III. Graph complexes

In this unit we shall study the fundamental groups for a special class of spaces which are built out of very simple pieces but turn out to be important in many branches of mathematics, and in some sense are "toy models" for the sorts of objects usually studied in algebraic and geometric topology. More precisely, these spaces (called finite graph complexes, edge-path graphs or more simply just *graphs*) are excellent test cases for applying the methods and results of this course.

Informally, a graph can be constructed by taking a finite collection of closed intervals and identifying their endpoints in a suitable fashion; following geometric intuition, the images of the intervals are called *edges* and the images of their endpoints are called *vertices*. Note that these are NOT graphs as defined and studied in coordinate geometry and calculus, but the name has stuck and become standard usage, both in mathematics and in its applications to numerous other subjects such as computer science, physics, chemistry, industrial engineering, the biological sciences and even to other areas of knowledge where it is useful to look at chains of relationships or passage from one state of a system to another. A fairly simple application of graph theory to a problem about relationships is given in the following online video:

`http://www.youtube.com/watch?v=b3lbjoiEAyA`

One of the main results in this unit is a complete description of the fundamental group of a connected finite graph using the numbers of edges and vertices. This result in turn leads to an algebraic theorem about free groups that is somewhat nonintuitive: If F is a free group on a finite number $n > 1$ of generators, then for each $m > n$ the group F contains a subgroup of finite index with m generators; in contrast, if A is a free abelian group on n generators, then every subgroup of A has at most n generators.

III.1 : Basic definitions

(**M**, §83; **H**, §1.A, Ch. 2 Introduction)

Since we have already described finite graphs intuitively, we shall proceed to the formal description.

Definition. A **finite edge-path graph complex** (more simply a *finite graph*) is a pair $(X, \mathcal{E})$ consisting of a compact Hausdorff space X and a finite family $\mathcal{E}$ of closed subsets with the following properties:

(1) Each subset $E \in \mathcal{E}$ is homeomorphic to the closed interval $[0, 1]$.

(2) The space X is the union of all the subspaces E in the family $\mathcal{E}$.

(3) If E_1 and E_2 are distinct subsets of $\mathcal{E}$, then either $E_1 \cap E_2$ is empty or else it is a single point corresponding to a vertex of each interval E_i.

COMMENTS ON THE DEFINITION. The endpoints of a set homeomorphic to $[0, 1]$ are topologically characterized by the fact that their complements are connected; for all other points, the complement has two components. As above, we shall say that a subset of $\mathcal{E}$ is an edge and an endpoint of an *edge* will be called a *vertex*.

The setting in Chapter 14 of Munkres is more general and includes examples where the set of edges is infinite but each vertex lies on only finitely many edges. We are restricting attention to examples with finitely many edges in order to simplify the discussion.

Examples. It is easy to draw many examples of graphs, and such drawings are extremely useful for understanding this concept. The file `graphpix1.pdf` contains a few examples, including some that will appear later in this course.

An alternate definition

Our definition of a graph assumes that two edges meet in just one endpoint, but in some situations it is convenient to consider examples for which the intersection of two edges is also allowed to be both vertices of the two edges as in the following illustration:

D

(Two vertices at the corners, two edges have these endpoints.)

We shall prove that every object of this more general type can be expressed as a graph in the sense of our definition.

LEMMA 1. *Let Γ be a system satisfying the conditions for an finite edge-vertex graph except that two edges may have both of their vertices, and let $\mathcal{E}$ be the collection of edges for this system. Then there is another family of closed subsets $\mathcal{E}'$ such that the following hold:*

(i) The family $\mathcal{E}'$ is a collection of edges for a graph structure on Γ.

(ii) Each element of $\mathcal{E}'$ is contained in a unique element of $\mathcal{E}$ such that one endpoint of $\mathcal{E}'$ is also an endpoint for $\mathcal{E}$ but another is not, and each edge in $\mathcal{E}$ is a union of two edges in $\mathcal{E}'$.

*(iii) The intersection of two distinct edges in $\mathcal{E}'$ is **at most** a common vertex.*

Proof. For each edge $E \in \mathcal{E}$, pick a point $b_E \in E$ that is not an endpoint. It follows that $E - \{b_E\}$ has two connected components, each of which contains exactly one endpoint of E. If x is an endpoint of E define the set $[x, E]$ to be the closure of the component of $E - \{b_E\}$ which contains x. If $\mathcal{E}'$ denotes the set of all such subsets $[x, E]$, then it follows immediately that $\mathcal{E}'$ has the properties stated in the lemma. Note that by construction the endpoints of a given edge $[x, E]$ are x and b_E.■

The family $\mathcal{E}'$ is frequently called the *derived* graph structure associated to $\mathcal{E}$.

As noted in one of the exercises, many examples of edge-vertex graphs are suggested by ordinary letters and numerals.

Subgraphs

Definition. Let $(X, \mathcal{E})$ be a finite edge-path graph. A *subgraph* $(X_0, \mathcal{E}_0)$ is given by a subfamily $\mathcal{E}_0 \subset \mathcal{E}$ such that X_0 is the union of all the edges in $\mathcal{E}_0$. It is said to be a *full subgraph* if two vertices $\mathbf{v}$ and $\mathbf{w}$ lie in X_0 and there is an edge $E \in \mathcal{E}$ joining them, then $E \in \mathcal{E}_0$.

PROPOSITION 2. *Let $(X, \mathcal{E})$ be a finite edge-path graph, and let $(X_0, \mathcal{E}_0)$ be a subgraph. Then the derived graph $(X_0, \mathcal{E}_0')$ is a full subgraph of $(X, \mathcal{E}')$.*

Proof. Suppose we are given an edge K in $(X, \mathcal{E}')$, so that its vertices must have the form $\mathbf{y}$, $\mathbf{m}_L$ where L is an edge in $(X, \mathcal{E})$ that has $\mathbf{y}$ as one of its endpoints and $\mathbf{m}_L$ is a non-vertex point in L. If both of these vertices belong to X_0, then the latter contains a point of L which is not a vertex, and since $(X_0, \mathcal{E}_0)$ is a subgraph it follows that L must be entirely contained in X_0. But this automatically implies that the edge in the derived complex with endpoints $\mathbf{y}$ and $\mathbf{m}_L$ must be contained in X_0.■

Connectedness

One immediate consequence of the definitions is that every point of a graph lies in the arc component of some vertex; specifically, if x lies on the edge E and the vertices of the latter are a and b, then x lies in the same arc component as both a and b. In fact, one can prove much stronger conclusions:

PROPOSITION 3. *If $(X, \mathcal{E})$ is a finite edge-path graph, then X is connected if and only if for each pair of distinct vertices $\mathbf{v}$ and $\mathbf{w}$ there is an* edge-path sequence *$E_1, \cdots, E_n$ such that $\mathbf{v}$ is one vertex of E_1, $\mathbf{w}$ is one vertex of E_n, for each k satisfying $1 < k \leq n$ the edges E_k and E_{k-1} have one vertex in common, and $\mathbf{v}$ and $\mathbf{w}$ are the "other" vertices of E_1 and E_n. Furthermore, X is a union of finitely many components, each of which is a full subgraph.*

IMPORTANT: In a general edge-path sequence defined as in the statement of the proposition, we do **NOT** make any assumptions about whether or not these two vertices are equal. If they are, then we shall say that the edge-path sequence is *closed* or that it is a *circuit* or **cycle**.

Proof. First of all, since every point lies on an edge, it follows that every point lie in the connected component of some vertex. In particular, there are only finitely many connected components. Define a binary relation on the set of vertices such that $\mathbf{v} \sim \mathbf{w}$ if and only if the two vertices are equal or there is an edge-path sequence as in the statement of the proposition. It is elementary to check that this is an equivalence relation, and that vertices in the same equivalence class determine the same connected component in X.

Given a vertex $\mathbf{v}$, let $Y_\mathbf{v}$ denote the union of all edges containing vertices which are equivalent to $\mathbf{v}$ in the sense of the preceding paragraph. If we choose one vertex $\mathbf{v}$ from each equivalence class, then we obtain a finite, pairwise disjoint family of closed connected subsets whose union is X, and it follows that these sets are must be the connected components of X. In fact, by construction each of these connected component is a full subgraph of $(X, \mathcal{E})$.■

Frequently it is convenient to look at edge-path sequences that are *minimal* or *simple* in the sense that one cannot easily extract shorter edge-path sequences from them. Here is a more precise formulation:

Definition. Let $E_1, \cdots, E_n$ be an edge-path sequence such that the vertices of E_i are $\mathbf{v}_{i-1}$ and $\mathbf{v}_i$. This sequence is said to be *reduced* if $\mathbf{v}_1, \cdots, \mathbf{v}_n$ are distinct and either $n \neq 2$ or else $\mathbf{v}_0 \neq \mathbf{v}_2$ (if $n = 2$ and $\mathbf{v}_0 = \mathbf{v}_2$, then the edge-path is just a sequence with $E_2 = E_1$, physically corresponding to going first along E_1 in one direction and then back in the opposite direction).

We then have the following result:

PROPOSITION 4. *If two distinct vertices $\mathbf{x}$ and $\mathbf{y}$ can be connected by an edge-path sequence, then they can be connected by a reduced sequence.*

Proof. Take a sequence with a minimum number of edges. We claim it is automatically reduced. If not, then there is a first vertex which is repeated, and a first time at which it is repeated. In other words, there is a minimal pair (i, j) such that $i < j$ and $\mathbf{v}_i = \mathbf{v}_j$, which means that if (p, q) is any other pair with this property we have $p \geq i$ and $q > j$. If we remove E_{i+1} through E_j from the edge-path sequence, we obtain a shorter sequence which joins the given two vertices.■

There may be several different reduced sequences joining a given pair of vertices. For example, take X to be a triangle graph in the plane whose vertices are the three noncollinear points **a**, **b** and **c**, and whose edges are the three line segments joining these pairs of points. Then **ab**, **bc** and **ac** are two reduced edge-path sequences joining **a** to **c**.

Definition. A circuit (or cycle) $E_1 \ , \cdots \ , E_n$ is called a a *simple circuit* or *simple cycle* if it is reduced.

COROLLARY 5. *Every simple circuit in a graph contains at least three edges.*■

Further topological properties of graphs

By definition and construction, a finite edge-path graph is compact Hausdorff, and in fact one can say considerably more:

PROPOSITION 6. *A finite edge-path graph is homeomorphic to a subset of* $\mathbb{R}^n$ *for some* n.

At the end of this section we shall prove that a graph is always homeomorphic to a subset of $\mathbb{R}^3$.

Proof. Suppose that the vertices are $\mathbf{v}_1, \ \cdots \ , \mathbf{v}_n$. Consider the graph in $\mathbb{R}^n$ whose vertices are the standard unit vectors $\mathbf{e}_i$ and whose edges are the closed line segments $A_{i,j}$ joining these vertices; the resulting compact subspace of $\mathbb{R}^n$ is a graph because two of these segments intersect in at most a common endpoint (use linear independence of the unit vectors to prove this). Define a continuous map f from original the graph X to the new graph Y such that if E is an edge with vertices $\mathbf{v}_i$ and $\mathbf{j}$ for $i < j$ and E is a homeomorphism from $[0, 1]$ to E such that $\mathbf{v}_i$ corresponds to 0 and $\mathbf{v}_j$ corresponds to 1, then $t \in [0, 1]$ is sent to

$$t\,\mathbf{e}_j \ + \ (1-t)\,\mathbf{e}_i$$

(since E is homeomorphic to $[0, 1]$ and endpoints are topologically characterized by the property that their complements are connected, it follows that either $\mathbf{v}_i$ corresponds to 0 and $\mathbf{v}_j$ corresponds to 1 or vice versa; in the second case, if we compose the original homeomorphism with the reflection on $[0, 1]$ sending s to $1 - s$, then we obtain a homeomorphism for which the first alternative holds).

It is a routine exercise to verify that f is continuous and 1–1, and therefore it maps X homeomorphically onto its image.■

The next result implies that a graph has a simply connected (universal) covering space.

PROPOSITION 7. *If* $(X, \mathcal{E})$ *is a finite edge-path graph and* $x \in X$, *then* x *has a (countable) neighborhood base of contractible open subsets.*

Proof. Suppose first that x is a vertex of X, and view X as a subset of $\mathbb{R}^n$ using the previous result. Define $\mathbf{OpenStar}\,(x, \mathcal{E})$ to be the complement of the set of all points on edges E' which do not have x as a vertex. The described set is the union of all E' which do not have x as a vertex and hence is closed, so its complement **OpenStar** $(x, \mathcal{E})$ must be open. For every ε such that $0 < \varepsilon < \sqrt{2}$ let

$$\mathbf{OpenStar}\,(\varepsilon; x, \mathcal{E})$$

denote the points in **OpenStar** $(x, \mathcal{E})$ whose distance from x is less than ε. Then it follows that there is a straight line homotopy from the identity on **OpenStar** $(\varepsilon; x, \mathcal{E})$ to the constant map with value x, and therefore every such neighborhood is contractible. Since X is presented as a metric space, it follows that a suitably chosen countable of this neighborhood family will be the desired countable neighborhood base of x.

Now suppose that x is not a vertex of X, so that there is a unique edge E containing x; by assumption x lies in the complement of the end points in E, and the corresponding subset of E is homeomorphic to the open interval $(0, 1)$. Since every point in $(0, 1)$ has a neighborhood base of contractible open subsets, the conclusion to the proposition will follow if we know that E − endpoints is open in X. The complement to this set is the set of all points that are either vertices of E or else lie on some edge other than E. This is a finite union of closed sets and hence closed, and therefore the set $E - \{\text{endpoints}\}$ must be open in X as desired.■

Addendum: Embedding graphs in $\mathbb{R}^3$

For many purposes it is enough to know that a connected graph is always topologically embeddable in some $\mathbb{R}^n$, but in some contexts it is useful to know the smallest n for which this is possible. The methods of point set topology show that a connected compact subset of $\mathbb{R}$ is an interval, and the final result of this section shows that for all other graphs the minimum value of n is 2 or 3.

THEOREM 8. *Let $(X, \mathcal{E})$ be a connected graph. Then there is a $1-1$ continuous mapping $\varphi : X \to \mathbb{R}^3$ such that every edge in $\mathcal{E}$ is mapped linearly to the closed segment joining the images of the vertices (in other words, the embedding is* rectilinear*).*

There are obvious examples for which we can take $n = 2$ (in particular, the complete graphs on three or four vertices), and later in this course we shall prove that $n = 3$ for two specific graphs; in other words, these graphs are not homeomorphic to subsets of $\mathbb{R}^2$ and are said to be *nonplanar*. A celebrated theorem due to K. Kuratowski characterizes all nonplanar connected graphs in terms of the two specific nonplanar examples that we shall analyze in Section VII.4.

Proof of Theorem 8. The proof is built around the following two key observations:

(A) *There is an infinite set of isolated points in $\mathbb{R}^3$ such that no four lie in a single plane.*

(B) *For the sequence of points in the preceding observation, the intersections of two line segments joining two pairs of vertices is at most a common vertex.*

Before proving these, we shall indicate how they yield the theorem. Given the connected graph $(X, \mathcal{E})$, define φ on vertices so that each vertex goes to a point in the sequence of (A) and the mapping is 1–1. Next, if K is an edge of $\mathcal{E}$ with vertices a and b, extend φ to K by sending the latter homeomorphically to the closed line segment joining $\varphi(a)$ to $\varphi(b)$. Construct such an extension for each edge in the graph. Then observation (B) implies that if $K \neq K'$ are distinct edges of $\mathcal{E}$, it follows that $\varphi[K] \cap \varphi[K']$ is at most a common vertex of the two line segments. Since the map on vertices is 1–1, it follows that this common vertex is a common vertex of K and K', and therefore the mapping φ will be continuous and 1–1, so that ϕ maps X homeomorphically to its image because X is compact.

Verification of observation (A). — Obviously $\mathbb{R}^3$ contains a set of 4 such points; for example, take $\mathbf{0}$ and the three standard unit vectors. Suppose we know that there is a set of n points, say

Y, satisfying the property in (A); by induction, it will suffice to find a similar subset containing Y with one additional point. Let

$$P_i \ , \qquad 1 \ \leq i \ \leq \ \binom{n}{3}$$

be the planes determined by triples of points in Y. Then $\mathbb{R}^3 - \cup_i P_i$ is an open dense subset. Form Y' by adding a point z in the complement of $\cup_i P_i$ such that $|z| \geq n+1$. If we continue to define points recursively in this manner, we obtain the sort of subset described in (A), and in fact is it a closed subset of $\mathbb{R}^3$.

Verification of observation (B). — There are two cases to consider, depending upon whether or not the two edges have a common vertex. Both rely upon the following elementary fact from linear algebra: *If* $\mathbf{a}$, $\mathbf{b}$, $\mathbf{c}$ *and* $\mathbf{d}$ *are nonplanar points in* $\mathbb{R}^3$, *then* $\mathbf{b}-\mathbf{a}$, $\mathbf{c}-\mathbf{a}$ *and* $\mathbf{d}-\mathbf{a}$ *are linearly independent.* (If the points were coplanar, then there would be a 2-dimensional vector subspace containing the three difference vectors.)

CASE 1. Suppose that the segments have one endpoint in common, so that the endpoints of one segment are $\mathbf{x}$ and $\mathbf{y}$ while the endpoints of the other are $\mathbf{y}$ and $\mathbf{z}$. We claim that they cannot have any other points in common. If there were such a point then there would be $s, t \in [0,1]$ such that

$$t\mathbf{x} \ + \ (1-t)\mathbf{y} \ \ = \ \ s\mathbf{z} \ + \ (1-s)\mathbf{y}$$

so that $t(\mathbf{x}-\mathbf{y}) = s(\mathbf{z}-\mathbf{y})$. Since there is a fourth point $\mathbf{w}$ such that $\mathbf{y}$, $\mathbf{x}$, $\mathbf{z}$ and $\mathbf{w}$ are nonplanar, it follows that $\mathbf{x}-\mathbf{y}$ and $\mathbf{z}-\mathbf{y}$ are linearly independent, and consequently the only way the displayed equation can be valid is if $s = t = 0$, in which case we have

$$t\mathbf{x} + (1-t)\mathbf{y} \ \ = \ \ \mathbf{y} \ \ = \ \ s\mathbf{z} + (1-s)\mathbf{y}$$

so that the common point must be the common endpoint of the two segments.

CASE 2. Suppose that the segments have no endpoints in common and that they join $\mathbf{y}$ to $\mathbf{x}$ and $\mathbf{z}$ to $\mathbf{w}$ respectively. As in the preceding case, if the segments did have some point in common, then there would be $s, t \in [0,1]$ such that

$$t\mathbf{x} \ + \ (1-t)\mathbf{y} \ \ = \ \ s\mathbf{w} \ + \ (1-s)\mathbf{z}$$

and if we subtract one side of this equation from the other we obtain the equation

$$t\mathbf{x} \ + \ (1-t)\mathbf{y} \ - \ s\mathbf{w} \ + \ (s-1)\mathbf{z} \ \ = \ \ \mathbf{0} \ .$$

By linear independence this equation is satisfied if and only if each of $t, (1-t), -s, (s-1)$ is zero. However, this is impossible because it implies both $t = 0$ and $t = 1$. Therefore the two segments cannot have any points in common. — As indicated previously, this completes the proof of Theorem 8.■

III.2 : Maximal trees

(**M**, §84; **H**, §1.A))

The graph $(X, \mathcal{E})$ is said to be a *tree* if for distinct vertices u and v in X there is a UNIQUE reduced edge path sequence

$$E_1, \ \cdots \ , E_r$$

such that u is the "initial" endpoint of E_1 and v is the "final" endpoint of E_r.

A reduced edge path sequence is sometimes called a **simple chain** in the graph.

One can visualize many examples of trees by looking at letters of the alphabet; examples include the letters

E, F, H, I, K, M, N, V, W, X, Y, Z.

The numerals 4 and 5 as depicted on a standard calculator display (with an open top on the 4) also correspond to examples of trees. On the other hand, the linear graphs corresponding to triangles, rectangles, pentagons, etc. are not trees. Other nonexamples include the letter A, the numeral 4 as depicted in print with a closed top, and the numerals 6 and 8 as depicted on a standard calculator display.

PROPOSITION 1. *Every tree has a vertex that lies on only one edge.*

Proof. If the tree has only one edge, then the result is trivial. Assume now that the tree $(X, \mathcal{E})$ has $m \geq 2$ edges. We shall assume further that every vertex of $(X, \mathcal{E})$ lies on at least two distinct edges and derive a contradiction.

Let A_1 be an edge of $(X, \mathcal{E})$, and let v_0 and v_1 be its edges. Let A_2 be a second edge which has v_1 as a vertex, and let v_2 be its other vertex. Continuing in this manner, we obtain an infinite sequence of edges A_n with vertices v_{n-1} and v_n such that $A_n \neq A_{n-1}$. Since there are only finitely many edges and vertices, there must be some first value of k such that $v_k = v_{k+j}$ for some $j > 0$ (in other words, there is a first repeated value in the sequence). By construction, we must have $j \geq 3$. By construction, we know that A_{k+j} defines a simple chain joining $v_k = v_{j+k}$ to v_{j+k-1} and similarly the sequence $A_{k+1}, \ \cdots \ A_{k+j-1}$ defines a simple chain joining these two vertices. Now the first simple chain consists of one edge, while the second consists of at least two because $k + j - 1 \geq k + 2 > k + 1$, and thus we have constructed two simple chains joining these vertices. Since we are assuming the graph is a tree, this is impossible, and therefore it follows that there must be some vertex which lies on only one edge.■

We shall also need the following companion result:

PROPOSITION 2. *Suppose $(X, \mathcal{E})$ is a tree and v_0 is a vertex which lies on only one edge, say E_0. Let $(X_0, \mathcal{E}_0)$ be the subgraph given by the union of all edges except E_0 (hence its vertices are all the vertices of the original graph except v_0). Then $(X_0, \mathcal{E}_0)$ is also a tree.*

Proof. Suppose that u and w are vertices of the subgraph and $A_1, \ \cdots \ , A_r$ is a simple chain connecting them. We claim that none of the edges A_i can be equal to E_0; if this is true then it will follow that the subgraph will be a tree (see the final step of the argument).

As usual, label the vertices of the edges A_i such that $u = a_0$, $w = a_r$, and the vertices of A_i are a_i and a_{i-1}. By hypothesis, $A_i \neq A_{i\pm1}$ for all i. Suppose that we have $A_j = E_0$ for some j; then either $v_0 = a_{i-1}$ or else $v_0 = a_i$. Let v_1 be the other vertex of E_0.

CASE 1: Suppose that $v_0 = a_{i-1}$. Since $a_0 = u \neq v_0$, it follows that $i > 0$. Since E_0 is the only edge containing the vertex v_0, it follows that $A_{i-1} = A_i$, with $v_1 = a_{i-2} = a_i$. This contradicts the definition of a simple chain, and hence we can conclude that $v_0 \neq a_{i-1}$. CASE 2. Suppose that $v_0 = a_i$. Since $a_r = w \neq v_0$, it follows that $i < r$. Since E_0 is the only edge containing the vertex v_0, it follows that $A_{i+1} = A_i$, with $v_1 = a_{i-1} = a_{i+1}$. This contradicts the definition of a simple chain, and hence we can conclude that $v_0 \neq a_i$. Combining these results, we can conclude that E_0 does not appear in the simple chain sequence $A_1, \cdots, A_r$, so that the latter is a simple chain in $(X_0.\mathcal{E}_0$. This simple chain is unique in the smaller complex by the uniqueness condition on the larger complex, and therefore the smaller complex must also be a tree.■

We are now ready to state one of the most important properties of trees:

THEOREM 3. *If $(T, \mathcal{E})$ is a tree and $\mathbf{v}$ is a vertex of this graph, then $\{\mathbf{v}\}$ is a strong deformation retract of X.*

Proof. This is trivial for graphs with one edge because they are homeomorphic to the unit interval. Suppose now that we know the result for trees with at most n edges, and suppose that $(T, \mathcal{E})$ has $n + 1$ edges.

By Lemma 84.2 we may write $T = T_0 \cup A$ where A is an edge and T_0 is a tree with n edges such that $A \cap T_0$ is a single vertex $\mathbf{w}$. Let $\mathbf{y}$ be the other vertex of A. The proof splits into cases depending upon whether or not the vertex $\mathbf{v}$ of T is equal to $\mathbf{y}$, $\mathbf{w}$ or some other vertex in T_0.

We shall need the following two results on strong deformation retracts; in both cases the proofs are elementary:

(1) *Suppose X is a union of two closed subspaces $A \cup B$, and let $A \cap B = C$. If C is a strong deformation retract of both A and B, then C is a strong deformation retract of X.*

(2) *Suppose X is a union of two closed subspaces $A \cup B$, and let $A \cap B = C$. If C is a strong deformation retract of B, then A is a strong deformation retract of X.*

Suppose first that the vertex is $\mathbf{w}$. Then $\{\mathbf{w}\}$ is a strong deformation retract of both A and T_0, so by the first statement above it is a strong deformation retract of their union, which is T.

Now suppose that the vertex is $\mathbf{y}$. Then the second statement above implies that A is a strong deformation retract of T. Since $\{\mathbf{y}\}$ is a strong deformation retract of A, it follows that $\{\mathbf{y}\}$ is also a strong deformation retract of T.

Finally, suppose that the vertex $\mathbf{v}$ lies in T_0 but is not $\mathbf{w}$. Another application of the second statement implies that T_0 is a strong deformation retract of T, and since $\{\mathbf{v}\}$ is a strong deformation retract of T_0, it follows that $\{\mathbf{v}\}$ is also a strong deformation retract of T.■

COROLLARY 4. *The fundamental group of a tree is trivial.*■

Definition. Let $(X, \mathcal{E})$ be a graph. A subgraph $M \subset X$ is a *maximal tree* in X if M is a tree and there is no tree M' in X which properly contains M.

It is fairly straightforward to show that maximal trees exist. First of all, X must contain subgraphs that are trees, for any subgraph consisting of a single edge is a tree. Because of this, it follows that there must be some tree in X with a maximal number of edges, and this will be a maximal tree.■

For the sake of completeness, we state the following elementary result:

LEMMA 5. *If $(X, \mathcal{E})$ is a graph with a maximal tree M and Y is a subgraph of X containing M, then M is a maximal tree in Y.*■

Finally, we shall need the following important property of maximal trees:

PROPOSITION 6. *If $(X, \mathcal{E})$ is a connected graph and $T \subset X$ is a maximal tree, then all the vertices of $(X, \mathcal{E})$ belong to T.*

Proof. As usual, assume the conclusion is false; then there is some vertex $v \notin T$. By connectedness there is an edge-path sequence joining v to some point in T, and among these sequences there is one $E_1, \cdots, E_n$ of minimum length. Since we have an edge-path sequence we can denote the vertices on the edges by $v_0, \cdots, v_n$ such that $v = v_0$, $v_n \in T$, and the edges of E_i are v_i and v_{i-1}. By the minimality of this sequence we know that $v_i \in T$ if and only if $i = n$.

Let $T_1 = T \cup E_n$. We claim that T_1 is also a tree. The key point in verifying this will be the following observation:

> If an edge-path sequence in T_1 contains E_n, then E_n is either the first edge or the last edge, and v_{n-1} is either the initial vertex or the final vertex.

This is true because the vertex v_{n-1} lies on E_n but not on any edges in T (if it did, then $v_{n-1} \in T$ and by our assumptions this is not the case). If E_n appeared in the middle of the sequence, the one of the two edges containing v_{n-1} would have to lie in T, and this would imply $v_{n-1} \in T$.

To prove that T_1 is a tree, consider an arbitrary pair of vertices w and w'. If they both lie in T, then there is a unique reduced edge-path in T joining them, and we claim that there is no other reduced edge-path in T_1 which joins them. Any such path would have to contain the edge E_n (the only edge not in T). Since a reduced edge-path containing E_n must start or end with E_n, such an edge-path cannot join two points in T. — Now consider reduced edge-path sequences joining v_{n-1} to some vertex w in T. Since $v_n \in T$, there is a unique edge-path sequence $K_1, \cdots, K_m$ joining v_n to w. If we insert E_n at the beginning of this sequence, we obtain a reduced edge-path sequence joining v_{n-1} to w in T_1. To see that this sequence is unique, note that every edge-path sequence joining v_{n-1} to w must start with E_n because no other edge in T_1 contains v_{n-1}. If we remove E_n from the sequence, we obtain a reduced edge-path sequence joining v_n to w, and since E_n does not appear in this sequence it must be an edge-path sequence in T. Therefore the sequence joining v_n to w must be the previously described edge-path sequence $K_1, \cdots, K_m$, and it follows that there is only one edge-path sequence in T_1 joining v_{n-1} to w.

The preceding shows that T_1 is a tree in X which properly contains T. Since T was assumed to be a maximal tree, this yields a contradiction, so our hypothesis about a vertex not in T must be false and hence T must contain all the vertices.■

III.3 : Fundamental groups of graphs

(**M**, §84; **H**, §1.A))

In this section we shall show that the fundamental group of a connected graph $(X, \mathcal{E})$ has a very simple description depending only upon the numbers of vertices and edges in $\mathcal{E}$.

We already know that the fundamental group of a tree is trivial, and the crucial step in proving the main result is to describe the fundamental group of a connected graph which is the union of a maximal tree and a single edge. A complete graph on three vertices has this property, and its fundamental group is infinite cyclic because such a graph is homeomorphic to S^1 (verify this), and the first result is a generalization of this fact to all graphs which are unions of a tree and a single edge.

PROPOSITION 1. *Suppose that the connected graph $(X, \mathcal{E})$ contains a maximal tree T such that X is the union of T with a single edge E^*. Then X is homotopy equivalent to S^1.*

The file `graphpix2.pdf` in the course directory contains a drawing and a simplified discussion of the main ideas in the proof.

Proof. Since T is a maximal tree, the vertices of E^* lie in T. If $\mathbf{a}$ and $\mathbf{b}$ are these vertices, then there is a reduced edge-path sequence $E_1, \cdots, E_n$ joining $\mathbf{a}$ to $\mathbf{b}$, and if we let Γ be the union of the $E - i$'s and E^*, it follows that Γ must be homeomorphic to S^1. By construction Γ determines a subgraph of X. For the sake of uniformity, set $\mathbf{v}_0 = \mathbf{a}$ and $\mathbf{v}_n = \mathbf{b}$.

We claim that Γ is a strong deformation retract of X. Let Y be the subgraph obtained by removing the edges E^* and E_i from $\mathcal{E}$, and for each i let Y_i be the component of $\mathbf{v}_i$. By our assumptions it follows that Y and the subgraphs Y_i are trees. It will suffice to prove that if $i \neq j$ then $\mathbf{v}_j \notin Y_i$, for then we have $Y_i \cap \Gamma = \{\mathbf{v}_i\}$ and we can repeatedly apply the criteria in the previous argument to show that Γ is a strong deformation retract of X.

Suppose now that $\mathbf{v}_j \notin Y_i$ for some $j \neq i$. Then there is some reduced edge-path sequence $F_1, \cdots, F_m$ joining $\mathbf{v}_i$ to $\mathbf{v}_j$ in Y_i. Since the vertices of the edges F_r contain at least one $\mathbf{v}_j$ other than $\mathbf{v}_i$, there is a first edge in the sequence F_s which contains such an edge, say $\mathbf{v}_k$. Of course, none of the edges F_r lies in Γ. However, we also know that there is a reduced edge path sequence in $\Gamma \cap T$ which joins $\mathbf{v}_j$ to $\mathbf{v}_k$, and we can merge this with the edge-path sequence $F_1, \cdots, F_s$ (whose edges lie in T but not Γ) to obtain a reduced cycle in T. Since T is a tree, this is a contradiction, and therefore we must have $Y_i \cap \Gamma = \{\mathbf{v}_i\}$. As noted before, this suffices to complete the proof.■

The preceding special case is a key step in proving the following general result:

THEOREM 2. *Let $(X, \mathcal{E})$ be a connected graph, let T be a maximal tree in X, and let p be a vertex of T. Then $\pi_1(X, x)$ is a free group on k generators, where k is the number of edges that are in X but not in T.*

Let T be a maximal tree in the connected graph X, and let $F_1, \cdots, F_b$ denote the edges of X which do not lie in T. Let $W \subset X$ be the open set obtained by deleting exactly one non-vertex point from each of the edges F_i, and let $U_j = W \cup F_j$. It then follows that each U_j is an open subset of X and if $i \neq j$ then $U_i \cap U_j = W$. Furthermore T is a strong deformation retract of W and for each subset of indices $i_1, \cdots, i_k$ the set $F_{i_1} \cup F_{i_k}$ is a strong deformation retract of $U_{i_1} \cup U_{i_k}$. In particular, we know that the sets W and U_i are all arcwise connected. By the preceding result we know that F_1 and U_1 are homotopy equivalent to S^1, and we claim by induction that the fundamental groups of $F_1 \cup \cdots \cup F_t$ and $U_1 \cup \cdots \cup U_t$ are free on t generators. For if the result is true for $t \geq 1$, then we have

$$\bigcup_{i \leq t+1} U_i = \left(\bigcup_{i \leq t} U_i \right) \cup U_{t+1} , \qquad W = \left(\bigcup_{i \leq t} U_i \right) \cap U_{t+1}$$

so that the Seifert-van Kampen Theorem implies that the fundamental group of $U_1 \cup \cdots \cup U_{t+1}$ is the free product of the fundamental groups of $U_1 \cup \cdots \cup U_t$ and U_{t+1}. By induction the group for the first space is free on t generators while the group for the second is infinite cyclic, and this completes the proof of the inductive step.■

The preceding results yield a few simple criteria for recognizing when a connected graph is a tree.

THEOREM 3. *If X is a connected graph, then the following are equivalent:*

(*i*) *X is a tree.*

(*ii*) *X is contractible.*

(*iii*) *X is simply connected.*

Proof. We already know that the first condition implies the second and the second implies the third, so it is only necessary to prove that (*iii*) implies (*i*). However, if T is a maximal tree in X and $T \neq X$, then we know that the fundamental group of X is a free group on k generators, where $k > 0$ is the number of edges which are in X but not in T. Therefore if X is simply connected we must have $T = X$.■

The Euler characteristic of a graph

If $(X, \mathcal{E})$ is a connected graph, then the preceding discussion shows that the fundamental group of X is a free group on a finite set of free generators. We would like to have a formula for the number of generators which can be read off immediately from the graph structure and does not require us to find an explicit maximal tree inside the graph.

Definition. The *Euler characteristic* of $(X, \mathcal{E})$ is the integer $\chi(X, \mathcal{E}) = v - e$, where e is the number of edges in the graph and v is the number of vertices.

If there is exactly one edge, then clearly $v = 2$, $e = 1$, and the Euler characteristic is equal to $1 = 2 - 1$. The first indication of the Euler characteristic's potential usefulness is an extension of this to arbitrary trees.

PROPOSITION 4. *If $(T, \mathcal{E})$ is a tree, then $\chi(T, \mathcal{E}) = 1$.*

Proof. Not surprisingly, this goes by induction on the number of edges. We already know the formula if there is one edge. As before, if we know the result for trees with n edges and T has $n+1$ edges we may write $T = T_0 \cup A$, where T_0 is a tree, A is a vertex, and their intersection is a single point. For each subgraph Y let $e(Y)$ and $v(Y)$ denote the numbers of edges and vertices in Y. Then we have $e(T) = e(T_0) + 1$, $v(T) = v(T_0) + 1$, and hence we also have

$$\chi(T) \quad = \quad v(T) - e(T) \quad = \quad [v(T_0) + 1] - [e(T_0) + 1] \quad = \quad v(T_0) - e(T_0) \quad = \quad 1$$

which is the formula we wanted to verify.■

THEOREM 5. *If $(X, \mathcal{E})$ is a connected graph, then the fundamental group of X is a free group on $1 - \chi(X, \mathcal{E})$ generators.*

Note that if G is a finite group on n generators, then there are exactly 2^n homomorphisms from G to $\mathbb{Z}_2$ (there are that many ways to define a function from the set of generators to $\mathbb{Z}_2$, and each such function extends uniquely to a group homomorphism). Therefore the number of free generators does not depend upon the choice of a generating set; more precisely, if m and n are positive integers such that G is free on sets of m and n generators, then $m = n$. — Similarly, a group G cannot be simultaneously free on both a finite and an infinite set of generators, for the number of homomorphisms into $\mathbb{Z}_2$ is finite if and only if the generating set is finite; finally, if α is a transfinite cardinal number, then a free group on a set of generators with cardinality α also has cardinality α (verify this), and therefore in all cases the cardinality of a set of free generators for a free group is independent of the choice of generators.

Proof of Theorem 5. We adopt the notational conventions in the preceding argument. Let T be a maximal tree in X, and suppose that there are k edges in X which are not in T, so that the fundamental group is free on k generators. By construction we know that $v(T) = v(X)$ and $e(X) = e(T) + k$, and by the preceding result we know that the Euler characteristic of T is 1. Therefore we have

$$\chi(X, \mathcal{E}) \quad = \quad v(X) - e(X) \quad = \quad v(T) - e(T) - k \quad = \quad 1 - k$$

so that $k = 1 - \chi(X, \mathcal{E})$ as required.∎

In the exercises we note that the theorem is also valid for the edge-path graphs as defined in the files for this course.

COROLLARY 6. *If two connected graphs X and X' are base point preservingly homotopy equivalent as topological spaces, then they have the same Euler characteristics.*

In particular, the corollary applies if X and X' are homeomorphic. For this reason we often suppress the edge decomposition and simply use $\chi(X)$ when writing the Euler characteristic.

Proof. If X and X' are homotopy equivalent, then their fundamental groups are isomorphic, and hence they are both free groups with the same numbers of generators. Since the Euler characteristics can be expressed as functions of these numbers of generators, it follows that the Euler characteristics of X and X' must be equal.∎

COROLLARY 7. *A connected graph X is a tree if and only if $\chi(X) = 1$.*

Proof. We know that $\chi(X) = 1$ if and only if X is simply connected.∎

REMARK. More generally, one has the following criteria for recognizing whether two connected graphs X and Y are homotopy equivalent:

(1) *The connected graphs X and Y are homotopy equivalent if and only if their fundamental groups are isomorphic.*∎

(2) *The connected graphs X and Y are homotopy equivalent if and only if their Euler characteristics are equal.*∎

The results of this course show that the fundamental groups are isomorphic if and only if the Euler characteristics are equal, so (2) will follow from (1). Proving the latter is not all that difficult, but we shall not give the details here.

III.4 : Finite coverings of graphs

(**M**, §§83, 85; **H**, §1.A)

As indicated at the beginning of this unit, we shall conclude our discussion of graphs with an application to some mildly counter-intuitive results on finite index subgroups of finitely generated free groups. Since each such group is the fundamental group of some graph X and finite index subgroups of $\pi_1(X)$ should correspond to finite coverings of X, we begin with an observation about finite covering spaces of graphs.

PROPOSITION 1. *Let $(X, \mathcal{E})$ be a connected finite graph, and suppose that $p : W \to X$ be a connected finite covering. Then there is a finite graph complex structure $\mathcal{E}'$ on W such that p maps each edge of $\mathcal{E}'$ homeomorphically to an edge of $\mathcal{E}$.*

Proof. Let k be the number of sheets in the covering space projection, so that each point in X has exactly k preimages in W. If E_α is an edge in $\mathcal{E}$ and $F_\alpha = p^{-1}[E_\alpha]$, then the restriction p_α of p to F_α defines an k-sheeted covering space projection over E_α; the space F_α is not necessarily connected, and in fact F_α is a finite disjoint union of connected covering spaces over X, with each such space corresponding to a component of F_α. Since E_α is simply connected, p_α must be a homeomorphism on each component of F_α, and it follows that F_α must be a union of pairwise disjoint compact subspaces $F_{\alpha,j}$ where $1 \leq j \leq k$ such that p_α maps each subspace homeomorphically onto E_α.

To show that the subsets $F_{\alpha,j}$ determine a graph structure on W, we need to look at the intersections

$$F_{\alpha,j} \ \cap \ F_{\beta,i}$$

where $(\alpha, j) \neq (\beta, i)$. If $\alpha = \beta$ then these intersections are empty by the reasoning in the preceding paragraph. If $\alpha \neq \beta$, then the relations

$$p\,[F_{\alpha,j} \ \cap \ F_{\beta,i}] \quad \subset \quad p\,[F_{\alpha,j}] \ \cap \ p\,[F_{\beta,i}] \quad = \quad E_\alpha \ \cap \ E_\beta$$

imply that $F_{\alpha,j} \ \cap \ F_{\beta,i}$ is empty if $E_\alpha \cap E_\beta$ is empty, while if the latter is not empty then the intersection is contained in the inverse image of the vertex common to $E_\alpha \cap E_\beta$. Since the restriction of p to each component of F_α is 1–1, it follows that $F_{\alpha,j} \cap F_{\beta,i}$ contains at most one point if $\alpha = \beta$, and if this happens then this point is a vertex of both $F_{\alpha,j}$ and $F_{\beta,i}$, proving that we have a graph structure on X.■.

We can use this result to prove the following purely algebraic result on subgroups of finite index:

PROPOSITION 2. *Let F be a free group on k generators, for some positive integer k, and let H be a subgroup of index n. Then H is a free group on $nk - n + 1$ generators.*

A standard result in algebra states that if M is a finitely generated free module on m generators over a principal ideal domain $\mathbb{D}$ and $N \subset M$ is a $\mathbb{D}$-submodule, then N is free on n generators for some $n \leq m$. In contrast, the result above says that a free subgroup of a free group may have more generators than the group containing it. After proving this result, we shall also describe an example to show that a finitely generated free group also contains a non-finitely generated free subgroup (which is not of finite index).

Proof. Let $(X, \mathcal{E})$ be a connected graph whose fundamental group is free on k generators; one method of constructing such a graph is to take edges A_i, B_i and C_i for $1 \leq i \leq k$, where the edges of A_i are x, p_i, and q_i, the edges of B_i are x, r_i, and s_i, and the edges of B_i are x, u_i, and v_i (topologically, X is a union of k circles such that each pair intersect at x and nowhere else). By the formula relating the number of generators for F and the Euler characteristic, we know that $k = 1 - \chi(X)$, or equivalently $\chi(X) = 1 - k$. Let Y be the connected covering space of X corresponding to the subgroup H. Then Y is a graph, and the fundamental group of Y is H, so that H is a free group.

We know that the number of free generators for H is given by $1 - \chi(Y)$, so it is only necessary to compute this Euler characteristic. Let e and v be the number of edges and vertices for $(X, \mathcal{E})$, so that $n = 1 - \chi(X)$, where $\chi(X) = v - e$. Since Y is an n-sheeted covering of X, if we take the associated edge decomposition of Y (such that each edge of Y maps homeomorphically to an edge of X) we see that the numbers of vertices and edges for Y are nv and ne respectively, so that

$$\chi(Y) \quad = \quad n \cdot \chi(X) \ .$$

Therefore the number of generators for the fundamental group of Y is given by

$$1 - \chi(Y) \quad = \quad 1 - n \cdot \chi(X) \quad = \quad 1 - n(1-k) \quad = \quad nk - n + 1$$

which is what we wanted to prove.■

COROLLARY 3. *If F is a free group on k generators for some $k \geq 2$, then F contains free subgroups on m generators for all $m \geq k$.*

QUESTION. Does this result extend to the case $k = 1$? Prove this or explain why it cannot be true.

Proof. It suffices to prove this result when $k = 2$ since F automatically contains a free subgroup with 2 generators (take a subset of some generating set for F).

Let X be a graph whose fundamental group is free on 2 generators u and v, and let Y_n be the n-sheeted covering space whose fundamental group is the (free) subgroup generated by by u and v^n for some $n \geq 2$. Then the fundamental group of Y_n is isomorphic to a free group on $n+1$ generators. It follows that for every positive integer $m \geq 3$ there is some n such that $\pi_1(Y_n)$ contains a free group on m generators.■

Example. The free group on two generators also contains a free subgroup with a countably infinite set of generators (hence the same is also true for every free group on more than two generators). Here is a sketch of the argument. Filling in the details is left to the reader as an exercise:

Let $X = S^1 \vee S^1$ with base point given by the common point of the two circles, and let u and v be free generators of $\pi_1(X)$ which are represented by the two circles. Let K denote the kernel of the homomorphism from $\pi_1(X)$ to $\mathbb{Z}$ which sends u to zero and v to a generator.

Let Y be the covering space of X whose fundamental group is isomorphic to K. It follows that Y is homeomorphic to a copy of the real line with a circle attached at each point $2q\pi$ where q runs through all integers (verify this!). An explicit model for Y is the set of all points $(x, y, z) \in \mathbb{R}^3$ such that either $(x, y) = (1, 0)$ (in other words, the line with parametric equations $(1, 0, t)$ for $t \in \mathbb{R}$) or $x^2 + y^2 = 1$ and $z = 2q\pi$ for some integer q. If we view X as the subset of $\mathbb{R}^2$ given by

$$\{ \ x^2 + y^2 \quad = \quad 1 \} \ \cup \ \{ \ (x-2)^2 + y^2 = 1 \ \}$$

then the covering space projection corresponds to the map sending (x, y, z) to (x, y) on the first piece and sending $(1, 0, t)$ to

$$\left(2 - \cos t, -\sin t\right)$$

on the second.

Let $A_m \subset Y$ be the set of all points such that $|z| \leq m$. Then A_m consists of a closed line segment with $2m + 1$ circles attached symmetrically with respect to the end points. It follows that $\pi_1(A_m)$ is a free group on $2m + 1$ generators, and the inclusion of $\pi_1(A_m)$ in $\pi_1(A_{m+1})$ is a 1–1 map sending the free generators of the first group to a subset of a set of free generators for the second.

Since every compact subset of Y is contained in some A_m, it follows that the fundamental group of Y is an increasing limit of the fundamental groups of the subspaces A_m. Since this limit is a free group on a countably infinite set of generators, it follows that $\pi_1(Y)$ must have the same property.■

III.∞ : Infinite graphs

(**M**, §§83–85; **H**, §1.B)

If we are given a connected graph $(X, \mathcal{E})$ with a nontrivial (hence infinite) fundamental group, then by Proposition 1.7 and Theorem I.3.1 we know that X has a simply connected covering space, and the methods of Corollary I.1.2 show that the inverse image of an edge is a disjoint union of edges and hence the universal simply connected covering space $\widetilde{X}$ has a decomposition into subsets homeomorphic to intervals. It is natural to think of this as an infinite graph complex, and results from Munkres and Hatcher provide a mathematically precise setting for doing so. One particularly significant application of infinite graphs is the following result (which was originally proved by algebraic methods):

Munkres, Theorem 85.1, p. 514. *If G is a free group and $H \subset G$ is a subgroup, then H is also free.*

As noted at the end of the preceding section, if G has a finite set of free generators and is not cyclic, then the number of generators for H can be any number between 1 and $\aleph_0$ (the cardinal number of the integers).

Other uses of infinite graphs are discussed in Section 1.B of Hatcher, particularly in the section on graphs of groups. Such constructions are often extremely valuable sources of geometric insights into a group's algebraic structure. The Wikipedia articles

`http://en.wikipedia.org/wiki/Geometric_group_theory`

`http://en.wikipedia.org/wiki/Group_cohomology`

contain more detailed additional information (and further links) concerning the ways in which topology and group theory — particularly infinite group theory — have interacted with each other in mathematical research during the past century.