is a closed map, then we can define its dual map M^* by interchanging faces and vertices. See Section 4.1 for some related duality notions for non-closed maps. A map is called a *cell-complex* if the intersection of any two faces, edges, or vertices

CAMBRIDGE

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

is a face, edge, vertex, or ∅. Maps with 2-gons are not cell-complexes; they are *CW-complexes* (see, for example, [Rot88]).

Denote by $\mathbb{S}^2$ the 2-dimensional sphere defined by $\{x = (x_1, x_2, x_3) \in \mathbb{R}^3 : x_1^2 + x_2^2 + x_3^2 = 1\}$. For a point A of $\mathbb{S}^2$, let A' be its opposite and the plane $\mathcal{H}_A$ be the plane orthogonal to AA' passing through A'. In the following, take $A = (0, 0, -1)$; then $\mathcal{H}_A$ is the set of all $x \in \mathbb{R}^3$ with $x_3 = 1$. If $B \in \mathbb{S}^2 - \{A\}$, then the intersection of the line AB with the plane $\mathcal{H}_A$ defines a point $f_A(B)$. This establishes a bijection, called a *Riemann map*, between $\mathbb{S}^2 - \{A\}$ and the plane $\mathcal{H}_A$; we can extend f_A to A by defining $f_A(A)$ to be the "point at infinity" of the plane $\mathcal{H}_A$.

Let G be a finite plane graph on $\mathbb{R}^2 \simeq \mathcal{H}_A$ and let $f_A^{-1}(G)$ denote its image in $\mathbb{S}^2$. The vertices of G correspond to points of the sphere $\mathbb{S}^2$, the edges of G correspond to non-intersecting curvilinear lines on $\mathbb{S}^2$ and the faces of G correspond to domains of $\mathbb{S}^2$ delimited by circuit of those lines. The exterior face of G corresponds to a domain of $\mathbb{S}^2$ containing A. Reciprocally, if we have a map M on the sphere, then we can find a point A, which does not belong to M and the corresponding plane $\mathcal{H}_A$, the image of M on $\mathcal{H}_A$ is a finite plane graph. So, by abuse of language, we will use the term "sphere" not only for the surface $\mathbb{S}^2$ but also for any combinatorial map on it, i.e. a finite plane graph.

A reader who is interested only in plane graphs, our main subject, can move now directly to Section 1.4. But, for full understanding of the toroidal case, we need maps in all their generality. For reference on Map Theory; see, for example, [BoLi95] and [MoTh01].

We will also work with maps having an infinite number of vertices, edges, and faces. The vertex-degrees will always be bounded by some constant; however, faces could have an infinity of edges. Amongst plane drawing of those maps, we will allow only *locally finite* ones, i.e. those admitting an embedding such that any bounded domain contains a finite number of vertices. Consider, for example, the map $\widetilde{\mathbb{Z}^2}$ obtained as the quotient map of square tiling $\mathbb{Z}^2$ by the translation operation $(x, y) \mapsto (x + 10, y)$. The map $\widetilde{Z^2}$ is an infinite cylinder made of consecutive rings of ten 4-gons. We can draw those rings concentrically on the plane, but the resulting plane graph will not be locally finite.

1.2.2 Orientability and classification of surfaces

A *flag of a map* is a triple (v, e, f), where v is a vertex contained in the edge e and e is contained in the face f. Given a flag $F = (v, e, f)$, the flags differing from F only in v, e or f are denoted by $\sigma_0(F)$, $\sigma_1(F)$ and $\sigma_2(F)$, respectively. $\sigma_0(F)$ and $\sigma_1(F)$ are always defined over the flag-set $\mathcal{F}(M)$ of M but $\sigma_2(F)$ is not always defined if M is not closed.

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

The map M is called *oriented* if there exist a bipartition $\mathcal{F}_1$, $\mathcal{F}_2$ of $\mathcal{F}(M)$ such that, for any $(v, e, f) \in \mathcal{F}_1$, the flag $\sigma_i(v, e, f)$, if it exists, belongs to $\mathcal{F}_2$. We will be almost exclusively concerned with oriented maps.

The notion of orientation is easy to define algebraically but difficult to visualize, because closed non-orientable maps cannot be represented by a picture. Fortunately, this is easier for maps with boundaries; see Figure 1.1 for a non-orientable map, called *Möbius strip*. The non-orientability can be seen in the following way: moving along one side of the strip and doing a full circuit, we arrives at the other side of the strip. All boundary edges of a Möbius strip belong to a unique cycle; after adding a face to this cycle, we obtain the *projective plane* $\mathbb{P}^2$. The projective plane can also be obtained by taking a map on the sphere (like Dodecahedron) and identifying the opposite vertices, edges, or faces, i.e. taking the *antipodal quotient*.

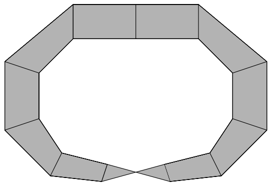

Figure 1.1 A Möbius strip

Given a surface S, we can add to it a *handle*:

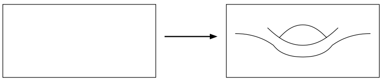

or a *cross-cap*.[1] The handle and cross-cap can be seen as a cylinder and a Möbius strip, respectively.

Consider now the classification of finite maps:

Theorem 1.2.1 *Any finite closed map is one of the following:*

1 *the sphere* $\mathbb{S}^2$ *(orientable),*
2 *the sphere* $\mathbb{S}^2$ *with* g *handles (orientable),*
3 *the sphere* $\mathbb{S}^2$ *with* g *cross-caps (non-orientable).*

[1] See, for example, http://mathworld.wolfram.com/Cross-Cap.html for pictures of cross-caps.

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

Theorem 1.2.1 is proved, for example, in [Mun71]. The number g above is called the *genus* of the map. All finite closed maps, that occur below, are:

1 the sphere $\mathbb{S}^2$ with $g = 0$ (orientable),
2 the *torus* $\mathbb{T}^2$ with $g = 1$ (orientable),
3 the *projective plane* $\mathbb{P}^2$ with $g = 1$ (non-orientable),
4 the *Klein bottle* $\mathbb{K}^2$ with $g = 2$ (non-orientable); one way to obtain the Klein bottle is to take the quotient of a torus $\mathbb{R}^2/\mathbb{Z}^2$ by the fixed-point-free automorphism $f(x, y) = (x + \frac{1}{2}, -y)$.

If M is a finite non-closed map, then we can add some faces along the boundary edges and obtain a closed map. So, finite non-closed maps are obtained by removing some faces of closed ones.

1.2.3 Fundamental groups, coverings, and quotient maps

Fix an orientation on every edge of a given map M and define the free group $G(\mathcal{M})$ with generators g_e indexed by the edge-set of M (see, for example, [Hum96] for relevant definitions in Group Theory). An *oriented path* $\mathcal{O}P = (v_1, v_2), \dots, v_m$ is a sequence of vertices with v_i adjacent to v_{i+1}. For an edge $e_i = (v_i, v_{i+1})$, denote by $g(v_i, v_{i+1})$ the group element g_{e_i} if e_i is oriented from v_i to v_{i+1}, and $g_{e_i}^{-1}$ otherwise. Associate to the oriented path the product $g(\mathcal{O}P) = g(v_1, v_2)g(v_2, v_3)\dots g(v_{m-1}, v_m)$.

Denote by $\mathcal{Z}_v(M)$ the set of all $g(\mathcal{O}P)$ with $\mathcal{O}P$ being the set of oriented closed paths starting and finishing at a given *base vertex* v. It is a group; reversing orientation corresponds to taking the inverse and product to concatenating closed paths. Given a face F of M, bounded by a circuit of vertices $(v_1, \dots, v_{|F|})$, and an oriented path $\mathcal{O}P$ from the vertex v to the vertex v_1, consider a group element $g(\mathcal{O}P)g(v_1, v_2, \dots, v_{|F|}, v_1)g(\mathcal{O}P)^{-1}$. Denote by $\mathcal{B}_v(M)$ the subgroup of $G(\mathcal{M})$ generated by all such elements. The *fundamental group* $\pi_1(M)$ is the quotient group of the group $\mathcal{Z}_v(M)$ by the normal subgroup $\mathcal{B}_v(M)$. Two oriented closed paths having a common vertex v are called *homotopic* if they correspond to the same element in the group $\pi_1(M)$. The group $\mathcal{B}_v(M)$ is the group of all elements homotopic to the *null path*, i.e. the path from v to v with 0 edges. If we replace the base vertex v by another base vertex w, then, for any oriented path $\mathcal{O}P$ from v to w, we have $\mathcal{Z}_v(M) = g\mathcal{Z}_w g^{-1}$ and $\mathcal{B}_v(M) = g\mathcal{B}_w g^{-1}$ with $g = g(\mathcal{O}P)$. So, the fundamental group depends on the base vertex, but only up to conjugacy. A map is called *simply connected* if $\pi_1(M)$ is trivial, i.e. every path is homotopic to the null path. This is equivalent to saying that every two paths with the same beginning and end can be continuously deformed one to the other.

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

See below three homotopic paths in the same map:

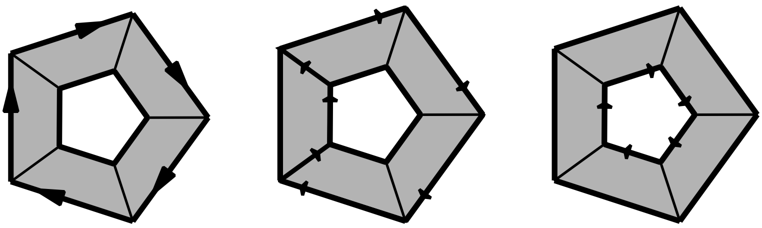

See below again three plane maps and a closed path represented in it:

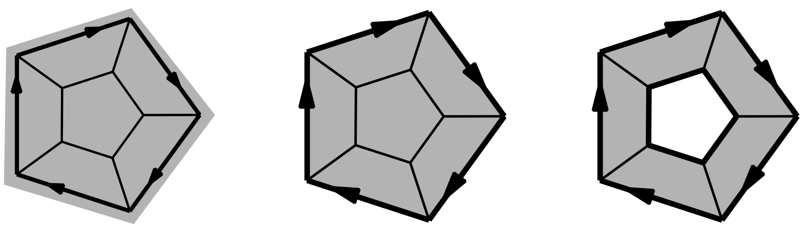

M_1 is simply connected M_2 is simply connected M_3 is not simply connected

In M_1 and M_2, the closed path is homotopic to a null path. In M_1, this cycle is the boundary of a face, while in M_2, the closed path is the boundary of all faces put together. More generally, a plane graph and a finite plane graph minus a face are simply connected. But the closed path in M_3 is not homotopic to a null path. Actually, this closed path is a generator of the fundamental group $\pi_1(M_3) \simeq \mathbb{Z}$.

Given two maps M and M', a *cell-homomorphism of maps* $\phi : M \to M'$ is a function that maps vertices, edges, and faces of M to the ones of M', while preserving the incidence relations. An *isomorphism* is a cell-homomorphism that is bijective. If $M = M'$, it is called an *automorphism*; the set of all automorphism of a map M is called the *symmetry group* of M. An automorphism f of a map M is called *fixed-point-free* if f is the identity or, for every vertex, edge, face of M, its image in f is different from it. If G is a group of fixed-point-free automorphisms of a map M, then M/G is the *quotient map* of M by G. Its vertices, edges, and faces are formed by orbits of vertices, edges, and faces of M (by G) with the incidence relations being induced by the ones of M.

The quotient of a map can be a map with loops and multiple edges. Consider, for example, the 4-valent plane tiling $\{4, 4\}$ (see Section 1.5) formed by 4-gons and the group $\mathbb{Z}^2$ acting by translations on it. There is one orbit of vertices, two orbits of edges, and one orbit of faces under $\mathbb{Z}^2$; so the quotient $\{4, 4\}/\mathbb{Z}^2$ is a torus represented by a single vertex and two loops.

For a vertex v (or edge e, or face f), the *standard neighborhood* $N(v)$ is the set of all vertices, edges, and faces incident to v. A *local isomorphism* is a continuous

© Cambridge University Press www.cambridge.org

Cambridge

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

mapping $\phi\colon M \to M'$ such that, for any vertex $v \in M$ (or edge, or face), the mapping from $N(v)$ to $N(\phi(v))$ is bijective. A *covering* is a local isomorphism such that for every vertex $v' \in M'$ (or edge, or face) and $w' \in N(v')$, if $\phi^{-1}(v') = (v_i)_{i \in I}$, we have an element $w_i \in N(v_i)$ such that $w_i \neq w_j$ if $i \neq j$ and $\phi^{-1}(w') = (w_i)_{i \in I}$.

If $\mathcal{O}P' = (v'_1, \dots, v'_m)$ is an oriented path in M', ϕ is a covering and v_1 is a vertex in M with $\phi(v_1) = v'_1$, then there exists a unique oriented path $\mathcal{O}P = (v_1, \dots v_m)$ in M such that $\phi(\mathcal{O}P) = \mathcal{O}P'$. A *deck automorphism* is an automorphism u of M such that $\phi \circ u = \phi$, u is necessarily fixed-point-free. If $v' \in M'$, then, for any two $v_1, v_2 \in \phi^{-1}(v')$, there exists a deck automorphism u such that $u(v_1) = v_2$.

Given a map M, its *universal cover* is a simply connected map $\widetilde{M}$ (unique up to isomorphism) with a covering $\phi : \widetilde{M} \to M$. The map $\widetilde{M}$ is finite if and only if M and $\pi_1(M)$ are finite. The fundamental group $\pi_1(M)$ is isomorphic to the group of deck automorphisms of ϕ. If H is a subgroup of a group G, then its *normalizer*, denoted by $N_G(H)$, is defined as:

$$N_G(H) = \{x \in G \ : \ xhx^{-1} \in H \text{ for all } h \in H\}.$$

The group $Aut(M)$ of automorphisms of M is identified with the quotient group:

$$N_{Aut(\widetilde{M})}(\pi_1(M))/\pi_1(M).$$

The simplest and most frequently used case is when M is a closed finite map on the sphere. In this case π_1 is trivial and we can represent the map nicely on the plane with a face chosen to be exterior. An infinite locally finite closed simply connected map can be represented on the plane. In this case, there is no exterior face and the map fills completely the plane.

A closed torus M can be represented as a 3-dimensional figure projected on to the plane, but this is not very practical. We represent its universal cover $\widetilde{M}$ as a plane having two periodicity directions, i.e. a 2-*periodic plane graph*. The group $\pi_1(M)$ is isomorphic to $\mathbb{Z}^2$ and it is represented on $\widetilde{M}$ as a group of translation symmetries.

By choosing a *finite index subgroup* H of the group G (i.e. such that there exist $g_1, \dots, g_m \in G$ with $G = \cup_i g_i H$) of deck transformations and taking the quotient, we can obtain a bigger torus; such tori have a translation subgroup, which is isomorphic to the quotient G/H.

On the other hand, given a torus with non-trivial translation group, there exists a unique *minimal torus* with the same universal cover and trivial translation subgroup. Those minimal tori correspond, in a one-to-one way, to periodic tilings of the plane.

1.2.4 Homology and Euler–Poincaré characteristic

Given a map M, assign an orientation on each of its edges and form a $\mathbb{Z}$-module $C_1(M)$ using this set of oriented edges as basis. The $\mathbb{Z}$-module $Z_1(M)$ is the submodule of $C_1(M)$ generated by the set of closed oriented cycles of M. Given any

CAMBRIDGE

Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

face of M, associate to it the set of incident edges in clockwise orientation; the generated $\mathbb{Z}$-module is denoted by $B_1(M)$. It is easy to see that $B_1(M)$ is a submodule of $Z_1(M)$.

The *homology group* $H_1(M)$ is the quotient of $Z_1(M)$ by its subgroup $B_1(M)$. Again, we refer to Algebraic Topology textbooks for details. If M is a torus, then $H_1(M)$ is isomorphic to $\pi_1(M)$.

If M is an orientable finite closed map, then $H_1(M)$ is isomorphic to $\mathbb{Z}^{2g}$, where g is the genus of M. The *Euler–Poincaré characteristic* of a finite map M is defined as $\chi(M) = v - e + f$ with v the number of vertices, e the number of edges, and f the number of faces.

Theorem 1.2.2 *For a finite closed map M of genus g it holds:*

(i) if M is orientable, then $\chi(M) = 2 - 2g$,
(ii) if M is non-orientable, then $\chi(M) = 2 - g$.

This theorem is the main reason why we are able to use topology in dimension two to derive non-trivial combinatorial results.

Theorem 1.2.3 *Let G be a k-valent closed map on a surface M; then:*

(i) the following Euler formula *is valid:*

$$\sum_{j\geq 2} p_j(2k - j(k-2)) = 2k\chi(M), \tag{1.1}$$

where p_i is the number of i-gonal faces.

(ii) If G has no 2-gonal faces, then $k \leq 5$ if M is a sphere and $k \leq 6$ if M is a torus.

Proof. (i) The relation $2e = kv$ allows us to rewrite the Euler–Poincaré characteristic as:

$$\chi(M) = \left(\frac{2}{k} - 1\right)e + \sum_{i\geq 2} p_i \ .$$

Using that $2e = \sum_{i\geq 2} ip_i$ in the above equation, yields the result.

If $j \geq 3$, then $2k - j(k-2) \leq 0$ for $k \geq 6$ and $2k - j(k-2) < 0$ for $k \geq 7$. Assertion (ii) is deduced by noticing that $\chi = 2, 0$ for sphere, and torus, respectively. □

1.3 Representation of maps

A *polytope* P is the convex hull of a finite set of points in $\mathbb{R}^n$; its *dimension* is the dimension of the smallest affine space containing it. We assume it to be full-dimensional. A linear inequality $f(x) \geq 0$ is called *valid* if it holds for all $x \in P$. A *face* of P is a set of the form $\{x \in P \ : \ f(x) = 0\}$ with $f \geq 0$ being a valid inequality.

© Cambridge University Press www.cambridge.org

CAMBRIDGE
Cambridge University Press
978-0-521-87307-9 - Geometry of Chemical Graphs: Polycycles and Two-faced Maps
Michel Deza and Mathieu Dutour Sikiric
Excerpt
More information

We will consider only 3-dimensional polytopes; they are called *polyhedra*. Their 0-dimensional faces are called *vertices* and the 2-dimensional faces are called just *faces*. Two vertices are called *adjacent* if there exist an *edge*, i.e. a 1-dimensional face containing both of them. The *skeleton* of a polyhedron is the graph formed by all its vertices with two vertices being *adjacent* if they share an edge. This graph is 3-connected and admits a plane embedding.

Given a polyhedron P, its skeleton $skel(P)$ is a planar graph. Furthermore, for any face F of P, we can draw $skel(P)$ on the plane so that F is the exterior face of the plane graph. Those drawings are called *Schlegel diagrams* (see, for example, [Zie95]). Steinitz proved that a finite graph is the skeleton of a polyhedron (and so, an infinity of polyhedra with the same skeleton) if and only if it is planar and 3-connected (see [Ste22], [Zie95, Chapter 4] and [Grü03] for a clarification of the history of this theorem).

A *Riemann surface* is a 2-dimensional compact differentiable surface, together with an infinitesimal element of length (see textbooks on differential and Riemannian geometry, for example, [Nak90]). The *curvature* $K(x)$ at a point x is the coefficient α in the expansion:

$$Vol(D(x,r)) = \pi r^2 - \alpha r^4 + o(r^4)$$

with $D(x,r)$ being the disc consisting of elements at distance at most r from x. The curvature of a Riemann surface S satisfies the *Gauss–Bonnet formula*:

$$\int_S K(x)dx = 2\pi(1-g).$$

All Riemann surfaces, considered in this section, will be of constant curvature. If a surface has constant curvature, then, for any two points x and y of it, there exist two neighborhoods N_x and N_y and a local isometry ϕ mapping x to y and N_x to N_y. Hence, Riemann surfaces of constant curvature do not have local invariants and the only invariants they have are global (see, for example, [Jos06]). For genus zero, the curvature has positive integral. Up to rescaling, we can assume that this curvature is 1. There is only one such Riemann surface: the sphere $\mathbb{S}^2$. For genus 1, the curvature has integral 0 and so it is 0. The Teichmüller space T_1 has dimension 2, which means that Riemann surfaces of genus 1 are parametrized by two real parameters. Geometrically, they are very easy to depict: take $\mathbb{R}^2$ and quotient it by a group $v_1\mathbb{Z} + v_2\mathbb{Z}$. For higher genus, the situation is much more complicated.

Given a map M, its *circle-packing representation* (see [Moh97]) is a set of disks on a Riemann surface Σ of constant curvature, one disk $D(v, r_v)$ for each vertex v of M, such that the following conditions are fulfilled:

1 the interior of disks are pairwise disjoint,
2 the disk $D(u, r_u)$, $D(v, r_v)$ touch if and only if uv is an edge of M.

© Cambridge University Press www.cambridge.org