Algebraic & Geometric Topology 9 (2009) 327–335

Graphs of subgroups of free groups

LARSEN LOUDER
D B MCREYNOLDS

We construct an efficient model for graphs of finitely generated subgroups of free groups. Using this we give a very short proof of Dicks's reformulation of the strengthened Hanna Neumann Conjecture as the Amalgamated Graph Conjecture. In addition, we answer a question of Culler and Shalen on ranks of intersections in free groups. The latter has also been done independently by R P Kent IV.

20E05

1 Introduction

One purpose of this article is to investigate the interplay between the join and intersection of a pair of finitely generated subgroups of a free group. Our main result, Theorem 2.4, is a minor generalization of the construction of the first author from [4] and produces a simple model for analyzing intersections and joins. We use this technique to give a quick proof of a theorem of Dicks [2]. Another application of Theorem 2.4 is an answer to an unpublished question of Culler and Shalen [1]. This has been done independently by Kent [3]. Explicitly, the result is the following theorem.

Theorem 1.1 *Let $G = H_1 *_M H_2$ be a graph of free groups such that each H_i has rank 2. If $G \twoheadrightarrow \mathbb{F}_3$ then M is cyclic or trivial.*

One can derive upper bounds on the rank of the intersection given lower bounds on the rank of the join. This has also been observed in the nice article of Kent [3], where some upper bounds are explicitly computed. The proof of Theorem 1.1 presented here differs only slightly from his. In the broadest terms, the two articles share with most papers in the subject an analysis of immersions of graphs, a method that dates back to Stallings [5]. Specifically, Kent uses directly the topological pushout of a pair of graphs along the core of their pullback, a graph which appears here as the underlying graph of a reduced graph of graphs.

Published: 23 February 2009 DOI: 10.2140/agt.2009.9.327

Acknowledgements The authors are very grateful to Richard Kent for many discussions on this topic, in particular those regarding Theorem 1.1. The first author thanks the California Institute of Technology for its hospitality during a visit when this work began. The second author would like to thank Ben Klaff for bringing the question of Culler and Shalen to his attention. Finally, many thanks to the referee for several useful comments and suggestions, especially a much simplified proof of Lemma 2.2.

Both authors were supported in part by NSF postdoctoral fellowships.

2 Graphs of graphs

A *graph of graphs* is a finite graph of spaces such that all vertex spaces are combinatorial graphs and all edge maps are embeddings. Below are some simple operations on graphs of graphs. All vertices and edges are indicated by lower case letters, and their associated spaces will be denoted by the corresponding letter in upper case. We will not keep track of orientation here despite its occasional importance—we trust the reader to sort out this simple matter when it arises. Let X be a graph of graphs with vertices v_i and edges e_j.

(M1) *Making vertex and edge spaces connected*: Let $V_{i,1},\dots,V_{i,n_i}$ be the connected components of the vertex space V_i associated to the vertex v_i of the underlying graph G and $E_{j,1},\dots,E_{j,m_i}$ be the connected components of the edge space E_j. We construct a new graph of graphs as follows. First, we build the underlying graph. For each i and j, we take a collection of vertices $v_{i,k}$ and edges $e_{j,l}$, one for each connected component of each vertex space and edge space, respectively. We label $v_{i,k}$ with $V_{i,k}$ and $e_{j,l}$ with $E_{j,l}$, and attach $e_{j,l}$ to $v_{i,k}$ if the image of $E_{j,l}$ in V_i is contained in $V_{i,k}$. The attaching maps for this graph of graphs are the obvious ones. If $e_{j,l}$ is adjacent to $v_{i,k}$, then we attach an end of $E_{j,l}\times I$ to $V_{i,k}$ by the inclusion map.

(M2) *Removing unnecessary vertices*: If V is a vertex space with exactly two incident edges such that both inclusions are isomorphisms, we remove v and regard the pair of incident edges as a single edge. If V has one incident edge and the inclusion is an isomorphism, we remove V and the incident edge.

(M3) *Removing isolated edges*: If a vertex space V has an edge e that is not the image of an edge from an incident edge space, we remove e from V.

(M4) *Collapsing free edges or vertices*: We call an edge e of a vertex space V *free* if it is the image of only one edge from the collection of incident edge spaces, say $e'\subset E$. In this case, we remove e and e' from V and E. If a vertex space V is a point and has only one incident edge space, we remove v and the incident edge.

A graph of graphs is *reduced* if any application of these operations leaves the space unchanged. Notice that any graph of graphs can be converted to a reduced graph of graphs by greedily applying (M1) through (M4). The remaining requisite operations on graphs of graphs are blow ups and blow downs at a vertex.

(M5) For a vertex space V, divide the incident edge spaces into two classes $E_1, \dots, E_n$ and $E_{n+1}, \dots, E_m$, and let V_1 (V_2, resp.) be the union of the images of E_i, $i \leq n$ ($i > n$, resp.). When $V_1 \cap V_2$ is nontrivial, we replace V by $V_1 \sqcup V_2$ and introduce a new edge $v_1 \cap v_2$ with the edge graph $V_1 \cap V_2$. Next, we attach E_i to V_1 for $i \leq n$, E_i to V_2 for $i > n$ and the newly introduced edge space $V_1 \cap V_2$ to V_1 and V_2 via the inclusion maps.

(M6) *Blow up*: We blow up a vertex by applying (M5). We pass to connected components of the newly created vertex and edge spaces via (M1). Finally, we pass to the associated reduced graph of graphs using (M2).

(M7) *Blow down*: Let E be an edge space of a graph of graphs. If the two embeddings of E have disjoint images, that is, $\iota(E) \cap \tau(E) = \varnothing$, then we remove the edge e of the underlying graph and identify the two endpoints of e. Finally, the graph carried by the new vertex is the one obtained by identifying the vertex space(s) at the ends of e by setting $\iota(f) = \tau(f)$, where f is either a vertex or an edge of E.

Remark Notice that if X has no free or isolated edges, then the space obtained by blowing up a vertex with an application of (M6) also has no free or isolated edges. Also, when V is connected, it follows that $V_1 \cap V_2$ is nontrivial and thus (M5) is applicable.

The *horizontal subgraph* of a graph of graphs is the graph obtained by restricting vertex and edge spaces to vertices. The *mid-graph* of a graph of graphs is the graph obtained by restricting vertex and edge spaces to midpoints of edges. These two subgraphs are denoted Γ_H and Γ_M, respectively. Note that neither of these graphs is necessarily connected. If X is reduced, then Γ_M and Γ_H do not have any valence one vertices. Conversely, if either one of them has a valence one vertex, then there must be a free edge or vertex in X. If there are isolated edges, then a component of Γ_M is a point. In particular, if X is reduced, then every component of Γ_M has nontrivial fundamental group.

Lemma 2.1 *Blowing up and blowing down are homotopy equivalences.*

This follows easily upon observing that if two of the edges introduced during a blowup are both adjacent to a vertex introduced during the blowup, then the images of the edge spaces they carry are disjoint. That the remaining moves, other than (M3), preserve the homotopy types of X, Γ_H and Γ_M, is clear. Note that (M3) only serves to remove trivial components of Γ_M.

It is important to know when to blow up X. The following lemma achieves this.

Lemma 2.2 *Let Δ be a connected graph with a collection $\mathcal{C} = \{\Delta_i\}_{i=1}^m$ of (not necessarily distinct) connected subgraphs. If every edge of Δ is contained in at least two Δ_i and $m > 3$, then after relabeling the Δ_j, there is a partition $\mathcal{C}_1 = \{\Delta_1, \dots, \Delta_n\}$ and $\mathcal{C}_2 = \{\Delta_{n+1}, \dots, \Delta_m\}$ of $\mathcal{C}$ such that at least two Δ_i in $\mathcal{C}_1$ intersect nontrivially and at least two Δ_i in $\mathcal{C}_2$ intersect nontrivially.*

Proof It suffices to find distinct $A, B, C, D \in \mathcal{C}$ such that $A \cap B \neq \varnothing$ and $C \cap D \neq \varnothing$. If all triple intersections are empty then $\mathcal{C}$ has at most two elements by connectivity of Δ. Let $A, B, C \in \mathcal{C}$ such that $A \cap B \cap C \neq \varnothing$. Since Δ is connected, there is some D meeting, again without loss, C. □

Remark Notice that if V is a vertex space of a reduced graph of graphs X with at least four incident edge spaces, then we can use Lemma 2.2 to ensure that (M5) is applicable.

Let X be a reduced graph of graphs such that all vertex and edge spaces are connected. The space X has an underlying graph that we shall denote by $\Gamma_U(X)$. Let $m(X)$ be the highest valence of vertex of $\Gamma_U(X)$, $n(x)$ the number of vertices with valence $m(X)$ and $\chi(\Gamma_U(X))$ the Euler characteristic of $\Gamma_U(X)$. The *complexity* of X is the lexicographically ordered 3–tuple

$$c(X) := (\chi(\Gamma_U(X)), m(X), n(X)).$$

We call a blowup of a vertex v using two sets of edge spaces satisfying Lemma 2.2 *nontrivial*. Our next lemma justifies this terminology.

Lemma 2.3 *Let X be reduced and $m(X) > 3$. If X' is obtained from X via a nontrivial application of* (M6) *to a vertex v with valence $m(X)$, then $c(X') < c(X)$.*

Proof Let $\{v_i\}$ be the vertices of X' introduced during a blow up of X at the vertex v. These vertices must have valence at least two, as otherwise X has a free edge and is not reduced. We assume contrary to the claim that $c(X) = c(X')$. If the Euler characteristics of the underlying graphs of X and X' are equal, then the subgraph B spanned by the edges associated to the connected components of $V_1 \cap V_2$ must be a tree. First observe that it is connected as otherwise V could not have been connected.

Second, if B is not a tree, then the Euler characteristic of the underlying graph must decrease. As B is a tree we have

$$1 - \frac{1}{2}\text{valence}(v) = \sum_i \left(1 - \frac{1}{2}\text{valence}(v_i)\right). \tag{1}$$

If both $m(X') = m(X)$ and $n(X') = n(X)$, then all but one of the vertices v_{i_0} has valence two since there are no valence one vertices making a positive contribution to the sum on the right hand side of (1). Therefore, every component of V_1 (the alternative is handled identically) is the image of exactly one incident edge space from one element of the partition of edges incident to v. However, this is impossible since the blowup X' was assumed to be nontrivial. □

We are now ready to state our main result.

Theorem 2.4 *Every graph of graphs X such that no connected component of Γ_M is a tree can be converted to a reduced graph of graphs X' all of whose vertex groups have valence three. There is a homotopy equivalence $(X', \Gamma'_H, \Gamma'_M) \to (X, \Gamma_H, \Gamma_M)$.*

The *corank* of a group G is the maximal rank of a free group that it maps onto and will be denoted by $\mathrm{cr}(G)$. Before proving Theorem 2.4, a few remarks are in order. First, observe that if X is reduced, then the natural map $\pi_1(X) \to \pi_1(\Gamma_U(X))$ is surjective. Second, the complexity of all graphs of graphs homotopy equivalent to X is bounded below by $(1 - \mathrm{cr}(\pi_1(X)), 3, 0)$. That said, we now give a proof of Theorem 2.4.

Proof of Theorem 2.4 First we apply (M4) until there are no free edges. This does not change the homotopy type of the triple (X, Γ_H, Γ_M). There are no isolated edges since each component of Γ_M is assumed to have nontrivial fundamental group. Next, we pass to connected components of edge and vertex spaces and then pass to the associated reduced graph of graphs by removing valence two vertex spaces. Let X be a reduced graph of graphs, and consider a sequence $\{X_i\}$ starting with X such that X_i is obtained from X_{i-1} by nontrivially blowing up a maximal valence vertex. Since all the X_i are homotopy equivalent and the maps $\pi_1(X_i) \to \pi_1(\Gamma_U(X_i))$, $i > 0$, are surjective, $c(X_i) \geq (1 - \mathrm{cr}(\pi_1(X)), 3, 0)$. According to Lemma 2.3, $c(X_i) > c(X_{i+1})$. Since the complexity is bounded below, for some n, X_n has only valence three vertices. □

A graph of graphs represents a graph of free groups when the ϵ–neighborhood of Γ_M is a product $I \times \Gamma_M$. In this case there are two natural immersions $\Gamma_M \to \Gamma_H$ in the sense of Stallings [5]. Moreover, there is an immersion $\Gamma_H \to \Gamma_U$. We say such a graph of graphs is *representing*. Conversely, suppose that $G = \Delta(H_1, \ldots, H_k, M_1, \cdots, M_l)$

is a graph of free groups with vertex groups H_i, edge groups M_j and that there is a map $\gamma\colon G \to \mathbb{F}$ which embeds each H_i. Let ι_j and τ_j be the two inclusion maps $M_j \hookrightarrow H_{\iota(j)}$ and $M_j \hookrightarrow H_{\tau(j)}$. Represent $\mathbb{F}$ as the fundamental group of a marked labeled graph R with one vertex, and find immersions of marked labeled graphs $\eta_i\colon \Gamma_{H_i} \to R$ representing $\gamma|_{H_i}$ and $\mu_j\colon \Gamma_{M_j} \to R$ representing $\gamma|_{M_j}$. We choose the notation Γ_{H_i} in anticipation of the fact that they are the connected components of the horizontal subgraph of the graph of graphs under construction.

The immersion μ_j factors through $\eta_{\iota(j)}$ and $\eta_{\tau(j)}$ via immersions $\iota_j\colon \Gamma_{M_j} \to \Gamma_{H_{\iota(j)}}$ and $\tau_j\colon \Gamma_{M_j} \to \Gamma_{H_{\tau(j)}}$. We construct a space X by taking the $\Gamma_{M_j} \times I$ as edge spaces, taking the Γ_{H_i} as vertex spaces and using as attaching maps $\iota_j\colon \Gamma_{M_j} \times \{0\} \to \Gamma_{H_{\iota(j)}}$ and $\tau_j\colon \Gamma_{M_j} \times \{1\} \to \Gamma_{H_{\tau(j)}}$. Let $\alpha_j\colon \Gamma_{M_j} \times I \to \Gamma_{M_j}$ be the projection to the first factor. Since $\eta_{\iota(j)} \circ \iota_j = \mu_j$ and $\eta_{\tau(j)} \circ \iota_j = \mu_j$ there is a well defined map $\pi\colon X \to R$ which restricts to η_i and agrees with $\mu_j \circ \alpha_j$.

We now endow X with the structure of a graph of graphs. Let b be the base point of R. Let $V = \pi^{-1}(b)$ and $E_l = \pi^{-1}(m_l)$, where m_l is the midpoint of an edge e_l of R. Each edge e_l of R induces two maps of E_l to V, each of which is an embedding. That these are embeddings can be seen as follows. If one fails to be injective on vertices of E_l, then some $\Gamma_{H_i} \to R$ is not an immersion. If it is injective on vertices but not on edges, then some $\Gamma_{M_j} \to R$ is not immersed. Thus, we may use this data to endow X with the structure of a graph of graphs. By Theorem 2.4, we can repeatedly blow up X until we produce a graph of graphs X' all of whose vertices have valence three. If (M3) is ever applied in the process, then it must be that some M_i was trivial.

Remark Let w be a vertex of a vertex space V of a graph of graphs X. It follows that w is a vertex of Γ_H and the valence of w in Γ_H is exactly the number of edge graphs incident to V whose images contain w. If X is reduced and V has valence three, then there must be a vertex of V which is contained in the image of all three incident edge graphs. Moreover, if Γ_M has a valence three vertex w, then the images of w in Γ_H must each have valence three in Γ_H.

Proof of Theorem 1.1 We begin by representing $G \twoheadrightarrow \mathbb{F}_3$ by a map from a graph of graphs X to a bouquet of three circles R. Note that since $G \twoheadrightarrow \mathbb{F}_3$, at least one of the maps $H_i \to \mathbb{F}_3$ is injective. If the other fails to be injective, the result is immediate. Consequently, we are reduced to the case when both are injective and thus the construction above can be implemented. By blowing X up, we may assume that X has only valence three vertices. The map $G \twoheadrightarrow \mathbb{F}_3$ factors through the map $G \cong \pi_1(X) \twoheadrightarrow \pi_1(\Gamma_U)$, and the rank of Γ_U must be either 3 or 4. If the latter holds, then M is trivial and the theorem holds. If Γ_U has rank 3, since all vertices of Γ_U are valence three, there must be exactly four. By the remark above, if Γ_M has a valence three vertex, then the map from the set of valence three vertices of Γ_H to the set of valence three vertices of Γ_U cannot be injective. Since the two components of Γ_H each have fundamental group $\mathbb{F}_2$, they have 2 valence three vertices apiece. However, this implies the contradiction that Γ_U has at most 3 valence three vertices. Thus, Γ_M has vertices of valence at most two and so has rank at most one, as claimed. □

Remark By [4], M is contained in the subgroup generated by a basis element in at least one of H_1 or H_2.

Other inequalities of this type are easily obtained through an analysis of a reduced valence three graph of graphs representing the intersection. In particular, special cases of the Hanna Neumann conjecture can be verified with this analysis. For explicit inequalities, we refer the reader to Kent [3] who has also derived them.

3 Intersections of subgroups of free groups

Let H_1 and H_2 be subgroups of a fixed free group $\mathbb{F}$. If $G = \Delta(H_1, H_2; M_j)$, a graph of free groups with two vertex groups $\{H_i\}$, edge groups $\{M_j\}$, with no monogons and a map $\pi\colon G \twoheadrightarrow \mathbb{F}$ embedding each of the factors H_i, then the vertex spaces of a graph of graphs X representing Δ are bipartite.

A graph of graphs is *simple-edged* if no vertex space has a bigon. To relate reduced graphs of graphs to intersections of free groups we need to understand what happens when a graph of graphs X as above is not simple-edged. Let p and q be the midpoints of a pair of offending edges, $\Gamma = \Gamma_U(X)$ the underlying graph of X, and give the edges of Γ distinct oriented labels. The labeling of Γ induces labelings of Γ_M and Γ_{H_i}. Let Γ'_M be the labeled graph obtained by identifying p and q. By folding the labeled graph Γ'_M (see for instance Stallings [5]), we obtain a labeled graph Γ_K with fundamental group $K = \pi_1(\Gamma_K, p)$. Folding endows Γ_K with a pair of immersions $\nu_i\colon \Gamma_K \to \Gamma_{H_i}$. In addition, there is an immersion $\eta_j\colon \Gamma_{M_j} \to \Gamma_K$ and the edge map $\Gamma_{M_j} \to \Gamma_{H_i}$ is just $\nu_i \circ \eta_j$.

We must consider two cases with regard to p and q after folding. Namely, the midpoints p and q are either in the same component of Γ_M or in distinct components of Γ_M. We address the latter first and assume, without loss of generality, that Γ_{M_1} and Γ_{M_2} are the components of Γ_M containing p and q. Compute the fundamental groups of Γ_{H_1} and Γ_{H_2} with respect to the images of p (which coincide with the images of q). From this we see that $\pi(H_1) \cap \pi(H_2)$ contains $\pi(M_1)$ and $\pi(M_2)$. If $\pi(M_1) \not< \pi(M_2)$ and $\pi(M_2) \not< \pi(M_1)$, then each inclusion $\pi(M_i) \hookrightarrow \pi(H_1) \cap \pi(H_2)$ is proper and the image of the fundamental group of Γ_K, computed with respect to the image of p, is precisely $\langle M_1, M_2\rangle$. If neither η_1 nor η_2 is an isomorphism of labeled graphs, then K properly contains M_1 and M_2. In the event we are in the first case, without loss of generality, we shall assume that p, q are contained in Γ_{M_1}. We identify the vertices p and q of Γ_M and then fold to obtain a labeled graph Γ_K. As before, the immersion $\Gamma_{M_1} \to \Gamma_{H_i}$ factors through the induced immersion $\Gamma_K \to \Gamma_{H_i}$. In this case $\Gamma_M \to \Gamma_K$ cannot be an isomorphism of graphs and $H_1 \cap H_2$ properly contains M_1.

Let X be a reduced simple-edged graph of graphs with underlying graph $\Gamma = \Gamma_U(X)$. Let $\mathcal{X}(\Gamma)$ be the collection of reduced simple-edged graphs of graphs with underlying graph Γ. If $X, X' \in \mathcal{X}(\Gamma)$, then $X \leq X'$ if there is a map of graphs of spaces $X \to X'$ such that all restrictions to vertex and edge spaces are embeddings. We can restrict to the subcollection $\mathcal{X}_X(\Gamma)$ such that for each $X' \in \mathcal{X}_X(\Gamma)$ there is a map $X \to X'$ and the map $\Gamma_H(X) \to \Gamma_H(X')$ is a graph isomorphism. Clearly $\mathcal{X}_X(\Gamma)$ has a maximal element Y. To link reduced simple-edged graphs of graphs to the strengthened Hanna Neumann conjecture, we only need to observe that since X is simple-edged, each component of $\Gamma_M(X)$ is an embedded subgraph of $\Gamma_M(Y)$ (ie the fundamental groups of components of $\Gamma_M(X)$ are free factors of the respective components of $\Gamma_M(Y)$).

The strengthened Hanna Neumann conjecture then implies that if G is as above and the associated graph of graphs is simple-edged, then

$$\chi(\Gamma_{H_1})\chi(\Gamma_{H_2}) + \chi(\Gamma_M) \geq 0.$$

The equivalence of the amalgamated graph conjecture and the strengthened Hanna Neumann conjecture of [2] follows immediately from the observation that the vertex and edge spaces of a representing simple-edged graph of graphs can be written as in the statement of Dicks' theorem. We leave the details of the construction of this correspondence to the reader, though we state a version of the equivalence for completeness.

Let X be a simple-edged reduced graph of graphs all of whose vertices are valence three that represents a homomorphism $\Delta(H_1, H_2, M_j) \to \mathbb{F}$. Let v_i be the vertices of $\Gamma_U(X)$, and for each i, let Δ_i be the intersection of the images of the three edge spaces incident to v_i. Finally, let Δ be the disjoint union of the Δ_i, $\Sigma_1 = \Delta \cap \Gamma_{H_1}$ and $\Sigma_2 = \Delta \cap \Gamma_{H_2}$ and μ be the number of edges in Δ.

Theorem 3.1

$$\chi(H_1)\chi(H_2) + \sum_i \chi(M_i) = \frac{1}{4}|\Sigma_1||\Sigma_2| - \frac{1}{2}\mu$$

The proof is straightforward. The number of valence three vertices of Γ_{H_i} is $|\Sigma_i|$, Γ_{H_i} has only valence two or valence three vertices, and the Euler characteristic of H_i is therefore $-\frac{1}{2}|\Sigma_i|$. The Euler characteristic of each Γ_{M_j} is computed in the same manner. μ is the number of valence three vertices of Γ_M. In this formulation, the amalgamated graph conjecture simply states that if one is given a reduced simple-edged representing graph of graphs whose horizontal graph has two components, then the right hand side of the above equality is nonnegative.

References

[1] **M Culler**, **P B Shalen**, *Four-free groups and hyperbolic geometry* `arXiv:0806.1188`

[2] **W Dicks**, *Equivalence of the strengthened Hanna Neumann conjecture and the amalgamated graph conjecture*, Invent. Math. 117 (1994) 373–389 MR1283723

[3] **R P Kent IV**, *Intersections and joins of free groups*, Algebr. Geom. Topol. 9 (2009) 305–325

[4] **L Louder**, *Krull dimension for limit groups III: Scott complexity and adjoining roots to finitely generated groups* `arXiv:math/0612222v3`

[5] **J R Stallings**, *Topology of finite graphs*, Invent. Math. 71 (1983) 551–565 MR695906

Department of Mathematics, University of Michigan
Ann Arbor, MI 48109-1043, USA

Department of Mathematics, University of Chicago
Chicago, IL 60637, USA

llouder@umich.edu, dmcreyn@math.uchicago.edu, dmcreyn@gmail.com

http://www.math.lsa.umich.edu/~llouder/,
http://www.math.uchicago.edu/~dmcreyn/

Received: 27 August 2008 Revised: 25 January 2009