GRAPH HOMOLOGY AND COHOMOLOGY

IAN DEWAN

Abstract. We provide an introduction to graph homology and cohomology, giving the definitions of four homology and cohomology theories of graphs, and proving some results on the relation between the combnatorial and topological properties of graphs and their homology. The only required background is undergraduate algebra.

Contents

1. Introduction

Graph theory traditionally is examined from a combinatorial perspective, treating problems in graph colourings, or finding subgraphs or cycles with certain properties while considering the graph only as a set of vertices and edges with an incidence relation. Graphs are also naturally treated topologically, where they are considered as one-dimensional simplical complexes, and problems like finding embeddings of graphs in other topological spaces can be considered.

Graph *homology* naturally combines these two perspectives by associating to a graph algebraic structures which can then be used to recover information about the combinarorial and topological structure of the graph. In addition to its application directly to graph theory it also has applications in other areas: for example in algebra, the moduli space of the skew-zigzag algebra associated with a graph

2010 *Mathematics Subject Classification.* 05C10, 05C25.

(introduced in [8]) has been shown recently (in [3]) to be related to the cohomology of the graph.

Although graph homology has appeared in the literature, there is little introductory material on it: this document aims to rectify this, by providing an introductory overview of graph homology accessible to a student with only undergraduate algebra (some sections also require basic topology and category theory). Section 2 provides the required background in graph theory, Sections 3 and 4 describe the vertex and edge spaces of a graph, and Section 5 introduces the concept of the homology of a chain complex in a general algebraic context. Sections 6–9 describe four homology and cohomology theories of graphs in detail, how the (co)homology can be obtained for several kinds of graphs, and its general properties. Sections 10 and 11 cover algebraic properties of homology, culminating in the Universal Coefficient Theorem, and the effect of base change on homology. Sections 12–14 cover some topological properties of graphs and how they relate to graph homology.

2. Graph theory background

A *graph* Γ is an ordered triple (V, E, ν), where V is a set of *vertices* (or *nodes*), and E a set of *edges*, which is disjoint from V, and ν is a function $E \to \mathcal{P}_2(V)$, where $\mathcal{P}_2(V) = \{\{u, v\} \mid u, v \in V\}$ is the set of one or two element subsets of V. If $\nu(e) = \{u, v\}$ for $e \in E$, we say that the edge e connects, or is an edge between, the vertices u and v; if $u = v$, then e is called a *loop.* If $f \in E$ is also an edge between u and v, then e and f are *parallel edges.* Note that this definition of a graph is often called a multigraph or pseudograph in other sources, which restrict graphs to have no loops or parallel edges; we call this restricted type of graph a *simple graph.*

Graphs are usually represented as diagrams, with a point for each node and an arc for each edge.

Example 2.1. The diagram below is the undirected graph $\Gamma = (V, E, \nu)$, with $V = \{a, b, c, d\}$, and $E = \{e, f, g, h\}$ and $\nu(e) = \{a, c\}$, $\nu(f) = \{b, c\}$, $\nu(g) = \{b, d\}$, and $\nu(h) = \{c, d\}$.

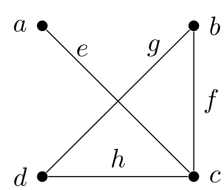

If u and v are vertices of a graph Γ, a u–v *walk* is an alternating sequence of vertices and edges of Γ $(u = a_1, e_1, a_2, e_2, \ldots, e_{n-1}, a_n = v)$ such that for $1 \leq i < n$, e_i is an edge connecting a_i and a_{i+1}. A u–v walk W is *closed* (or a *cycle*) if $a_1 = a_n$, a *path* if no vertex is repeated, and a *simple cycle* if it is closed and no vertex but $a_1 = a_n$ is repeated. The *length* of W, $\ell(W)$, is the number of edges it contains, which is always $n - 1$.

For every vertex u in Γ, there is a *trivial walk* of length zero, denoted (u), which is a path and a cycle, and is not to be confused with a walk (of length one) passing through a loop on u. Given walks $W = (a_1, e_1, \ldots, e_{n-1}, a_n)$ and $V = (b_1, f_1, \ldots, b_{m-1}, b_m)$, where $a_n = b_1$, there is a *concatenation* $W + V = (a_1, e_1, \ldots, e_{n-1}, a_n = b_1, f_1, \ldots, f_{m-1}, b_m)$. Moreover, we have that $\ell(W + V) =$

$\ell(W) + \ell(V)$. Concatenation is associative, and concatenation with a trivial walk (when defined) leaves a walk unchanged.

Proposition 2.2. *Let Γ be a graph, and u and v vertices of Γ. If there is a u–v walk W in Γ, then there is a u–v path which is a subsequence of W.*

Proof. Let $P = (a_1, e_1, \ldots, e_{n-1}, a_n)$ be a u–v walk of minimal length which is a subsequence of W; we show P is a path. Suppose P were not a path: then there would be $1 \leq i < j \leq n$ such that $a_i = a_j$. We can divide P into three components, $P_1 = (a_1, e_1, \ldots, e_{i-1}, a_i)$, $P_2 = (a_i, e_i, \ldots, e_{j-1}, a_j)$, and $P_3 = (a_j, e_j, \ldots, e_{n-1}, a_n)$. Then $P = P_1 + P_2 + P_3$. Since $i < j$, $\ell(P_2) > 0$. Since $a_i = a_j$, $P_1 + P_3$ exists, and $\ell(P_1 + P_3) = \ell(P_1) + \ell(P_3) < \ell(P_1) + \ell(P_2) + \ell(P_3) = \ell(P)$. Then $P_1 + P_2$ is a u–v walk, a subsequence of W, and of length less than P. Contradiction. Therefore P is a path. □

In a graph $\Gamma = (V, E)$, the *degree* of a vertex v is the number of edges incident on that vertex; loops on v are counted twice.

Let $\Gamma = (V, E, \nu)$ and $\Delta = (V', E', \nu')$ be graphs. If $V' \subseteq V$, $E' \subseteq E$, and ν' is equal to ν on its domain, then Δ is a *subgraph* of Γ. Let $f_V : V \to V'$ and $f_E : E \to E'$, and suppose that for every $e \in E$, $\nu'(f_E(e)) = f_V(\nu(e))$; then $f = (f_V, f_E)$ is a *graph homomorphism,* and we write $f : \Gamma \to \Delta$. We can compose graph homomorphisms by composing the underlying functions, and the composition of graph homomorphisms is a homomorphism. The identity functions on V and E give rise to a identity homomorphism $\mathrm{id}_\Gamma : \Gamma \to \Gamma$: therefore there is a category Graph with graphs as objects, and graph homomorphisms as morphisms. The isomorphisms of this category are the f such that the underlying functions are bijective and their inverses also form graph homomorphism; they are called *graph isomorphisms.* There is a full subcategory $\mathsf{FinGraph}$ the objects of which are the graphs with a finite number of vertices and edges.

Two vertices u and v of a graph are *connected* if there exists a u-v walk; connectedness is an equivalence relation on vertices, and the subgraphs determined by the equivalence classes are called *connected components* of the graph. A graph is connected if it has only one connected component. A graph with no simple cycles is a *forest,* or *acyclic,* and a connected forest is a *tree.* Note that every forest must be a simple graph, since loops and parallel edges always give rise to cycles.

Proposition 2.3 ([4, Th. 1.5.1])**.** *Let T be a graph. Then the following conditions are equivalent:*

(1) *T is a tree;*
(2) *There is a unique path between any two vertices in T;*
(3) *T is minimally connected, i.e. removing any edge from T will leave an unconnected graph;*
(4) *T is maximally acyclic, i.e. the graph produced by adding any edge to T contains a cycle.*

Corollary 2.4 ([4, Cor. 1.5.2])**.** *The vertices of a finite tree can always be enumerated $v_1, v_2, \ldots, v_n$ so that for every v_i, $i \geq 2$, there is exactly one edge between v_i and some element of $\{v_1, v_2, \ldots, v_{i-1}\}$.*

Corollary 2.5 ([4, Cor. 1.5.3])**.** *A connected graph with n vertices is a tree if and only if it has $n - 1$ edges.*

A *spanning tree* is a subgraph of a graph Γ which is a tree and contains all the vertices of Γ. If T is a spanning tree, any edges of Γ not in T are called *chords* of T in Γ.

Proposition 2.6 ([4, Prop. 1.5.6])**.** *Every connected graph has a spanning tree.*

Proposition 2.6 applies to all graphs, and in the case of infinite graphs is equivalent to the axiom of choice.

It will often be natural to want to give edges a direction, for which we introduce the notion of an *oriented edge:* an oriented edge in a graph $\Gamma = (V, E, \nu)$ is a triple $(e, s, t) \in E \times V \times V$ such that $\nu(e) = \{s, t\}$. We say the vertex s is the source and t the target of the oriented edge. Clearly there are two possible oriented edges for any edge (except for loops, where there is only one): an *orientation* of a graph is a selection of one oriented edge for each edge of the graph.

3. The cycle space

Let R be a principal ideal domain.

Definition 3.1 (vertex space)**.** Let $\Gamma = (V, E, \nu)$ be a graph. The *vertex space* of Γ, written $\mathcal{V}_R(\Gamma)$, is the free R-module over V.

We will often drop the subscript when the ring is unimportant or clear from context. The vertices of Γ form a basis for $\mathcal{V}(\Gamma)$ by construction. There is a similar concept for edges:

Definition 3.2 (edge space)**.** Let Γ be a graph. Then the *edge space* of Γ, $\mathcal{E}_R(\Gamma)$, is the R-module generated by the set of oriented edges of Γ subject to the relations

$$(e, s, t) + (e, t, s) = 0$$

for every oriented edge (e, s, t) in Γ.

As before, we often drop the subscript R. Note that $\mathcal{E}(\Gamma)$ is free and isomorphic to the free R-module over E: this isomorphism is not natural, but every choice of orientation of the edges of Γ gives such an isomorphism.

Our definition of the edge space implicitly makes it a quotient of the free R-module over the set of oriented edges of Γ, O. This means that there is a projection $\pi\colon O \to \mathcal{E}(\Gamma)$ allowing us to treat any oriented edge as an element of $\mathcal{E}(\Gamma)$, and the morphism $f\colon O \to \mathcal{V}(\Gamma)$, which maps an oriented edge (e, s, t) to $t - s$ and extends to all of O by linearity, factors uniquely through π, giving

$$\partial\colon \mathcal{E}(\Gamma) \to \mathcal{V}(\Gamma)$$

such that $\partial \circ \pi = f$: we call this ∂ the *boundary operator.*

Definition 3.3 (cycle space)**.** The *cycle space* of Γ is $\mathcal{C}(\Gamma) = \ker \partial$.

We can associate to any walk $W = (u_1, e_2, \dots, e_{n-1}, u_n)$ in Γ an element of $\mathcal{E}(\Gamma)$, $\sigma(W) = \sum_{i=1}^{n-1}(e_i, u_i, u_{i+1})$. We then have that $\partial(\sigma(W)) = \partial\left(\sum_{i=1}^{n-1}(e_i, u_i, u_{i+1})\right) = \sum_{i=1}^{n-1} f((e_i, u_i, u_{i+1})) = \sum_{i=1}^{n-1}(u_{i+1} - u_i) = u_n - u_1$: thus if W is closed, $\partial(\sigma(W)) = 0$.

Lemma 3.4. *Let $T = (V, E, \nu)$ be a tree with $|V| = n$, $E = \{e_1, \dots, e_{n-1}\}$, and $a_1, \dots, a_{n-1} \in R$. Suppose we select an orientation of T that assigns to each edge e_i an oriented edge (e_i, s_i, t_i): then if $\sum_{i=1}^{n-1} a_i(t_i - s_i) = 0$, $a_1 = \dots = a_{n-1} = 0$.*

Proof. The proof is by induction on the number n of vertices in T. If $n < 2$, there are no edges in the graph, and trivially all the non-existent a_i are zero. Suppose that the lemma is true for trees of size $n-1$. By Corollary 2.4, the vertices can be numbered $v_1, \dots, v_n$ so there is exactly one edge between v_n and an element of $\{v_1, \dots, v_{n-1}\}$; assume, without loss of generality, that edge is e_{n-1}. Since $\{v_1, \dots, v_n\} \subseteq \mathcal{V}(T)$ is linearly independent, the coefficient of v_n in the sum $0 = \sum_{i=1}^{n-1} a_i(t_i - s_i)$ must be 0: but since e_{n-1} is the only edge incident on v_n, the coefficient of v_n is either a_{n-1} or $-a_{n-1}$, depending on the orientation of e_{n-1}, and therefore $a_{n-1} = 0$. If we remove v_n and e_{n-1} from T, the result is a tree of $n-1$ vertices, and $\sum_{i=1}^{n-2} a_i(t_i - s_i) = 0$. Then by the induction hypothesis, $a_1 = \cdots = a_{n-1} = 0$. □

Theorem 3.5 ([4, Theo. 1.9.6])**.** *Let $\Gamma = (V, E, \nu)$ be a graph with n vertices, m edges, and c connected components. Then* $\operatorname{rank} \mathcal{C}(\Gamma) = m - n + c$.

Proof. Fix an orientation of Γ: this allows us to consider edges of Γ as elements of $\mathcal{E}(\Gamma)$. Each connected component is a connected subgraph, and by Proposition 2.6 has a spanning tree. Thus there is a family of trees $T_1 = (V_1, E_1, \nu_1), \dots, T_c = (V_c, E_c, \nu_c)$ with V the disjoint union of the V_i. Consider $B = \partial\left(\bigcup_{i=1}^{c} E_i\right)$: we will show this is a basis for $\operatorname{im} \partial$. As a consequence of Lemma 3.4, $\partial(E_i)$ is linearly independent for all i, and since there can be no vertices in common between edges in distinct connected components, the union B must be linearly independent.

To show B spans $\operatorname{im} \partial$, it is sufficient to show that the image of any edge under ∂ is a linear combination of elements of B. Evidently this is the case for all elements of E_i for some i. Let f be a chord of T_i in a connected component of Γ. Since T_i is a tree and therefore maximally acyclic, the graph $(V_i, E_i \cup \{f\})$ must have a cycle C_f, of which f must be one of the edges. Then $\partial(\sigma(C_f)) = 0$, and $\partial(f) = \partial(f - \sigma(C_f))$. But $f - \sigma(C_f)$ is a sum of (negations of) elements of E_T, and so $\partial(f)$ is a linear combination of elements of B.

Thus B is a basis for $\operatorname{im} \partial$, and $|B| = \sum_{i=1}^{c} |E_i| = \sum_{i=1}^{c}(|V_i| - 1) = \left(\sum_{i=1}^{c} |V_i|\right) - c = n - c$. Since R is a PID, by the rank-nullity theorem, $\operatorname{rank} \mathcal{C}(\Gamma) = \operatorname{rank} \ker \partial = \operatorname{rank} \mathcal{E}(\Gamma) - \operatorname{rank} \operatorname{im} \partial = m - n + c$. □

4. Infinite graphs and the cycle space

The edge and vertex spaces are free R-modules, and therefore contain only sums of finitely many vertices or edges. In the case of finite graphs, these are the only such sums that are possible, but for infinite graphs we might wish to consider sums that contain an infinite number of vertices or edges. We therefore define a new vertex space:

Definition 4.1 (locally finite vertex space)**.** Let $\Gamma = (V, E, \nu)$ be a graph. The *locally finite vertex space, $\mathcal{V}_R^\infty(\Gamma)$*, is the R-module

$$R^V = \prod_{v \in V} R.$$

The elements of the R-module R^V can be treated as functions $V \to R$, with addition and scalar multiplication defined pointwise, or as possibly infinite formal linear combinations $\sum_{v \in V} a_v v$ of vertices. There is also a new version of the edge space.

Definition 4.2 (locally finite edge space)**.** Let Γ be a graph, and O the set of oriented edges of Γ. The R-module $R^O = \prod_{(e,s,t)\in O} R$ has a submodule L consisting of all formal sums $\sum_{o\in O} a_o o$ such that for any vertex v in Γ, a_o is zero for all but finitely many oriented edges o incident on v. There is in turn a submodule $Q \subseteq L$ comprised of all formal sums $\sum_{o\in O} a_o o \in L$ such that for every oriented edge (e, s, t) of Γ, $a_{(e,s,t)} = a_{(e,t,s)}$. The *locally finite edge space* is the R-module $\mathcal{E}_R^\infty(\Gamma) = L/Q$.

If each vertex of Γ has only finitely many edges incident on it, then $L = R^O$ and $\mathcal{E}^\infty(\Gamma)$ is isomorphic to the space $R^E = \prod_{e\in E} R$; the isomorphism is not canonical, but every choice of orientation of Γ gives such an isomorphism. Such a graph is called *locally finite.*

If a graph Γ is finite, $\mathcal{V}_R(\Gamma)$ and $\mathcal{V}_R^\infty(\Gamma)$ are isomorphic, as are the edge spaces, but in the infinite case, this is not necessarily true. In fact, the locally finite vertex and edge spaces may not even be free R-modules. Consider the graph

given by $\Gamma_1 = (\mathbb{N}, \mathbb{N}, n \mapsto \{n, n+1\})$. Then $\mathcal{V}_\mathbb{Z}^\infty(\Gamma_1) = \mathbb{Z}^\mathbb{N}$ (also called the Baer-Specker group) is not free [2, Cor. 12.5].

We define a boundary operator $\partial\colon \mathcal{E}^\infty(\Gamma) \to \mathcal{V}^\infty(\Gamma)$ as before: there is a homomorphism $f\colon L \to \mathcal{V}^\infty(\Gamma)$, which maps an oriented edge (e, s, t) to $t - s$ and extends to the entire space by linearity. This homomorphism factors uniquely through $\pi\colon L \to \mathcal{E}^\infty(\Gamma)$, and ∂ is the unique homomorphism satisfying $f = \partial \circ \pi$. This explains why we needed to add the extra step of restricting consideration to L instead of R^O in the definition of $\mathcal{E}^\infty(\Gamma)$: the corresponding homomorphism $f\colon R^O \to \mathcal{V}^\infty(\Gamma)$ is not well defined for non-locally finite graphs. For example take the "infinite star" $\Gamma_2 = (\mathbb{N} \cup \{\bullet\}, \mathbb{N}, n \mapsto \{\bullet, n\})$: the sum $\sum_{n=1}^\infty (n, \bullet, n)$ is an element of R^O, but the coefficient of the vertex $\bullet$ in $f\left(\sum_{n=1}^\infty (n, \bullet, n)\right)$ is infinite. Restricting $\mathcal{E}^\infty(\Gamma)$ so only finitely many edges incident on one vertex have non-zero coefficient allows f to be well-defined, and is why this edge space is called locally finite. For a graph Γ, the *locally finite cycle space* is, as before, $\mathcal{C}_R^\infty(\Gamma) = \ker \partial$.

Remark 4.0.1. Another way of dealing with infinite sums would be to consider the boundary operator to be only defined for locally finite graphs.

The name "cycle space" is perhaps a misnomer in the infinite case: take for example the graph

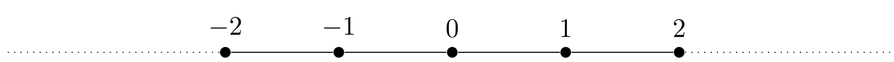

given by $\Gamma_3 = (\mathbb{Z}, \mathbb{Z}, n \mapsto \{n, n+1\})$. Then $\sum_{n\in\mathbb{Z}} (n, n, n+1)$ is an element of $\mathcal{E}^\infty(\Gamma_3)$, and since $\partial\left(\sum_{n\in\mathbb{Z}} (n, n, n+1)\right) = 0$, it is also an element of $\mathcal{C}^\infty(\Gamma_3)$, despite not being a cycle or sum of cycles in the traditional sense (in fact, Γ_3 is acyclic).

5. Homology and cohomology

Let R be a commutative ring, and suppose we have a sequence of R-modules M_i and homomorphisms d_i

$$\cdots \xrightarrow{d_3} M_2 \xrightarrow{d_2} M_1 \xrightarrow{d_1} M_0 \xrightarrow{d_0} M_{-1} \xrightarrow{d_{-1}} \cdots$$

such that $\forall i,\ d_i \circ d_{i+1} = 0$. Such a sequence is called a *chain complex of R-modules,* and denoted $(M_\bullet, d_\bullet)$, or simply $M_\bullet$. Because the composition of adjacent homomorphisms is trivial, $\operatorname{im} d_{i+1} \subseteq \ker d_i$. Therefore, for each module in the chain there is a quotient

$$H_i(M_\bullet) = \frac{\ker d_i}{\operatorname{im} d_{i+1}},$$

called the *i-th homology* of the complex.

Similarly to a chain complex, a *cochain complex* $(M^\bullet, d^\bullet)$ is a sequence of R-modules and homomorphisms

$$\cdots \xleftarrow{d^2} M^2 \xleftarrow{d^1} M^1 \xleftarrow{d^0} M^0 \xleftarrow{d^{-1}} M^{-1} \xleftarrow{d^{-2}} \cdots$$

such that $\forall i,\ d^i \circ d^{i-1} = 0$. Then $\operatorname{im} d^{i-1} \subseteq \ker d^i$, and we define the *$i$-th cohomology*

$$H^i(M^\bullet) = \frac{\ker d^i}{\operatorname{im} d^{i-1}}.$$

Chain complexes and cochain complexes are in fact precisely the same thing: a cochain complex can be turned into a chain complex by setting $M_i = M^{-i}$ and $d_i = d^{-i}$, and the ith homology module of the new chain complex is the $-i$th cohomology module of the cochain complex. Thus although the results in the remainder of this section are stated for chain complexes and homology, they also apply to cochain complexes and cohomology.

Chain complexes form a category $\mathsf{Ch}(R\text{-}\mathsf{Mod})$, in which an object is a chain complex of R-modules, and a morphism between $(M_\bullet, d_\bullet)$ and $(N_\bullet, e_\bullet)$ is a sequence of R-module homomorphisms $h_i \colon M_i \to N_i$ such that the diagram

$$\begin{array}{ccccccccccc}
\cdots & \xrightarrow{d_3} & M_2 & \xrightarrow{d_2} & M_1 & \xrightarrow{d_1} & M_0 & \xrightarrow{d_0} & M_{-1} & \xrightarrow{d_{-1}} & \cdots \\
 & & \downarrow{\scriptstyle h_2} & & \downarrow{\scriptstyle h_1} & & \downarrow{\scriptstyle h_0} & & \downarrow{\scriptstyle h_{-1}} & & \\
\cdots & \xrightarrow{e_3} & N_2 & \xrightarrow{e_2} & N_1 & \xrightarrow{e_1} & N_0 & \xrightarrow{e_0} & N_{-1} & \xrightarrow{e_{-1}} & \cdots
\end{array}$$

commutes. Composition and the identities are derived from composition and identity in R-Mod in the obvious way. Note that this means that a morphism in $\mathsf{Ch}(R\text{-}\mathsf{Mod})$ is an isomorphism if and only if the h_i are isomorphisms in R-Mod.

Lemma 5.1. *The category $\mathsf{Ch}(R\text{-}\mathsf{Mod})$ has products and coproducts; specifically, given a family of chain complexes $(M^\alpha_\bullet, d^\alpha_\bullet)$ indexed by $\alpha \in A$, the coproduct is the chain complex*

$$\cdots \xrightarrow{\oplus d_3^\alpha} \bigoplus_{\alpha\in A} M_2^\alpha \xrightarrow{\oplus d_2^\alpha} \bigoplus_{\alpha\in A} M_1^\alpha \xrightarrow{\oplus d_1^\alpha} \bigoplus_{\alpha\in A} M_0^\alpha \xrightarrow{\oplus d_0^\alpha} \bigoplus_{\alpha\in A} M_{-1}^\alpha \xrightarrow{\oplus d_{-1}^\alpha} \cdots$$

and the product the chain complex

$$\cdots \xrightarrow{\Pi d_3^\alpha} \prod_{\alpha\in A} M_2^\alpha \xrightarrow{\Pi d_2^\alpha} \prod_{\alpha\in A} M_1^\alpha \xrightarrow{\Pi d_1^\alpha} \prod_{\alpha\in A} M_0^\alpha \xrightarrow{\Pi d_0^\alpha} \prod_{\alpha\in A} M_{-1}^\alpha \xrightarrow{\Pi d_{-1}^\alpha} \cdots.$$

Proof. We will give the proof for the case of products; the proof for coproducts is analogous. Let the chain complex $P_\bullet$ be the candidate product in the statement of the lemma, and for all $a \in A$ define the morphism of chain complexes $\pi^a_\bullet \colon P_\bullet \to M^a_\bullet$ by letting π^a_i be the canonical surjection $\prod_{\alpha\in A} M_i^\alpha \to M_i^a$.

Let $(N_\bullet, e_\bullet)$ be a chain complex, and for every $\alpha \in A$, let $f^\alpha_\bullet \colon N_\bullet \to M^\alpha_\bullet$ be a morphism of chain complexes. Then for every i, there is a unique module homomorphism $\sigma_i \colon N_i \to \prod_{\alpha \in A} M_i^\alpha = P_i$ such that for all $\alpha \in A$, $f_i^\alpha = \pi_i^\alpha \circ \sigma_i$. Then $\sigma_\bullet$ is a morphism of chain complexes such that for all $\alpha \in A$, the diagram

$$\begin{array}{ccccccccc}
\cdots & \xrightarrow{e_{i+2}} & N_{i+1} & \xrightarrow{e_{i+1}} & N_i & \xrightarrow{e_i} & N_{i-1} & \xrightarrow{e_{i-1}} & \cdots \\
 & f_{i+1}^\alpha & \downarrow \sigma_{i+1} & f_i^\alpha & \downarrow \sigma_i & f_{i-1}^\alpha & \downarrow \sigma_{i-1} & & \\
\cdots & \xrightarrow{\Pi d_{i+2}^\alpha} & P_{i+1} & \xrightarrow{\Pi d_{i+1}^\alpha} & P_i^\alpha & \xrightarrow{\Pi d_i^\alpha} & P_{i-1} & \xrightarrow{\Pi d_{i-1}^\alpha} & \cdots \\
 & & \downarrow \pi_{i+1}^\alpha & & \downarrow \pi_i^\alpha & & \downarrow \pi_{i-1}^\alpha & & \\
\cdots & \xrightarrow{d_{i+2}^\alpha} & M_{i+1}^\alpha & \xrightarrow{d_{i+1}^\alpha} & M_i^\alpha & \xrightarrow{d_i^\alpha} & M_{i-1}^\alpha & \xrightarrow{d_{i-1}^\alpha} & \cdots
\end{array}$$

commutes, and thus $f_\bullet^\alpha = \pi_\bullet^\alpha \circ \sigma_\bullet$, and the uniqueness of $\sigma_\bullet$ is given by the uniqueness of the individual σ_i. □

Theorem 5.2 ([1, Lem. IX.3.4])**.** *For all $i \in \mathbb{Z}$, the mapping*

$$M_\bullet \mapsto H_i(M_\bullet)$$

defines a functor $\mathsf{Ch}(R\text{-}\mathsf{Mod}) \to R\text{-}\mathsf{Mod}$.

Proof. We need only show that this mapping induces a functorial mapping of morphisms between the two categories. Let $h_\bullet \colon M_\bullet \to N_\bullet$ be a morphism in $\mathsf{Ch}(R\text{-}\mathsf{Mod})$, and fix i. Consider the diagram

$$\begin{array}{ccccc}
M_{i+1} & \xrightarrow{d_{i+1}^M} & M_i & \xrightarrow{d_i^M} & M_{i-1} \\
\downarrow h_{i-1} & & \downarrow h_i & & \downarrow h_{i+1} \\
N_{i+1} & \xrightarrow{d_{i+1}^N} & N_i & \xrightarrow{d_i^N} & N_{i-1}
\end{array}$$

which commutes, since $h_\bullet$ is a homomorphism of chain complexes. Then $\ker(d_i^N \circ h_i) \subseteq \ker d_i^M$ (since $d_i^N \circ h_i = h_{i+1} \circ d_i^M$), and therefore $h_i(\ker d_i^M) \subseteq \ker d_i^N$. So h_i can be restricted to a homomorphism $\ker d_i^M \to \ker d_i^N$, and composing with the projection $\ker d_i^N \to \frac{\ker d_i^N}{\operatorname{im} d_{i-1}^N} = H_i(N_\bullet)$ gives a homomorphism $g \colon \ker d_i^M \to H_i(N_\bullet)$.

Again by commutativity of the above diagram, $h_i(\operatorname{im} d_{i+1}^M) \subseteq \operatorname{im} d_{i+1}^N$, and $g(\operatorname{im} d_{i+1}^M) = \{0\}$. Therefore h_i induces (via g) a homomorphism $H_i(M_\bullet) \to H_i(N_\bullet)$. □

Corollary 5.3. *If $M_\bullet \cong N_\bullet$ as chain complexes, then for all i,*

$$H_i(M_\bullet) \cong H_i(N_\bullet)$$

Lemma 5.4. *The functor $H_i \colon$ $\mathsf{Ch}(R\text{-}\mathsf{Mod}) \to R\text{-}\mathsf{Mod}$ preserves products and coproducts.*

Proof. Let $M_\bullet^\alpha$ be a family of chain complexes indexed by $\alpha \in A$. Consider the coproduct in $\mathsf{Ch}(R\text{-}\mathsf{Mod})$,

$$\cdots \xrightarrow{\oplus d_{i+2}^\alpha} \bigoplus_{\alpha \in A} M_{i+1}^\alpha \xrightarrow{\oplus d_{i+1}^\alpha} \bigoplus_{\alpha \in A} M_i^\alpha \xrightarrow{\oplus d_i^\alpha} \bigoplus_{\alpha \in A} M_{i-1}^\alpha \xrightarrow{\oplus d_{i-1}^\alpha} \cdots.$$

Let $x \in \bigoplus_{\alpha\in A} M_i^\alpha$. Then there exist $x^\alpha \in M_i^\alpha$ such that $x = \sum_{\alpha\in A} x^\alpha$ (and all but finitely many of the x^α are 0). By construction,

$$\oplus d_i^\alpha(x) = \oplus d_i^\alpha\left(\sum_{\alpha\in A} x^\alpha\right) = \sum_{\alpha\in A} d_i^\alpha(x^\alpha),$$

and since each of the $d_i^\alpha(x^\alpha)$ is in a different M_{i-1}^α, it follows that $\oplus d_i^\alpha(x) = 0$ if and only if $\forall\alpha \in A. d_i^\alpha(x^\alpha) = 0$. Thus $\ker \oplus d_i^\alpha = \bigoplus_{\alpha\in A} \ker d_i^\alpha$.

Similarly, if $y \in \operatorname{im} \oplus d_i^\alpha$, then there are $x^\alpha \in M_{i+1}^\alpha$ (all but finitely many equal to zero) such that

$$y = \oplus d_{i+1}^\alpha\left(\sum_{\alpha\in A} x^\alpha\right) = \sum_{\alpha\in A} d_{i+1}^\alpha(x^\alpha),$$

and so $\operatorname{im} \oplus d_{i+1}^\alpha = \bigoplus_{\alpha\in A} \operatorname{im} d_{i+1}^\alpha$. Therefore,

$$H_i\left(\coprod_{\alpha\in A} M_\bullet^\alpha\right) = \frac{\ker \oplus d_i^\alpha}{\operatorname{im} \oplus d_{i+1}^\alpha} = \frac{\bigoplus_{\alpha\in A} \ker d_i^\alpha}{\bigoplus_{\alpha\in A} \operatorname{im} d_{i+1}^\alpha} = \bigoplus_{\alpha\in A} \frac{\ker d_i^\alpha}{\operatorname{im} d_{i+1}^\alpha} = \bigoplus_{\alpha\in A} H_i(M_\bullet^\alpha)$$

A similar argument shows that the functor preserves products. □

6. Homology of graphs

For any graph Γ, we can construct a very simple chain complex $\Gamma_\bullet$ on any principal ideal domain R:

$$\cdots \xrightarrow{d_3=0} 0 \xrightarrow{d_2=0} \mathcal{E}_R(\Gamma) \xrightarrow{d_1=\partial} \mathcal{V}_R(\Gamma) \xrightarrow{d_0=0} 0 \xrightarrow{d_{-1}=0} \cdots,$$

where the only non-trivial homomorphism is the boundary operator, and the only non-trivial modules are the edge and vertex spaces. The homology of this chain complex is called the *ordinary homology* of the graph. All of the homology modules are trivial except possibly $H_1(\Gamma_\bullet)$ and $H_0(\Gamma_\bullet)$. Since the image of $d_2 = 0$ is trivial, $H_1(\Gamma_\bullet) = \ker \partial = \mathcal{C}_R(\Gamma)$. In the case of $H_0(\Gamma_\bullet)$ we have the following lemma:

Lemma 6.1. *For any graph Γ with a finite number of vertices and c connected components,* $\operatorname{rank} H_0(\Gamma_\bullet) = c$.

Proof. Since d_0 is trivial, $\ker d_0 = \mathcal{V}(\Gamma)$, and thus $H_0(\Gamma_\bullet) = \frac{\mathcal{V}(\Gamma)}{\operatorname{im}\partial}$. As seen in the proof of Theorem 3.5, $\operatorname{rank} \operatorname{im} \partial = n - c$, where n is the number of vertices of Γ and c the number of connected components. Thus $\operatorname{rank} H_0(\Gamma_\bullet) = \operatorname{rank} \mathcal{V}(\Gamma) - \operatorname{rank} \operatorname{im} \partial = c$. □

Allowing for infinite graphs, Theorem 3.5 and Lemma 6.1 generalize to the following two theorems.

Theorem 6.2 ([5, Prop. 2.7.1])**.** *Let Γ be a connected graph with at least one vertex. Then* $\operatorname{rank} H_0(\Gamma_\bullet) = 1$.

Proof. Consider the homomorphism $f\colon \mathcal{V}(\Gamma) \to R$ which maps every vertex of Γ to 1. Since Γ has at least one vertex, f is surjective. If (e, s, t) is a directed edge of Γ, then $f(\partial((e, s, t))) = f(t - s) = 0$, and since directed edges span $\mathcal{E}(\Gamma)$, it follows that $\operatorname{im} \partial \subseteq \ker f$. Suppose a linear combination $x = \sum_{i=1}^n a_i v_i$ of vertices is an element of $\ker f$. Then $\sum_{i=1}^n a_i = 0$, and therefore $a_1 = -\sum_{i=2}^n a_i$. Since Γ is connected, there exists a v_1–v_i walk W_i for $1 < i \leq n$, and an element of $\mathcal{E}(\Gamma)$, $\sigma(W_i)$ corresponding to W_i, with $\partial(\sigma(W_i)) = v_i - v_1$. Thus $x = \sum_{i=1}^n a_i v_i =$

$\left(-\sum_{i=2}^n a_i\right) v_1 + \sum_{i=2}^n a_i v_i = \sum_{i=2}^n a_i(v_i - v_1) = \sum_{i=2}^n a_i \partial(\sigma(W_i)) \in \operatorname{im}\partial$. Thus $\ker f = \operatorname{im}\partial$, and there is an exact sequence

$$\mathcal{E}(\Gamma) \xrightarrow{\partial} \mathcal{V}(\Gamma) \xrightarrow{f} R \to 0,$$

making $H_0(\Gamma_\bullet) = \frac{\mathcal{V}(\Gamma)}{\operatorname{im}\partial} \cong R$. □

Theorem 6.3 ([5, Prop. 2.7.2])**.** *Let Γ be a graph, and $\{\Gamma^\alpha \mid \alpha \in A\}$ its set of connected components. Then for all n,*

$$H_n(\Gamma_\bullet) \cong \bigoplus_{\alpha \in A} H_n(\Gamma^\alpha_\bullet)$$

Proof. We will show that $\Gamma_\bullet \cong \coprod_{\alpha\in A} \Gamma^\alpha_\bullet$ as chain complexes; then by Corollary 5.3 and Lemma 5.4,

$$H_i(\Gamma_\bullet) \cong H_i\left(\coprod_{\alpha\in A} \Gamma^\alpha_\bullet\right) \cong \bigoplus_{\alpha\in A} H_i(\Gamma^\alpha_\bullet).$$

Every edge in Γ appears in exactly one connected component; this implies that there is a function $\varphi_E \colon \bigoplus_{\alpha\in A} \mathcal{E}(\Gamma^\alpha) \to \mathcal{E}(\Gamma)$ that maps an oriented edge (e, s, t) of $\mathcal{E}(\Gamma^\alpha)$ to the same oriented edge (e, s, t) in $\mathcal{E}(\Gamma)$, and extends over the whole space by linearity, and φ_E is an isomorphism. Each vertex in Γ is also in only one connected component, so there is a similarly constructed isomorphism $\varphi_V \colon \bigoplus_{\alpha\in A} \mathcal{V}(\Gamma^\alpha) \to \mathcal{V}(\Gamma)$. Each connected component of Γ has a boundary operator $\partial^\alpha \colon \mathcal{E}(\Gamma^\alpha) \to \mathcal{V}(\Gamma^\alpha)$; these can be combined to form $\oplus\partial^\alpha \colon \bigoplus_{\alpha\in A} \mathcal{E}(\Gamma^\alpha) \to \bigoplus_{\alpha\in A} \mathcal{V}(\Gamma^\alpha)$, which maps elements of each Γ^α by ∂^α, and extends to the whole direct sum by linearity.

We show that the following diagram

$$\begin{array}{ccc} \bigoplus_{\alpha\in A} \mathcal{E}(\Gamma^\alpha) & \xrightarrow{\oplus\partial^\alpha} & \bigoplus_{\alpha\in A} \mathcal{V}(\Gamma^\alpha) \\ \downarrow{\scriptstyle\varphi_E} & & \downarrow{\scriptstyle\varphi_V} \\ \mathcal{E}(\Gamma) & \xrightarrow{\partial} & \mathcal{V}(\Gamma) \end{array}$$

commutes. Let (e, s, t) be an oriented edge of Γ: it falls in exactly one connected component Γ^{α_0}. Then

$$\begin{aligned} \varphi_V(\oplus\partial^\alpha((e,s,t))) &= \varphi_V(\partial^{\alpha_0}((e,s,t))) \\ &= \varphi_V(t) - \varphi_V(s) \\ &= t - s \\ &= \partial((e,s,t)) \\ &= \partial(\varphi_E((e,s,t))) \end{aligned}$$

and since oriented edges span $\bigoplus_{\alpha\in A} \mathcal{E}(\Gamma^\alpha)$, $\varphi_V \circ \oplus\partial^\alpha = \partial \circ \varphi_E$.

Since φ_E and φ_V are isomorphisms, the diagram shows that $\Gamma_\bullet$ is isomorphic to the complex $\coprod_{\alpha\in A} \Gamma^\alpha_\bullet$. □

By replacing each of the modules by their locally finite analogue, we obtain the locally finite chain complex $\Gamma^\infty_\bullet$

$$\cdots \xrightarrow{d_3=0} 0 \xrightarrow{d_2=0} \mathcal{E}^\infty_R(\Gamma) \xrightarrow{d_1=\partial} \mathcal{V}^\infty_R(\Gamma) \xrightarrow{d_0=0} 0 \xrightarrow{d_{-1}=0} \cdots,$$

with homology modules $H_1(\Gamma^\infty_\bullet) = \mathcal{C}^\infty(\Gamma)$ and $H_0(\Gamma^\infty_\bullet)$; these are called the *locally finite homology* of the graph.

For finite graphs, the locally finite homology is the same as the ordinary homology, since the locally finite vertex and edge spaces are isomorphic to the ordinary vertex and edge spaces. For infinite graphs, the locally finite versions of Lemma 6.2 and Theorem 6.3 are:

Lemma 6.4 ([5, Prop. 11.1.3])**.** *Let Γ be an infinite and locally finite graph. If Γ is connected, then $H_0(\Gamma^\infty_\bullet) = 0$.*

Proof. Let $\Gamma = (V, E, \nu)$ be infinite, locally finite, and connected. We define the distance between any two vertices u and v of Γ to be the length of the shortest path between them. Then the distance is a metric on V. For a vertex v, $B_n(v)$ will denote the *closed* ball of radius n around v in the distance metric. If Δ is a subgraph of Γ, the complementary components of Δ are the connected components of Γ without the vertices of Δ and the edges incident on a vertex of Δ. If Δ is finite, at least one of its complementary components is unbounded in the distance metric.

Let y be an arbitrary vertex of Γ, and let Y_n be the union of $B_n(y)$ and the bounded complementary components of $B_n(y)$, for $n \in \mathbb{N}$. Then each Y_n is bounded (and thus finite) and connected. Pick a strictly increasing sequence of natural numbers $\{n_k\}_{k=1}^\infty$ such that $Y_{n_k} \subseteq B_{n_{k+1}}(y) \subseteq Y_{n_{k+1}}$. We say a vertex v has index n_k for the least n_k such that v is in Y_{n_k}, and the index of an edge is the minimum of the indices of its endpoints. Suppose that v is a vertex of index n_k; then there is a path from v to a vertex of index n_{k+1}. Repeating this indefinitely gives an infinite path, and this a sequence of oriented edges $\{e_j\}_{j=1}^\infty$, that starts at v and such that, for each m, all but finitely many e_j are not in Y_{n_m}. Then let $c_v = \sum_{j=1}^\infty e_j$. Clearly $\partial(c_v) = v$; moreover, for any edge e of Γ, there are only finitely many v such that an orientation of e is in c_v, since such an orientation appears in c_v only if the index of v is less than or equal to the index of e. Thus for any $d = \sum_{\alpha \in A} m_\alpha v_\alpha \in \mathcal{V}^\infty(\Gamma)$, $\partial\left(\sum_{\alpha \in A} m_\alpha c_{v_\alpha}\right) = \sum_{\alpha \in A} m_\alpha \partial(c_{v_\alpha}) = d$, and thus $\partial\colon \mathcal{E}^\infty(\Gamma) \to \mathcal{V}^\infty(\Gamma)$ is surjective, and thus $\frac{\mathcal{V}^\infty(\Gamma)}{\operatorname{im} \partial} = 0$. □

Note that the above lemma requires that the graph be locally finite: if it is not, a finite subgraph could have no complementary components that are unbounded, and therefore there would be no infinite sequences of edges that pass through each vertex only finitely many times. For example, the infinite star mentioned before, $\Gamma = (\mathbb{N} \cup \{\bullet\}, \mathbb{N}, n \mapsto \{n, \bullet\})$, is infinite and connected, but not locally finite, and $\operatorname{im} \partial$ is a strict subset of $\mathcal{V}^\infty(\Gamma)$ (for example, $\bullet \notin \operatorname{im} \partial$), meaning $H_0(\Gamma^\infty_\bullet) \neq 0$.

Theorem 6.5 ([5, Prop. 11.1.5])**.** *Let Γ be a graph, and $\{\Gamma_\alpha \mid \alpha \in A\}$ its set of connected components. Then for all n,*

$$H_n(\Gamma^\infty_\bullet) \cong \prod_{\alpha \in A} H_n({\Gamma_\alpha}^\infty_\bullet)$$

Proof. We will show that $\Gamma^\infty_\bullet \cong \prod_{\alpha \in A} {\Gamma_\alpha}^\infty_\bullet$ as chain complexes; then by Corollary 5.3 and Lemma 5.4,

$$H_i(\Gamma^\infty_\bullet) \cong H_i\left(\prod_{\alpha \in A} {\Gamma_\alpha}^\infty_\bullet\right) \cong \prod_{\alpha \in A} H_i({\Gamma_\alpha}^\infty_\bullet).$$

We can consider $x \in \prod_{\alpha \in A} \mathcal{E}^\infty(\Gamma_\alpha)$ as a possibly infinite sum $\sum_{\alpha \in A} x_\alpha$, where $x_\alpha \in \mathcal{E}^\infty(\Gamma_\alpha)$, and an element of $\mathcal{E}^\infty(\Delta)$ (Δ equal to either Γ or Γ_α for some α) as a possibly infinite sum $\sum_{e \in E} a_e e$, where E is the set of edges of Δ, and $a_e \in R$ for all e. Each edge e of Γ lies in exactly one connected component Γ_{α_e}. Then we can define a function $\varphi_E^\infty \colon \prod_{\alpha \in A} \mathcal{E}^\infty(\Gamma_\alpha) \to \mathcal{E}^\infty(\Gamma)$ by

$$\sum_{\alpha \in A} \sum_{e \in E_\alpha} a_e e \mapsto \sum_{e \in E} a_e e$$

Then φ_E^∞ is evidently linear, and is an isomorphism. We can similarly construct an isomorphism $\varphi_V^\infty \colon \prod_{\alpha \in A} \mathcal{V}^\infty(\Gamma_\alpha) \to \mathcal{V}^\infty(\Gamma)$.

We will show that the following diagram

$$\begin{array}{ccc} \prod_{\alpha \in A} \mathcal{E}^\infty(\Gamma_\alpha) & \xrightarrow{\Pi \partial_\alpha} & \prod_{\alpha \in A} \mathcal{V}^\infty(\Gamma_\alpha) \\ \downarrow \varphi_E^\infty & & \downarrow \varphi_V^\infty \\ \mathcal{E}^\infty(\Gamma) & \xrightarrow{\partial} & \mathcal{V}^\infty(\Gamma) \end{array}$$

commutes, which implies that $\Gamma_\bullet^\infty \cong \prod_{\alpha \in A} \Gamma_{\alpha}{}_\bullet^\infty$. Let $x = \sum_{\alpha \in A} \sum_{(s,t) \in E_\alpha} a_{(s,t)}(s,t) \in \prod_{\alpha \in A} \mathcal{E}^\infty(\Gamma_\alpha)$. Then

$$\begin{aligned} (\varphi_V^\infty \circ \Pi \partial_\alpha)(x) &= \varphi_V^\infty \left(\sum_{\alpha \in A} \sum_{(s,t) \in E_\alpha} a_{(s,t)} \partial_\alpha((s,t)) \right) \\ &= \sum_{(s,t) \in E} a_{(s,t)}(t - s) \\ &= \sum_{(s,t) \in E} a_{(s,t)} \partial((s,t)) \\ &= \partial \left(\varphi_E^\infty \left(\sum_{\alpha \in A} \sum_{(s,t) \in E_\alpha} a_{(s,t)}(s,t) \right) \right) \\ &= (\partial \circ \varphi_E^\infty)(x) \quad \square \end{aligned}$$

We have the following invariants of graphs derived from the homology:

Definition 6.6 (Betti numbers)**.** For any $n \in \mathbb{N}$, the *nth Betti number* of a graph Γ is $\operatorname{rank} H_n(\Gamma_\bullet)$. All but the first and zeroth Betti numbers are zero. The first Betti number is also called the *cyclomatic number.*

Definition 6.7 (Euler characteristic)**.** The *Euler characteristic* of a graph Γ is the alternating sum of the Betti numbers; i.e. $\chi(\Gamma) = \operatorname{rank} H_0(\Gamma_\bullet) - \operatorname{rank} H_1(\Gamma_\bullet)$, if both ranks exist and are finite. It is equal to (for finite graphs) the number of vertices minus the number of edges, which follows from Theorem 3.5 and Lemma 6.1.

7. Dual modules

Given a module M over a ring R, there is a natural way of deriving a new module from M, called its *dual.*

Definition 7.1 (dual module)**.** Let R be a commutative ring; the *dual module* of an R-module M is the module $M^\vee = \mathrm{Hom}_R(M, R)$, the set of R-module homomorphisms from M to R with addition and scalar multiplication defined pointwise. A homomorphism of R-modules $f: M \to N$ induces a *dual homomorphism* $f^\vee : N^\vee \to M^\vee$, $h \mapsto h \circ f$.

In particular, the dual of a zero module is a zero module, and of a zero homomorphism is a zero homomorphism. In addition, the dual of a direct sum of modules is the product of the duals of the modules. In particular, we have that

Lemma 7.2 ([1, Cor. VIII.5.7])**.** *Let S be a set, $R^{\oplus S}$ the free R-module over S, and R^S the module of all functions $S \to R$ with addition and scalar multiplication defined point wise. Then*

$$(R^{\oplus S})^\vee \cong R^S$$

Proof. Let $\phi\colon (R^{\oplus S})^\vee \to R^S$ be defined by $\phi(g)(s) = g(s)$. This is clearly linear. Suppose $g \in (R^{\oplus S})^\vee$ such that $\phi(g) = 0$; then for all $s \in S$, $g(s) = 0$. Since S is a basis for $R^{\oplus S}$, it follows that $g = 0$, and $\ker \phi = \{0\}$. Thus ϕ is injective. Let $f: S \to R$; since S is a basis for $R^{\oplus R}$, f can be extended (uniquely, in fact) to a homomorphism $g_f : R^{\oplus S} \to R$. Then $g_f \in (R^{\oplus S})^\vee$, and $\phi(g_f) = F$. Thus ϕ is surjective. ∎

If S is finite, then $R^S \cong R^{\oplus S}$, and so by the above lemma free modules of finite rank are isomorphic to their duals. In the specific case of the vertex and edge spaces, we have the following:

Corollary 7.3. *Let Γ be a graph. Then $\mathcal{V}(\Gamma)^\vee \cong \mathcal{V}^\infty(\Gamma)$, and if Γ is locally finite, $\mathcal{E}(\Gamma)^\vee \cong \mathcal{E}^\infty(\Gamma)$.*

8. Cohomology of graphs

For any graph Γ there is a cochain complex $\Gamma^\bullet$ defined in terms of the duals of $\Gamma_\bullet$. We define $\Gamma^\bullet$ to be the cochain complex

$$\cdots \xleftarrow{d^3=0} 0 \xleftarrow{d^2=0} \mathcal{E}_R(\Gamma)^\vee \xleftarrow{d^1=\partial^\vee} \mathcal{V}_R(\Gamma)^\vee \xleftarrow{d^0=0} 0 \xleftarrow{d^{-1}=0} \cdots$$

which is defined by taking the dual at each point of the chain complex $\Gamma_\bullet$. The cohomology modules of the above complex are called the *ordinary cohomology,* and the homomorphism $\partial^\vee$ the *coboundary operator.* Intuitively, the coboundary operator takes every vertex to the sum of all of the edges incident on that vertex. We obtain the following equivalents to Theorem 3.5, Lemma 6.2, and Theorem 6.3, respectively:

Lemma 8.1 ([5, Prop. 12.1.2])**.** *Let Γ be connected. Then* $\operatorname{rank} H^0(\Gamma^\bullet) = 1$.

Proof. Consider an element of $\mathcal{V}(\Gamma)^\vee$ as a homomorphism $f: \mathcal{V}(\Gamma) \to R$, and let u and v be two vertices in Γ. Since Γ is connected, there is a u–v path ($u = a_1, e_1, \ldots, e_{n-1}, a_n = v$) in Γ. Suppose that $f \in \ker \partial^\vee$; we will show by induction on n that $f(u) = f(v)$. If $n = 1$, u and v are the same vertex. Suppose $n > 1$; by the induction hypothesis, $f(u) = f(a_{n-1})$: then there is an edge e_{n-1} from a_{n-1} to v, and $\partial(e_{n-1}) = v - a_{n-1}$. Thus $\partial^\vee(f)(e_{n-1}) = (f \circ \partial)(e_{n-1}) = f(v - a_{n-1}) = f(v) - f(a_{n-1})$; but since $f \in \ker \partial^\vee$, $\partial^\vee(f)(e_{n-1}) = 0$, and $f(v) = f(a_{n-1}) = f(u)$. Since the vertices form a basis for $\mathcal{V}(\Gamma)$, this means that $\ker \partial^\vee$ is generated by the homomorphism that maps every vertex to 1, and $\operatorname{rank} H^0(\Gamma^\bullet) = 1$. ∎

Lemma 8.2. *Let* Γ *be a finite connected graph, with* n *vertices and* m *edges. Then* $\operatorname{rank} H^1(\Gamma^\bullet) = m - n + 1$.

Proof. Since Γ is finite, the duals of the vertex and edge spaces are isomorphic to the vertex and edge spaces themselves, and thus $\operatorname{rank} \mathcal{V}(\Gamma)^\vee = n$ and $\operatorname{rank} \mathcal{E}(\Gamma)^\vee = m$. By Lemma 8.1, $\operatorname{rank} \ker \partial^\vee = 1$. Thus by the rank-nullity theorem, $\operatorname{rank} \operatorname{im} \partial^\vee = \operatorname{rank} \mathcal{V}(\Gamma)^\vee - \operatorname{rank} \ker \partial^\vee = n - 1$, and $\operatorname{rank} H^1(\Gamma^\bullet) = \operatorname{rank} \frac{\mathcal{E}(\Gamma)^\vee}{\operatorname{im} \partial^\vee} = \operatorname{rank} \mathcal{E}(\Gamma)^\vee - \operatorname{rank} \operatorname{im} \partial^\vee = m - n + 1$. □

Theorem 8.3 ([5, Prop. 12.1.3])**.** *Let* Γ *be a graph, and* $\{\Gamma_\alpha \mid \alpha \in A\}$ *its connected components. Then for all* n*,*

$$H^n(\Gamma^\bullet) \cong \prod_{\alpha \in A} H^n(\Gamma_\alpha^\bullet)$$

Proof. We will show that $\Gamma^\bullet \cong \prod_{\alpha \in A} \Gamma_\alpha^\bullet$ as chain complexes; then by Corollary 5.3 and Lemma 5.4,

$$H^n(\Gamma^\bullet) \cong H^n\left(\prod_{\alpha \in A} \Gamma_\alpha^\bullet\right) \cong \prod_{\alpha \in A} H^n(\Gamma_\alpha^\bullet).$$

Since no edges or vertices are in more than one connected component, a homomorphism with domain $\mathcal{E}(\Gamma)$ or $\mathcal{V}(\Gamma)$ can be uniquely divided into homomorphisms with domains $\mathcal{E}(\Gamma_\alpha)$ and $\mathcal{V}(\Gamma_\alpha)$ for all $\alpha \in A$; let $\varphi_E \colon \prod_{\alpha \in A} \mathcal{E}(\Gamma_\alpha)^\vee \to \mathcal{E}(\Gamma)^\vee$ and $\varphi_V \colon \prod_{\alpha \in A} \mathcal{V}(\Gamma_\alpha)^\vee \to \mathcal{V}(\Gamma)^\vee$ be the isomorphisms produced by this division. We will show that the following diagram

$$\begin{array}{ccc} \prod_{\alpha \in A} \mathcal{E}(\Gamma_\alpha)^\vee & \xleftarrow{\Pi \partial_\alpha^\vee} & \prod_{\alpha \in A} \mathcal{V}(\Gamma_\alpha)^\vee \\ \downarrow \varphi_E & & \downarrow \varphi_V \\ \mathcal{E}(\Gamma)^\vee & \xleftarrow[\partial^\vee]{} & \mathcal{V}(\Gamma)^\vee \end{array}$$

commutes, which implies that $\Gamma^\bullet \cong \prod_{\alpha \in A} \Gamma_\alpha^\bullet$. Let $x \in \prod_{\alpha \in A} \mathcal{V}(\Gamma_\alpha)^\vee$ be the product of homomorphisms $f_\alpha \colon \mathcal{V}(\Gamma_\alpha) \to R$ for all $\alpha \in A$. Then for every edge e of Γ, supposing the connected component containing e is Γ_{α_e},

$$\begin{aligned} (\varphi_E \circ \Pi \partial_\alpha^\vee)(x)(e) &= \left(\sum_{\alpha \in A} \varphi_E(\partial_\alpha^\vee(f_\alpha))\right)(e) \\ &= \sum_{\alpha \in A} \varphi_E(f_\alpha \circ \partial_\alpha)(e) \\ &= f_{\alpha_e}(\partial_{\alpha_e}(e)) \\ &= f_{\alpha_e}(\partial(e)) \\ &= \partial^\vee(\varphi_V(f_{\alpha_e}))(e) \\ &= \partial^\vee\left(\varphi_V\left(\sum_{\alpha \in A} f_\alpha\right)\right)(e) \\ &= (\partial^\vee \circ \varphi_V)(x)(e) \end{aligned}$$

And since edges form a basis for $\mathcal{E}(\Gamma)$, $(\varphi_E \circ \Pi \partial_\alpha^\vee)(x) = (\partial^\vee \circ \varphi_V)(x)$ for all such x. □

There is another cohomology which can be expressed in terms of duals, but in this case of duals that are finite. Let $\mathcal{E}_R(\Gamma)^f$ be the R-module of homomorphisms $\mathcal{E}_R(\Gamma) \to R$ that are zero for all but finitely many edges, and $\mathcal{V}_R(\Gamma)^f$ the R-module of homomorphisms $\mathcal{V}_R(\Gamma) \to R$ that are zero for all but finitely many vertices. Obviously $\mathcal{V}(\Gamma)^f \subseteq \mathcal{V}(\Gamma)^\vee$ and $\mathcal{E}(\Gamma)^f \subseteq \mathcal{E}(\Gamma)^\vee$, and the restriction of $\partial^\vee$ to $\mathcal{V}(\Gamma)^\vee$ has image contained in $\mathcal{E}(\Gamma)^f$. There is a cochain complex $\Gamma_f^\bullet$:

$$\cdots \xleftarrow{d^3=0} 0 \xleftarrow{d^2=0} \mathcal{E}_R(\Gamma)^f \xleftarrow{d^1=\partial^\vee} \mathcal{V}_R(\Gamma)^f \xleftarrow{d^0=0} 0 \xleftarrow{d^{-1}=0} \cdots .$$

The cohomology of this complex is called the *cohomology with compact support,* since the elements of $\mathcal{E}(\Gamma)^f$ or $\mathcal{V}(\Gamma)^f$ are those homomorphisms with support on the edges or vertices of a finite (and therefore compact) subgraph of Γ.

For finite graphs, the cohomology with compact support is the same as the ordinary cohomology, since every homomorphism from the vertex or edge space to R is zero for all but finitely many vertices or edges, and thus the finite duals are equal to the normal duals. For infinite graphs, we obtain the following equivalents to results for the ordinary cohomology:

Lemma 8.4 ([5, Prop. 12.1.2])**.** *Let* Γ *be infinite and connected. Then* $\operatorname{rank} H^0(\Gamma_f^\bullet) = 0$.

Proof. An element $f \in \mathcal{V}(\Gamma)^f$ is a homomorphism $\mathcal{V}(\Gamma) \to R$, where $f(v)$ is zero for all but finitely many vertices v. Let u and v be vertices of Γ. Since Γ is connected there is a u–v path in Γ, and by the same argument as in Lemma 8.1, if $f \in \ker \partial^\vee$, $f(u) = f(v)$. So if $f \in \ker \partial^\vee$, the value of f is the same for all vertices, and since there are infinitely many vertices and finitely many non-zero $f(v)$, they must all be zero, and $f = 0$. Thus $\operatorname{rank} H^0(\Gamma_f^\bullet) = \operatorname{rank} \ker \partial^\vee = \operatorname{rank}\{0\} = 0$. □

Theorem 8.5 ([5, Prop. 12.1.3])**.** *Let* Γ *be a graph, and* $\{\Gamma_\alpha \mid \alpha \in A\}$ *its connected components. Then for all* n*,*

$$H^n(\Gamma_f^\bullet) \cong \bigoplus_{\alpha \in A} H^n(\Gamma_{\alpha}{}_f^\bullet)$$

9. Homology and cohomology as functors

Homology and cohomology define functors from the category Graph (or $\mathsf{FinGraph}$) into the category of R-modules, R-Mod. The functors are covariant for homology and contravariant for cohomology.

Lemma 9.1. *The mapping from a graph to its vertex or edge space is a functor* $\mathsf{Graph} \to R\text{-}\mathsf{Mod}$.

Proof. Since a graph homomorphism is just a function from the vertex set of one graph to another that satisfies certain properties, there is a forgetful functor $\mathscr{V}\colon \mathsf{Graph} \to \mathsf{Set}$ that maps graphs to their vertex sets and graph homomorphisms to their underlying function. There is also a free functor $\mathscr{F}_R\colon \mathsf{Set} \to R\text{-}\mathsf{Mod}$ that maps a set to a free R-module with that set as a basis. Then $\mathscr{F}_R \circ \mathscr{V}\colon \mathsf{Graph} \to R\text{-}\mathsf{Mod}$ is a functor that maps a graph Γ to the free R-module over the vertex set of Γ, that is to $\mathcal{V}(\Gamma)$.

There is another functor $\mathscr{E}\colon \mathsf{Graph} \to \mathsf{Set}$, which maps a graph to its edge set, and a graph homomorphism f to the function between edge sets $f_E\colon E \to E'$, $\{u, v\} \mapsto \{f(u), f(v)\}$ (such an edge must exist in E', by the definition of a graph homomorphism, and must be unique, since there are no parallel edges). Then $\mathscr{E}$ preserves the identity and respects composition, and is thus a functor. Then $\mathscr{F}_R \circ \mathscr{E} \colon \mathsf{Graph} \to R\text{-}\mathsf{Mod}$ is a functor that maps a graph Γ to the free R-module over its edge set, that is to $\mathcal{E}(\Gamma)$. □

Lemma 9.2. *The boundary operator ∂ is a natural transformation from the edge to the vertex space.*

Proof. We need only to show that for any graphs Γ and Δ and a graph homomorphism $f \colon \Gamma \to \Delta$, the diagram

$$\begin{array}{ccc} \mathcal{E}(\Gamma) & \xrightarrow{\mathcal{E}(f)} & \mathcal{E}(\Delta) \\ \downarrow{\scriptstyle\partial} & & \downarrow{\scriptstyle\partial} \\ \mathcal{V}(\Gamma) & \xrightarrow{\mathcal{V}(f)} & \mathcal{V}(\Delta) \end{array}$$

commutes. Let (e, s, t) be an oriented edge of Γ: we will show that $(\mathcal{V}(f) \circ \partial)((e,s,t)) = (\partial \circ \mathcal{E}(f))((e,s,t))$. Since the oriented edges span $\mathcal{E}(\Gamma)$, this shows that $\mathcal{V}(f) \circ \partial = \partial \circ \mathcal{E}(f)$.

$$\begin{aligned} (\mathcal{V}(f) \circ \partial)((e,s,t)) &= \mathcal{V}(f)(\partial((e,s,t))) \\ &= \mathcal{V}(f)(t) - \mathcal{V}(f)(s) \\ &= f(t) - f(s) \\ &= \partial((f(e), f(s), f(t))) \\ &= \partial(\mathcal{E}(f)((e,s,t))) \\ &= (\partial \circ \mathcal{E}(f))((e,s,t)) \quad \square \end{aligned}$$

Rotating the diagram in the above proof and adding some zero modules to each end gives

$$\begin{array}{ccccccccccc} \cdots & \xrightarrow{0} & 0 & \xrightarrow{0} & \mathcal{E}(\Gamma) & \xrightarrow{\partial} & \mathcal{V}(\Gamma) & \xrightarrow{0} & 0 & \xrightarrow{0} & \cdots \\ & & \downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \mathcal{E}(f)} & & \downarrow{\scriptstyle \mathcal{V}(f)} & & \downarrow{\scriptstyle \mathrm{id}} & & \\ \cdots & \xrightarrow{0} & 0 & \xrightarrow{0} & \mathcal{E}(\Delta) & \xrightarrow{\partial} & \mathcal{V}(\Delta) & \xrightarrow{0} & 0 & \xrightarrow{0} & \cdots \end{array}$$

which is clearly a chain map from $\Gamma_\bullet$ to $\Delta_\bullet$; thus we have a new functor from the category of graphs to the category of R-chain complexes, mapping Γ to $\Gamma_\bullet$. Theorem 5.2 shows that the mapping from a chain complex to its nth homology module is a functor, and so, by composition, we have the following.

Proposition 9.3. *For any i, the mapping*

$$\Gamma \mapsto H_i(\Gamma_\bullet)$$

defines a covariant functor $\mathsf{Graph} \to R\text{-}\mathsf{Mod}$.

In the case of cohomology, we have the following.

Lemma 9.4. *The mapping from a graph to the dual of its vertex or edge space is a contravariant functor* $\mathsf{Graph} \to R\text{-}\mathsf{Mod}$.

Lemma 9.5. *The coboundary operator $\partial^\vee$ is a natural transformation from the dual of the vertex space to the dual of the edge space.*

These follow from their equivalents for homology by composing the functors from Lemma 9.1 with the contravariant functor $\mathscr{D}\colon R\text{-}\mathsf{Mod} \to R\text{-}\mathsf{Mod}$ that maps modules to their duals. For any graph homomorphism $f\colon \Gamma \to \Delta$, we then have the cochain map from $\Delta^\bullet$ to $\Gamma^\bullet$

$$\begin{array}{ccccccccccc} \cdots & \xleftarrow[0]{} & 0 & \xleftarrow[0]{} & \mathcal{E}(\Gamma)^\vee & \xleftarrow[\partial^\vee]{} & \mathcal{V}(\Gamma)^\vee & \xleftarrow[0]{} & 0 & \xleftarrow[0]{} & \cdots \\ & & \uparrow{\scriptstyle \mathrm{id}} & & \uparrow{\scriptstyle \mathcal{E}(f)^\vee} & & \uparrow{\scriptstyle \mathcal{V}(f)^\vee} & & \uparrow{\scriptstyle \mathrm{id}} & & \\ \cdots & \xleftarrow[0]{} & 0 & \xleftarrow[0]{} & \mathcal{E}(\Delta)^\vee & \xleftarrow[\partial^\vee]{} & \mathcal{V}(\Delta)^\vee & \xleftarrow[0]{} & 0 & \xleftarrow[0]{} & \cdots \end{array}$$

and by the same argument as for homology, there is a contravariant functor that maps a graph Γ to $\Gamma^\bullet$. Then by Theorem 5.2,

Proposition 9.6. *For any i, the mapping*

$$\Gamma \mapsto H^i(\Gamma^\bullet)$$

defines a contravariant functor $\mathsf{Graph} \to R\text{-}\mathsf{Mod}$.

10. The Universal Coefficient Theorem

Two important additive bifunctors $R\text{-}\mathsf{Mod} \times R\text{-}\mathsf{Mod} \to R\text{-}\mathsf{Mod}$ are the tensor product and the Hom functor. Neither is an exact functor, although the tensor product is right-exact in both arguments and the Hom functor left-exact in both arguments. An important class of R-modules are those for which these functors are exact: an R-module N is *flat* if the functor $M \mapsto M \otimes_R N$ (or equivalently $M \mapsto N \otimes_R M$, since the tensor product is commutative up to natural isomorphism) is exact, *projective* if the functor $M \mapsto \mathrm{Hom}_R(N, M)$ is exact, and *injective* if the (contravariant) functor $M \mapsto \mathrm{Hom}_R(M, N)$ is exact.

Since tensoring with R is the identity, R is flat over itself, and since the tensor product distributes over direct sums, direct sums of modules are flat if *and only if* the summands are flat: this implies that all free modules are flat. A characterization of projective modules is that P is projective if and only if there is a module Q such that $P \oplus Q$ is free ([1, Prop. VIII.6.4]); this implies that free modules and direct sums of projective modules are projective and that all projective modules are flat. However, there are projective modules that are not free: consider $\mathbb{Z}/2\mathbb{Z}$ as a $\mathbb{Z}/6\mathbb{Z}$-module; since $\mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/3\mathbb{Z} = \mathbb{Z}/6\mathbb{Z}$, $\mathbb{Z}/2\mathbb{Z}$ is projective, but it cannot be free, since it has fewer elements than $\mathbb{Z}/6\mathbb{Z}$ (in a PID, however, all projective modules are free). Similarly, $\mathbb{Q}$ is flat as a $\mathbb{Z}$-module, but not projective. The $\mathbb{Z}$-modules $\mathbb{Z}/n\mathbb{Z}$ ($n > 0$) are examples of non-flat modules: the sequence

$$0 \to \mathbb{Z} \xrightarrow{\times n} \mathbb{Z} \to \mathbb{Z}/n\mathbb{Z} \to 0$$

is exact, but when tensored with $\mathbb{Z}/n\mathbb{Z}$,

$$0 \to \mathbb{Z}/n\mathbb{Z} \xrightarrow{\times n} \mathbb{Z}/n\mathbb{Z} \to \mathbb{Z}/n\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/n\mathbb{Z} \to 0$$

the sequence is no longer exact, since the second homomorphism is now the zero morphism, and not injective (this is a specific case of the general result that flat modules over an integral domain must be torsion-free).

Injective modules are more difficult to characterize. For example, free modules are not necessarily injective: $\mathbb{Z}$ is not injective over itself, since the exact sequence

$$0 \to \mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \to \mathbb{Z}/2\mathbb{Z} \to 0$$

is mapped by the contravariant functor $N \mapsto \operatorname{Hom}_{\mathbb{Z}}(N, \mathbb{Z})$ to

$$0 \to 0 \to \mathbb{Z}^\vee \xrightarrow{\gamma} \mathbb{Z}^\vee \to 0$$

where γ, the dual of multiplication by two, is not surjective, and thus the sequence is not exact. One characterization of injective modules is Baer's criterion: an R-module Q is injective if and only if every homomorphism $I \to Q$, where I is a nonzero ideal of R, extends to a homomorphism $R \to Q$ ([1, Thm. VIII.6.6]). A direct product of modules is injective if and only if the factors are injective. Injective modules over an integral domain are always *divisible* (that is the homomorphism defined by multiplication by any non-zero-divisor in R is surjective), and in a PID, all divisible modules are injective.

Every R-module M is a quotient of a projective module; that is there is an exact sequence

$$P \to M \to 0$$

where P is projective (in fact, P can be assumed to be free: in the very worst case, we can always take $P = \bigoplus_{m \in M} R$). By repeating this construction indefinitely, we can show that every module M has a *projective resolution,* that is a semi-infinite exact sequence

$$\cdots \to P_2 \to P_1 \to P_0 \to M \to 0$$

with all the P_i projective. Similarly, every module M is a submodule of an injective module; that is there is an exact sequence

$$0 \to M \to Q$$

where Q is injective ([1, Cor. VIII.6.12]). By repeating this construction indefinitely, we obtain an *injective resolution* of M, that is a semi-infinite exact sequence

$$0 \to M \to Q_0 \to Q_1 \to Q_2 \to \cdots$$

with all the Q_i injective.

Suppose we have an additive functor $\mathscr{F}$ which is left- or right-exact. It is natural to ask if $\mathscr{F}$ is exact, and if not how "far" it is from being exact: a way to measure the inexactness of $\mathscr{F}$ are the *derived functors* of $\mathscr{F}$.

Definition 10.1 (derived functors)**.** Let $\mathscr{F} \colon R\text{-}\mathsf{Mod} \to S\text{-}\mathsf{Mod}$ be a covariant additive functor and M an R-module.

The module M has a projective resolution in terms of projective modules P_i, and

$$\cdots \to P_2 \to P_1 \to P_0 \to 0 \to 0 \to 0 \to \cdots$$

is a chain complex $P_\bullet$ which is exact everywhere but possibly at P_0. Then since $\mathscr{F}$ is additive,

$$\cdots \to \mathscr{F}(P_2) \to \mathscr{F}(P_1) \to \mathscr{F}(P_0) \to 0 \to 0 \to 0 \to \cdots$$

is a chain complex $\mathscr{F}P_\bullet$, and the *nth left-derived functor of* $\mathscr{F}$, $\mathscr{L}_n\mathscr{F} \colon R\text{-}\mathsf{Mod} \to S\text{-}\mathsf{Mod}$ is the functor that maps M to $H_n(\mathscr{F}P_\bullet)$.

Similarly, M has an injective resolution in terms of injective modules Q_i, and there is a cochain complex $Q_\bullet$

$$\cdots \leftarrow Q_2 \leftarrow Q_1 \leftarrow Q_0 \leftarrow 0 \leftarrow 0 \leftarrow 0 \leftarrow \cdots$$

and a cochain complex of S-modules $\mathscr{F}Q_\bullet$. The *nth right-derived functor of* $\mathscr{F}$, $\mathscr{R}^n\mathscr{F} \colon R\text{-}\mathsf{Mod} \to S\text{-}\mathsf{Mod}$, is the functor that maps M to $H^n(\mathscr{F}Q_\bullet)$.

There is a similar construction for contravariant functors, which differs only in that additive contravariant functors map chain complexes to cochain complexes and vice versa. It turns out that the derived functors do not depend (up to natural isomorphism) on the choice of resolution ([1, Ex. IX.7.6]). If $\mathscr{F}$ is right-exact, then $\mathscr{L}_0\mathscr{F}$ is naturally isomorphic to $\mathscr{F}$, and if it is left-exact, $\mathscr{R}^0\mathscr{F}$ is naturally isomorphic to $\mathscr{F}$ ([1, Prop. IX.7.13]). We obtain from the derived functor the following long exact sequence:

Lemma 10.2 ([1, Thm. IX.7.12])**.** *Let $\mathscr{F}\colon R\text{-}\mathsf{Mod} \to S\text{-}\mathsf{Mod}$ be an additive functor, and*

$$0 \to L \to M \to N \to 0$$

an exact sequence of R-modules. There are homomorphisms δ_i making the following sequence exact:

$$\begin{array}{l}\cdots \longrightarrow \mathscr{L}_2\mathscr{F}(L) \longrightarrow \mathscr{L}_2\mathscr{F}(M) \longrightarrow \mathscr{L}_2\mathscr{F}(N) \xrightarrow{\delta_2} \\ \mathscr{L}_1\mathscr{F}(L) \longrightarrow \mathscr{L}_1\mathscr{F}(M) \longrightarrow \mathscr{L}_1\mathscr{F}(N) \xrightarrow{\delta_1} \\ \mathscr{L}_0\mathscr{F}(L) \longrightarrow \mathscr{L}_0\mathscr{F}(M) \longrightarrow \mathscr{L}_0\mathscr{F}(N) \longrightarrow 0,\end{array}$$

and homomorphisms δ^i making the following sequence exact:

$$\begin{array}{l}0 \longrightarrow \mathscr{R}^0\mathscr{F}(L) \longrightarrow \mathscr{R}^0\mathscr{F}(M) \longrightarrow \mathscr{R}^0\mathscr{F}(N) \xrightarrow{\delta^1} \\ \mathscr{R}^1\mathscr{F}(L) \longrightarrow \mathscr{R}^1\mathscr{F}(M) \longrightarrow \mathscr{R}^1\mathscr{F}(N) \xrightarrow{\delta^2} \\ \mathscr{R}^2\mathscr{F}(L) \longrightarrow \mathscr{R}^2\mathscr{F}(M) \longrightarrow \mathscr{R}^2\mathscr{F}(N) \longrightarrow \cdots .\end{array}$$

If the original functor $\mathscr{F}$ is right-exact, the first diagram above becomes

$$\begin{array}{l}\cdots \longrightarrow \mathscr{L}_2\mathscr{F}(L) \longrightarrow \mathscr{L}_2\mathscr{F}(M) \longrightarrow \mathscr{L}_2\mathscr{F}(N) \xrightarrow{\delta_2} \\ \mathscr{L}_1\mathscr{F}(L) \longrightarrow \mathscr{L}_1\mathscr{F}(M) \longrightarrow \mathscr{L}_1\mathscr{F}(N) \xrightarrow{\delta_1} \\ \mathscr{F}(L) \longrightarrow \mathscr{F}(M) \longrightarrow \mathscr{F}(N) \longrightarrow 0,\end{array}$$

and similarly for the second diagram if the functor is left-exact. This explains how the derived functors measure the failure of an additive functor to be exact: if $\mathscr{F}$ is right exact, it maps the original exact sequence to

$$\mathscr{F}(L) \to \mathscr{F}(M) \to \mathscr{F}(N) \to 0,$$

where we have lost the 0 on the left to the (possible) lack of left-exactness, and we can extend this to the long exact sequence of derived functors. If $\mathscr{F}$ is left-exact, then the left derived functors $\mathscr{L}_i\mathscr{F}$ are all the zero functor for $i > 0$, and if it is right-exact then the right derived functors $\mathscr{R}^i\mathscr{F}$ are all the zero functor for $i > 0$. Conversely, if the first left derived functor is zero, then $\mathscr{F}$ is right-exact, and if the first right derived functor is zero, $\mathscr{F}$ is left-exact.

In the specific case of the tensor product, we call the derived functors the *Tor functors*. Let N be an R-module: then the functor $M \mapsto \operatorname{Tor}_i^R(M, N)$ is the ith left derived functor of the tensor product $M \mapsto M \otimes_R N$. This leads to an obvious question: since the tensor product is a bifunctor, why do we take the derived functor only of the tensor product on the right? In fact, it turns out that Tor is *balanced*: that is, the module obtained by applying the ith left derived functor of the functor $L \mapsto M \otimes_R L$ to N is isomorphic to the module obtained by applying the ith left derived functor of $L \mapsto L \otimes_R N$ to M for all R-modules M and N ([1, Thm. IX.8.13]), and we call them both $\operatorname{Tor}_i^R(M, N)$. Since R is commutative, $M \otimes_R N \cong N \otimes_R M$, and from this we can conclude that $\operatorname{Tor}_i^R(M, N) \cong \operatorname{Tor}_i^R(N, M)$. If M or N is flat, the tensor product is exact, and $\operatorname{Tor}_i^R(M, N) = 0$ for $i > 0$.

In the case of the Hom functor, we call the ith right derived functor of the functor $N \mapsto \operatorname{Hom}_R(M, N)$ (which is left-exact) $\operatorname{Ext}_R^i(M, N)$. Like Tor, Ext balances, and we get the same modules by applying the right derived functors of the contravariant left-exact functor $M \mapsto \operatorname{Hom}_R(M, N)$ ([1, Thm. IX.8.14]). The functor $\operatorname{Ext}_R^i(M, N)$ is, like Hom, contravariant in the first argument and covariant in the second. A module P is projective if and only if for all modules M, $\operatorname{Ext}_R^1(P, M) = 0$, and a module Q is injective if and only if for all modules M, $\operatorname{Ext}_R^1(M, Q) = 0$.

The objective of this excursion into homological algebra is to obtain the following result:

Theorem 10.3 (Universal Coefficient Theorem in Homology [7, Thm. V.2.5])**.** *Let $C_\bullet$ be a chain complex of flat modules over a principal ideal domain R, A an R-module, and $D_\bullet$ the chain complex with $D_n = A \otimes_R C_n$ and boundary operators $\mathrm{id}_A \otimes d_i$. Then there is, for all n, a short exact sequence*

$$0 \to A \otimes_R H_n(C_\bullet) \to H_n(D_\bullet) \to \operatorname{Tor}_1^R(A, H_{n-1}(C_\bullet)) \to 0,$$

which splits.

There is an equivalent result for cohomology:

Theorem 10.4 (Universal Coefficient Theorem in Cohomology [7, Thm. V.3.3])**.** *Let $C_\bullet$ be a chain complex of free modules over a principal ideal domain R, B an R-module, and $D^\bullet$ the cochain complex with $D^n = \operatorname{Hom}_R(C_n, B)$ and coboundary operators $f \mapsto f \circ d_i$. Then there is, for all n, a short exact sequence*

$$0 \to \operatorname{Ext}_R^1(H_{n-1}(C_\bullet), B) \to H^n(D^\bullet) \to \operatorname{Hom}_R(H_n(C_\bullet), B) \to 0,$$

which splits.

The name "universal coefficient theorem" comes from the fact that $A \otimes_R H_n(C_\bullet)$ and $\operatorname{Hom}_R(H_n(C_\bullet), B)$ may be regarded as the homology module (resp. cohomology module) of $C_\bullet$ with coefficients in the module A (resp. B), and the theorem shows that the homology and cohomology with coefficients in a module is determined by the homology with coefficients in the ring.

11. Base changes and homology

Let R and S be rings, and $f\colon R \to S$ a ring homomorphism; it is natural to ask whether f induces a mapping between modules over R and S. The answer is yes; in fact, there are three different ways to do so. The first is called *restriction of scalars*, and maps S-modules to R-modules. We turn an S-module M into an R-module M', where M' has the same underlying abelian group as M, and the action of $r \in R$ on $m \in M$ is given by

$$rm = f(r)m.$$

If f is an injection, R is (isomorphic to) a sub-ring of S, and the new module simply restricts the scalars by which we can multiply to those in R, hence the name. The map is denoted f_*, and is an additive functor $S\text{-}\mathsf{Mod} \to R\text{-}\mathsf{Mod}$; however it will not be of much importance, and is included merely for completeness.

The second way of obtaining a functor from f is called *extension of scalars,* which is a functor $R\text{-}\mathsf{Mod} \to S\text{-}\mathsf{Mod}$. The homomorphism f allows us to give S a natural R-module structure (the action of $r \in R$ is multiplication by $f(r)$). Then for an R-module M, we define $f^*(M) = S \otimes_R M$. This module has a natural S-module structure, given by, for all $s, x \in S$, $m \in M$,

$$s(x \otimes m) = (sx) \otimes m.$$

The third way to obtain a functor from f is *coextension of scalars,* which is a functor $R\text{-}\mathsf{Mod} \to S\text{-}\mathsf{Mod}$. It associates any R-module M with the module $f^!(M) = \mathrm{Hom}_R(S, M)$, which can be given an S-module structure by, for all $h \in f^!(M)$, $s, x \in S$,

$$(sh)(x) = h(sx).$$

Proposition 11.1 ([1, § VIII.3.3])**.** *For any ring homomorphism $f\colon R \to S$, $f^*, f^!\colon R\text{-}\mathsf{Mod} \to S\text{-}\mathsf{Mod}$ are additive functors.*

The functor f^* maps an R-module homomorphism h to the S-module homomorphism $\mathrm{id}_S \otimes h$, and $f^!$ maps h to $f \mapsto h \circ f$.

There are adjunctions between the functors:

Proposition 11.2 ([1, Prop. VIII.3.6])**.** *For any ring homomorphism $f\colon R \to S$, $f^* \dashv f_* \dashv f^!$; i.e. f^* is left-adjoint to f_* and f_* left-adjoint to $f^!$.*

Extension of scalars preserves the ordinary edge and vertex spaces:

Lemma 11.3. *Let $f\colon R \to S$ be a ring homomorphism, and M a free R-module with a basis $B \subseteq M$. Then $f^*(M)$ is a free S-module with a basis $B' = \{1_S \otimes b \mid b \in B\}$.*

Proof. Since M is free with basis B, $M \cong \bigoplus_{b\in B} Rb$. Then since the tensor product distributes over direct sums, $f^*(M) = S \otimes_R M \cong \bigoplus_{b\in B} S \otimes_R Rb$, and evidently $S \otimes_R Rb \cong S$ as a S module (since $s \otimes rb = f(r)s(1_S \otimes b)$, every element of $S \otimes_R Rb$ is a multiple of $1_S \otimes b$), so $f^*(M)$ is free and generated by B'. □

The above lemma implies that the images under f^* of the ordinary vertex and edges spaces over R are isomorphic to the vertex and edge spaces over S. The same is not true, however, of the locally finite edge and vertex spaces. Let Γ be a locally finite graph with countably infinitely many edges $e_1, e_2, \ldots$; then $\mathcal{E}_R^\infty(\Gamma) \cong \prod_{i\in\mathbb{N}} R$. Let $f\colon \mathbb{Z} \to \mathbb{Q}$ be the unique ring homomorphism from $\mathbb{Z}$ to $\mathbb{Q}$. Then

$f^*(\mathbb{Z}) = \mathbb{Q} \otimes_{\mathbb{Z}} \mathbb{Z} \cong \mathbb{Q}$. If $\mathcal{E}^{\infty}_{\mathbb{Q}}(\Gamma) \cong f^*(\mathcal{E}^{\infty}_{\mathbb{Z}}(\Gamma))$, then $\prod_{i\in\mathbb{N}} \mathbb{Q} \otimes_{\mathbb{Z}} \mathbb{Z} = \prod_{i\in\mathbb{N}} \mathbb{Q} \cong f^*(\prod_{i\in\mathbb{N}} \mathbb{Z}) = \mathbb{Q} \otimes_{\mathbb{Z}} \prod_{i\in\mathbb{N}} \mathbb{Z}$, which is not the case.

Let $f\colon R \to S$ be a ring homomorphism, and $(M_\bullet, d_\bullet)$ a chain complex of R-modules. Then for all i, $f^*(M_i)$ is an S-module, and $f^*(d_i)\colon f^*(M_i) \to f^*(M_{i-1})$ a homomorphism of S-modules; in addition, $f^*(d_i) \circ f^*(d_{i+1}) = f^*(d_i \circ d_{i+1}) = f^*(0) = 0$, since f^* is additive. Thus $(f^*(M_\bullet), f^*(d_\bullet))$ is a chain complex of S-modules.

Lemma 11.4. *The following diagram commutes:*

$$\begin{array}{ccc} f^*(\mathcal{E}_R(\Gamma)) & \xrightarrow{f^*(\partial)} & f^*(\mathcal{V}_R(\Gamma)) \\ \downarrow{\scriptstyle \psi_E} & & \downarrow{\scriptstyle \psi_V} \\ \mathcal{E}_S(\Gamma)) & \xrightarrow{\quad\partial\quad} & \mathcal{V}_S(\Gamma) \end{array}$$

where ψ_V and ψ_E are the isomorphisms that identify a vertex (resp. edge) x in $\mathcal{V}_S(\Gamma)$ (resp. $\mathcal{E}_S(\Gamma)$) with $1_S \otimes x$ in $f^*(\mathcal{V}_R(\Gamma))$ (resp. $f^*(\mathcal{E}_R(\Gamma))$).

Proof. Let (e, s, t) be a directed edge of Γ; then $1_S \otimes (e, s, t)$ is an element of $f^*(\mathcal{E}_R(\Gamma))$, and

$$\psi_V(f^*(\partial)(1_S \otimes (e, s, t))) = \psi_V(1_S \otimes \partial((e, s, t))) = \partial((e, s, t)) = \partial(\psi_V(1_S \otimes (e, s, t))).$$

Then since the elements $1_S \otimes (s, t)$ form a basis for $f^*(\mathcal{E}_R(\Gamma))$, $\psi_V \circ f^*(\partial) = \partial \circ \psi_E$. ☐

This implies that the image of the chain complex of R-modules $\Gamma_\bullet$ under f^* is isomorphic as a chain complex to the chain complex of S-modules $\Gamma_\bullet$.

Theorem 11.5. *Let $f\colon R \to S$ be a ring homomorphism, and Γ a graph. Let $\Gamma^R_\bullet$ be the ordinary chain complex of Γ over R, and $\Gamma^S_\bullet$ the ordinary chain complex over S. Then for all i,*

$$S \otimes_R H_i(\Gamma^R_\bullet) = f^*(H_i(\Gamma^R_\bullet)) \cong H_i(\Gamma^S_\bullet).$$

Proof. The only cases to check are $n = 0$ and $n = 1$, since all the other homology modules are trivial, and f^*, being exact, maps trivial modules to trivial modules.

Since $\Gamma^S_\bullet$ is isomorphic to the image of $\Gamma^R_\bullet$ under f^*, their homology modules are also isomorphic (Corollary 5.3). Thus by the Universal Coefficient Theorem (Theorem 10.3), the sequence

$$0 \to f^*(H_n(\Gamma^R_\bullet)) \to H_n(\Gamma^S_\bullet) \to \mathrm{Tor}^R_1(S, H_{n-1}(\Gamma^R_\bullet)) \to 0$$

is exact for all n, and if $\mathrm{Tor}^R_1(S, H_{n-1}(\Gamma^R_\bullet)) = 0$, then $f^*(H_n(\Gamma^R_\bullet)) \cong H_n(\Gamma^S_\bullet)$, and the proof is complete.

If $n = 0$, then $H_{n-1}(\Gamma^R_\bullet) = 0$, and thus $\mathrm{Tor}^R_1(S, H_{n-1}(\Gamma^R_\bullet)) = 0$. By Theorem 6.2 and Theorem 6.3, $H_0(\Gamma^R_\bullet)$ is always free and therefore flat, and thus if $n = 1$, $\mathrm{Tor}^R_1(S, H_{n-1}(\Gamma^R_\bullet)) = 0$. ☐

Because there is a unique ring homomorphism from $\mathbb{Z}$ to any other ring, this theorem entails that knowing the homology over $\mathbb{Z}$ is sufficient to determine it over any ring.

12. Topologically equivalent graphs

The homology of a graph is a purely topological property of the graph, which means that graphs that are topologically equivalent have the same homology. In order to prove this statement, we need to have a better understanding of the topological structure of graphs. Graphs have an "obvious" representation as a topological space, where the vertices are points and the edges arcs. The following definition formalizes this idea.

Definition 12.1 (topological realization)**.** Let $\Gamma = (V, E, \nu)$ be a graph. A *topological realization* of Γ is a Hausdorff space X such that there exists a partition P of X and a bijection $\lambda\colon V \cup E \to P$ such that

(1) for all $v \in V$, $\lambda(v)$ is a singleton,
(2) for all $e \in E$, $\lambda(e)$ is homeomorphic to $(0,1)$,
(3) for all $e \in E$, there exists a continuous function $f\colon [0,1] \to X$ such that the restriction of f to $(0,1)$ is a homeomorphism onto $\lambda(e)$ and $f(\{0,1\}) = \lambda(\nu(e))$, and
(4) $A \subseteq X$ is closed if and only if for all $e \in E$, $A \cap \overline{\lambda(e)}$ is closed.

This definition is essentially a special case of the definition of a CW-complex. together with a bijection associating specific cells with the vertices and edges of the graph. The next few propositions, therefore, formalize the notion that a graph is, topologically, a one-dimensional CW (or simplicial) complex.

Proposition 12.2. *A locally finite graph has a topological realization.*

Proof. Let $\Gamma = (V, E, \nu)$ be a locally finite graph, and select an orientation for Γ. Let

$$X = V \cup ((0,1) \times E)$$

and

$$\mathcal{B} = \{(a,b) \times \{e\} \mid 0 \le a < b \le 1, e \in E\} \cup \{B_{r,v} \mid 0 < r < 1, v \in V\},$$

where $B_{r,v}$ is the union of $\{v\}$, $(0,r) \times \{e\}$ for all edges e starting at v under the given orientation, and $(1-r,1) \times \{e\}$ for all edges e ending at v under the given orientation. Then $\mathcal{B}$ is clearly a basis for a topology on X. We will show that X with the topology generated by $\mathcal{B}$ is a topological realization of Γ.

First we define P and λ. The elements of the partition P are singletons $\{v\}$, for all $v \in V$, and the sets $(0,1) \times \{e\}$, for all $e \in E$. Then

$$\lambda(x) = \begin{cases} \{x\} & \text{if } x \in V \\ (0,1) \times \{x\} & \text{if } x \in E \end{cases}$$

obviously defines a bijection $\lambda\colon V \cup E \to P$.

Let $x, y \in X$, $x \neq y$. There are several possible cases:

- if $x = (r, e)$ and $y = (s, d)$ for $0 < r, s < 1$ and edges e and d, then
  - if e and d are distinct edges, $(0,1) \times \{e\}$ and $(0,1) \times \{d\}$ are disjoint neighbourhoods of x and y respectively, and
  - if $e = d$, then $(r - \rho, r + \rho) \times \{e\}$ and $(s - \rho, r + \rho) \times \{e\}$, where $\rho = |r - s|/2 > 0$, are disjoint neighbourhoods of x and y respectively;
- if x and y are vertices, $B_{1/2,x}$ and $B_{1/2,y}$ are disjoint neighbourhoods; and

- if x is a vertex and $y = (r, e)$, then either $B_{r/2,x}$ and $(r/2, 1) \times \{e\}$ or $B_{(1+r)/2,x}$ and $(0, (1 + r)/2) \times \{e\}$ are disjoint neighbourhoods of x and y respectively.

Thus X is Hausdorff.

Now we demonstrate the conditions in Definition 12.1:

(1) This is trivial from the definition of λ.
(2) Let e be an edge of Γ; then $\lambda(e) = (0, 1) \times \{e\}$. Let $f : \lambda(e) \to (0, 1)$, $(r, e) \mapsto r$; obviously f is a bijection. An open set in $(0, 1)$ is a union of open intervals (a, b), $0 < a, b < 1$, and an open set in $\lambda(e)$ is a union of elements of $\mathcal{B}$: that is of sets $(a, b) \times \{e\}$, since the other type of basis element contains a vertex. Then f and f^{-1} take basis elements to basis elements, and thus f is a homeomorphism.
(3) Let e be an edge of Γ which starts at a vertex u and ends at v in the given orientation, and define $f : [0, 1] \to X$ by

$$f(x) = \begin{cases} u & \text{if } x = 0, \\ v & \text{if } x = 1, \\ (x, e) & \text{otherwise.} \end{cases}$$

Then f is clearly continuous, the restriction of f to $(0, 1)$ is simply the homeomorphism from the previous part, and $f(\{0, 1\}) = \lambda(\nu(e))$.
(4) The only if case is true for any topological space. Suppose that $A \subseteq X$, and for all $e \in E$, $A \cap \overline{\lambda(e)}$ is closed. Let $x \in X$ be a limit point of A. If x is a point (r, e), then $\lambda(e) = (0, 1) \times \{e\}$ is a neighbourhood of x, and thus contains points of A. Thus x is a limit point of $A \cap \overline{\lambda(E)}$, and since that set is closed, $x \in A$. If x is a vertex, let B be the union of the $\overline{\lambda(e)}$ for all e incident on x: since there are only finitely many such e, B is a finite union of closed sets, and is closed. B contains a neighbourhood of x, and therefore contains points of A: then x is a limit point of $B \cap A$, and since that set is a finite union of closed sets of the form $\overline{\lambda(E)} \cap A$, it is closed, and $x \in B \cap A \subseteq A$. Thus A is closed.

Therefore, X is a topological realization of Γ. □

Proposition 12.3. *Let $\Gamma = (V, E, \nu)$ be a graph, and X a topological realization of Γ. Then a topological space Y is homeomorphic to X if and only if Y is also a topological realization of Γ.*

Proof. Suppose Y is a topological realization of Γ; we will show it is homeomorphic to X. Let P and λ be a partition and bijection making X a topological realization, and Q and κ a partition and bijection making Y a topological realization. For every edge $e \in E$, there is a homeomorphism $g_e : \lambda(e) \to (0, 1)$, and a homeomorphism $h_e : (0, 1) \to \kappa(e)$. Let $f : X \to Y$ be defined by

$$f(x) = \begin{cases} \text{the only element of } \kappa(v) & \text{if } x \in \lambda(v) \text{ for some } v \in V, \\ h_e(g_e(x)) & \text{if } x \in \lambda(e) \text{ for some } e \in E. \end{cases}$$

That f is well-defined follows from the fact that P is a partition and λ a bijection; that it is a bijection follows from the fact that κ and $h_e \circ g_e$ (for all e) are bijections and P a partition. It therefore remains only to show that f and f^{-1} are continuous.

Since for every $e \in E$, $h_e \circ g_e \colon \lambda(e) \to \kappa(e)$ is a homeomorphism, f and f^{-1} are continuous when restricted to $\lambda(e)$ or $\kappa(e)$ respectively.

Conversely, suppose there exists $f \colon X \to Y$ a homeomorphism; we will show that Y is a topological realization of Γ. Let P and λ be a partition and bijection making X a topological realization. Since f is a bijection, there is a partition Q of Y induced by P, the elements of which are the images of the elements of P under f. Let $\kappa \colon V \cup E \to Q$ be defined by $x \mapsto f(\lambda(x))$; then κ is a bijection. That Y is Hausdorff follows from being homeomorphic to X. It remains only to check the numbered conditions from Definition 12.1.

(1) Since $\lambda(v)$ is a singleton and f a bijection, $\kappa(v)$ must also be a singleton.
(2) If $g \colon (0,1) \to \lambda(e)$ is a homeomorphism, then the composition of g with the restriction of f to $\lambda(e)$ is a homeomorphism $(0,1) \to \kappa(e)$.
(3) There is a continuous function $h \colon [0,1] \to X$ satisfying the required conditions; $f \circ h \colon [0,1] \to Y$ is also continuous, since the restriction of h to $(0,1)$ is a homeomorphism, the restriction of $f \circ h$ is also, and $(f \circ h)(\{0,1\}) = f(\lambda(\nu(e))) = \kappa(\nu(e))$.
(4) The only if portion is true for all topological spaces. Let $A \subseteq Y$ such that $A \cap \overline{\kappa(e)}$ is closed for all $e \in E$. Then for all $e \in E$, $f^{-1}(A \cap \overline{\kappa(e)}) = f^{-1}(A) \cap f^{-1}(\overline{\kappa(e)}) = f^{-1}(A) \cap \overline{\lambda(e)}$ is closed, and therefore $f^{-1}(A)$ is closed, and since f is a homeomorphism, A is closed.

Thus Y is a topological realization of Γ. □

We say that two graphs are *topologically equivalent* if their topological realizations are homeomorphic. If two graphs are topologically equivalent, then they have the same topological realizations. One way of producing a topologically equivalent graph from an existing graph is to subdivide edges of the graph: that is to replace an edge e from u to v with n edges $e_1, e_2, \ldots, e_n$ such that e_1 is from u to a new vertex w_1, e_2 from w_1 to w_2, etc., up to e_n from w_{n-1} to v. We will show that this not only produces topologically equivalent graphs, but that it is in a sense the only way to produce topologically equivalent graphs: any two topologically equivalent graphs have a third graph, which can be produced from each by subdivision of edges, and which is topologically equivalent to both. Showing that topologically equivalent graphs have the same homology is then as simple as showing that subdivision of edges preserves the homology.

Lemma 12.4. *Let $\Gamma = (V, E, \nu)$ be a graph, and X a topological realization of Γ, with a bijection $\lambda \colon V \cup E \to P$. Then for all edges $e \in E$,*

(1) $\overline{\lambda(e)} = \lambda(e) \cup \lambda(\nu(e))$*, and*
(2) $\lambda(e)$ *is open.*

Proof. Let $e \in E$, and $f \colon [0,1] \to X$ the function required by condition (3) of Definition 12.1. Since $[0,1]$ is compact and f continuous, $f([0,1]) = \lambda(e) \cup \lambda(\nu(e))$ is compact, and thus (since X is Hausdorff) closed; therefore $\overline{\lambda(e)} \subseteq \lambda(e) \cup \lambda(\nu(e))$. Since $\lambda(e)$ must be a subset of its closure, it remains only to show that the elements of $\nu(e)$ are limit points of $\lambda(e)$. This is the case, since if $u \in \nu(e)$ and U is a neighbourhood of $\lambda(u)$, $f^{-1}(U)$ is non-empty and open in $[0,1]$, and it cannot be a subset of $\{0,1\}$, because that would make $[0,1]$ disconnected; thus U contains a point of $\lambda(e)$ other than u.

To show $\lambda(e)$ is open, we will show $X \setminus \lambda(e)$ is closed. Let $e' \in E$, and consider $\overline{\lambda(e')} \cap (X \setminus \lambda(e)) = \overline{\lambda(e')} \setminus \lambda(e)$. If $e' \neq e$, this set is, by the previous part, $\overline{\lambda(e')}$, and if $e = e'$ it is $\lambda(\nu(e))$, a set of one or two points. Either way, it is closed, and by condition (4) of Definition 12.1, $X \setminus \lambda(e)$ is closed. □

Lemma 12.5. *Let P and Q be two partitions of a set X, and R their coarsest common refinement. Then every element of R is the intersection of an element of P and an element of Q.*

Proof. Let $r \in R$: since R is a refinement of P and Q, there is $p \in P$ and $q \in Q$ such that $r \subseteq p \cap q$. Let $s \in R$ be such that $s \cap p \cap q \neq \varnothing$; then since R is a refinement, $s \subseteq p \cap q$. Then it must be the case that $r = s$, since otherwise we could replace r and s by $r \cup s$ and have a coarser refinement. Thus r is the only element of R such that $r \cap p \cap q \neq \varnothing$, and since every element of X is in some element of R, $r = p \cap q$. □

Theorem 12.6. *Let Γ and Δ be topologically equivalent graphs. Then there is a graph Λ which is isomorphic to graphs obtained from Γ and Δ by subdividing edges into finitely many new edges.*

Proof. First we note that if Γ and Δ both contain a connected component that consists entirely of a single vertex and a single loop on that vertex, we can obtain a Λ by removing that component from both Γ and Δ, applying the theorem, and then adding such a component to the resulting graph. If Γ and Δ contain multiple such components, we can remove sets of such components of equal cardinality from Γ and Δ. Thus we can assume, without loss of generality, that Γ and Δ do not both contain such a component.

Suppose Γ and Δ are topologically equivalent. Let X be a topological realization of Γ: then since X is homeomorphic to some topological realization of Δ, by Proposition 12.3 it is also a topological realization of Δ. Let P and λ be a partition and bijection making X a topological realization of Γ, and Q and κ a partition and bijection making X a topological realization of Δ. Then P and Q have a coarsest common refinement R. Every element of R is a non-empty intersection of an element of P and an element of Q, by Lemma 12.5: if either of those elements is a singleton, the intersection is as well.

Suppose $p \in P$ and $q \in Q$ are not singletons and have a non-empty intersection; they must be the image of edges in the topological realization, and there exists a continuous function $f\colon [0,1] \to X$ which when restricted to $(0,1)$ is a homeomorphism onto p. By Lemma 12.4, p and q are open, and therefore so is $f^{-1}(p \cap q)$; we will show it is connected, and therefore an open interval. Suppose it were not connected: then it would have at least two connected components, which would be open intervals (a,b) and (c,d), with $0 \leq a < b \leq c < d \leq 1$ and $\{a,b,c,d\} \cap f^{-1}(p \cap q) = \varnothing$. Then $\{a,b,c,d\} \subseteq \overline{f^{-1}(p \cap q)}$, and thus $f(\{a,b,c,d\}) \subseteq \overline{p \cap q} \subseteq \overline{q}$, and $f(\{a,b,c,d\}) \cap q = f(\{a,d\} \cap \{0,1\})$. But by Lemma 12.4, $1 \leq |\overline{q} \setminus q| \leq 2$, and this can only happen if $a = 0$, $b = c$, $d = 1$, and f maps 0 and 1 to the same point. Then both $f(0) = f(1)$ and $f(b) = f(c)$ are the realization of a vertex, of Γ and of Δ respectively, and the only edge incident on them is the one realized by p and by q respectively (if there were another edge in Γ, then $f(0)$ would be a limit point of the image of that edge under λ, but since q is open, $f(0)$ has a neighbourhood in q, which would contain no points of those edges; for the same reason, there can be no other edge in Δ). But then $\overline{p} = \overline{q}$ is the realization of connected components of both Γ and Δ with a single vertex and a single loop incident on that vertex, and we supposed that at least one of Γ and Δ contains no such component. Thus there exists $0 \leq x < y \leq 1$ such that $f^{-1}(p \cap q) = (x, y)$, $f(\{x, y\}) \cap p \cap q = \varnothing$, and $f(\{x, y\}) \subseteq \overline{p \cap q}$. There exists a function $g\colon [0,1] \to [0,1]$, $z \mapsto (y-x)z + x$, which is continuous, injective, and has image $[x, y]$, and therefore $f \circ g\colon [0,1] \to X$ is continuous, homeomorphically maps $(0,1)$ to $p \cap q$, and the singletons $\{(f \circ g)(0)\}$ and $\{f \circ g)(1)\}$ are elements of R. Such a function $f \circ g$ exists for every $r \in R$ which is not a singleton, and we will call it f_r.

Let $p \in P$: we will show that there are only finitely many elements of R that are subsets of p. It is sufficient to show that there are only finitely many singleton elements of R which are subsets of p. Each such an element must also be a singleton element of Q; let S be the union of all such singletons. For any edge e of Δ, $S \cap \overline{\kappa(e)}$ has at most two elements, and since X is Hausdorff, is closed; then by condition (4) of Definition 12.1, S is closed. Therefore $f^{-1}(S) \subseteq [0,1]$ is closed. Suppose S is infinite: then $f^{-1}(S)$ is too, since $S \subseteq p \subseteq f([0,1])$, and since $[0,1]$ is compact, $f^{-1}(S)$ has a limit point $x \in [0,1]$, and since $f^{-1}(S)$ is closed, $x \in f^{-1}(S)$. Then $f(x) \in S$ is a limit point of S, and $\{f(x)\}$ is the realization of a vertex of Δ: since $f(x) \in P$, there are exactly two edges incident on that vertex, e and d, and $\kappa(e) \cup \{f(x)\} \cup \kappa(d)$ is an open neighbourhood of $f(x)$. But this set contains no elements of S other than $f(x)$, which contradicts our conclusion that $f(x)$ is a limit point of S: thus S must be finite. By the same argument, we can show that there are only finitely many elements of R that are subsets if any $q \in Q$.

Let $V = \{r \in R \mid r \text{ a singleton}\}$ and $E = R \setminus V$; then there exists a function $\nu\colon E \to \mathcal{P}_2(V)$ that maps $r \in E$ to $\{\{f_r(0)\}, \{f_r(1)\}\}$; then $\Lambda = (V, E, \nu)$ is a graph, and X, together with the partition R, the identity function $V \cup E \to R$, and the functions f_r satisfy the first three conditions of Definition 12.1. Suppose that $A \subseteq X$ is such that for all $e \in E$, $A \cap \overline{e}$ is closed. Let e' be an edge of Γ; then by the previous paragraph, $\overline{\lambda(e')} = \bigcup_{i=1}^{n} \overline{e_i}$ for $e_i \in E$. Thus $A \cap \overline{\lambda(e')} = \bigcup_{i=1}^{n} A \cap \overline{e_i}$ is the union of finitely many closed sets, and thus closed. Then since e' was any arbitrary edge of Γ, A is closed, and X is a topological realization of Λ.

It remains only to show that Λ can be obtained from Γ and from Δ by subdividing edges finitely many times, which follows from the fact that R is a refinement of P and of Q and the fact that each element of P or of Q contains only finitely many elements of R. □

We now show that subdivision preserves the homology.

Lemma 12.7 ([6, Cor. 5.7])**.** *Let Γ be a graph, and Δ a graph obtained from Γ by subdividing any number of edges into* two *new edges each. Then for all n,*

$$H_n(\Gamma_\bullet) \cong H_n(\Delta_\bullet).$$

Proof. Let $\alpha_E\colon \mathcal{E}(\Gamma) \to \mathcal{E}(\Delta)$ be defined on edges of Γ by

$$\alpha_E(e) = \begin{cases} e & \text{if } e \text{ is not subdivided in } \Delta \\ e_1 + e_2 & \text{if } e \text{ is subdivided into } e_1 \text{ and } e_2, \end{cases}$$

and extended to the whole space by linearity. Let $\alpha_V\colon \mathcal{V}(\Gamma) \to \mathcal{V}(\Delta)$ be the inclusion map (since the set of vertices of Γ is a subset of the set of vertices of Δ). It is clear that both α_E and α_V are injective. We will show the following diagram commutes:

$$\begin{array}{ccc} \mathcal{E}(\Gamma) & \xrightarrow{\partial_\Gamma} & \mathcal{V}(\Gamma) \\ \Big\downarrow{\scriptstyle \alpha_E} & & \Big\downarrow{\scriptstyle \alpha_V} \\ \mathcal{E}(\Delta) & \xrightarrow{\partial_\Delta} & \mathcal{V}(\Delta). \end{array}$$

Let e be an edge of Γ from u to v. If e is not subdivided in Δ, then $(\alpha_V \circ \partial_\Gamma)(e) = v - u = (\partial_\Delta \circ \alpha_E)(e)$. If e is subdivided into e_1 from u to w and e_2 from w to v, then $\partial_\Delta(\alpha_E(e)) = \partial_\Delta(e_1 + e_2) = (w - u) + (v - w) = v - u = (\alpha_V \circ \partial_\Gamma)(e)$. Thus the diagram commutes.

We will show that $\alpha_E(\ker \partial_\Gamma) = \ker \partial_\Delta$; since α_E is injective, this implies that the restriction of α_E to $\ker \partial_\Gamma$ is an isomorphism onto $\ker \partial_\Delta$, and $H_1(\Gamma_\bullet) = \ker \partial_\Gamma = \ker \partial_\Delta = H_1(\Delta_\bullet)$. One direction, $\alpha_E(\ker \partial_\Gamma) \subseteq \ker \partial_\Delta$, is clear from the commutativity of the above diagram. Let $x \in \ker \partial_\Delta$. If e_1 from u to w and e_2 from w to v are two edges in Δ produced by subdividing an edge e from u to v in Γ, then the coefficients of e_1 and e_2 in x must be equal, because otherwise the coefficient of w in $\partial_\Delta(x)$ would be non-zero (since e_1 and e_2 are the only edges incident on w). Thus there is an element $x' \in \mathcal{E}(\Gamma)$ with the coefficient of an edge e equal to its coefficient in x, if e is undivided, or the coefficients of its subdivisions in x if it was divided; this means that $\alpha_E(x') = x$, and by the commutativity of the above diagram, $x' \in \ker \partial_\Gamma$. Thus $\ker \partial_\Delta = \alpha_E(\ker \partial_\Gamma)$.

By the universal property of quotients, there exists a homomorphism ϕ making the following diagram commute:

$$\begin{array}{ccccc} \mathcal{V}(\Gamma) & \xrightarrow{\alpha_V} & \mathcal{V}(\Delta) & \xrightarrow{\pi_\Delta} & H_0(\Delta_\bullet) \\ & \searrow^{\pi_\Gamma} & & \nearrow_{\phi} & \\ & & H_0(\Gamma_\bullet) & & \end{array}$$

where the π_Γ and π_Δ are the canonical surjections, provided that $\ker(\pi_\Delta \circ \alpha_V) = \alpha_V^{-1}(\operatorname{im} \partial_\Delta) \subseteq \operatorname{im} \partial_\Gamma$. To show this, suppose that $x \in \alpha_V^{-1}(\operatorname{im} \partial_\Delta)$: then there exists $y \in \mathcal{E}(\Delta)$ such that $\alpha_V(x) = \partial_\Delta(y)$. This implies that the coefficient of $\partial_\Delta(y)$ on any of the vertices added by a subdivision is zero, and by the same argument as in the previous paragraph, there exists $z \in \mathcal{E}(\Gamma)$ such that $\alpha_E(z) = y$. Then $\partial_\Delta(\alpha_E(z)) = \partial_\Delta(y)$, and by commutativity of the above diagram, $\alpha_V(\partial_\Gamma(z)) = \partial_\Delta(y) = \alpha_v(x)$, and by injectivity of α_V, $x = \partial_\Gamma(z)$. Thus $x \in \operatorname{im} \partial_\Gamma$. In addition to implying the existence of ϕ, this also shows it is injective. Suppose $x \in \ker \phi$; then there exists $y \in \mathcal{V}(\Gamma)$ such that $\pi_\Gamma(y) = x$, and $0 = (\phi \circ \pi_\Gamma)(y) = (\pi_\Delta \circ \alpha_V)(y)$. Thus $y \in \ker(\pi_\Delta \circ \alpha_V) \subseteq \operatorname{im} \partial_\Gamma$, and therefore $x = \pi_\Gamma(y) = 0$. Thus $\ker \phi = \{0\}$.

To show that ϕ is surjective, it is sufficient to show that the homomorphism $\pi_\Delta \circ \alpha_\Gamma$ is surjective. To show this, it is in turn sufficient to show that $\pi_\Delta(v) \in \operatorname{im}(\pi_\Delta \circ \alpha_V)$ for all vertices v of Δ, since the vertices are a basis of $\mathcal{V}(\Delta)$. If v is also a vertex of Γ, then clearly this is the case, and if v was added during the subdivision of some edge, then there is a vertex u of Γ such that there is an edge in Δ from u to v, and $\pi_\Delta(v) = \pi_\Delta(v - u + u) = \pi_\Delta(v - u) + \pi_\Delta(u) = \pi_\Delta(u) \in \operatorname{im}(\pi_\Delta \circ \alpha_V)$. Thus ϕ is an isomorphism, and $H_0(\Gamma_\bullet) \cong H_0(\Delta_\bullet)$. □

Subdividing an edge into finitely many edges is simply equivalent to performing the subdivision into two edges recursively. From the above lemma and induction we obtain the following result:

Corollary 12.8. *Let Γ be a graph, and Δ a graph which can be obtained from Γ by subdividing edges of Γ, each into finitely many new edges. Then for all n,*

$$H_n(\Gamma_\bullet) \cong H_n(\Delta_\bullet)$$

Then from Corollary 12.8 and Theorem 12.6 we obtain the following result.

Theorem 12.9 ([6, Thm. 5.13])**.** *Topologically equivalent graphs have isomorphic homology modules.*

13. Dual graphs, graph embeddings, and homology

For any graph Γ, there are two natural notions of a *dual graph*, one combinatorial, and one topological. The combinatorial dual has a vertex for every edge of Γ, and in which there is an edge between vertices if the correspond to edges incident on a common vertex. Formally, we define

Definition 13.1 (dual graph 1)**.** Let $\Gamma = (V, E, \nu)$ be a simple graph. The *dual graph* of Γ is the graph $\Gamma^* = (E, \{\{e, f\} \in \mathcal{P}_2(E) \mid \nu(e) \cap \nu(f) \neq \varnothing\}, \iota)$, where ι is the inclusion map.

The dual is defined here only for simple graphs because it is not clear what the correct definition is for graphs that have loops and multiple edges: should there be exactly one edge between every pair of vertices that correspond to edges incident on a common vertex in the original graph, or should there be one edge between them for every vertex in common?

It is clear that if Γ is connected and has more than one vertex, then so is Γ^*, since a path $(a_1, e_1, \ldots, e_{n-1}, a_n)$ in Γ provides a path from e_1 to e_{n-1} in Γ^*. If Γ has only one vertex, then its dual is the empty graph. The dual of a disconnected graph is the disjoint union of the duals of its components. Then since by Theorems 6.2 and 6.3 $H_0(\Gamma_\bullet)$ is the free module generated by the connected components of Γ, $H_0(\Gamma^*_\bullet)$ is the free module generated by the components with more then one vertex.

There is not such a simple relationship between the first homology modules. For any tree T, $H_1(T_\bullet) \cong 0$. However, the first homology module of the dual can be any free module of finite rank.

Proposition 13.2. *For all $n \in \mathbb{N}$, there exists a finite tree T such that* $\operatorname{rank} H_1(T^*_\bullet) = n$.

Proof. The proof is by induction on n. If $n = 0$, then we can take $T = (\{u, v\}, \{e\}, e \mapsto \{u, v\})$, which has dual $T^* = (\{e\}, \varnothing, 0)$; then $H_1(T^*_\bullet) = 0$.

Suppose $n > 0$. By the induction hypothesis, there exists a tree S such that $H_1(S^*_\bullet) = n - 1$. By Corollary 2.4, the vertices of S can be ordered $v_1, v_2, \ldots, v_k$ such that there is exactly one edge connecting v_k to the rest of the tree. We obtain T by adding two new vertices v_{k+1} and v_{k+2} to S and edges connecting both new vertices to v_k. Then T^* is obtained from S^* by adding two new vertices, for the two new edges of T, which are connected to each other and to the vertex corresponding to the one edge connecting v_k to the rest of the tree, since all of those edges are incident on v_k. This adds precisely one new cycle to the dual graph, and so $\operatorname{rank} H_1(T^*_\bullet) = \operatorname{rank} H_1(S^*_\bullet) + 1 = n$. □

The trees constructed in the proof of the above proposition have the general form

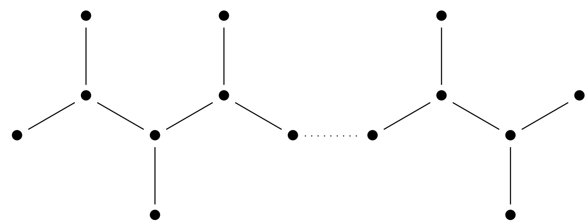

where there are n vertices of degree 3: the dual graph then looks like

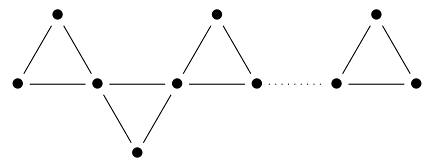

with n cycles joined at their vertices.

The topological notion of a dual graph only applies to certain graphs and relies on the idea of embeddings of graphs. An *embedding* of a graph Γ into a topological space X is a subset of X that is a topological realization of Γ. If E is such an embedding into X, the connected components of $X \setminus E$ are called the *faces* of the embedding. A graph is *planar* if it has an embedding into $\mathbb{R}^2$.

For the purposes of this section, we will consider only embeddings that are *nice.* A nice embedding into a 2-manifold X is one in which the realization of every edge lies in the boundary of one or two faces, is disjoint from the boundary of every other face, and for every face F of the embedding there is a continuous function $f\colon \overline{D} \to X$ from the closed disk to X satisfying

(1) $f(\overline{D}) = \overline{F}$,
(2) the restriction of f to the open disk is a homeomorphism onto F,
(3) the realization of a vertex in the boundary of F has a finite preimage, and
(4) the realization of a point in an edge in the boundary of F has a preimage consisting of one point if that edge is in the boundary of multiple faces, and two points if it is in the boundary only of F.

Our new notion of a dual graph will then be defined for any graph that can be nicely embedded into a 2-manifold, and we will generally focus on graphs that are planar.

Definition 13.3 (dual graph)**.** Let $\Gamma = (V, E, \nu)$ be a graph, X a 2-manifold, and E an embedding of Γ into X. The *dual graph of* Γ *with respect to* E is the graph $\Gamma^\dagger = (V', E, \nu')$, where V' is the set of faces of E, and ν' maps an edge to the set of faces that have a limit point in its realization.

Note that $\Gamma^\dagger$ does not just depend on Γ and the manifold in which it is embedded: it also depends on the choice of embedding. Consider the graph $\Gamma = (\mathbb{Z}, \mathbb{Z}, n \mapsto \{n, n+1\})$, which is an infinite line: it can be embedded into $\mathbb{R}^2$ in an obvious way, mapping a vertex n to the point $(n, 0)$ and an edge n to the set $[n, n+1] \times \{0\}$, but it can also be embedded by mapping a vertex n to $(a_n, 0)$ and an edge n to the set $[a_n, a_{n+1}] \times \{0\}$, where $a_n = 1 - \sum_{i=0}^{|n|} 2^{-i}$ for positive n, and $a_{-n} = -a_n$. For the first embedding, $\Gamma^\dagger$ has two vertices, but under the second it has only one.

Lemma 13.4. *Let E be a nice embedding of a graph Γ with at least one vertex into a 2-manifold, and F a face of E. Then there exists a cycle in Γ the edges* *and vertices of which are precisely those with realizations contained in the boundary of F.*

Proof. Since the embedding is nice, there exists a continuous surjective function $f\colon S^1 \to \partial F \subseteq E$; thus ∂F is compact. Since E is Hausdorff, this means that ∂F is closed, and if e is an edge with realization contained in ∂F, then the realizations of the endpoints of e are also in ∂F. For every vertex v with realization in ∂F, let $U_v \subseteq E$ be the union of the realization of v and the realizations of all the edges incident on it; then $\{U_v \mid v \text{ a vertex with realization in } \partial F\}$ is a cover of ∂F, and has a finite subcover. But each vertex is contained in only one of the U_v, and thus the only subcover is the cover itself. Thus there are only finitely many vertices in ∂F, and since each vertex has finite preimage, there are only finitely many points in S^1 mapped to vertices. The points between any two of these points must map to an edge in the boundary of F, by the continuity of f. Thus if we pick one vertex to start with, and one of its preimages in S^1, and then list vertices and edges in the order that one encounters their preimages while passing around S^1 clockwise, we obtain a cycle in Γ, which must contain all and only the edges and vertices wih realization in the boundary of F because f is surjective. ☐

Proposition 13.5. *Let E be a nice embedding of a locally finite connected graph $\Gamma = (V, E, \nu)$ with at least one vertex into a 2-manifold X. Then there is an embedding of $\Gamma^\dagger$ into X with respect to which $\Gamma^{\dagger\dagger} \cong \Gamma$.*

Proof. For a face F of E, let $f_F\colon \overline{D} \to X$ be the function required by the niceness of the embedding, and pick a point x_e in the realization of every edge e. By the definition of a nice embedding, the realization of every edge is in the boundary of 1 or 2 faces: if it is in the boundary of two faces F and G, then every point of it is the image of one point under f_F and of one point under f_G, and if it is in the boundary of one face F then it is the image of two points under f_F. Let

$$M = \bigcup_{F \text{ a face}} \{f_F(S) \mid S \text{ a line segment from } (0,0) \text{ to } x \text{ where } f_F(x) = x_e \text{ for some edge } e\};$$

we will show that M is an embedding of $\Gamma^\dagger$. The vertices of $\Gamma^\dagger$ are the faces of E: each face F is realized by the point $f_F((0,0))$. The edges of $\Gamma^\dagger$ are the edges of Γ: for each edge e between faces F and G, M contains a path from $f_F((0,0))$ to x_e, and from x_e to $f_G((0,0))$: then the concatenation of these paths realizes e (that the interior of the path is homeomorphic to $(0,1)$ follows from the fact that f_F and f_G are homeomorphisms). The realizations of edges and vertices clearly partition M. It remains only to show that a subset of M that meets the closure of every edge in a closed set in closed. Let $Q \subseteq M$ be such a set, and x a limit point of Q. If x is the realization of a vertex, then there are finitely many edges incident on that vertex, with realizations $A_1, A_2, \ldots, A_n$, and the set $\bigcup_{i=1}^n \overline{A_i} \cap Q$ is closed: but x must be a limit point of this set, since it has a neighbourhood (in M) contained entirely in $\bigcup_{i=1}^n \overline{A_i}$ and is a limit point of Q, and thus x is contained in this set, and therefore in Q. If x is contained in the realization A of an edge, then A is a neighbourhood of x in M, and must contain points of Q: thus x is a limit point of $\overline{A} \cap Q$, and since this set is closed, $x \in \overline{A} \cap Q \subseteq Q$. Thus M is an embedding of $\Gamma^\dagger$ into X.

It remains to show that $\Gamma^{\dagger\dagger} \cong \Gamma$. The edge sets of the two graphs are identical. If F is a face of M, then by Lemma 13.4, its boundary is the realization of a cycle of $\Gamma^\dagger$, and contains the realization of some edge; since a realization of an edge in M intersects exactly one realization of an edge in E in exactly one point, one vertex of the edge must be in F. Conversely, suppose there were more than one vertex of Γ with realization in F. There cannot be any edge of Γ which meets F and does not pass through the boundary of F (since every edge of Γ has an edge of $\Gamma^\dagger$ intersecting it); thus any edges incident on the vertices in F pass through edges in the boundary of F, and by Lemma 13.4 there are only finitely many of them, and they can be ordered. Since the realizations of edges in E can never cross, it is possible to find a path in $\overline{D}$ connecting the preimages under f_F of two vertices of $\Gamma^\dagger$ in the boundary of F, while never intersecting $f_F^{-1}(E)$ or $f_F^{-1}(M)$: but this means that the two vertices lie in the same face of E, and by construction, there are no two such vertices. Thus each face of M has exactly one vertex of E in it, and there is a bijection between the vertex sets of $\Gamma^{\dagger\dagger}$ and Γ. It remains only to show that and edge of Γ connects two vertices of Γ if and only if the equivalent edge of $\Gamma^{\dagger\dagger}$ connects the equivalent vertices of $\Gamma^{\dagger\dagger}$, or equivalently, an edge of Γ connects two vertices if and only if its realization intersects the realization of and edge of $\Gamma^{\dagger\dagger}$ contained in the boundary of the faces containing the same vertices; □

14. Simply connected covers of graphs

In this section, we look at one last topologically motivated property of graphs, that of simply connected covers. A *covering* of a graph Γ is a graph Δ together with a graph homomorphism $f\colon \Delta \to \Gamma$ that is surjective on both vertices and edges and that preserves the degree of vertices. Being simply connected is a topological property, which in the case of graphs reduces to being acyclic.

Proposition 14.1. *Every connected graph has a covering by a tree.*

Proof. Let Γ be a connected graph, and let Γ' be the graph obtained from Γ by replacing every loop e in Γ by two loops e_1 and e_2 on the same vertex as e. To every edge of Γ' we associate an edge of Γ, which we call the *underlying edge:* for an edge e of Γ', the underlying edge is e itself if e is not a loop, and the loop e replaces if it is a loop. If Γ is empty, then it is itself a tree, and covers itself by the identity map; assume therefore that Γ (and therefore Γ') has a vertex v_0. Let P be the set of walks in Γ' starting at v_0 and never using the same edge twice in a row, let $Q = \{(p, q) \in P \mid p \text{ is a prefix of } q \text{ and } \ell(q) = \ell(p) + 1\}$, and let $\Delta = (P, Q, (p, q) \mapsto \{p, q\})$. We will show that Δ is a tree covering Γ.

First we show that Δ is a tree by showing there exists a unique path between any two vertices. A walk in Δ is equivalent to a sequence of operations, each of which is either adding an additional vertex and edge to the end of a walk or removing the final vertex and edge from a walk. Let p and q be vertices in Δ: they have a longest common prefix (which consists of at least the trivial path on v_0). The next vertex in p after this prefix must be deleted at some point, since it's not in q: thus any path from p to q must start with enough deletions to end up at the longest common prefix (if we allowed any additions, we would need to delete them again eventually to get to the prefix, and the walk would not be a path). Next the path must add any vertices needed to get to q, and these can be the only operations in the path (if we deleted any vertices in the prefix or the we added to get to q we will need to add them again eventually, and if we add any vertices not in q, we will need to delete them again eventually, and in either case the walk would not be a path). Thus there is a path from p to q, and it is unique.

Now we show Δ covers Γ. Let $f\colon \Delta \to \Gamma$ be the graph homomorphism that maps vertices in Δ to their terminal vertex, and edges in Δ to the underlying edge of their last edge. Then f is clearly a homomorphism, and because Γ is connected, Γ' is also, and there is a path from v_0 to any vertex, which makes f surjective on vertices. Similarly, f is surjective on edges, since there is always a path from v_0 to one endpoint of the edge, which can be extended to a walk ending in the edge. Moreover, f preserves the degree of vertices: if a vertex of Δ is a walk ending at a vertex v of Γ, then it has one incident edge for the edge leading into it, and one for each edge leading out, which does not include the edge leading in, since the walks represented by vertices in Δ never use the same edge twice in a row. Since loops in Γ are duplicated in Γ', they are counted twice in the degree of corresponding vertex of Δ. Thus a vertex of Δ has the same degree as the vertex f maps it to, and f and Δ cover Γ. $\square$

If Γ is a tree, the covering graph constructed in Proposition 14.1 is isomorphic to Γ, but in the case of other graphs, the tree will be considerably larger, and in fact infinite. For example, the cover of a cycle of any number of vertices (including a single vertex with a loop) is an infinite line, and the cover of the graph

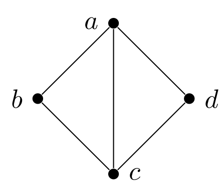

is the tree

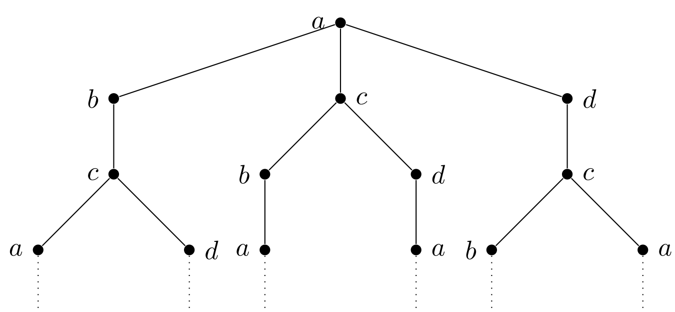

where vertices are labelled with the vertex they are mapped to by the covering.

References

1. Paolo Aluffi, *Algebra: chapter 0*, Graduate Studies in Mathematics, vol. 104, American Mathematical Society, Providence, RI, 2009. MR 2527940
2. Reinhold Baer, *Abelian groups without elements of finite order*, Duke Math. J. **3** (1937), no. 1, 68–122. MR 1545974
3. Chad Couture, *Skew-zigzag algebras*, SIGMA Symmetry Integrability Geom. Methods Appl. **12** (2016), 062, 19 pages. MR 3514943
4. Reinhard Diestel, *Graph theory*, fourth ed., Graduate Texts in Mathematics, vol. 173, Springer, Heidelberg, 2010. MR 2744811
5. Ross Geoghegan, *Topological methods in group theory*, Graduate Texts in Mathematics, vol. 243, Springer, New York, 2008. MR 2365352
6. Peter Giblin, *Graphs, surfaces and homology*, third ed., Cambridge University Press, Cambridge, 2010. MR 2722281

7. Peter John Hilton and Urs Stammbach, *A course in homological algebra*, Springer-Verlag, New York-Berlin, 1971, Graduate Texts in Mathematics, Vol. 4. MR 0346025
8. Ruth Stella Huerfano and Mikhail Khovanov, *A category for the adjoint representation*, J. Algebra **246** (2001), no. 2, 514–542. MR 1872113