Topology of Finite Graphs

John R. Stallings

University of California, Berkeley, Department of Mathematics, Berkeley, CA 94720/USA

Introduction

This paper derives from a course in group theory which I gave at Berkeley in 1982. I wanted to prove the standard theorems on free groups, and discovered that, after a few preliminaries, the notion of "locally injective" map (or "immersion") of graphs was very useful. This enables one to see, in an effective, easy, algorithmic way just what happens with finitely generated free groups. One can understand in this way (1) Howson's theorem that if A and B are finitely generated subgroups of a free group, then $A \cap B$ is finitely generated, and (2) M. Hall's theorem that finitely generated subgroups of free groups are closed in the profinite topology.

During this course, S.M. Gersten came up with a simple proof of H. Neumann's inequality on the ranks in Howson's theorem. One of the ideas in Gersten's proof was to use core-graphs (graphs with no trees hanging on). Subsequently, I found that some consequences of a paper of Greenberg's could be proved using core-graphs and "covering translations" of immersions; the most striking such result is that if A and B are finitely generated subgroups of a free group and if $A \cap B$ is of finite index in both A and B, then $A \cap B$ is of finite index in $A \vee B$, the subgroup generated by $A \cup B$.

1. The Category of Graphs

1.1. By graph and map of graphs, I mean something purely combinatorial or algebraic. Pictures can be drawn, but one has to understand that maps are rigid and not just continuous, maps do not collapse edges or wrap edges around several edges. The formulation below is due to Serre [9].

Specifically, a *graph* Γ consists of two sets E and V, and two functions $E \to E$ and $E \to V$: For each $e \in E$, there is an element $\bar{e} \in E$, and an element $\iota(e) \in V$.

The rules to be satisfied are just these: $\bar{\bar{e}} = e$ and $\bar{e} \neq e$. An $e \in E$ is a *directed edge* of Γ; $\bar{e}$ is the *reverse* of e. The elements of V are called *vertices* of Γ; $\iota(e)$ is the *initial vertex* of the edge e. We define the *terminal vertex* of e to be $\tau(e) = \iota(\bar{e})$.

An *orientation* of Γ consists of a choice of exactly one edge in each pair $\{e, \bar{e}\}$. Another way to say this: the group Z_2 acts freely on the edges of Γ, and an orientation is a choice of a representative in each orbit.

A *map of graphs* $f: \Gamma \to \Delta$ consists of a pair of functions, edges to edges, vertices to vertices, preserving the structure. Given f, an orientation of the target Δ determines a unique orientation of Γ which is preserved by f, called the orientation of Γ induced by f.

Thus, graphs and their maps form a category in which various categorial concepts can be discussed. There are two functors, "edges" and "vertices," from graphs to the category of sets; I call a categorical construction *obvious* if it is preserved by these two functors. Thus, monomorphism, epimorphism, direct limit, product, coproduct, pullback, pushout are all obvious. However, some of these deserve more than a word.

1.2. Let us consider *pushouts*. In the category of sets, pushouts have a mysterious quality, being defined in terms of equivalence relations *generated by* certain binary relations. For example, if G is a group containing subgroups A and B, and if $C = A \vee B$ is the subgroup generated by $A \cup B$, and if G/X denotes the set of left cosets, then

$$
\begin{array}{ccc}
G & \longrightarrow & G/A \\
\downarrow & & \downarrow \\
G/B & \longrightarrow & G/C
\end{array}
$$

is a pushout diagram.

In graphs, pushouts do not always exist. The reason is that a pushout of free Z_2-sets may not be *free*. The necessary and sufficient condition for the pushout of a pair of maps of graphs, $\alpha_1: \Gamma \to \Delta_1$ and $\alpha_2: \Gamma \to \Delta_2$, to exist is that there exist orientations of $\Gamma, \Delta_1, \Delta_2$ which are preserved by α_1 and α_2. When this condition is satisfied, the pushout is "obvious."

1.3. *Pullbacks* in the category of graphs always exist and are "obvious" and are easy to construct:

Let $\beta_1: \Gamma_1 \to \Delta$, $\beta_2: \Gamma_2 \to \Delta$ be maps of graphs. Define Γ_3 to have
vertex-set $= \{(u, v) | u$ a vertex of Γ_1, v a vertex of Γ_2, $\beta_1(u) = \beta_2(v)\}$
edge-set similarly

and so on. There will be maps of Γ_3 to Γ_1 and Γ_2 which form, with β_1 and β_2, a pullback diagram.

This pullback is a sub-graph of the *product* $\Gamma_1 \times \Gamma_2$, which is an "obvious" construction. A curious fact is, however, that products seem to be quite useless from the point of view of group theory, whereas pullbacks are important.

2. Paths

2.1. A *path* p in Γ, of *length* $n = |p|$, with *initial vertex* u and *terminal vertex* v, is an n-tuple of edges of Γ, $p = e_1 e_2 \ldots e_n$, such that for $i = 1, \ldots, n-1$, we have $\tau(e_i) = \iota(e_{i+1})$, and such that $u = \iota(e_1)$ and $v = \tau(e_n)$. For $n = 0$, given any vertex v, there is a unique path Λ_v of length 0 whose initial and terminal vertices coincide and are equal to v. Another way to say this: The standard arc of length n, A_n, can be described as the interval $[0, n]$ subdivided at the integer points; then our path p is a map of graphs $p: A_n \to \Gamma$ such that $p(0) = u$, $p(n) = v$.

A path p is called a *circuit* if its initial and terminal vertices coincide.

If p and q are paths in Γ and the terminal vertex of p equals the initial vertex of q, they may be *concatenated* to form a path pq with $|pq| = |p| + |q|$, whose initial vertex is that of p and whose terminal vertex is that of q.

The set of all paths in Γ with the operation of concatenation is a small category, denoted $P(\Gamma)$. It can be thought of as the category generated by Γ. A map of graphs $f: \Gamma \to \Delta$ induces a length-preserving homomorphism (functor) denoted by the same symbol $f: P(\Gamma) \to P(\Delta)$.

2.2. A *round-trip* is a path of the form $e\bar{e}$. If a path p contains two adjacent edges forming a round-trip, then by deleting that round-trip we get a path p' with the same initial and terminal vertices as p, and with $|p'| = |p| - 2$. We write $p \searrow p'$, and think of p' as obtained from p by an *elementary reduction*.

The equivalence relation on $P(\Gamma)$ generated by $\searrow$ is called *homotopy* and is denoted by $\sim$. Concatenation of paths is compatible with homotopy, and thus the set of $\sim$-classes of $P(\Gamma)$ forms a small category denoted by $\pi(\Gamma)$. Each element of $\pi(\Gamma)$ has an inverse: Let $[p]$ be the homotopy class of p; if Λ_v is the path of length 0 at v, let $\bar{\Lambda}_v = \Lambda_v$; if $p = qe$, where e is an edge, define recursively $\bar{p} = \bar{e}\bar{q}$; then $[p]^{-1} = [\bar{p}]$. This justifies calling $\pi(\Gamma)$ the *fundamental groupoid* of Γ. The set of elements of $\pi(\Gamma)$ starting and ending at a fixed vertex v forms a group, $\pi_1(\Gamma, v)$, the *fundamental group* of Γ based at v.

Given a map $f: \Gamma \to \Delta$ there are defined homomorphisms, denoted by the same symbol,

$$f: \pi(\Gamma) \to \pi(\Delta),$$
$$f: \pi_1(\Gamma, v) \to \pi_1(\Delta, f(v)).$$

2.3. The case of a graph Γ with just one vertex is classical: $P(\Gamma)$ is just the free monoid on E; if $\mathcal{O}$ is an orientation of Γ, then $\pi(\Gamma) = \pi_1(\Gamma, v)$ is the free group on $\mathcal{O}$.

We can define, in general Δ, a *reduced path* to be a path in Δ containing no round-trip. Every path is clearly homotopic to some reduced path.

In the case of a one-vertex graph Γ, the classical theory shows that each homotopy class of paths contains a unique reduced path. This can be extended easily (see **5.2**) to general graphs.

2.4. A graph is *connected* if any pair of vertices is joined by some path. In general, a graph is the disjoint union (coproduct) of its connected components.

A graph is a *forest* if the only reduced circuits have length 0. A *tree* is a connected forest.

If u and v are vertices in a tree T, then there is a unique reduced path in T starting at u and ending at v. This path is denoted $[u, v]_T$.

We now list a few classical exercises, some more difficult than others:

(a) *Every graph Γ contains a maximal forest.*

(b) *Every maximal forest in Γ contains all the vertices of Γ.*

(c) *If Γ is connected, then every maximal forest in Γ is a tree.*

(d) *Let v be a vertex of a connected graph Γ, containing a maximal tree T. Let $\mathcal{O}$ be an orientation of Γ; for each $e \in \mathcal{O} - T$, let*

$$p_e = [v, \iota(e)]_T \, e \, [\tau(e), v]_T.$$

Then $\pi_1(\Gamma, v)$ is free on $\mathcal{O} - T$, the element of $\pi_1(\Gamma, v)$ corresponding to e being the homotopy class of p_e.

3. Stars

3.1. If v is a vertex of the graph Γ, the *star* of v in Γ is the set of edges of Γ:

$$St(v, \Gamma) = \{e \in E \mid \iota(e) = v\}.$$

The cardinality of $St(v, \Gamma)$ is called the *valence* of v in Γ.

A map of graphs $f: \Gamma \to \Delta$, yields, for each vertex v of Γ, a function

$$f_v: St(v, \Gamma) \to St(f(v), \Delta).$$

If, for each vertex v of Γ, f_v is injective, we call f an *immersion*. If each f_v is bijective, we call f a *covering*. If each f_v is surjective, we say f is *locally surjective*.

For example, a reduced path of length n in Γ is exactly the same as an immersion from the standard arc of length n to Γ.

3.2. A pair of edges (e_1, e_2) of Γ is said to be *admissible* if $\iota(e_1) = \iota(e_2)$ and $e_1 \neq \bar{e}_2$. In this case, we can identify $\tau(e_1)$ to $\tau(e_2)$, e_1 to e_2, $\bar{e}_1$ to $\bar{e}_2$, to obtain a graph denoted by $\Gamma/[e_1 = e_2]$, which we call the result of *folding* (e_1, e_2) in Γ. This is a particularly simple instance of the pushout construction:

The condition $e_1 \neq \bar{e}_2$ is the orientation condition discussed in **1.2.**

3.3. If $f: \Gamma \to \Delta$ is a map of graphs, and (e_1, e_2) is a pair of edges of Γ such that $\iota(e_1) = \iota(e_2)$ and $f(e_1) = f(e_2)$, then (e_1, e_2) is an admissible pair of Γ-edges, and we say that f *folds* that admissible pair. In this case, f factors through $\Gamma/[e_1 = e_2]$.

Any map which is not an immersion folds some admissible pair (e_1, e_2) nontrivially – that is, $e_1 \neq e_2$.

Thus if Γ is a finite graph and $f: \Gamma \to \Delta$ is a map of graphs, we can find a finite sequence of foldings: $\Gamma = \Gamma_0 \to \Gamma_1 \to \Gamma_2 \ldots \to \Gamma_n$ and an immersion $\Gamma_n \to \Delta$, so that the composition of the immersion and the sequence of foldings is equal to f. [The sequence of foldings is not unique, but the final immersion *is* unique.]

4. Coverings

4.1. The theory of coverings of graphs is almost completely analogous to the topological theory of covering spaces. The bijectivity of star-maps easily implies:

(a) (Path-lifting): *If* $f: \Gamma \to \Delta$ *is a covering,* v *a vertex of* Γ, p *a path in* Δ *with initial vertex* $f(v)$*, then there exists a unique path* $\tilde{p}$ *in* Γ *with initial vertex* v *such that* $f\tilde{p} = p$.

(b) (Homotopy-lifting): *In* (a)*, if* p *is a round-trip, then* $\tilde{p}$ *is a round-trip. Hence if* $p \sim q$*, then* $\tilde{p} \sim \tilde{q}$.

These are the major lemmas from which we can prove the standard facts, namely:

(c) (General lifting):

If $f: \Gamma \to \Delta$ *is a covering,* $g: \Theta \to \Delta$ *a map of graphs, and if* Θ *is connected, and if* u, v *are vertices of* Γ, Θ *such that* $f(u) = g(v)$*, then: there exists* $\tilde{g}: \Theta \to \Gamma$ *such that* $\tilde{g}(v) = u$ *and* $f\tilde{g} = g$*, if and only if* $g\pi_1(\Theta, v) \subset f\pi_1(\Gamma, u)$*; and if* $\tilde{g}$ *exists, it is unique.*

(d) (Existence of coverings): *If* $f: \Gamma \to \Delta$ *is a covering and* u *a vertex of* Γ, *then*

$$f: \pi_1(\Gamma, u) \to \pi_1(\Delta, f(u))$$

is injective.

If Δ *is connected,* v *a vertex,* $S \subset \pi_1(\Delta, v)$ *a subgroup, then there exists a covering* $f: \Gamma \to \Delta$ *where* Γ *is connected, with vertex* u*, such that* $f(u) = v$ *and* $f\pi_1(\Gamma, u) = S$*. Any two such coverings are isomorphic (in the standard sense). The index of* S *in* $\pi_1(\Delta, v)$ *is the cardinality of* $f^{-1}(v)$.

4.2. If G is a group, a *G-graph* Γ is a graph together with an action of G on the left on Γ by maps of graphs, such that, for all $g \in G$ and every edge e, $ge \neq \bar{e}$. In this case, we can define a quotient graph Γ/G and a quotient map of graphs $\Gamma \to \Gamma/G$. It is easy to show, in general, that $\Gamma \to \Gamma/G$ is locally surjective.

We say that G *acts freely* on Γ, when, whenever v is a vertex of Γ, $g \in G$, and $gv = v$, then $g = 1$, the identity element of G. In this case $\Gamma \to \Gamma/G$ is an immersion, and hence is a covering.

The *universal cover* $f: \tilde{\Gamma} \to \Gamma$, of a connected graph Γ, is a covering with $\tilde{\Gamma}$ connected and $\pi_1(\tilde{\Gamma})$ trivial. For v a Γ-vertex, $G = \pi_1(\Gamma, v)$ acts freely, by "covering translations", on $\tilde{\Gamma}$, and f is isomorphic to the quotient map $\tilde{\Gamma} \to \tilde{\Gamma}/G$.

4.3. If A, B are subgroups of a group G, their *join* $A \vee B$ is the subgroup generated by $A \cup B$.

Theorem (Pushout represents join). *Suppose that*

$$
\begin{array}{ccc}
\Gamma & \xrightarrow{\ \alpha_1\ } & \varDelta_1 \\
{\scriptstyle \alpha_2}\Big\downarrow & & \Big\downarrow{\scriptstyle \beta_1} \\
\varDelta_2 & \xrightarrow[\ \beta_2\ ]{} & \Theta
\end{array}
$$

is a pushout diagram, and that $\Gamma, \varDelta_1, \varDelta_2$ *are connected. Let* u *be a vertex of* Γ*; call the images of* u *in* $\varDelta_1, \varDelta_2, \Theta$, *respectively,* v_1, v_2, w. *Then*

$$(*) \qquad\qquad \pi_1(\Theta, w) = \beta_1 \pi_1(\varDelta_1, v_1) \vee \beta_2 \pi_1(\varDelta_2, v_2).$$

Proof. Let $f: \tilde{\Theta} \to \Theta$ be the covering of Θ corresponding to the subgroup on the right side of $(*)$, which covering exists by **4.1(d)**. By **4.1(c)**, there are liftings $\tilde{\beta}_1: \varDelta_1 \to \tilde{\Theta}, \tilde{\beta}_2: \varDelta_2 \to \tilde{\Theta}$, preserving base-point. Then $\tilde{\beta}_1 \alpha_1$ and $\tilde{\beta}_2 \alpha_2 : \Gamma \to \tilde{\Theta}$ are liftings of $\beta_1 \alpha_1 = \beta_2 \alpha_2$, and so, by the uniqueness part of **4.1(c)**, $\tilde{\beta}_1 \alpha_1 = \tilde{\beta}_2 \alpha_2$. Using the fact that our diagram is a pushout diagram, there is then a map $\Theta \to \tilde{\Theta}$ which is a cross-section of $f: \tilde{\Theta} \to \Theta$. This implies that f is an isomorphism of graphs, and that says that $(*)$ is true.

4.4. Corollary. *If* (e_1, e_2) *is an admissible pair of edges in a connected graph* Γ, *then the folding map* $\Gamma \to \Gamma/[e_1 = e_2]$ *is surjective on fundamental groups.*

Proof. The pushout diagram for a folding (see **3.2**) has the upper left corner connected and the lower left corner has trivial fundamental group. Then an application of **4.3** proves this result.

Of course, a more explicit result can be proved. If $\tau(e_1) \neq \tau(e_2)$, the folding map is a π_1-isomorphism. If $e_1 \neq e_2$ and $\tau(e_1) = \tau(e_2)$, the folding map exactly kills one element of a particular free basis of the fundamental group of Γ.

5. Immersions

5.1. Immersions have some of the properties of coverings. Generally, liftings may not exist, but if they do, they are unique. Immersions represent subgroups more efficiently than do coverings. We now present a series of exercises on immersions, pointing out any tricky points.

5.1. (a) (Preservation of reduced paths). *If* $f: \Gamma \to \varDelta$ *is an immersion of graphs, and* p *is a reduced path in* Γ, *then* fp *is a reduced path in* $\varDelta$.

This is, in fact, a special case of the fact that the composition of two immersions is an immersion.

(b) (Uniqueness of path-lifting). *If* $f: \Gamma \to \varDelta$ *is an immersion, and* p, q *are paths in* Γ *having the same initial vertex, and if* $fp = fq$, *then* $p = q$. [Induction on $|p|$.]

(c) (Uniqueness of lifting). *If $f: \Gamma \to \Delta$ is an immersion, Θ a connected graph, and $g_1, g_2: \Theta \to \Gamma$ are maps of graphs such that $fg_1 = fg_2$, and if there is a vertex v of Θ such that $g_1(v) = g_2(v)$, then $g_1 = g_2$.*

This follows from **(b)**.

5.2. Proposition. (Uniqueness of reduced paths). *If Γ is any graph, and if p and q are reduced, homotopic paths in Γ, then $p = q$.*

Proof. Let Δ be the result of identifying all vertices of Γ to one vertex. The identification map $f: \Gamma \to \Delta$ is bijective on edges, and therefore is an immersion. Then since $p \sim q$, we have $fp \sim fq$. Since p and q are reduced, by **5.1(a)** both fp and fq are reduced. By the theory of free groups, **2.3**, it follows that $fp = fq$. By **5.1(b)** and the fact, implied by $p \sim q$, that p and q have the same initial vertex, we see that $p = q$.

(Note that this proof appeals to the exterior mathematical world, the word-problem in free groups. A proof within graph-theory is possible, using the universal cover $\tilde{\Gamma}$ and facts about trees, such as the uniqueness of reduced paths in a tree; however, this seems excessively complex.)

5.3. Proposition (Injectivity of π_1). *If $f: \Gamma \to \Delta$ is an immersion, and v is a vertex of Γ, then*
$$f: \pi_1(\Gamma, v) \to \pi_1(\Delta, f(v))$$
is injective.

Proof. Let $\alpha \in \pi_1(\Gamma, v)$, $\alpha \neq 1$. Then α is represented by a circuit p based at v, with p reduced and $|p| \geq 1$. By **5.1(a)**, fp is reduced. Since $|fp| = |p| \geq 1$, by **5.2** fp is not homotopic to a path of length 0, and so $f\alpha \neq 1$.

5.4. Algorithm (Finitely generated subgroups). *Given a finite set of elements $\{\alpha_1, \ldots, \alpha_n\} \subset \pi_1(\Delta, u)$, there is an algorithm that represents the subgroup S of $\pi_1(\Delta, u)$ generated by $\{\alpha_1, \ldots, \alpha_n\}$, by an immersion $f: \Gamma \to \Delta$, as follows:*

Represent α_i by a circuit p_i based at u. Let Γ_1 be the disjoint union of n standard arcs $B_1, \ldots, B_n$, the length of B_i being $|p_i|$. Map Γ_1 into Δ by $p_1 \cup \ldots \cup p_n$. Identify all the initial and terminal vertices of all B_i to a single vertex v, obtaining Γ_2 and a map $f_2: \Gamma_2 \to \Delta$. In other words, Γ_2 is a wedge of n subdivided circles, mapping the i^{th} circle by p_i.

Then $f_2 \pi_1(\Gamma_2, v) = S$.

By **3.3**, f_2 can be factored through a series of folds and an immersion:

$$\underbrace{\Gamma_2 \to \Gamma_3 \to \ldots \to \Gamma_k}_{\text{folds}} \underbrace{\xrightarrow{\ f_k\ }}_{\text{immersion}} \Delta$$

By **4.4**, each fold is surjective on π_1, and so, letting w be the image of v in Γ_k,
$$f_k \pi_1(\Gamma_k, w) = S.$$

Thus f_k is the desired immersion.

Now, by choosing a maximal tree in Γ_k, we can find a free basis of $\pi_1(\Gamma_k, w)$, which, by **5.3**, yields a free basis of S.

We can see easily whether the folds in this procedure are isomorphisms on π_1 or not (remark after the proof of **4.4**); this will decide if $\{\alpha_1, \ldots, \alpha_n\}$ is a free basis of S. As we shall see later in **7.6**, it is easy to see whether or not S has finite index in $\pi_1(\varDelta, u)$ in case $\varDelta$ has only one vertex, and it is easy to determine what that finite index is.

5.5. Theorem (Pullback of immersions represents intersection). *Let*

$$
\begin{array}{ccc}
\Gamma_3 & \xrightarrow{\ g_1\ } & \Gamma_1 \\
{\scriptstyle g_2}\big\downarrow & & \big\downarrow{\scriptstyle f_1} \\
\Gamma_2 & \xrightarrow{\ f_2\ } & \varDelta
\end{array}
$$

be a pullback diagram of graphs. Suppose that f_1 and f_2 are immersions. Let v_1, v_2 be vertices in Γ_1, Γ_2 such that $f_1(v_1) = f_2(v_2) = w$; let v_3 be the corresponding vertex in Γ_3. Define $f_3 = f_1 g_1 = f_2 g_2 : \Gamma_3 \to \varDelta$, and define

$$S_i = f_i \pi_1(\Gamma_i, v_i) \quad \text{for } i = 1, 2, 3.$$

(These are subgroups of $\pi_1(\varDelta, w)$.)
    Then
$$S_3 = S_1 \cap S_2.$$

Proof. It is clear from $f_1 g_1 = f_2 g_2$ that S_3 is contained in $S_1 \cap S_2$. To show the reverse inclusion, let $\alpha \in S_1 \cap S_2$. Then there are reduced circuits p_1, p_2 in Γ_1, Γ_2 based at v_1, v_2 such that $f_1 p_1$ and $f_2 p_2$ belong to the homotopy class α. By **5.1(a)**, both $f_1 p_1$ and $f_2 p_2$ are reduced paths, and so by **5.2**, $f_1 p_1 = f_2 p_2$. By the pullback property, there exists a path p_3 in Γ_3 such that $p_1 = g_1 p_3$ and $p_2 = g_2 p_3$, and p_3 is a circuit based at v_3. Thus $f_3 p_3$ represents an element of S_3, and represents α.

5.6. Corollary (Howson's theorem [5]). *If S_1 and S_2 are finitely generated subgroups of a free group F, then $S_1 \cap S_2$ is finitely generated (and a free basis of $S_1 \cap S_2$ can be determined by an easy algorithm).*

Proof. Represent F as $\pi_1(\varDelta)$, where $\varDelta$ is a graph with one vertex. Using **5.4**, represent S_1 and S_2 by immersions $f_1 : \Gamma_1 \to \varDelta$, $f_2 : \Gamma_2 \to \varDelta$, where Γ_1 and Γ_2 are connected finite graphs. Construct the pullback Γ_3 by **1.3**; clearly, Γ_3 is a finite graph, and by **5.5**, a component of Γ_3 (containing v_3) represents $S_1 \cap S_2$. It is easy to check that $f_3 : \Gamma_3 \to \varDelta$ is an immersion, and therefore a finite free basis of S_3 can be determined by choosing a maximal tree in that component of Γ_3.

5.7. (a) As Gersten has shown, [2], a careful check of Euler characteristics proves H. Neumann's inequality [8] on the rank of $S_1 \cap S_2$: If $S_1 \cap S_2$ is non-trivial, then
$$r(S_1 \cap S_2) - 1 \leq 2 \cdot (r(S_1) - 1) \cdot (r(S_2) - 1).$$

H. Neumann's conjecture that "2" can be replaced by "1" seems to be a much more difficult combinatorial problem. Perhaps these techniques may clarify this question.

(b) The fact that the pullback Γ_3 has only a finite number of components leads to this curious result: If S_1 and S_2 are finitely generated subgroups of a free group F, then, as α varies over F, the subgroups $\{S_1 \cap \alpha S_2 \alpha^{-1}\}$ belong to only a finite number of conjugacy classes of subgroups of F. [To see this, $\alpha S_2 \alpha^{-1}$ is represented by Γ_2 with perhaps an arc attached by one end-vertex, and a change of basepoint v_2. The pullback then is changed by perhaps adding several arcs, each attached by one end-vertex or not at all attached. Any component of the pullback, representing $S_1 \cap \alpha S_2 \alpha^{-1}$, is either an arc, representing $\{1\}$, or a component of the original pullback with some hairs growing on it; in the latter case the represented subgroup is conjugate to the group represented by a component of the original pullback Γ_3.] This theorem is due to Imrich [6].

5.8. For further reference, we discuss translations of immersions.

A *translation* of a map of graphs $f: \Gamma \to \Delta$ is a map $g: \Gamma \to \Gamma$ which is an isomorphism of graphs and for which $fg = f$. The set of all translations of f forms a group $G(f)$ which acts on Γ, and f factors through the quotient $\Gamma \to \Gamma / G(f)$.

Proposition. *If $f: \Gamma \to \Delta$ is an immersion, and Γ is connected, then $G(f)$ acts freely on Γ.*

Proof. This is a consequence of the definition of "acts freely" in **4.2**, and the uniqueness of lifting, **5.1(c)**.

6. Marshall Hall's Theorem

6.1. Theorem. *Let $f: \Gamma \to \Delta$ be an immersion of graphs. Suppose that Δ has only one vertex, and that Γ has only finitely many vertices. Then there exists a graph Γ' containing Γ, such that $\Gamma' - \Gamma$ consists only of edges, and there exists a map $f': \Gamma' \to \Delta$ extending f, such that f' is a covering. (And, if Δ is a finite graph, the proof shows how to effectively construct and enumerate all such extensions Γ', f'.)*

Proof. Let V be the set of vertices of Γ. Choose an orientation $\mathcal{O}$ of Δ. For each $e \in \mathcal{O}$, define

$$R_e = \{(u, v) \in V \times V \mid \text{ there exists an edge } e_1 \text{ of}$$

$$\Gamma \text{ such that } f(e_1) = e, \quad \iota(e_1) = u, \quad \tau(e_1) = v\}.$$

Because f is an immersion, the two coordinate projections $R_e \to V$ are injective. Thus R_e is a bijection of one subset of V onto another; since V is finite, that bijection can be extended to all of V, so that there exists $S_e \subset V \times V$ such that $R_e \subset S_e$ and S_e is a bijection $V \to V$.

Using a choice of S_e for all $e \in \mathcal{O}$, we construct Γ' and f', as follows.

The set of vertices of Γ' is V.

The set of edges of Γ' is E', defined thus:

$$E' = \{(u, v, e) \mid u, v \in V, e \text{ and edge of } \Delta; \quad \text{if } e \in \mathcal{O}, \quad \text{then}$$

$$(u, v) \in S_e; \quad \text{if } \bar{e} \in \mathcal{O}, \quad \text{then } (v, u) \in S_{\bar{e}}\}.$$

Define, for $\varepsilon = (u, v, e) \in E'$, $\bar{\varepsilon} = (v, u, \bar{e})$, and $\iota(\varepsilon) = u$.

Define $f': \Gamma' \to \Delta$ by taking V to the unique vertex of Δ, and $f'(\varepsilon) = e$, for $\varepsilon = (u, v, e) \in E'$.

Embed Γ into Γ' by a map $a: \Gamma \to \Gamma'$ as follows: For a vertex v, $a(v) = v$; for an edge e_1 of Γ, $a(e_1) = (\iota(e_1), \tau(e_1), f(e_1))$.

The proof that this all works is left as an exercise.

6.2. Example: Let F be free on $\{x, y\}$. Then there exist exactly 36 subgroups S of index 5 in F for which the following form a system of representatives of the right cosets of S in F: $\{1, x, xy, xyx^{-1}, xyx^{-1}y^{-1}\}$.

To see this, represent F as the fundamental group of Δ:

Map the standard arc Γ of length 4 into Δ thus:

This is an immersion $f: \Gamma \to \Delta$. For the edge x, the partial bijection R_x says: $0 \to 1$, $3 \to 2$; while R_y says: $1 \to 2$, $4 \to 3$. Each of these can be extended to total bijections S_x and S_y in $3! = 6$ ways, and so there are $6 \cdot 6$ possible constructions of Γ', each yielding (base-point $= 0$) one of the subgroups under consideration.

6.3. Corollary (Hall [4], Burns [1]). *Let* $\alpha_1, \ldots, \alpha_k$, $\beta_1, \ldots, \beta_\ell$ *be elements of a free group F. Let S be the subgroup of F generated by* $\{\alpha_1, \ldots, \alpha_k\}$. *Suppose that* $\beta_i \notin S$ *for* $i = 1, \ldots, \ell$. *Then there exists a subgroup S' of finite index in F, such that* $S \subset S'$, $\beta_i \notin S'$ *for* $i = 1, \ldots, l$, *and there exists a free basis of S' having a subset which is a free basis of S.*

Proof. Represent F as $\pi_1(\Delta)$, where Δ has only one vertex. Let Γ_1 be a wedge of circles and arcs, subdivided appropriately, and $f_1: \Gamma_1 \to \Delta$ a map, so that the circles in Γ_1, under f_1, represent α_j and the arcs in Γ_1 represent β_i. Thus, with appropriate base-point v_1,

$$f_1 \pi_1(\Gamma_1, v_1) = S.$$

Since Γ_1 is a finite connected graph, by **3.3** f_1 can be factored through a series of folds and an immersion $f: \Gamma \to \Delta$. By **4.4**, $f\pi_1(\Gamma, v) = S$, and since $\beta_i \notin S$, the image of the i^{th} arc of Γ_1 in Γ is *not* a circuit.

Now apply **6.1** to $f: \Gamma \to \Delta$, extending f to a covering $f': \Gamma' \to \Delta$, without adding new vertices. We define $S' = f' \pi_1(\Gamma', v)$. The index of S' in F is the number of vertices of Γ. Clearly $S \subset S'$; the paths in Γ which represent β_i – the images of the arcs of Γ_1 – are not circuits, and so $\beta_1 \notin S'$. A maximal tree T of Γ is also a maximal tree of Γ'; using this, we find a free basis of $\pi_1(\Gamma', v)$ of which a subset is a free basis of $\pi_1(\Gamma, v)$.

7. Core Graphs

7.1. Every connected graph with non-trivial fundamental group contains a core where the fundamental group is concentrated, and the original graph consists of this core with various trees hanging on.

In the case of a *finite graph* Γ, we can obtain the core of Γ by a process of clipping exterior hairs one by one. Thus, if Γ has no vertex of valence 1, Γ is its own core. If Γ has a vertex v such that there is a unique edge e with $v = \iota(e)$, then by removing v, e, and $\bar{e}$, we obtain a smaller graph Γ'. A finite number of such changes yields the core of Γ.

It is more elegant, however, to give an intrinsic definition.

7.2. A *cyclically reduced circuit* in a graph Γ is a circuit $p = e_1 e_2 \dots e_n$, which is reduced as a path, and for which $e_1 \neq \bar{e}_n$. This is the same as an immersion of a subdivided circle into Γ.

A graph Γ is said to be a *core-graph* if Γ is connected, has at least one edge, and every edge belongs to at least one cyclically reduced circuit.

If Γ is a connected graph with non-trivial fundamental group, an *essential edge* of Γ is an edge belonging to some cyclically reduced circuit. The *core of* Γ consists of all essential edges of Γ and all initial vertices of essential edges. We sometimes refer to the process of finding the core of Γ as "shaving off trees."

7.3. Some elementary exercises:

(a) *A finite connected graph Γ is a core-graph if and only if the valence of each vertex is at least two.* Thus the process of clipping off hairs described in **7.1**, yields the core of any finite connected graph with non-trivial fundamental group.

(b) *If Γ is a connected graph with non-trivial fundamental group, and Γ' is the core of Γ, then Γ' is a core-graph. If v is a vertex of Γ', then the inclusion $\pi_1(\Gamma', v) \to \pi_1(\Gamma, v)$ is an isomorphism.*

7.4. If S is a subgroup of a group G, we say that $S \subset G$ satisfies the *Burnside condition* when, for every $g \in G$, there exists some positive integer n, such that $g^n \in S$. Clearly if S satisfies the Burnside condition in G, then so does any conjugate subgroup.

7.5. (a) *Let $f : \Gamma \to \Delta$ be a finite-sheeted covering of connected graphs, v a vertex of Γ. Then $f\pi_1(\Gamma, v) \subset \pi_1(\Delta, f(v))$ satisfies the Burnside condition.*

(b) *Conversely: Let $f : \Gamma \to \Delta$ be an immersion of connected graphs. Suppose that Δ is a core-graph; v a vertex of Γ, $f\pi_1(\Gamma, v) \subset \pi_1(\Delta, f(v))$ satisfies the Burnside condition. Then f is a covering.*

Proof. **(a)** is clear, because subgroups of finite index satisfy the Burnside condition.

To prove **(b)**, we must show f is locally surjective. Let w be a vertex of Γ and e an edge of Δ with $f(w) = \iota(e)$. Since the Burnside condition is preserved under conjugation, and Γ is connected, then $f\pi_1(\Gamma, w) \subset \pi_1(\Delta, f(w))$ satisfies the Burnside condition. Since Δ is a core-graph, there is a cyclically reduced circuit

p in Δ whose first term is e. By the Burnside condition, there exists n such that the homotopy class of p^n belongs to $f\pi_1(\Gamma, w)$. Since p is cyclically reduced, p^n is reduced. There is a reduced circuit q of Γ based at w, such that $fq \sim p^n$; fq is reduced since f is an immersion (**5.1(a)**), and so $fq = p^n$ by **5.2**. Then the first term of q is an edge e_1 of Γ with $\iota(e_1) = w$ and $f(e_1) = e$.

7.6. Remark. This proposition finishes the algorithmic determination, promised in **5.4**, of whether a given finitely generated subgroup of a free group is of finite index. (We could also have used Marshall Hall's theorem to see this.) Represent the subgroup by an immersion $f \colon \Gamma \to \Delta$, where Γ is finite and Δ is a one-vertex graph (and hence a core-graph); if f is a covering, then the subgroup is of finite index equal to the number of vertices of Γ. If f is not a covering, then **7.5** shows the subgroup does not satisfy the Burnside condition, and so is of infinite index.

7.7. We can now sketch Gersten's proof of H. Neumann's bound on the rank of $S_1 \cap S_2$ (see **5.7**). We represent the ambient free group by the graph

$$\Delta = \; \begin{array}{c}\bullet\!\bullet\end{array}$$

We represent S_1 and S_2 by immersions $\Gamma_1, \Gamma_2 \to \Delta$, and the intersection $S_1 \cap S_2$ and by a component of the pullback Γ_3. If $S_1 \cap S_2$ is non-trivial, we can conjugate the situation within $\pi(\Delta)$ so as to bring the base-point of Γ_3 into the core Γ_3' of the component representing $S_1 \cap S_2$. Shaving the trees off of Γ_1 and Γ_2 to get their cores Γ_1' and Γ_2', we get a commutative square: (Note that under an immersion $A \to B$, the core of A maps into the core of B.)

$$
\begin{array}{ccc}
\Gamma_3' & \longrightarrow & \Gamma_1' \\
\downarrow{\scriptstyle f_3} & \searrow & \downarrow{\scriptstyle f_1} \\
\Gamma_2' & \underset{\scriptstyle f_2}{\longrightarrow} & \Delta
\end{array}
$$

Here, f_1 and f_2 are immersions, Γ_3' is a subgraph of their pullback, and all these graphs are finite core-graphs.

Let $\rho(X)$ denote the number of vertices of valence 3 in X. An elementary computation using Euler characteristics shows that, for a finite core-graph X in which all vertices have valence ≤ 3, $\rho(X) = 2(r(\pi_1(X)) - 1)$, where $r(G)$ is the rank of the free group G.

From the pullback nature of the diagram, we easily conclude that

$$\rho(\Gamma_3') \leqq \rho(\Gamma_1') \cdot \rho(\Gamma_2').$$

Putting these facts together with **7.3(b)**, we get the H. Neumann bound.

7.8. Greenberg [3] proved some theorems on Fuchsian groups, from which one can deduce the following remarkable result, which I shall prove using graph-theory.

Theorem. *Let* S_1 *and* S_2 *be finitely generated subgroups of a free group* F. *Suppose that* $S_1 \cap S_2$ *is of finite index both in* S_1 *and in* S_2. *Then* $S_1 \cap S_2$ *is of finite index in the join* $S_1 \vee S_2$.

Proof. As in the sketch of Gersten's proof of H. Neumann's inequality in **7.7**, we can represent the situation by a commutative square

In this picture, f_1 and f_2 are immersions, Γ_3 is a subgraph of the pullback; Γ_1, Γ_2, and Γ_3 are finite core-graphs; $f_3 = f_1 g_2 = f_2 g_2$; v_3 is a vertex of Γ_3, v_1 and v_2 are the images of v_3 in Γ_1 and Γ_2; $w = f_1(v_1) = f_2(v_2) = f_3(v_3)$. The group $F = \pi_1(\Delta, w)$, and

$$S_i = f_i(\Gamma_i, v_i) \quad \text{for } i = 1, 2, 3,$$
$$S_3 = S_1 \cap S_2.$$

By **7.5(b)**, since S_3 is of finite index in both S_1 and S_2, it follows that both g_1 and g_2 are coverings.

Let $r: \tilde{\Gamma}_3 \to \Gamma_3$ be the universal covering. Splicing this into our diagram:

We see that $\tilde{g}_1$ and $\tilde{g}_2$ are themselves universal coverings, and that the group $G(\tilde{g}_1)$ of covering translations of $\tilde{g}_1$ can be identified with S_1, and likewise $G(\tilde{g}_2) = S_2$.

Now f_3 is an immersion, and so the composition $\tilde{\Gamma}_3 \xrightarrow{r} \Gamma_3 \xrightarrow{f_3} \Delta$ is an immersion, denoted by $h = f_3 r$. Consider $\sigma \in G(\tilde{g}_1)$. Then $\sigma: \tilde{\Gamma}_3 \to \tilde{\Gamma}_3$ is an automorphism of graphs such that $\tilde{g}_1 \sigma = \tilde{g}_1$; that is, $g_1 r\sigma = g_1 r$; hence $f_1 g_1 r\sigma = f_1 g_1 r$, or $f_3 r\sigma = f_3 r$, or $h\sigma = h$. Thus σ is a translation of the immersion h. Similarly, the covering translations of $\tilde{g}_2$ are translations of h.

Let K be the group of translations of h generated by $G(\tilde{g}_1) \cup G(\tilde{g}_2)$. Then $\tilde{\Gamma}_3/K$ fits into a commutative diagram:

$$
\begin{array}{ccc}
\tilde{\Gamma}_3 & \xrightarrow{\ \tilde{g}_1\ } & \Gamma_1 \\
\tilde{g}_2 \downarrow & \ \ t_3 \searrow & t_1 \downarrow \quad \searrow f_1 \\
\Gamma_2 & \xrightarrow[\ t_2\ ]{} & \tilde{\Gamma}_3/K \\
& \searrow f_2 & \quad \searrow s \\
& & \Delta
\end{array}
$$

(This construction is a device to obtain the pushout $\tilde{\Gamma}_3/K$ of g_1 and g_2.)

Now K acts freely on $\tilde{\Gamma}_3$, by **5.8**. Thus t_3 is a covering. Since $\tilde{g}_1$ is a covering, it follows that t_1 is a covering. (Similarly t_2 is a covering). Because Γ_1 is a finite graph, t_1 is a finite-sheeted covering. This shows that S_1 is of finite index in $s\pi_1(\tilde{\Gamma}_3/K)$, which contains both S_1 and S_2; and so S_1 is of finite index in $S_1 \vee S_2$. With the fact that $S_1 \cap S_2$ is of finite index in S_1, this shows that $S_1 \cap S_2$ is of finite index in $S_1 \vee S_2$.

8. Comments

8.1. The argument in **7.8** can be done within the topological category. There is an esoteric theorem, which may be of some interest. By an *immersion*, we would mean a locally injective, continuous function. A *translation* σ of $f: A \to B$ would be a homeomorphism $A \to A$ such that $f\sigma = f$. The analogue of **5.8** would say that if A is a connected Hausdorff space and $f: A \to B$ is an immersion, then the group of translations of f acts freely, properly discontinuously on A. The analogue of the argument of **7.8** would say:

Let

$$
\begin{array}{ccc}
A & \xrightarrow{\ g_1\ } & B_1 \\
g_2 \downarrow & & \downarrow f_1 \\
B_2 & \xrightarrow[\ f_2\ ]{} & C
\end{array}
$$

be a commutative diagram; f_1, f_2 immersions; A a subspace of the pullback; g_1, g_2 coverings (a covering would be a local product with discrete fibre).

Let

$$
\begin{array}{ccc}
A & \xrightarrow{\ g_1\ } & B_1 \\
g_2 \downarrow & & \downarrow h_1 \\
B_2 & \xrightarrow[\ h_2\ ]{} & D
\end{array}
$$

be the topological pushout. Then if A, B_1, B_2 are connected, and A is Hausdorff, locally path-connected, and semi-locally-1-connected, then h_1 and h_2 are coverings.

(I am not sure I have all the topological hypotheses exactly right.)

8.2. An automorphism of a free group can be described by means of a subdivision of a graph and a sequence of foldings. I wonder if this would clarify various things about automorphisms, such as McCool's work [7].

8.3. The categorical viewpoint seems to be very fruitful in this subject. I wonder if "difference kernels" or equalizers, of injective homomorphisms of free groups have a graph-theoretic interpretation. For instance, I might generalize a conjecture of Scott's as follows: Suppose $\alpha, \beta: A \to B$ are homomorphisms of finitely generated free groups, and that α is a monomorphism. Define

$$D(\alpha, \beta) = \{x \in A \mid \alpha(x) = \beta(x)\}.$$

The conjecture is that $D(\alpha, \beta)$ is finitely generated. Perhaps the conjecture should be for the case that both α and β are monic. It is easily seen to be false when neither is assumed to be injective.

References

1. Burns, R.G.: A note on free groups. Proc. Amer. Math. Soc. **23**, 14–17 (1969)
2. Gersten, S.M.: Intersections of finitely generated subgroups of free groups and resolutions of graphs. Invent. Math. **71** (1983)
3. Greenberg, L.: Discrete groups of motions. Canad. J. Math. **12**, 414–425 (1960)
4. Hall, M., Jr.: Coset representations in free groups. Trans. Amer. Math. Soc. **67**, 421–432 (1949)
5. Howson, A.G.: On the intersection of finitely generated free groups. J. London Math. Soc. **29**, 428–434 (1954)
6. Imrich, W.: Subgroup theorems and graphs. Combinatorial Mathematics V. Lecture Notes Vol. 622. Berlin-Heidelberg-New York: Springer 1977
7. McCool, J.: Some finitely presented subgroups of the automorphism group of a free group. J. Algebra **35**, 205–213 (1975)
8. Neumann, H.: On the intersection of finitely generated free groups. Publ. Math. Debrecen **4**, 36–39 (1956); Addendum **5**, 128 (1957)
9. Serre, J-P.: Arbres, amalgames SL_2. Astérisque No. 46, Société Math. de France (1977)