Folding graphs and applications, d'après Stallings

Mladen Bestvina

Fall 2001 Class Notes, updated 2010

1 Folding and applications

A *graph* is a 1-dimensional cell complex. Thus we can have more than one edge with the same pair of endpoints, and we can have edges that are loops. We will usually label the *oriented* edges by $a, b, \cdots$, and when we want the same edge with the opposite orientation we will put a "bar" over the label, or else capitalize the letter. E.g. $\overline{a} = A$ and $\overline{\overline{a}} = a$. The initial vertex of an edge e is denoted $\iota(e)$ and the terminal vertex is $\tau(e)$. Thus $\tau(e) = \iota(\overline{e})$.

More formally (Serre), a graph is a quadruple $(V, E, \overline{\ }, \iota)$ where V and E are sets (of vertices and oriented edges, respectively), $\overline{\ } : E \to E$ is a free involution (i.e. $\overline{e} \neq e$ and $\overline{\overline{e}} = e$), and $\iota : E \to V$ is a function.

A *morphism* of graphs is a cellular map that sends each open edge homeomorphically onto an open edge. Changing such a map by a homeomorphism isotopic to the identity rel vertices will be regarded as the same morphism. On the formal side, a morphism $(V, E, \overline{\ }, \iota) \to (V', E', \overline{\ }, \iota)$ is a pair of maps $V \to V'$, $E \to E'$ that commute with $\overline{\ }$ and ι.

Definition 1. A morphism $f : G \to G'$ is an *immersion* if it is locally injective (i.e. each point has a neighborhood on which f is injective). Since it is always injective on open edges, we only need to check this on the vertices. Serre would give the definition as:

$$\iota(e_1) = \iota(e_2) \ \&\ f(e_1) = f(e_2) \Rightarrow e_1 = e_2.$$

E.g. covering maps are immersions, and so are inclusions of subgraphs. Compositions of immersions are immersions.

Exercise 2. Suppose $f : X \to Y$ is an immersion between finite graphs. Show that it is possible to attach finitely many 0- and 1-cells to X to create a graph $\tilde{X}$ and to extend f to $\tilde{f} : \tilde{X} \to Y$ so that $\tilde{f}$ is a covering map. Conclude that every immersion between finite graphs is the composition of an embedding and a covering map. Moreover, if Y has only one vertex, one can construct $\tilde{X}$ by attaching 1-cells only.

Definition 3. An *edge-path* in a graph G is a sequence of edges $e_1, e_2, \cdots, e_k$ such that $\tau(e_i) = \iota(e_{i+1})$. An edge-path can be thought of as a morphism $I \to G$ where I is a graph homeomorphic to $[0,1]$ (provided with an orientation) or a point (sometimes it's convenient to allow empty paths as well).

An edge-path is *tight* or *reduced* if $I \to G$ is an immersion, or equivalently, no two consecutive edges are of the form $e, \overline{e}$. An *elementary homotopy* is the move that deletes such consecutive edges (with the understanding that if $e, \overline{e}$ is the entire path, the new path is the initial point), or inserts them.

Exercise 4. If two edge-paths are related by an elementary homotopy, then they are homotopic rel endpoints (with the understanding that if one path is a point, then the statement means that the other is homotopic to constant rel endpoints).

Exercise 5. Every edge-path can be transformed to a reduced edge-path by applying a sequence of elementary homotopies. This process is called *tightening*.

OK, this is too easy! Here is a slightly more challenging problem:

Exercise 6. If two reduced paths are homotopic rel endpoints, then they are equal. In particular, two edge-paths are related by a sequence of elementary homotopies iff they are homotopic rel endpoints. Hint: Universal cover. Define when a graph is a tree, show that the universal cover of a connected graph is a tree, and finally show that reduced edge-paths in a tree are embedded and they are determined by their endpoints.

Exercise 7. A nontrivial reduced edge-path that starts and ends at the base vertex represents a nontrivial element of π_1. An immersion between connected graphs is injective in π_1.

Definition 8. Suppose e_1, e_2 are two edges of G with $e_1 \neq e_2$ and $e_1 \neq \overline{e_2}$ and with $\iota(e_1) = \iota(e_2)$. Form a new graph G' by identifying e_1 with e_2 and identifying $\tau(e_1)$ with $\tau(e_2)$ (these may already be identified!). The quotient map $q : G \to G'$ is a morphism of graphs, and we call it a *fold*. A formal definition is left to the reader. See Figure 1

Exercise 9. Show that folds of the second kind are epimorphisms in π_1.

Theorem 10 (Stallings, 1983)**.** *Every morphism $G \to G'$ of finite graphs factors as*

$$G = G_0 \to G_1 \to G_2 \to \cdots \to G_k \to G'$$

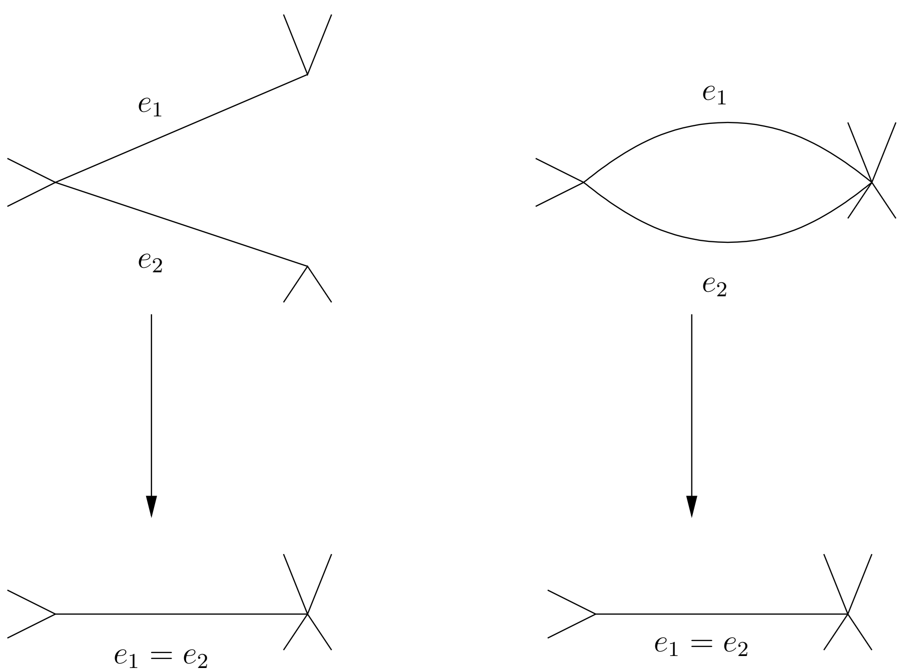

Figure 1: Two kinds of folds. The first is a homotopy equivalence. The second is not.

where the last map $G_k \to G'$ is an immersion and all the other maps in the sequence are folds. Moreover, such a factorization can be found by a very fast algorithm.

OK, so this is all kindergarten. Still, the beauty of it is that it can be used to give simple proofs of classical theorems about free groups. Just remember the proof that every subgroup of a free group is free!

2 Subgroups, cores

Start by considering the following problem:

> *Given $x_1, x_2, \cdots, x_k \in F_n$, find a basis of the subgroup $H =< x_1, x_2, \cdots, x_k >$ generated by them.*

We will consider the following example:
$F_2 =< a, b >$, free group of rank 2, $H =< a^3b, \overline{a}bab, a^2\overline{b}a >$.

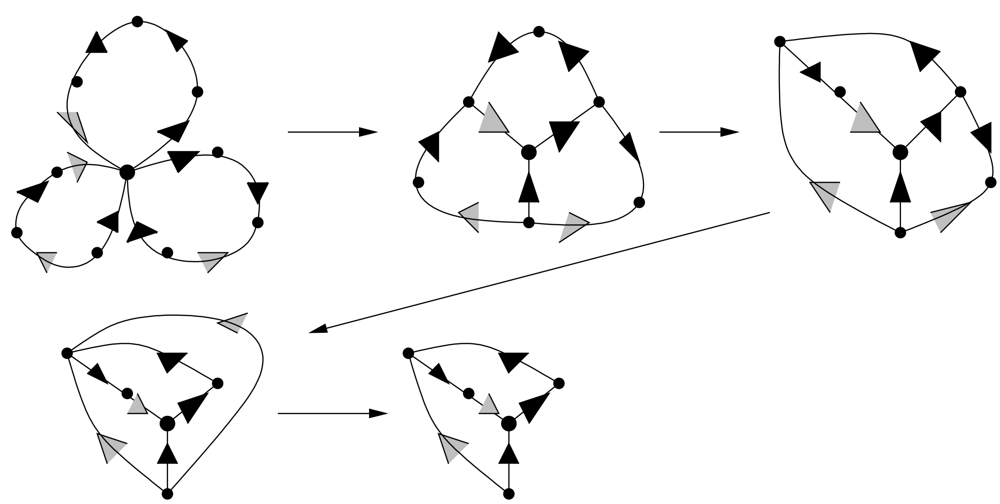

Figure 2: Solid arrows represent a, and the shaded arrows are b. The basepoint is denoted by the large dot.

We start with the map from the bouquet of 3 circles to the bouquet of two circles representing the 3 generators of H. Such a map is described by a labeling of a subdivision of the bouquet of 3 circles. Then we find the factorization as in Theorem 10.

The first map represents 3 simultaneous folds. The others are single folds, and the last fold is of the second kind. Keep in mind that each labeled graph G_i represents a morphism $g_i : G_i \to Y$ to the bouquet of two circles Y. Note that folds commute with the g_i's, so that the image in π_1 of all g_i's is the same. But for $i = 1$ it is H. From the last graph we read off a basis for H: $a^3b, \overline{a}bab$.

Question 11. Does the ultimate graph depend on how we chose to fold?

Of course not! Why? It is precisely the *core* of the covering space Y_H of Y corresponding to H. By the core I mean any of the following equivalent items:

- the largest connected finite subgraph that contains the base vertex and has no valence 1 vertices [1] except possibly the base vertex.
- the smallest subgraph that contains the base vertex and to which Y_H deformation retracts.
- the union of the images of all reduced edge-paths that begin and end at the base vertex.
- the union of the images of the finitely many reduced edge-paths that represent generators of H.

Exercise 12. Prove that the above descriptions are equivalent.

Exercise 13. Show that Y_H can be constructed from the core by attaching (infinite, usually) trees to the vertices.

Hints for Exercises 12 and 13: First observe that the third and the fourth bullet are equivalent, by arguing that the image of a path representing the product of two elements is contained in the union of the images of paths representing the factors. Call this the core. Then argue the property from Exercise 13 (if it fails, find an element whose edge-path is not contained in the core). Then argue the first two bullets.

So you could construct the whole covering space by adding a bunch of trees to the last graph. For practice, draw a few edges of these trees. The core is thus canonically associated to the subgroup H (given a fixed basis of the underlying free group). It is the topological representative of H. So when someone tells you: "Let H be a finitely generated subgroup of a free group F_n..." you should immediately visualize an immersion $G \to Y$ as above.

[1] The valence of a vertex v is the cardinality of $\iota^{-1}(v)$.

Exercise 14. Suppose H is a finitely generated normal subgroup of F_n. Show that either H has finite index in F_n or $H = \{1\}$.

Exercise 15. Can you always compute the normalizer

$$N(H) = \{\gamma \in F_n | \gamma H \gamma^{-1} = H\}?$$

What can you say about the index $[N(H) : H]$? (Answer: It is always finite, and bounded by the number of vertices in the graph representing H. Recall that $N(H)/H$ is the deck group.)

Question 16. Given $w \in F_n$, can you algorithmically decide whether $w \in H$.

Yes! If the reduced word w can be realized by an immersed path in the core that begins and ends at the base vertex, then $w \in H$. Otherwise $w \notin H$. E.g., in our example $a^2\bar{b}a \in H$ while $b \notin H$ (it cannot be drawn from the base vertex) and $a \notin H$ (it can be drawn, but doesn't close up).

Question 17. Given $w \in F_n$, can you algorithmically decide whether w is conjugate to an element of H.

Yes! First replace w by its conjugate if necessary so that it is cyclically reduced (e.g. $ababa^{-1}$ would be replaced by bab). If w can be realized by an immersed path in the core that begins and ends at the same (but arbitrary) vertex, then w can be conjugated into H. Otherwise, the answer is no. (The reason for passing to the cyclically reduced word is to guarantee that any lift is contained in the core.) conjugacy class of w can be thought of as an immersion $S^1 \to Y$ and the question is whether this lifts to the core. This is a finite check.

Now some exercises for you.

Exercise 18. Can you tell if H is normal in F_n?

Exercise 19. Given a homomorphism $h : F_n \to F_m$, can you tell if h is injective, surjective, bijective? Answer: Injective iff there are no folds of the second kind. Surjective iff the last map is a homeomorphism.

Exercise 20. Show that for every homomorphism $h : F_n \to F_m$ there is a free factorization $F_n = A * B$ such that h kills A and is injective on B.

Exercise 21. Show that for every finitely generated $H \subset F_n$ there is a subgroup $H' \subset F_n$ such that $H \subset H'$, H is a free factor in H', and H' has finite index in F_n. This is called Marshall Hall's theorem. You can find H' algorithmically. Do it for H in our example. Hint: Add some edges to G to turn an immersion $G \to Y$ into a covering map.

Question 22. Can you decide if H has finite index in F_n? Hint: Yes! The core must be a covering.

Remark 23. The same argument shows more: If $H \subset F_n$ has the property that for every $w \in F_n$ there is $k = k(w) > 0$ such that $w^k \in H$, then H has finite index in F_n. This is not true for arbitrary groups. For example, there exist infinite groups (even finitely generated!) where every element has finite order. The most famous such family of groups are called Burnside groups.

Exercise 24. This is a generalization of Exercise 21. Show that for every finitely generated subgroup $H < F_n$ and any $g_1, \cdots, g_k \in F_n \setminus H$ there is a subgroup H' such that $H < H' < F_n$, H is a free factor in H', H' has finite index in F_n and $g_1, \cdots, g_k \notin H'$. In particular, F_n satisfies LERF.

Hints, discussion: A group G satisfies LERF (locally extended residually finite) if for every f.g. $H < G$ and every $g \in G \setminus H$ there is a subgroup of finite index H' in G such that $H < H'$ and $g \notin H$ (why is one g here enough?). G is residually finite when this property holds for $H = \{1\}$. One of the issues in 3-manifold topology is whether (say hyperbolic) 3-manifold fundamental groups are LERF. This would allow the construction of certain covering spaces where immersed surfaces lift to embeddings. Now for the hint: Start with the wedge of circles and arcs representing H and the g_i's. Fold (what else?) until the gadget immerses. Note that the endpoints of the arcs are different from the basepoint. Complete to a covering map. Done.

3 Nielsen's generators of $Aut(F_n)$

The group of all automorphisms of $F_n =< a_1, \cdots, a_n >$ is denoted $Aut(F_n)$. It is perhaps surprising that this group is finitely generated.

Theorem 25 (Nielsen)**.** *The following 3 types of automorphisms σ generate $Aut(F_n)$.*

(1) *(Permutation of basis elements) σ induces a permutation of $a_1, \cdots, a_n$.*

(2) *(Sign changes) σ sends each a_i either to a_i or to a_i^{-1}.*

(3) *(Change of maximal tree) There is a special basis element a_i that is sent to $a_i^{\pm 1}$. Every other basis element a_j is sent to one of: a_j, $a_i^{\pm 1}a_j$, $a_ja_i^{\pm 1}$, or $a_i^{\pm 1}a_ja_i^{\mp 1}$.*

Discussion: The first two types generate a finite subgroup, called the *signed permutation group*, of order $2^n n!$. There is quite a bit of redundancy among these generators.

Exercise 26. Using Nielsen's theorem show that $Aut(F_n)$ is generated by the following 4 automorphisms:

$\sigma_1(a_i) = a_{i+1}$ (indices are taken cyclically).
$\sigma_2(a_1) = a_2, \sigma_2(a_2) = a_1, \sigma_2(a_i) = a_i$ for $i > 2$.
$\sigma_3(a_1) = a_1^{-1}, \sigma_3(a_i) = a_i$ for $i > 1$.
$\sigma_4(a_1) = a_1 a_2, \sigma_4(a_i) = a_i$ for $i > 1$.

Bernhard Neumann found a two-element generating set for $Aut(F_n)$, $n \geq 4$.

Why is the last type called "change of maximal tree"? To identify $\pi_1(G)$ with F_n we need some additional data, most important of which is a maximal tree $T \subset G$. We also choose orientations of the edges outside T and we choose a bijection between the edges and the bases elements of F_n. Now every closed edge-path based at the base vertex determines an element of F_n by reading off the labels of the edges outside T. Changing the bijection corresponds to type (1) automorphisms, and changing orientations corresponds to type (2). Now consider changing T: Let e be an edge outside T. Then $T \cup e$ is homotopy equivalent to S^1, and there is a cycle of edges that contains e. Let f be an edge of T contained in this cycle, so that $T' = T \cup e \setminus f$ is a maximal tree. We can also assume that e and f are compatibly oriented, so that they determine an orientation of the circle. The T-rule assigns to e the closed path that from the base vertex travels along T to the circle, then around the circle, and back to the base vertex. The word read off is f, i.e. $f = e$. Now take some e_i that is outside both T and T' and represents a basis element with both choices of trees. The T-rule assigns to e_i a path that can be schematically viewed as "Y" with e_i connecting the two endpoints at the top and the base vertex at the bottom. If f is not in Y, then $e_i \mapsto e_i$. If f is in the vertical arc pointing up (we may assume this by changing orientations) then $e_i \mapsto f e_i f^{-1}$. If f is in the north-west arc, then we have $e_i \mapsto f e_i$, and the north-east case reads $e_i \mapsto e_i f^{-1}$. Thus we have a type (3) automorphisms. See an example in Figure 3.

Exercise 27. If T and T' are two maximal trees, there is a sequence $T = T_0, T_1, \cdots, T_k = T'$ of maximal trees such that any two consecutive trees differ in only one edge, as above.

Exercise 28. This is a bit more ambitious. Consider the simplicial complex whose vertices are non-closed edges of G, and a collection of edges spans a simplex if their union is a forest. Draw some examples. Can you make a conjecture about the homotopy type of the complex?

Proof of Theorem 25. Let $\alpha : F_n \to F_n$ be an automorphism. Let R be

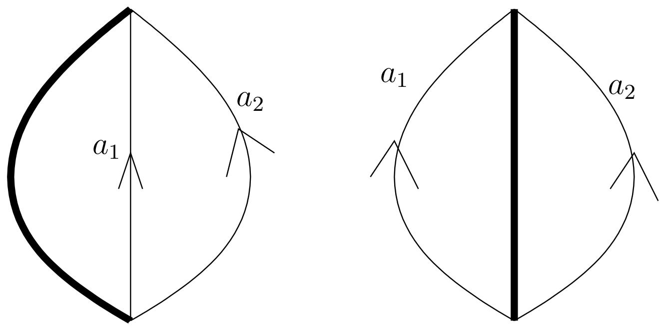

Figure 3: Thick lines represent maximal trees. The base vertex is at the bottom. To find the associated automorphism $\sigma :< a_1, a_2 >\to< a_1, a_2 >$, compute $\sigma(a_1)$ by first drawing the corresponding loop on the left graph (up the middle, down on left), translating the loop to the right graph, and reading the labels, to get $\sigma(a_1) = a_1^{-1}$. Similarly, $\sigma(a_2) = a_2 a_1^{-1}$ (up on the right, down on the left).

the rose corresponding to $a_1, \cdots, a_n$ and let X be the rose with subdivided edges representing α. So the first loop of X represents $\alpha(a_1)$ and it is subdivided into $length(\alpha(a_1))$ edges labeled by $a_1, \cdots, a_n$ and their inverses, and reading off the word $\alpha(a_1)$. There is an induced map $\rho : X \to R$. Now factor this map into the composition of folds. All folds are of the first type and the last map is a homeomorphism.

Identify $\pi_1(X)$ with F_n using a maximal tree, and identify $\pi_1(R)$ with F_n. If the orientations are chosen correctly, the map induced by ρ is α. First suppose that the first fold involves two embedded edges. Change the maximal tree if necessary to arrange that the first fold takes place in T. Then the image of T is a maximal tree and the induced homomorphism is identity. Then change the tree again etc. If a fold involves an embedded edge and a loop, arrange that the embedded edge is in T – the induced homomorphism is of type (3). At the end we have a homeomorphism and the induced map in π_1 is a signed permutation. □

Exercise 29. Let $F_2 =< x, y >$. Show that every automorphism $F_2 \to F_2$ sends the commutator $[x, y] = xyx^{-1}y^{-1}$ to an element that is conjugate to either $[x, y]$ or to $[x, y]^{-1}$. Hint: Check that this holds for the Nielsen generators. Info: There is nothing comparable in higher rank.

Exercise 30. It is a fact from surface topology that a homotopy equivalence $X \to Y$ between two compact connected surfaces which induces a homeomorphism $\partial X \to \partial Y$ is homotopic to a homeomorphism $X \to Y$. Using this fact and the previous exercise, show that for every automorphism $h : F_2 \to F_2$ there is a homeomorphism of the torus with a disk cut out that in π_1 induces an automorphism conjugate to h (we don't have basepoints, so getting h up to conjugacy is the best we can hope for). Hint: Torus with a disk cut out is obtained from the wedge of two circles by gluing an annulus along one boundary circle and with the commutator as attaching map. Extend the map to the annulus....

Exercise 31. Find explicit homeomorphisms that correspond to the Nielsen generators. That's another way to prove the statement of Exercise 30. Hint: Isometries (with respect to *some* metric) or Dehn twists.

4 Intersections

Given two subgroups H_1, H_2 how do you compute their intersection? Suppose $f_i : G_i \to Y$ is an immersion that represents H_i $(i = 1, 2)$. The *pull-back* of f_1, f_2 is a pair $g_1 : G_0 \to G_1$ and $g_2 : G_0 \to G_2$ of morphisms of graphs such that $f_1 g_1 = f_2 g_2$ and characterized by the universal property: For every pair $\phi_1 : G' \to G_1, \phi_2 : G' \to G_2$ of morphisms such that $f_1\phi_1 = f_2\phi_2$ there is a *unique* morphism $\phi : G' \to G_0$ such that $g_1\phi = \phi_1$ and $g_2\phi = \phi_2$.

$$\begin{tikzcd}
G' \arrow[dr, dashed, "\phi"] \arrow[drr, "\phi_1"] \arrow[ddr, "\phi_2"'] & & \\
& G_0 \arrow[r, "g_1"'] \arrow[d, "g_2"] & G_1 \arrow[d, "f_1"] \\
& G_2 \arrow[r, "f_2"'] & Y
\end{tikzcd}$$

From the universal property it is easy to see that pull-backs are unique (if they exist). The existence is shown by the usual direct construction: A vertex of G_0 is a pair (v_1, v_2) with v_i a vertex of G_i such that $f_1(v_1) = f_2(v_2)$. Similarly, an (oriented) edge of G_0 s a pair (e_1, e_2) with e_i an edge of G_i and with $f_1(e_1) = f_2(e_2)$. The involution is defined by $\overline{(e_1, e_2)} = (\overline{e_1}, \overline{e_2})$ and

$\iota(e_1, e_2) = (\iota(e_1), \iota(e_2))$. Maps $g_1 : G_0 \to G_1$ and $g_2 : G_0 \to G_2$ are induced by the projections to the first and second coordinates.

Exercise 32. Check that this is really a graph and that it satisfies the definition of the pull-back.

If v_i is the base vertex of G_i that maps to the base vertex of Y, we declare (v_1, v_2) the base vertex of G_0.

Proposition 33. *If f_1, f_2 are immersions, so are g_1, g_2. In that case, the image in π_1 of $f_1 g_1 = f_2 g_2 : G_0 \to Y$ equals the intersection of the images in π_1 of f_1 and f_2.*

Proof. Exercise. Use the universal property for G' an arc. $\square$

As an example, consider $H_1, H_2 \subset F_2 =< a, b >$ given by $H_1 =< a^2, ba >$, $H_2 =< ba, b^3 a\bar{b}a >$. First, using folding, find the core representatives $f_i : G_i \to Y$ of H_i. Then construct the pull-back. See Figure 4.

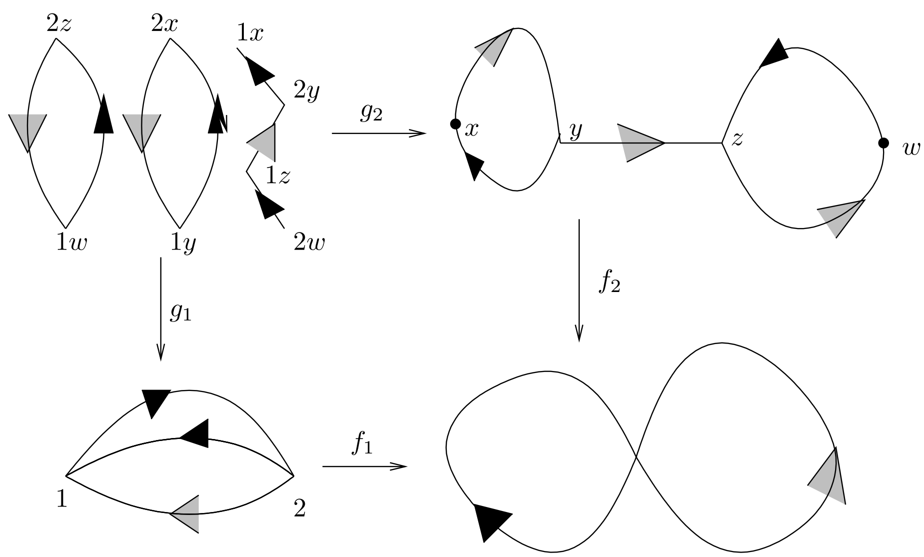

The base vertices are 2, x, and $2x$. We read off that $H_1 \cap H_2 =< ba >$.

What is the significance of the other circle in the pull-back? Well, if our base vertices were 1, w, and $1w$ we would be looking at the other circle. The change of base vertices corresponds to conjugating the subgroups. Specifically, we see that $aH_1\bar{a} \cap a\bar{b}aH\bar{a}b\bar{a} =< ab >$. This is because a is a path in G_1 from 1 to 2, and $a\bar{b}a$ is a path in G_2 from w to x.

What if we used a different path between the vertices, e.g. $\bar{b}$ from 1 to 2? Well, we would get the same conjugate: $\bar{b}H_1b = aH_1\bar{a}$, since $ba \in H_1$.

Here is the theorem that can be proved using pull-backs.

Theorem 34. *Suppose H_1, H_2 are two (free) subgroups of F_n of finite rank.*

(1) *The intersection $H_1 \cap H_2$ has finite rank.*

(2) *There are only finitely many double cosets $H_1gH_2 \in H_1\backslash G/H_2$ such that $g^{-1}H_1g \cap H_2 \neq \{1\}$.*

(1) This is known as Howson's theorem.

Regarding (2), note that if $\gamma = h_1gh_2$ for some $h_1 \in H_1$ and $h_2 \in H_2$ then $\gamma^{-1}H_1\gamma \cap H_2 = h_2^{-1}(g^{-1}H_1g \cap H_2)h_2$. Furthermore, simultaneous conjugation does not change the rank of the intersection: $g^{-1}H_1g \cap g^{-1}H_2g = g^{-1}(H_1 \cap H_2)g$. To disallow this simultaneous conjugation, we chose to keep H_2 fixed, and further, we consider only representatives of double cosets.

Proof. (1) follows from the fact that pull-backs of finite graphs are finite.

(2) Let $G_i \to Y$ be immersions representing H_i ($i = 1, 2$) and let $\tilde{G}_i$ be the covering space of Y corresponding to the subgroup H_i. Thus G_i is the core of $\tilde{G}_i$. Let $\tilde{G}_0$ be the pull-back of $\tilde{G}_1 \to Y$ and $\tilde{G}_2 \to Y$. Thus $G_0 \subset \tilde{G}_0$. Note that $\tilde{G}_0$ may have (and usually does) infinitely many components. However, using the universal property of pull-backs, it is easy to check that every component of $\tilde{G}_0$ that is not contractible intersects G_0 nontrivially, and in fact any immersed loop in $\tilde{G}_0$ is contained in G_0. In particular, only finitely many components of $\tilde{G}_0$ are not trees, and each of these components has finite rank.

Let H_1' and H_2' be two conjugates of H_1 and H_2, say by w_1 and w_2 respectively. Developing w_i in $\tilde{G}_i$ from the basepoint produces a new basepoint x_i. The component of $\tilde{G}_0$ that contains (x_1, x_2) computes $H_1' \cap H_2'$. It remains to observe that if (x_1, x_2) and (x_1', x_2') obtained from (w_1', w_2') belong to the same component of $\tilde{G}_0$, then there is $w \in F_n$ (coming from a path in $\tilde{G}_0$ joining (x_1, x_2) with (x_1', x_2')) such that

$$h_i := w_iww_i'^{-1} \in H_i \qquad i = 1, 2$$

which implies that

$$w_1w_2^{-1} = (h_1w_1'w^{-1})(ww_2'^{-1}h_2^{-1}) = h_1w_1'w_2'^{-1}h_2^{-1}$$

and $w_1'w_2'^{-1}$ are in the same double coset. Thus the pair $(H_1^{w_1}, H_2^{w_2})$ is related to $(H_1^{w_1'}, H_2^{w_2'})$ by simultaneous conjugation and then changing the double coset representative. $\square$

It is an interesting question to estimate the rank of $H_1 \cap H_2$ in terms of the ranks of H_1 and H_2. We see immediately that the ambient free group F_n plays no role since we could embed it into another free group, e.g. F_2. To compute the rank of a graph, we make use of the following fact.

Exercise 35. Let G be a finite graph. Then the quantity

$$-\frac{1}{2}V_1 + \frac{1}{2}V_3 + V_4 + \frac{3}{2}V_5 + \cdots$$

equals the negative of the Euler characteristic of G, i.e. $rank(G) - 1$ when G is connected.

Now consider the pull-back diagram corresponding to H_1, H_2. For simplicity, we take the ambient group to be F_2, so all vertices have valence ≤ 4. Denote by V_3, V_4 the numbers of valence 3 and 4 vertices in G_1 and by V_3', V_4' the analogous numbers for G_2. Consider the core of the component of the pull-back that contains the base vertex. How many vertices of what valence can we have there? Well, the worst case is that we have V_4V_4' vertices of valence 4 and $V_3V_3' + V_3V_4' + V_4V_3'$ vertices of valence 3. We need to estimate the quantity

$$V_4V_4' + \frac{1}{2}(V_3V_3' + V_3V_4' + V_4V_3').$$

It is clearly estimated above by

$$2(V_4 + \frac{1}{2}V_3)(V_4' + \frac{1}{2}V_3').$$

So we proved:

Theorem 36. *If* $H_1 \cap H_2 \neq \{1\}$ *then*

$$rank(H_1 \cap H_2) - 1 \leq 2(rank(H_1) - 1)(rank(H_2) - 1).$$

This theorem is known as the Hanna Neumann inequality. In the same paper, she states the Hanna Neumann Conjecture, that the factor 2 can be dropped. In fact, the above proof shows that the sum

$$\sum (rank(H_1' \cap H_2) - 1)$$

over all distinct conjugates H_1' of H_1 with $H_1' \cap H_2 \neq \{1\}$ is bounded by $2(rank H_1 - 1)(rank H_2 - 1)$. It was conjectured by Walter Neumann (Hanna's son) that 2 can be removed in this stronger assertion. Update 2014: The more general conjecture was proved independently by Joel Friedman and by Igor Mineyev.

To finish, another exercise.

Exercise 37. A subgroup H of a group Γ is said to be *malnormal* if

$$\gamma \in \Gamma, h \in H, \gamma h \gamma^{-1} \in H \Rightarrow \gamma \in H \text{ or } h = 1.$$

Show that $H \subset F_n$ is malnormal iff the pull-back diagram for the pair (H, H) has at most one component that is not a tree. E.g. $< a, ba\overline{b} >$ is not malnormal.

5 Todd-Coxeter

Algorithmic questions for general groups tend to be unsolvable. Let

$$< a, \cdots, b | r, \cdots, s >$$

be a finite presentation of a group G. There are three famous questions, due to Dehn, that one tries to solve:

1. (The Word Problem) Is there an algorithm that for a given a word w in $a, \overline{a}, \cdots, b, \overline{b}$ decides whether w is trivial in G?

2. (The Conjugacy Problem) Is there an algorithm that for two words w_1, w_2 decides whether they are conjugate in G?

3. (The Isomorphism Problem) Is there an algorithm that decides whether a given finite presentation determines the trivial group? More generally, is there an algorithm that decides whether two given finite presentations determine isomorphic groups?

Novikov (1955) constructed an example of a f.p. group with unsolvable word problem. Then Adyan (1955) proved the following decisive result: Consider a property P of groups with the following features:

a) If $G \cong G'$ and G has P, so does G',

b) If H is a subgroup of G and G has P, so does H,

c) The free group F_n does not satisfy P for some n, and

d) There exists a group that satisfies P.

Then there is no algorithm that decides whether a given finite presentation determines a group that satisfies P.

For example "being the trivial group", or "being abelian" etc. are such properties.

However, there is a "process" that can be applied to a finite presentation, that attempts to construct the Cayley graph of the group. It will terminate iff the group is finite. Even when the group is infinite and "geometric" it will give better and better approximations to the Cayley graph. There is a lot of current research that attempts to understand "geometric" classes of groups, e.g. hyperbolic groups (Gromov, 1985), CAT(0) groups etc.

The process, known as the Todd-Coxeter process, can be described in terms of folding (Stallings-Wolf, 1987).

First, recall what the Cayley graph is. Let $F = \pi_1(Y)$ be the free group on $a, \cdots, b$, the fundamental group of the rose Y. Let R be the smallest normal subgroup of F that contains $r, \cdots, s$. Thus $G = F/R$, by definition. Let X be the covering space of Y corresponding to $R \subset F$. Then X is the Cayley graph of G. This is a regular covering space and the deck group can be identified with G. After choosing a base vertex, labeled $1 \in G$, the vertex set is in a natural 1-1 correspondence with the elements of G: for a vertex v choose a path α from 1 to v and project it to Y. The projection is a loop and it determines an element of G. Furthermore, the (oriented) edges of X have labels $a, \overline{a}, \cdots, b, \overline{b}$, and if an edge labeled a goes from a vertex labeled g to the vertex labeled g', then $ga = g'$. This gives another (more standard) description of the Cayley graph.

Here is then the procedure:

1. Start with the rose Y with oriented petals labeled $a, \cdots, b$.

2. Let Z be the rose with one petal for each relation $r, \cdots, s$. The petal labeled r (say) is subdivided into $length(r)$ edges that are labeled according to r (and similarly for other relations). The labeling defines a morphism $f : Z \to Y$. Now run the folding algorithm to replace f by an immersion $f_1 : Z_1 \to Y$. The image in π_1, call it R_1, did not change.

3. Take two copies of Z_1 and identify a pair of vertices. Convert the resulting morphism to an immersion $f_2 : Z_2 \to Y$ with the image R_2 in π_1. Note that R_2 is generated by the union of two conjugates of R_1.

4. Continue this process indefinitely.

The aim is to make $Z_i \to Y$ look more and more like a covering map in a larger and larger neighborhood of 1. When G is finite, the procedure can be made to terminate (with clever choices of vertices that are identified) with a regular covering space, which is the Cayley graph of G.

The process can be generalized to the situation where H is a subgroup of G and the covering space being built corresponds to the group generated by $R \cup H$. The procedure will terminate when H has finite index in G. This is called the Todd-Coxeter coset enumeration process. There are programs on the internet, notably GAP, that will perform this process free of charge.