BASS-SERRE THEORY AND COMPLEXES OF GROUPS

1. Free groups, free products

Just a summary for now. See [Ser03, I.§3] for some of this.

- Three points of view: Universal, Typographical, Topological. (Exercise: these all describe the same thing.)
- Exercise: Every element of a free group is represented by a unique reduced word. Every element of a free product as well (though you have to interpret "reduced word" to allow "letters" which are arbitrary non-trivial elements of the free factors).
- A group is free if and only if it acts freely on a tree.

2. Amalgams

Given a diagram $\mathcal{D}: \ A \xleftarrow{\alpha} C \xrightarrow{\beta} B$ the colimit $\varinjlim(\mathcal{D})$ (or *pushout*) exists and is usually written $A *_C B$ (suppressing the maps α and β from the notation). It comes with canonical maps $A \to A *_C B$ and $B \to A *_C B$ making a commutative square

$$\begin{CD} C @>{\beta}>> B \\ @V{\alpha}VV @VVV \\ A @>>> A *_C B \end{CD} \tag{1}$$

which satisfies a universal property: For any homomorphisms $\phi\colon A \to G$, $\psi\colon B \to G$ so that $\phi\alpha = \psi\beta$, there is a unique homomorphism from $A *_C B$ to G making the following diagram commute:

(2)

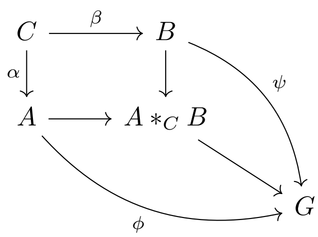

One way to define the group is as follows:

Definition 2.1. The amalgam $A *_C B$ is the quotient of $A * B$ by the normal subgroup generated by $\{\alpha(c)\beta(c)^{-1} \mid c \in C\}$.

There is also a topological interpretation:

Definition 2.2. For $F \in \{A, B, C\}$ let K_F be a $K(F, 1)$, and let $f_A\colon K_C \to K_A$ and $f_B\colon K_C \to K_B$ induce the maps α, β, respectively. Let $Y = K_A \sqcup K_C \times [0,1] \sqcup K_B / \sim$ where $(x, 0) \sim f_A(x)$ and $(x, 1) \sim f_B(x)$ for $x \in K_C$. We can also define $A *_C B$ to be $\pi_1(Y)$.

Date: September 14, 2022.

The equivalence of these follows from van Kampen's Theorem (see [Hat02, 1.20]).

The space Y defined in the last paragraph is a special case of something called a *graph of spaces*, in this case with two *vertex spaces* (K_A and K_B) and one *edge space* (K_C). We will see more graphs of spaces below.

The key result we proved was about *normal forms*. To define normal form we first choose right transversals T_A for $C < A$ and T_B for $C < A$. This means that T_A is a subset of A containing one element of each right coset Cg in A. We also require that $C \cap T_A = \{1\}$. Similar statements apply to T_B. A normal form for $g \in A *_C B$ is a tuple $(c; x_1, \ldots, x_n)$ so that

(1) $c \in C$;
(2) $x_i \in T_A \cup T_B - \{1\}$ for each i;
(3) If $x_i \in T_A$ then $x_{i+1} \in T_B$ and vice versa; and
(4) $cx_1 \cdots x_n = g$.

In class we proved:

Theorem 2.3. *Each $g \in A *_C B$ has a unique normal form.*

A more general result (where possibly more than two subgroups are amalgamated along a common subgroup) is proved in [Ser03, I.§1]. Various corollaries can also be found there, and we talked about a few of them.

3. HNN extensions

(Here we followed Section I.1.4 of Serre pretty closely, so I won't reproduce it here.)

4. Trees and amalgams

(See also [Ser03, I.4].) There is a third important way to think of amalgams; namely, as groups which act on a tree in a certain way. In the following, we write $A *_C B$ for the amalgam to indicate we are thinking of the maps α and β as inclusions.

Theorem 4.1. *The following are equivalent.*

(1) *$G \cong A *_C B$.*
(2) *There is an action of G on a (bipartite) tree T with strict fundamental domain consisting of a single edge e with endpoints v and w, and a commutative diagram*

$$(*)\qquad \begin{array}{ccccc} A & \longleftarrow & C & \longrightarrow & B \\ \downarrow & & \downarrow & & \downarrow \\ \mathrm{Stab}(v) & \longleftarrow & \mathrm{Stab}(e) & \longrightarrow & \mathrm{Stab}(w) \end{array}$$

with each vertical map an isomorphism.

Proof. (2) $\implies$ (1)**:** The diagram ($*$) together with the universal property of pushouts gives a unique homomorphism $\phi\colon A *_C B \to G$ which extends the given maps. We first show that ϕ is surjective (using that T is connected) and then show that ϕ is injective (using that T is simply connected).

Suppose that g is not in the image of ϕ. Then $g \notin \mathrm{Stab}(e)$. Choose g so that the distance from ge to e in T is minimized. Let e_1 be the first edge on a path from ge to e. Then $e_1 = ge_2$ for some e_2 adjacent to e. We claim that there is an

$h \in \mathrm{Stab}(v) \cup \mathrm{Stab}(w)$ so that $he = e_2$. Given the claim, we obtain a contradiction: since $ghe = e_1$ is closer to e than ge is we must have gh in the image of ϕ, implying that g is in the image of ϕ (since h obviously is).

To prove the claim, note that since the action of G on the set of edges is transitive, there must be *some* $h \in G$ so that $he = e_2$. But this h preserves the bipartite structure on T and so must fix $e \cap e_2$, which is either v or w.

Now we prove injectivity. Suppose that $g \in A *_C B$. Then g has a normal form $(c; x_1, \ldots, x_n)$ as in Theorem 2.3. We build an immersed path in T of length equal to $n+1$ whose first edge is e and whose last edge is ge. This will show (since T contains no circuits) that $ge = e$ if and only if $g \in C$. So if $g \in \ker \phi$, we have $g \in C$. The map ϕ is injective on C by assumption, so we must have $g = 1$.

It remains to build the path. If the normal form is $(c;)$, the path consists of only the edge $e_0 = e$. So suppose inductively we have found an immersed path $(e_0, \ldots, e_{n-1})$ starting with e and ending at $e_{n-1} = cx_1 \cdots x_{n-1}e$. Since $x_n \in (A \cup B) - C$, the edge $x_n e$ is adjacent to but not equal to e. This implies that $e_n = ge = (cx_1 \cdots x_{n-1})x_n e$ is adjacent to but not equal to e. The only worry is that $e_n = e_{n-2}$. But this cannot happen: if $x_n \in A$, then $e_n \cap e_{n-1}$ is in the orbit Gv, whereas $e_{n-1} \cap e_{n-2}$ is in the orbit Gw, and if $x_n \in B$ the reverse is true.

(1) $\Longrightarrow$ (2): There are multiple approaches to this implication. In class Monday I sketched two. The first was essentially Serre's approach (see [Ser03, I.4]). The second was more topological, constructing the tree as a quotient of the universal cover of the graph of spaces described in Definition 2.2. A third approach is given as a special case of Theorem **??** below. ∎

5. Graphs of groups and small categories

Let Γ be a 1–complex. We denote by $\underline{\Gamma}$ the opposite poset of cells of Γ. This is a small category so that each edge object is the source of two outgoing arrows pointed at vertex objects and one identity arrow. Each vertex object is the source only of an identity arrow. This is quite a boring category, as any composition ab is equal either to a or b!

We denote by $\underline{\mathrm{Grp}}^{1-1}$ the category of groups and monomorphisms.

Definition 5.1 (Graph of Groups)**.** Let Γ be a connected 1–complex. A *graph of groups* with underlying graph Γ is a functor $G\colon \underline{\Gamma} \to \underline{\mathrm{Grp}}^{1-1}$.

To each graph of groups we associate a small category:

Definition 5.2 (The category $\mathcal{C}G$)**.** Let $G\colon \underline{\Gamma} \to \underline{\mathrm{Grp}}^{1-1}$ be a graph of groups. The category $\mathcal{C}G$ has objects $\mathrm{Obj}(\underline{\Gamma})$ and morphisms of the form (g, a) where a is an arrow of $\underline{\Gamma}$ and $g \in G(t(a))$. For an arrow a, write ψ_a for $G(a)$. The composition law is

$$(g, a) \circ (h, b) = (g\psi_a(h), ab)$$

defined whenever $t(b) = i(a)$.

In a category, a *composable chain of n arrows* is an n–tuple of arrows $(a_1, \ldots, a_n)$ so that the composition $a_1 \cdots a_n$ is defined. For any small category $\mathcal{C}$ we can construct a simplicial complex called its *realization* $\mathsf{R}(\mathcal{C})$. This complex has vertex set equal to $\mathrm{Obj}(\mathcal{C})$, an edge for every arrow, and an n–simplex for every composable chain of n arrows. More explicitly, we have $\partial a = t(a) - i(a)$ for every arrow a, and

$$\begin{aligned}\partial(a_1, \dots, a_n) =& (a_1, \dots, a_{n-1}) + \sum_{i=1}^{n-1} (-1)^{n+i}(a_1, \dots, a_i a_{i+1}, \dots, a_n) \\ &+ (-1)^n (a_2, \dots, a_n)\end{aligned}$$

Definition 5.3 (Fundamental group)**.** If G is a graph of groups, the *fundamental group of* G is the fundamental group of the complex $\mathsf{R}(\mathcal{C}G)$, and written $\pi_1(G)$ (or $\pi_1(G, v)$ if we want to specify a base-point).

Example 5.4. If Γ consists of a single vertex v with no edges, then a graph of groups G on Γ is just a single group $A = G(v)$. The category $\mathcal{C}G$ has a single object, one arrow for each group element, and composition is given by the group law. One then has $\pi_1(\mathsf{R}(\mathcal{C}G)) \cong A$ in a canonical way. In fact (exercise) $\mathsf{R}(\mathcal{C}G)$ is a $K(A, 1)$.

Example 5.5. If Γ is a single edge e joining vertices v and w, then $\underline{\Gamma}$ is the category

$$v \xleftarrow{\;a\;} e \xrightarrow{\;b\;} w$$

(suppressing identity arrows). So a graph of groups is a diagram of monomorphisms

$$A \xleftarrow{\;\alpha\;} C \xrightarrow{\;\beta\;} B$$

where $A = G(v)$, $B = G(w)$ and so on. For each $c \in C$ there is a commutative diagram of arrows:

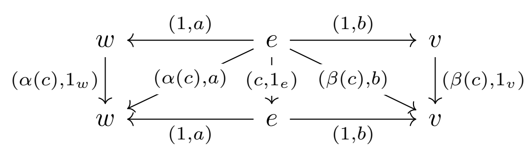

(There are a lot of "1"s in this diagram. The convention used here is that an undecorated 1 is the identity element of some *group*, whereas an *identity arrow* coming from $\underline{\Gamma}$ is given a subscript.) In the realization, this gives an annulus between the loop coming from $\alpha(c) \in A$ to the loop coming from $\beta(c) \in B$. In fact the realization is an example of the complex Y discussed in Definition 2.2, so this shows that in this case the fundamental group of the graph of groups is the amalgam.

6. The graph of groups coming from a tree action

We assume that G acts on the tree T *without inversions.* This means that for each edge e of T with endpoints v, w, we have

$$\mathrm{Stab}(e) = \mathrm{Stab}(v) \cap \mathrm{Stab}(w).$$

In this case the quotient $\Gamma = {}_{G}\backslash^{T}$ is a graph with vertices (resp., edges) in one to one correspondence with G–orbits of vertex (resp., edge) of T. This will be the underlying graph of the graph of groups structure.

We must define a functor from $\underline{\Gamma}$ to $\underline{\mathrm{Grp}}^{1-1}$. To do so we must make some choices. Namely, for each object o of $\underline{\Gamma}$ we choose an orbit representative $\tilde{o}$ in $\underline{T}$. We set $G(o) = \mathrm{Stab}(\tilde{o})$. To define the monomorphisms, we note that each arrow a of $\underline{\Gamma}$ lifts to a unique arrow $\tilde{a}$ with $i(\tilde{a}) = \widetilde{i(a)}$. This arrow may not point at $\widetilde{t(a)}$ but we may choose a group element h_a so that $h_a(t(\tilde{a})) = \widetilde{t(a)}$. The monomorphism ψ_a is given by conjugation: $\psi_a(g) = h_a g h_a^{-1}$. (We could also write this as $\psi_a = \mathrm{Ad}(h_a)|_{\mathrm{Stab}(\widetilde{i(a)})}$.)

We made a number of choices but these don't matter too much.

Definition 6.1 (Isomorphism of graphs of groups)**.** Let (G, ψ_a) and (H, ϕ_a) give two different graphs of groups with the same underlying graph Γ. An *isomorphism* from $G(\Gamma)$ to $H(\Gamma)$ is a collection of isomorphisms $f_o \colon G(o) \to H(o)$ so that for any nontrivial arrow $e \xrightarrow{a} v$, the diagram

$$\begin{array}{ccc} G(e) & \xrightarrow{\psi_a} & G(v) \\ \downarrow{\scriptstyle f_e} & & \downarrow{\scriptstyle f_v} \\ H(e) & \xrightarrow{\phi_a} & H(v) \end{array}$$

commutes up to a conjugacy in $H(v)$.

Note that the conjugating element of $H(v)$ may depend on the arrow.

Lemma 6.2. *Different choices in the definition of the graph of groups at the beginning of this section lead to isomorphic graphs of groups.*

Proof. Exercise. □

Later in these notes we will see how to reverse this process, and turn a graph of groups into a tree action.

7. Covers of categories

7.1. **Coverings.**

Definition 7.1. A *covering* of categories is a functor $f \colon \tilde{\mathcal{C}} \to \mathcal{C}$ so that for each object v of $\tilde{\mathcal{C}}$, the restriction of f to the collection of arrows with source or target v is a bijection. Given such a covering, we may say the *label* of an arrow $a \in \tilde{\mathcal{C}}$ is the arrow $f(a)$.

Remark 7.2. Any arrow in a cover is determined by its source and label.

Any functor induces a simplicial map on realizations. A covering of categories induces a map in which 2–simplices (compositions) are easily seen to lift. Indeed, one has the following.

Lemma 7.3. *Let $f \colon \tilde{\mathcal{C}} \to \mathcal{C}$ be a covering of categories. Then the induced map on realizations is a covering map.*

Proof. Exercise. □

7.2. **Paths and homotopies.** We would like to define the *universal cover* of a category in terms of homotopy classes of paths, just as we do in the topological setting. To do so we need to say what a path in a category is and what a homotopy is. Essentially, we will define $\mathcal{C}$–paths as combinatorial paths in the 1–skeleton of $\mathsf{R}(\mathcal{C})$, and homotopies of $\mathcal{C}$–paths as (combinatorial) homotopies in the 2–skeleton, rel endpoints. Each such homotopy can be written as a sequence of *elementary homotopies*, which are of three types:

(1) Insert or delete a backtrack.

(2) Homotop a concatenation of arrows $b \cdot a$ to ab across a triangle.
(3) Insert an identity arrow (going forwards or backwards).

More formally, we define an *oriented edge* of $\mathcal{C}$ to be a symbol a^+ or a^- where a is an arrow of $\mathcal{C}$. We have the following (at first counterintuitive) conventions for the source $i(e)$ and target $t(e)$ of an oriented edge e.

$$i(a^-) = i(a), \quad t(a^-) = t(a), \quad \text{and} \quad i(a^+) = t(a), \quad t(a^+) = i(a).$$

This choice is made so that the concatenation $a^+ \cdot b^+$ is $(ab)^+$ whenever the composition ab is defined.

A *path in* $\mathcal{C}$ of length n is a concatenation $p = e_1 \cdots e_n$ of oriented edges so that $t(e_i) = i(e_{i+1})$ for all $1 \leq i < n$, and has $i(p) = i(e_1), t(p) = t(e_n)$. A path p of length 0 is a choice of object v and has $i(p) = t(p) = v$.

Definition 7.4. An *elementary homotopy* of paths is one of the following:

(1) $p \cdot q \simeq p \cdot 1_v^{\pm} q$ when $t(p) = v$.
(2) $p \cdot q \simeq p \cdot a^+ \cdot a^- \cdot q$ if $t(p) = t(a)$ and $p \cdot q \simeq p \cdot a^- \cdot a^+ \cdot q$ if $t(p) = i(a)$.
(3) $p \cdot a^+ \cdot b^+ \cdot q \simeq p \cdot (ab)^+ \cdot q$ or $p \cdot b^- \cdot a^- \cdot q \simeq p \cdot (ab)^- \cdot q$ whenever both sides are defined.

7.3. **The universal cover.** We can now define the universal cover of a connected category.

Definition 7.5. Let $\mathcal{C}$ be a connected category (meaning any two objects are connected by a $\mathcal{C}$–path), and fix and object v in $\mathcal{C}$. We will define a category $\tilde{\mathcal{C}}$ and a covering $\pi\colon\ \tilde{\mathcal{C}} \to \mathcal{C}$.

The objects of the universal cover $\tilde{\mathcal{C}}$ are homotopy classes of $\mathcal{C}$–paths starting at v. If $[p]$ and $[p \cdot a^-]$ are both objects of $\tilde{\mathcal{C}}$, where a is an arrow of $\mathcal{C}$, then $\tilde{\mathcal{C}}$ will contain a unique arrow $([p], a)$ from $[p]$ to $[p \cdot a^-]$ so that $\pi([p], a) = a$. We define composition by

$$([p \cdot b^-], a)([p], b) = ([p], ab)$$

whenever $t(p) = i(b)$ and ab is defined in $\mathcal{C}$.

Lemma 7.6. *π is a covering of categories, and the realization of $\tilde{\mathcal{C}}$ is simply connected.*

Proof. Exercise. □

In the topological category, we have an action of the fundamental group of a space on its universal cover. The same happens here.

Lemma 7.7. *If $\mathcal{C}$ is a connected category, and v an object, then $\pi_1(\mathcal{C}, v)$ acts on $\tilde{\mathcal{C}}$ with quotient $\mathcal{C}$. The action is by pre-concatenation of paths; if σ is a $\mathcal{C}$–loop based at v, and p a path starting at v, then $[\sigma] \cdot [p] = [\sigma \cdot p]$.*

Proof. Exercise. □

8. The Bass-Serre Tree as a quotient category

Given any small category $\mathcal{C}$ one can attempt to form a quotient with objects isomorphism classes of objects of $\mathcal{C}$ and with the following equivalence on arrows:

$$a \sim b \quad \text{iff} \quad a = ybx \quad \text{for} \quad y, x \quad \text{invertible}$$

Call this set of objects and arrows the *bleaching* of $\mathcal{C}$, and write it as $\mathsf{Bl}(\mathcal{C})$. To make it a category we should define composition. The obvious thing to do is to define composition by $[a][b] = [ab]$ whenever $[ab]$ is defined. Unfortunately the composition will not be well-defined in general. But sometimes it is. For example if $\mathcal{C}$ is a group, then the quotient will be a category containing a single object and its identity arrow.

Lemma 8.1. *If $\mathcal{C} = \mathcal{C}G(\Gamma)$ for some graph of groups G, then $\mathsf{Bl}(\mathcal{C}G(\Gamma))$ is a category isomorphic to $\underline{\Gamma}$.*

Proof. Exercise. □

We next deal with covers. We'll need the following.

Lemma 8.2. *Let $\pi\colon \tilde{\mathcal{C}} \to \mathcal{C}$ be a covering of categories, and let a be an arrow of $\tilde{\mathcal{C}}$. Then a is invertible if and only if $\pi(a)$ is invertible.*

Proof. Exercise. □

A key fact about a category $\mathcal{C}G(\Gamma)$ for a graph of groups is that *vertex* and *edge* objects can be distinguished from one another by whether they are the source of any non-invertible arrows. Arrows (g, a) are invertible if and only if $a = 1_o$ for some o.

Proposition 8.3. *Let $\pi\colon \tilde{\mathcal{C}} \to \mathcal{C}G(\Gamma)$ be a covering of categories. Then $\mathsf{Bl}(\tilde{\mathcal{C}})$ is equal to $\underline{\Lambda}$ for some graph Λ and π descends to a surjective graph morphism $\Lambda \to \Gamma$.*

Proof. (Sketch) We first note that if $\mathcal{D} \to \mathcal{E}$ is any covering of connected categories, then there is an induced map $\mathsf{Bl}(()\mathcal{D}) \to \mathsf{Bl}(()\mathcal{E})$ which is surjective. (This follows from Lemma 8.2.) So the main thing we have to show is that $\mathsf{Bl}(()\tilde{\mathcal{C}})$ is equal to $\underline{\Lambda}$ for some 1–complex Λ. This amounts to showing two things: First, if $\pi(o)$ is a vertex object of $\mathcal{C}G(\Gamma)$ then the equivalence class of o in $\mathsf{Bl}(()\tilde{\mathcal{C}})$ is the source of no non-trivial arrows; such o will be the vertex objects of $\underline{\Lambda}$. Second, if $\pi(o)$ is an edge object, then the equivalence class of o in $\mathsf{Bl}(()\tilde{\mathcal{C}})$ is the source of exactly two non-trivial arrows, each of which points to a vertex object.

Let $\tilde{v}$ satisfy $\pi(\tilde{v}) = v$ for a vertex object v. Every arrow with $i(a) = \tilde{v}$ is invertible, by Lemma 8.2. So $[\tilde{v}]$ only has an identity arrow in $\mathsf{Bl}(\tilde{\mathcal{C}})$.

If $\tilde{e}$ satisfies $\pi(\tilde{e}) = e$ for an edge object e, then there are non-invertible arrows with label (g, a) where $i(a) = \tilde{e}$ and a is one of the two non-invertible arrows of $\underline{\Gamma}$ with source e. Fixing a with $t(a) = v$ and the equivalence class of $\tilde{e}$, these arrows are all equivalent. Indeed, and arrow labeled (g, a) is equivalent to one labeled $(1, a)$, as the following commutative diagram with labels from $\mathcal{C}G(\Gamma)$ shows:

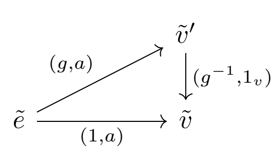

And any two arrows labeled $(1, a)$ with equivalent sources are equivalent by the following diagram (the objects $\tilde{v}'$ may be different for these two diagrams).

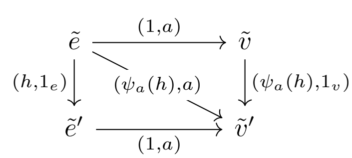

On the other hand, if a and b are the two outgoing arrows from e in $\underline{\Gamma}$, then they are not equivalent to one another, since they are distinguished in the quotient $\mathsf{Bl}(\tilde{\mathcal{C}}) \to \underline{\Gamma}$ (see Lemma 8.1). □

Proposition 8.4 (The Bass-Serre Tree)**.** *If $\pi\colon \tilde{\mathcal{C}} \to \mathcal{C}G(\Gamma)$ is the universal cover of the category associated to a graph of groups, then $\mathsf{Bl}(\tilde{\mathcal{C}})$ is a tree.*

Proof. Proposition 8.3 implies that $\mathsf{Bl}(\tilde{\mathcal{C}}) = \underline{\Lambda}$ for a 1–complex Λ. We must show that Λ has no circuit. Consider a shortest circuit c in Λ

$$a_1^+ \cdot b_1^- \cdots a_n^+ \cdot b_n^-$$

connecting edge objects $e_i = i(a_i) = i(b_i)$ and vertex objects $v_i = t(b_i) = i(a_{i+1})$ (understanding indices mod n). This circuit can be lifted to a circuit

$$\tilde{c} = \tilde{a}_1^+ \cdot \tilde{b}_1^- \cdots \tilde{a}_n^+ \cdot \tilde{b}_n^-$$

in $\tilde{\mathcal{C}}$. (Note that $\mathsf{Bl}\colon \tilde{\mathcal{C}} \to \underline{\Lambda}$ is *not* a covering so by "lifting" we just mean finding a section of Bl over the cycle.) By Lemma 7.6, $\tilde{\mathcal{C}}$ is simply connected. Using simplicial approximation, we can represent a null-homotopy by a cellulated disk D (like a *van Kampen diagram*), where each cell gives some elementary homotopy. In terms of the description in Definition 7.4, a 2–cell representing an elementary homotopy of type (n) is an n–gon. (In fact 2–gons are unnecessary – can you see why?) The faces of the cell are oriented and labeled with arrows of $\mathcal{C}G(\Gamma)$. Notice that if any side of a 2–cell is labeled by an arrow (g, a) for a non-invertible, there is exactly one other side of that 2–cell labeled (g', a) for some $g' \in G(t(a))$; if there is a third side it is labeled by an invertible arrow. In particular, the sides with a label (g, a) for some g fit together into "a–corridors". Supposing that (g_0, a) is the label of $\tilde{a}_1$, we see that there must be a subdiagram of D which is such a corridor, and which can only end at another arrow on the boundary of D. See Figure 8.

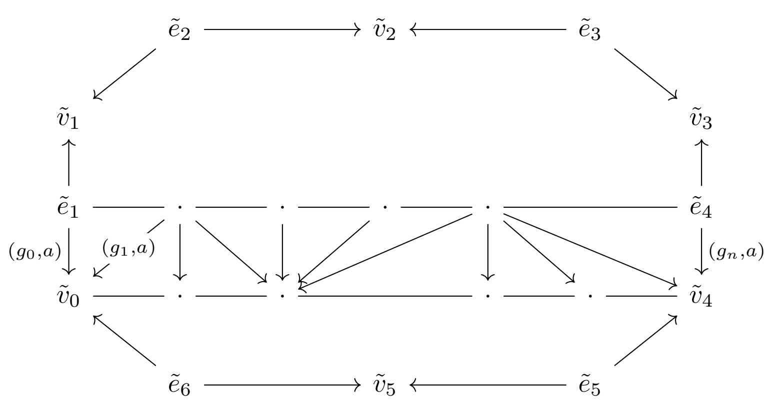

FIGURE 1. An example of an "a–corridor"; the arrows on the boundary of the corridor are invertible arrows so could point either left or right. The rest of the disk is filled in with triangles (and 1–gons and 2–gons) which don't interact with the corridor.

But this corridor furnishes a proof that the beginning and ending arrows are equivalent in $\mathsf{Bl}(\mathcal{C})$, contradicting the assumption that c is embedded.

□

9. The fundamental theorem of Bass-Serre Theory

Our goal in this section is to prove what could be called the Fundamental Theorem of Bass-Serre Theory (FTBS for short). Informally, this says that the processes of turning a tree action into a graph of groups and vice versa are inverses of one another. We'll state this as two separate theorems: Theorem 9.1 and 9.5.

9.1. From a graph of groups to a tree and back again.

This subsection is devoted to proving the following.

Theorem 9.1 (FTBS, Part 1)**.** *Let $H(\Gamma)$ be a graph of groups, and let v be a vertex of Γ. Then there is an action of $\pi_1(H(\Gamma), v)$ on a tree T so that the graph of groups associated to the action is isomorphic to $H(\Gamma)$.*

If a is an arrow of $\underline{\Gamma}$, we will use ϕ_a to denote the monomorphism $H(a)$.

The tree T has been defined in the last section; it is the 1–complex whose associated category $\underline{T}$ is the bleaching of the universal cover of $\mathcal{C}H(\Gamma)$. The action of $\pi_1(H, v)$ on the universal cover $\tilde{C}$ of $\mathcal{C}H(\Gamma)$ preserves the equivalences defining the bleaching, so the action descends to $\underline{T}$, and hence to T.

Lemma 9.2. *The quotient $\pi_1(H(\Gamma))\backslash T$ can be canonically identified with Γ.*

Proof. Exercise. □

We proceed to define a graph of groups $G(\Gamma)$ from the action of $\pi_1(H(\Gamma), v)$ on T. We will use ψ_a to denote the monomorphism $G(a)$.

Following the recipe in Section 6, we should choose arbitrary lifts $\tilde{o}$ in T of all the objects o of $\underline{\Gamma}$. But different choices will give isomorphic graphs of groups, and if we choose them in a slightly more systematic way things will work out more nicely. Namely, for each object o in Γ, we choose a path $p(o)$ in $\underline{\Gamma}$ from v to o. By abuse of notation this can also be regarded as a path in $\mathcal{C}H(\Gamma)$, replacing any arrow $a^{\pm}$ with the corresponding arrow $(1, a)^{\pm}$. The path $p(o)$ thus gives a point $[p(o)] \in \tilde{C}$, and we set $\tilde{o}$ equal to the equivalence class of $[p(o)]$ in $\underline{T} = \mathsf{Bl}(\tilde{C})$.

We define $G(o) = \mathrm{Stab}(\tilde{o})$.

Definition 9.3. If $p = a_1^{\epsilon_1} \cdots a_n^{\epsilon_n}$ is a path in a category, then $\overline{p} = a_n^{-\epsilon_n} \cdots a_1^{-\epsilon_1}$. Note that $p \cdot \overline{p}$ is always homotopic to a constant path at $i(p)$.

Lemma 9.4. $G(o) = \{[p(o) \cdot (g, 1_o)^+ \cdot \overline{p(o)}]\}$

Proof. Exercise. □

Next we choose the elements h_a carefully. If a is an arrow from e to w then we claim we can choose

$$h_a = [p(w) \cdot (1, a)^+ \cdot \overline{p(e)}].$$

Indeed, the arrow $\tilde{a}$ in $\underline{\mathrm{T}}$ is the equivalence class of the arrow $([p(e)], (1, a))$; the target of $([p(e)], (1, a))$ is $[p(e) \cdot (1, a)^-]$. (See Definition 7.5 for this notation.) We see $h_a[p(e)] = [p(w) \cdot (1, a)^+]$ so the target of $h_a([p(e)], (1, a))$ is $[p(w)]$. Passing to $\underline{T}$, we have verified that the target of $h_a\tilde{a}$ is $\tilde{w}$.

The inclusion map $\psi_a \colon G(e) \to G(w)$ given by h_a can thus be written:

$$\begin{aligned} \psi_a([p(e) \cdot (g, 1_e)^+ \cdot \overline{p(e)}]) &= [p(w) \cdot (1, a)^+ \cdot (g, 1_e)^+ \cdot (1, a)^- \cdot \overline{p(w)}] \\ &= [p(w) \cdot (\phi_a(g), 1_w)^+ \cdot \overline{p(w)}] \end{aligned} \tag{3}$$

This completes the description of $G(\Gamma)$. We now need to show that $G(\Gamma)$ is isomorphic to $H(\Gamma)$. To do so, we need to define maps $f_o\colon H(o) \to G(o)$ and conjugating elements $g_a \in G(t(a))$ for each object o and each arrow a of $\underline{\Gamma}$. We will take all the g_a to be *identity* elements, and define, for $g \in H(o)$,

$$f_o(g) = [p(o) \cdot (g, 1_o)^+ \cdot \overline{p(o)}].$$

By Lemma 9.4, each f_o is an isomorphism. It remains to check that squares of the form

$$\begin{array}{ccc} H(e) & \xrightarrow{\phi_a} & H(w) \\ \downarrow f_e & & \downarrow f_w \\ G(e) & \xrightarrow{\psi_a} & G(w) \end{array}$$

commute. For each $g \in H(e)$ we have, using equation (3)

$$\begin{aligned} f_w\phi_a(g) &= [p(w) \cdot (\phi_a(g), 1_w)^+ \cdot \overline{p(w)}] \\ &= \psi_a([p(e)) \cdot (g, 1_e)^+ \cdot \overline{p(e)}]) \\ &= \psi_a f_e(g). \end{aligned}$$

This concludes the proof of Theorem 9.1.

9.2. **From a tree action to a graph of groups and back again.** We now prove the other direction, namely:

Theorem 9.5 (FTBS, Part 2)**.** *Suppose that G acts on a tree T without inversions, and with quotient Γ. Let $H(\Gamma)$ be the associated graph of groups. Then for any vertex v in Γ, there is an isomorphism $f\colon \pi_1(H(\Gamma), v) \to G$ and an f–equivariant isomorphism from the Bass-Serre tree of $H(\Gamma)$ to T.*

References

[Hat02] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[Ser03] Jean-Pierre Serre. *Trees.* Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.