Michigan Math. J. 53 (2005)

Magnus Intersections in One-Relator Products

James Howie

1. Introduction

Recall that a group G is *locally indicable* if each of its nontrivial, finitely generated subgroups admits an infinite cyclic homomorphic image. A *one-relator product* of groups A_λ $(\lambda \in \Lambda)$ is a quotient

$$G = (*_{\lambda\in\Lambda} A_\lambda)/\langle\langle R\rangle\rangle$$

of their free product by the normal closure of a word R, which is called the *relator,* and is assumed not to be conjugate to an element of one of the A_λ. Let $A_\Lambda := *_{\lambda\in\Lambda} A_\lambda$. Then, by the Freiheitssatz for locally indicable groups [3] (see Theorem 2.1 to follow), a free factor $A_M := *_{\mu\in M} A_\mu$ of A_Λ embeds in G provided that R is not conjugate in A_Λ to an element of A_M. The image of this embedding is called the *Magnus subgroup* corresponding to the subset $M \subset \Lambda$.

The purpose of this paper is to examine the intersection of two Magnus subgroups of a one-relator product of locally indicable groups. Suppose M and N are two subsets of Λ. Then clearly the Magnus subgroup $A_{M\cap N}$ is contained in the intersection of Magnus subgroups $A_M \cap A_N$. In almost all cases it turns out that

$$A_M \cap A_N = A_{M\cap N},$$

but it is easy to construct examples where this equation fails.

In the special case where the A_λ are all infinite cyclic (so that G is a one-relator group), Collins [5] has proved the following result.

Theorem A. *Let A_M and A_N be Magnus subgroups of a one-relator group $G = A_\Lambda/\langle\langle R\rangle\rangle$. Then*

$$A_M \cap A_N = A_{M\cap N} * I,$$

where I is a free group of rank 0 *or* 1.

Both possibilities occur, but the most usual situation is that I has rank 0. If I has rank 1 then we say that A_M and A_N have *exceptional intersection.* In this case, nontrivial elements of I are called *exceptional elements.*

The purpose of this paper is twofold: to generalize Theorem A to the case of arbitrary locally indicable factors; and to investigate precisely under what circumstances an exceptional intersection can occur.

Received July 9, 2004. Revision received February 2, 2005.

We prove the following results. The first is a generalization of a result of Newman [11] for the case of a one-relator group with torsion.

THEOREM B. *Let G be a one-relator product of locally indicable groups in which the relator is a proper power. Then Magnus subgroups of G have no exceptional intersection.*

Next is our main result.

THEOREM C. *Let $G = A_\Lambda/\langle\langle R\rangle\rangle$ be a one-relator product of locally indicable groups, and let A_M and A_N be Magnus subgroups that have exceptional intersection in G. Then one of the following is true.*

(1) *Some conjugate R' of R is contained in a subgroup of A_Λ of the form $D * E$, where D and E are finitely generated subgroups of A_M and A_N (respectively) and where D^{ab} and E^{ab} have torsion-free rank 1. In this case*

$$A_M \cap A_N = A_{M\cap N} * I,$$

where I is the intersection of D and E in $(D * E)/\langle\langle R'\rangle\rangle$.

(2) *Some conjugate R' of R is contained in a subgroup of A_Λ of the form $D * \langle xy\rangle$, where D is a finitely generated subgroup of $A_{M\cap N}$ with D^{ab} of torsion-free rank 1 and where $x \in A_M$ and $y \in A_N$. In this case, either*

$$A_M \cap A_N = A_{M\cap N} * \langle x\rangle$$

or

$$A_M \cap A_N = A_{M\cap N} * x^{-1}Ix,$$

where I is the intersection of D and $(xy)D(xy)^{-1}$ in $(D * \langle xy\rangle)/\langle\langle R'\rangle\rangle$.

As a consequence, we can derive the following analogue of Theorem A.

THEOREM D. *Let $G = A_\Lambda/\langle\langle R\rangle\rangle$ be a one-relator product of locally indicable groups, and let A_M and A_N be Magnus subgroups. Then*

$$A_M \cap A_N = A_{M\cap N} * I,$$

where I is a free group of rank 0 or 1.

In the special case where the A_λ are all cyclic, we can use the classification of exceptional intersections given in Theorem C to obtain a solution to the algorithmic problem of recognizing when two Magnus subgroups admit an exceptional intersection and, if so, of finding a generating word for the exceptional factor I.

THEOREM E. *There is an algorithm for determining and computing exceptional intersections of Magnus subgroups in one-relator groups, in the following sense. Suppose we are given a finite presentation $\langle a_1, \dots, a_k, b_1, \dots, b_l, c_1, \dots, c_m \mid R = 1\rangle$ of a one-relator group G, where the relator R contains at least one letter $a_i^{\pm 1}$ and at least one letter $c_i^{\pm 1}$. Then the algorithm decides whether or not the intersection of the Magnus subgroups $M_1 = \langle a_1, \dots, a_k, b_1, \dots, b_l\rangle$ and*

$M_2 = \langle b_1, \ldots, b_l, c_1, \ldots, c_m \rangle$ *in* G *contains exceptional elements. If so, the algorithm finds a word* u *in the generators* $\{a_1, \ldots, a_k, b_1, \ldots, b_l\}$, *and a word* v *in the generators* $\{b_1, \ldots, b_l, c_1, \ldots, c_m\}$, *such that* $u = v$ *in* G *and* $\{b_1, \ldots, b_l, u\}$ *is a free basis for the intersection of* M_1 *and* M_2 *in* G.

There are two obstructions to extending this theorem to the general locally indicable case. One is the standard counting argument: there are uncountably many locally indicable groups but only countably many algorithms. The second is the absence of a suitable analogue—for more general one-relator products of locally indicable groups—of a key ingredient of our algorithm: the Baumslag–Taylor algorithm [1] for identifying the center of a one-relator group.

In practice, we restrict attention to the case where $|\Lambda| = 3$, $\Lambda = M \cup N$, and $|M| = |N| = 2$. The general case can readily be seen to reduce to this special case.

The remainder of the paper is organized as follows. In Section 2 we review a number of known results and standard techniques that we will use in the sequel. In Section 3 we focus attention on *minimal intersection van Kampen diagrams,* which arise whenever an exceptional intersection occurs, and prove a number of structural properties of such diagrams. In Section 4, the technique of towers is applied to our intersection diagrams in order to show that an exceptional intersection can occur only when the relator word R lies in a rank-2 subgroup of the free product $*_\Lambda A_\lambda$ of a certain form. In Section 5 we classify all such rank-2 subgroups and deduce the main results. Finally, in Section 6 we restrict our attention to the classical case of a one-relator group and prove that the problem of recognizing and finding a generator for an exceptional intersection subgroup is algorithmically soluble, following the guideline of the Baumslag–Taylor algorithm [1] for finding the center of a one-relator group.

Acknowledgment. It is a pleasure to acknowledge the help and encouragement of Don Collins toward this work, which grew out of conversations with him about [5]. In particular, the proofs in Section 5 would have been much longer and less elegant without his insights.

2. Preliminaries

2.1. Properties of Locally Indicable Groups

Throughout this paper, we will make extensive use of the following facts about one-relator products of locally indicable groups.

Theorem 2.1 (Freiheitssatz) [3]. *Let* $G = (A * B)/\langle\langle R \rangle\rangle$ *be a one-relator product of locally indicable groups* A, B *such that the relator* R *is not conjugate in* $A * B$ *to an element of* $A \cup B$. *Then the natural maps* $A \to G$ *and* $B \to G$ *are injective.*

Theorem 2.2 [6]. *Let* $G = (A * B)/\langle\langle R \rangle\rangle$ *be a one-relator product of locally indicable groups* A, B *such that the relator* R *is neither conjugate in* $A * B$ *to an element of* $A \cup B$ *nor a proper power in* $A * B$. *Then* G *is locally indicable.*

THEOREM 2.3 [6]. *Let A, B be locally indicable groups, let $R = a_1 b_1 \cdots a_k b_k \in A * B$ with $a_i \in A \setminus \{1\}$ and $b_i \in B \setminus \{1\}$, and let $G = (A * B)/\langle\langle R \rangle\rangle$ be the corresponding one-relator product. Then the $2k$ initial segments $S_{2i-1} = a_1 b_1 \cdots b_{i-1} a_i$ and $S_{2i} = a_1 b_1 \cdots a_i b_i$ $(i = 1, \ldots, k)$ represent pairwise distinct elements of G.*

THEOREM 2.4 [4]. *Locally indicable groups are right orderable.*

Recall that a *right ordering* on a group G is a linear ordering $<$ on G such that $x < y \Rightarrow xg < yg$ for all $x, y, g \in G$. A group G is *right orderable* if there exists a right ordering on G. In general, a right ordering is not unique. (For example, the abelian group $\mathbb{Z}^2$ has uncountably many right orderings.) At one point in our deliberations it will be important that we have some freedom in the choice of right ordering.

A particular type of right ordering arises from a homomorphism $\phi\colon G \to H$, where H and $\mathrm{Ker}(\phi)$ are equipped with right orderings (both denoted $<$). We define a right ordering on G lexicographically: $x < y$ if either (i) $\phi(x) < \phi(y)$ in H or (ii) $\phi(x) = \phi(y)$ and $xy^{-1} < 1$ in $\mathrm{Ker}(\phi)$.

We say that the resulting right ordering on G is *dominated* by the right ordering on H via the homomorphism ϕ. In particular, if G is a right orderable group and if $\phi\colon G \to H$ is any homomorphism to a group with a given right ordering, then any right ordering on G restricts to one on $\mathrm{Ker}(\phi)$ and so there exists a right ordering on G that is dominated by the right ordering on H via ϕ.

2.2. *Van Kampen Diagrams*

Recall that one may associate, to any group presentation $\mathcal{P}$, a 2-dimensional CW-complex $X = X(\mathcal{P})$ (usually with a single 0-cell) whose fundamental group $\pi_1(X)$ is isomorphic to the group $G = G(\mathcal{P})$ defined by the presentation. Recall also [9] that a *van Kampen diagram* over $\mathcal{P}$ consists of a simply connected, compact, planar 2-complex Δ and a combinatorial (or cellular) map $f\colon \Delta \to X(\mathcal{P})$. The complement of Δ in the plane is an open annulus, one of whose ends describes a closed edge-path P_0 in the 1-skeleton of Δ, called the *boundary path*. The image of P_0 in the 1-skeleton $\Delta^{(1)}$ of Δ is called the *boundary* of Δ, denoted $\partial\Delta$. In the 1-skeleton $X^{(1)}$ of X, the image $f(P)$ of any path P in the $\Delta^{(1)}$ can be expressed as a word $\lambda(P)$ in the generators of the presentation $\mathcal{P}$, called the *label* of P. In particular, $\lambda(P_0)$ is called the *boundary label* of the van Kampen diagram. Each 2-cell (or *region*) α of Δ is an open disc whose unique end describes a closed edge-path P_α in $\Delta^{(1)}$. We denote the image of P_α in $\Delta^{(1)}$ by $\partial\alpha$ and call it the *boundary* of α. We also refer to $\lambda(P_\alpha)$ as the *label* of α. Since f is a cellular map, it follows that $f(\alpha)$ is one of the 2-cells of X, so $\lambda(P_\alpha)$ is a conjugate of one of the relators of the presentation $\mathcal{P}$ or of the inverse of a relator.

Note that the paths P_0 and P_α are not entirely well-defined. They depend, in general, on a choice of starting point, which we call the *base-points* of the diagram as a whole and of α, respectively, and on a direction of travel. In this paper we will make the convention that the base-points of α are chosen such that $\lambda(\alpha)$ is equal to a relator or to the inverse of a relator, for each region α, and that the direction of travel is clockwise around Δ and α (with respect to some fixed orientation of the plane). We will also assume that $\lambda(\alpha)$ is always a cyclically reduced word.

Since Δ is simply connected, it follows that the boundary path P_0 of the diagram is null-homotopic in Δ and so $\lambda(P_0)$ represents the identity element of $\pi_1(X) = G$. Conversely, given any word W in the generators of $\mathcal{P}$ that represents the identity element of G, one can find a van Kampen diagram whose boundary label (with respect to some choice of base-point) is equal to W.

Finally, suppose that the boundary paths of two regions α and β in Δ have a point v in common, and suppose that the cyclic conjugates of their labels as read from v are mutually inverse words in the generators of $\mathcal{P}$. Then there is a process of *cancellation* by which α and β can be removed from Δ, and their boundaries sewn together, to produce an amended diagram with the same boundary label but fewer regions.

For details on these matters and for more about van Kampen diagrams in the general context, see [9].

2.3. Intersection Diagrams

Now let us turn to the situation considered in this paper. We have a one-relator product $G = (A * B * C)/\langle\langle R\rangle\rangle$ of locally indicable groups A, B, and C, and we wish to study the intersection in G of the Magnus subgroups $A * B$ and $B * C$. Choose presentations $\mathcal{P}_A, \mathcal{P}_B, \mathcal{P}_C$ for A, B, C (respectively) satisfying the following assumption.

Assumption A. Each letter of the word R is (uniquely) either a generator or the inverse of a generator in one of the presentations $\mathcal{P}_A, \mathcal{P}_B, \mathcal{P}_C$.

It is easy to check that presentations satisfying Assumption A exist. (For uniqueness, this requires the fact that locally indicable groups have no 2-torsion.) Now we can take the disjoint union $\hat{\mathcal{P}}$ as a presentation for the free product $A * B * C$. The resulting complex $X(\hat{\mathcal{P}})$ is just the one-point union of the 2-complexes $X(\mathcal{P}_A), X(\mathcal{P}_B), X(\mathcal{P}_C)$. However, we choose a slightly different topological model $\hat{X}$ for $\hat{\mathcal{P}}$: namely, the disjoint union of $X(\mathcal{P}_A)$, $X(\mathcal{P}_B)$, and $X(\mathcal{P}_C)$ together with two oriented 1-cells, labeled a and c, that we use to join the base-point of $X(\mathcal{P}_B)$ to those of $X(\mathcal{P}_A)$ and $X(\mathcal{P}_C)$, respectively. Note that the 1-cells a and c form a maximal tree T in the 1-skeleton of $\hat{X}$ and that the quotient space $\hat{X}/T$ is exactly $X(\hat{\mathcal{P}})$. (An algebraic analogue of this model is to consider R as a word in $(aAa^{-1}) * B * (cCc^{-1}) \subset A * B * C * \langle a, c\rangle$.)

We then use Assumption A to obtain a canonical choice of a cyclically reduced word in the generators of $\hat{\mathcal{P}}$ representing $R \in A * B * C$—namely, the one in which each letter of R is represented by a generator or the inverse of a generator. We add this canonical word as an additional defining relator to $\hat{\mathcal{P}}$ in order to obtain a presentation $\mathcal{P}$ of G. (We will abuse notation by using R for this additional relator of $\mathcal{P}$ as well as for the original element of $A * B * C$.) Note that this presentation satisfies the following.

PROPERTY B. Suppose that two subwords U, V of R or R^{-1} represent the same element of $A * B * C$. Then U and V are identically equal as words in the generators of $\mathcal{P}$.

To form our 2-complex model for $\mathcal{P}$, we attach a 2-cell to $\hat{X}$ along the edge-path representing R. Notice that this edge-path will involve the edges a and c as well as those labeled by generators of $\mathcal{P}$. We call the resulting 2-complex X.

Suppose that there are elements $u \in (A * B) \setminus B$ and $v \in (B * C) \setminus B$ such that $u = v$ in G. Then we may choose words U, V in the generators of $\mathcal{P}$ that represent u and v (respectively) as well as a van Kampen diagram $f \colon \Delta \to X(\mathcal{P})$ with boundary label UV^{-1}. We will call such a van Kampen diagram an *intersection van Kampen diagram* and the equation $u = v$ an *intersection equation*. We will denote the images of the paths U and V in $\partial\Delta$ as $\partial_+\Delta$ and $\partial_-\Delta$, respectively.

Note that, by Theorem 2.3 and the fact that Δ is simply connected, the boundary path of any region α of Δ is a simple closed path in $\Delta^{(1)}$, so that the closure of α is a closed disc.

If, amongst all intersection van Kampen diagrams, for all intersection equations $u = v$ (with $u \in (A * B) \setminus B$ and $v \in (B * C) \setminus B$) the number of regions in Δ is smallest possible, and if the total number of cells is smallest possible subject to the minimality of the number of regions, then we will call $f \colon \Delta \to X$ a *minimal intersection van Kampen diagram* and $u = v$ a *minimal intersection equation*.

There are four different types of 2-cells in the 2-complex $X = X(\mathcal{P})$: those coming from $X = X(\mathcal{P}_A)$, $X = X(\mathcal{P}_B)$, and $X = X(\mathcal{P}_C)$; and the single 2-cell corresponding to the additional relator R. We refer to the regions of a van Kampen diagram $f \colon \Delta \to X$ as A-, B-, C-, and R-regions according to the type of their images in X. The edges $f^{-1}(a)$ and $f^{-1}(c)$ of Δ will be called a- and c-edges, respectively. Note that only R-regions have a- or c-edges in their boundaries.

2.4. *Ordering of Regions*

Let $G = (A * B * C)/\langle\langle R\rangle\rangle$ be a one-relator product of locally indicable groups A, B, C, where R is not conjugate to an element of A, B, or C. Then we may write (uniquely) $R = \bar{R}^m$ for some $m \geq 1$ and some $\bar{R} \in A * B * C$ that is not a proper power.

By Theorem 2.2, the quotient group $\bar{G} = (A * B * C)\langle\langle \bar{R}\rangle\rangle$ of G is locally indicable and hence, by Theorem 2.4, right orderable. Choose a right ordering $<$ of $\bar{G}$, which we will regard as fixed for the rest of this section (although later we require the freedom to vary the right ordering). We use $<$ to define a partial ordering of the R-regions in a van Kampen diagram $f \colon \Delta \to X$ as follows.

Suppose that α, β are two R-regions of Δ. Choose an edge-path $Q = Q(\alpha, \beta)$ in $\Delta^{(1)}$ from the base-point of α to the base-point of β. We say that $\alpha < \beta$ if $\lambda(Q) > 1$. (Note that this definition is independent of the choice of the path Q, since Δ is simply connected.) An R-region of Δ is called *minimal* or *maximal* if it is minimal or maximal with respect to this partial ordering of R-regions. It is called *locally minimal* (resp., *locally maximal*) if it is no greater than (resp., no less than) any other R-region with which it shares an a- or c-edge. Clearly, minimal (resp. maximal) R-regions are also locally minimal (resp. locally maximal).

Without loss of generality, we assume that the attaching path in X of the 2-cell corresponding to R begins and ends at the base-point for $X(\mathcal{P}_B)$, so that we may think of R and $\bar{R}$ as cyclically reduced words in $(aAa^{-1}) * B * (cCc^{-1})$ (with $R = \bar{R}^m$).

The occurrences of a in $\bar{R}$ correspond to initial segments of $\bar{R}$ considered as a word in $(aAa^{-1}) * (B * cCc^{-1})$. By Theorem 2.3, these segments represent pairwise distinct elements of the right-ordered group $\bar{G}$. We denote by $a_{\min}$ and $a_{\max}$ the occurrences of a in $\bar{R}$ corresponding to the least and greatest of the initial segments of $\bar{G}$ (with respect to the chosen right ordering $<$); we define occurrences $c_{\min}$ and $c_{\max}$ of c in R in an analogous way.

Since $R = \bar{R}^m$, we can naturally regard R as containing m occurrences each of $a_{\min}$, $a_{\max}$, $c_{\min}$, and $c_{\max}$. These occur in a pattern that is repeated m times. In particular, for example, the m occurrences of $a_{\max}$ and $c_{\max}$ in R alternate.

In our van Kampen diagram $f\colon \Delta \to X$, each region has precisely m edges of each type $a_{\min}$, $a_{\max}$, $c_{\min}$, and $c_{\max}$, corresponding to its labeling $R^{\pm 1} = \bar{R}^{\pm m}$.

LEMMA 2.5. *Let α and β be regions of the minimal intersection van Kampen diagram Δ. If some $a_{\max}$- or $c_{\max}$-edge of α is also an edge of β, or if some $a_{\min}$- or $c_{\min}$-edge of β is also an edge of α, then $\alpha < \beta$.*

Proof. We assume that an $a_{\max}$-edge of α is also an edge of β. The other cases are similar. There is a path in $\partial\alpha$ from the base-point of α to the start of the given $a_{\max}$ edge that has label $\bar{R}^i S_M$ for some $i = 0, \dots, m-1$, where S_M is the greatest initial segment of $\bar{R}$ (with respect to the right ordering of $\bar{G}$). There is a path in $\partial\beta$ from the base-point of β to the start of this same a-edge that has label $\bar{R}^j S_p$ for some j, p, where S_p is also an initial segment of $\bar{R}$.

Note that $p \neq M$, since otherwise α and β would cancel, contradicting the minimality of Δ. Hence $S_p < S_M$. The concatenation of these two paths is a path Q, from the base-point of α to the base-point of β, whose label $\lambda(Q)$ is equal in $\bar{G}$ to $S_M S_p^{-1} > 1$. Hence $\alpha < \beta$ as claimed. □

COROLLARY 2.6. *Let $f\colon \Delta \to X$ be a minimal intersection van Kampen diagram, and let α be a locally minimal (resp. locally maximal) R-region of Δ. Then each $a_{\min}$- (resp. $a_{\max}$-) edge of α belongs to $\partial_+\Delta$ and each $c_{\min}$- (resp. $c_{\max}$-) edge of α belongs to $\partial_-\Delta$.*

Proof. We assume that α is locally maximal; the other case is similar. Let e be one of the $a_{\max}$-edges of α. If e is also an edge of another region β of Δ, then $\alpha < \beta$ by Lemma 2.5. But this contradicts local maximality of α, so $e \in \partial\Delta$. But $e \notin \partial_-\Delta$, since e is an a-edge, so $e \in \partial_+\Delta$. Similarly, every $c_{\max}$-edge of α belongs to $\partial_-\Delta$. □

This allows us to prove the first of our main results.

Proof of Theorem B. In the setup we have described, we need to show that the intersection of $(A * B)$ and $(B * C)$ in G is equal to B if $m \geq 2$. If not, then there exists a minimal intersection van Kampen diagram $f\colon \Delta \to X$. Clearly Δ must contain at least one R-region, for otherwise the intersection equation given by $\partial\Delta$ is a consequence of the relators of $\mathcal{P}_A$, $\mathcal{P}_B$, and $\mathcal{P}_C$. But then the intersection of $(A * B)$ and $(B * C)$ in $A * B * C$ would be strictly greater then B, which is absurd. Since Δ contains at least one R-region, it contains a maximal one, α say. Then α is also locally maximal. By Corollary 2.6, every $a_{\max}$-edge of α belongs to $\partial_+\Delta$ and every $c_{\max}$-edge of α belongs to $\partial_-\Delta$. But the $m \geq 2$ $a_{\max}$-edges of α alternate with the m $c_{\max}$-edges, so this is impossible. This contradiction completes the proof. □

2.5. Gurevich Regions

Let $X = X(\mathcal{P})$ be the 2-complex model for $G = (A * B * C)/\langle\langle R\rangle\rangle$ constructed in Section 2.2. Suppose that $\phi\colon \Delta \to X$ is a van Kampen diagram and that α is a region of Δ. A *Gurevich* path for α is an edge-path Q in $\partial\alpha \cap \partial\Delta$ that contains at least one edge labeled $a^{\pm 1}$ and at least one edge labeled $c^{\pm 1}$. A *Gurevich region* of Δ is a region (necessarily an R-region) admitting a Gurevich path. See Figure 2.1.

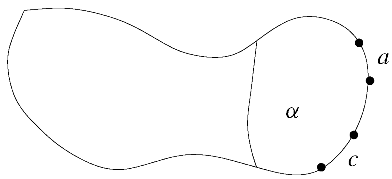

Figure 2.1

The following is essentially a special case of a result from [8].

THEOREM 2.7. *Let $\phi\colon \Delta \to X$ be a minimal van Kampen diagram with no boundary B-regions. Then the following statements hold.*

1. *If Δ has precisely one R-region, then that R-region is a Gurevich region.*
2. *If Δ has more than one R-region, then Δ contains at least two Gurevich regions.*

Sketch of Proof. In the language of [8], let $\mathcal{K}$ be the (staggered) generalized 2-complex whose 1-skeleton is the graph of groups (with trivial edge groups)

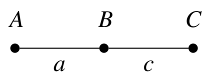

and a single 2-cell $c(R)$. We can amend $\phi\colon \Delta \to X$ to a diagram $\bar{\Delta}$ over $\mathcal{K}$ by absorbing any A-, B-, or C-regions into neighboring R-regions. The regions of $\bar{\Delta}$ are then in one-to-one correspondence with R-regions of Δ. By [8, Thm. 3.1], either $\bar{\Delta}$ has precisely one region and that region admits a Gurevich path, or there are at least two regions of $\bar{\Delta}$ that admit Gurevich paths. It therefore suffices to show that, if a region $\bar{\alpha}$ of $\bar{\Delta}$ admits a Gurevich path, then so does the corresponding region α of Δ.

Let Q be a Gurevich path of shortest length in $\partial\bar{\alpha}$. Then (with a suitable choice of orientation) Q has label $a^{-1}b_1 \cdots b_k c$ for some $b_1, \ldots, b_k \in B$. The first and last edges of Q are necessarily in $\partial\Delta$. If the other edges of Q are not contained in $\partial\Delta$ then they, together with a path in $\partial\Delta$, bound a union of B-regions that become absorbed into $\bar{\alpha}$ in the passage from Δ to $\bar{\Delta}$. But the latter possibility contradicts the condition that Δ has no boundary B-regions. □

3. Analysis of Minimal Intersection Diagrams

LEMMA 3.1. *Let Δ be a minimal intersection van Kampen diagram over the one-relator product $G = (A * B * C)/\langle\langle R\rangle\rangle$ of locally indicable groups A, B, and C. Suppose $\partial\Delta$ is the union of two paths $\partial_+\Delta$ and $\partial_-\Delta$, with a common initial point p_0 and a common terminal point p_1, labeled $u \in (A * B) \setminus B$ and $v \in (B * C) \setminus B$ (respectively) and such that $u \cdot v^{-1}$ is the boundary label of Δ.*

1. *Δ is a topological disc, or (equivalently) $\partial\Delta$ is a simple closed path.*
2. *Each boundary region of Δ is an R-region.*
3. *If Δ has more than one R-region, then it has precisely two Gurevich R-regions. The boundary of one of these contains p_0 and the boundary of the other contains p_1.*
4. *Each Gurevich region α of Δ is locally minimal or locally maximal. More precisely, each α admits a Gurevich path containing either the $a_{\max}$- and $c_{\max}$-edges of α or its $a_{\min}$- and $c_{\min}$-edges.*
5. *If Δ has a locally minimal (resp. maximal) Gurevich region, then any locally minimal (resp. maximal) region of Δ is a Gurevich region.*

Proof. 1. Suppose first that $\partial_+\Delta$ is not a simple path. Then there is a closed subpath, labeled (say) $u' \in A * B$, that bounds a subdiagram Δ' of Δ; hence $u' = 1$ in G. By the Freiheitssatz (Theorem 2.1), $u' = 1$ in $A * B$, so we can delete Δ' from Δ to get a smaller diagram with the same boundary label (evaluated in $A * B * C$), contrary to the minimality hypothesis.

Hence we may assume that $\partial_+\Delta$ is a simple path, and similarly $\partial_-\Delta$. Next suppose that these paths have an intermediate vertex in common. Then $u = u_1u_2$ and $v = v_1v_2$ in such a way that $u_iv_i^{-1}$ is the boundary label of a subdiagram Δ_i of Δ for $i = 1, 2$. Clearly, each Δ_i has strictly fewer cells than Δ, while at least one of the intersection equations $u_i = v_i$ is nontrivial (since the equation $u = v$ is nontrivial), contradicting the minimality of Δ once again.

2. Suppose that some boundary region α is not an R-region. Clearly α cannot be the whole of Δ. Hence $\partial\alpha$ contains an arc γ that meets $\partial\Delta$ only in its endpoints. See Figure 3.1.

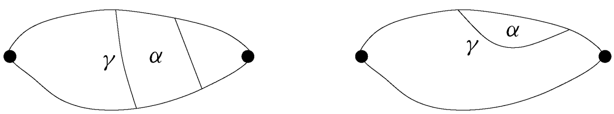

Figure 3.1

Since α is not an R-region, $\lambda(\gamma)$ is a word in the generators of $\mathcal{P}_A$, or of $\mathcal{P}_B$, or of $\mathcal{P}_C$. Cutting Δ along γ yields two smaller diagrams Δ_1 and Δ_2, at least one of which is a nontrivial intersection diagram, contradicting the minimality of Δ.

3. Since Δ has no boundary B-regions, Δ has at least two Gurevich regions by Theorem 2.7. But any Gurevich path in $\partial\Delta$ contains edges with labels a and c. Such a path intersects both $\partial_+\Delta$ and $\partial_-\Delta$ and so must contain (in its interior) one of p_0, p_1. The result follows.

4. If Δ has only one R-region, then it is simultaneously a minimal, maximal, and Gurevich region, so there is nothing to prove. We will therefore assume that Δ has at least two R-regions.

By definition, the $a_{\max}$-edge of any locally maximal R-region of Δ belongs to $\partial_+\Delta$, as does the $a_{\min}$-edge of any locally minimal R-region. Consider the collection $e_1, \dots, e_k$ of edges of this type (in the order in which they occur in $\partial_+\Delta$). Since Δ has at least one minimal and one maximal R-region, it follows that $k \geq 2$. We will show that e_1 and e_k belong to the two Gurevich regions.

Consider the R-region α to which e_1 belongs. This region is either locally minimal or locally maximal, so it meets both parts of $\partial\Delta$. We assume that α is locally maximal; the other case is similar.

Let P denote the subpath of $\partial\Delta$ whose first and last edges are (respectively) the $a_{\max}$- and $c_{\max}$-edges of α and that passes through p_0. If $P \subset \partial\alpha$, then P is a Gurevich path containing the $a_{\max}$- and $c_{\max}$-edges of α, as required.

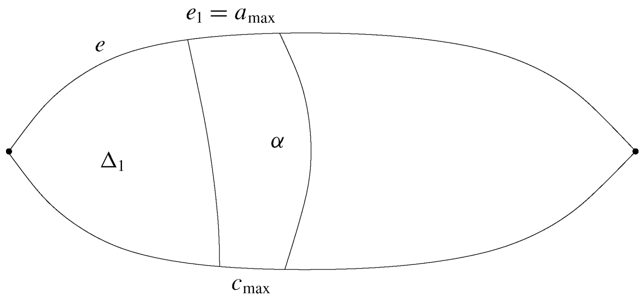

Figure 3.2

If $P \not\subset \partial\alpha$, choose an edge e of P with $e \notin \partial\alpha$, remove α and the interior of $\partial\alpha \cap \partial\Delta$ from Δ, and consider the component Δ_1 of what remains that contains e. See Figure 3.2. Now Δ_1 contains at least one R-region (e.g., the boundary region that meets e). Hence it contains an R-region β that is minimal in Δ_1. The only region of $\Delta \setminus \Delta_1$ with which β can share an a- or c-edge is α, and if this happens then $\beta < \alpha$ by local maximality of α. Combining this with the minimality of β in Δ_1, it follows that β is locally minimal in Δ. But then, by Corollary 2.6, the $a_{\min}$-edge of β lies in $\partial_+\Delta$ between p_0 and e_1, a contradiction.

We have shown (a) that the region containing e_1 is the Gurevich region containing p_0 and (b) that this region admits a Gurevich path containing either its

$a_{\max}$- and $c_{\max}$-edges or its $a_{\min}$- and $c_{\min}$-edges. A similar argument applies to the region containing e_k.

5. Suppose Δ has two locally minimal R-regions α and β, and suppose that α is a Gurevich region. Then, by part 4 of the Lemma, α admits a Gurevich path $P \subset \partial\alpha$ containing both the $a_{\min}$- and $c_{\min}$-edges of α. Let Q be the (oriented) subpath of P whose first edge is the $a_{\min}$-edge of α and whose last edge is the $c_{\min}$-edge of α. Then the label $w = \lambda(Q)$ of Q belongs to $(A * B) \cdot (B * C)$.

There is a path $Q' \in \partial\beta$ whose label $\lambda(Q')$ is also equal to w and whose first and last edges are the $a_{\min}$- and $c_{\min}$-edges of β, respectively. In particular, Q' joins the two parts of $\partial\Delta$. Let Q'' be a subpath of Q' that joins the two parts of $\partial\Delta$ and is minimal with respect to that property; then its label $\lambda(Q'')$ also belongs to $(A * B) \cdot (B * C)$. If Q'' consists of a single point, then β is a Gurevich R-region, as claimed.

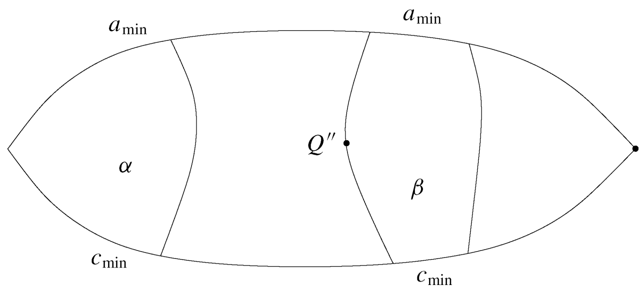

Figure 3.3

Otherwise, cutting along Q'' divides Δ into two subdiagrams Δ_1 and Δ_2, each with boundary label in $(A * B) \cdot (B * C)$ and so each representing an exceptional intersection equation. At least one of these equations is nontrivial, contradicting the minimality of Δ. See Figure 3.3. □

THEOREM 3.2. *Let Δ be a minimal intersection van Kampen diagram. Then Δ has a unique minimal R-region and a unique maximal R-region, each of which is a Gurevich R-region.*

Proof. In the free product $A * B * C$ we have $(A * B) \cap (B * C) = B$, so Δ must contain at least one R-region and hence, by Theorem 2.7, at least one Gurevich R-region. If Δ has a single R-region then there is nothing to prove. From now on we will assume that Δ has more than one R-region and hence that it has precisely two Gurevich R-regions; by Lemma 3.1, each of these is either locally minimal or locally maximal.

Suppose that one of the Gurevich R-regions is locally minimal and the other is locally maximal. By Lemma 3.1, any locally maximal or locally minimal R-region is a Gurevich R-region, so there is precisely one locally minimal R-region and precisely one locally maximal R-region. Since minimal (resp. maximal) R-regions are locally minimal (resp. locally maximal), the result follows in this case.

Next suppose that the two Gurevich R-regions of Δ are locally minimal (and that neither is locally maximal). We will show that this is impossible by deriving a contradiction. By symmetry, the same argument will apply to the case where both Gurevich R-regions are locally maximal and not locally minimal.

By Lemma 3.1, only the Gurevich R-regions are locally minimal. Moreover, Δ has at least one maximal R-region, say α. Since α is maximal, it is locally maximal and hence not a Gurevich R-region. See Figure 3.4.

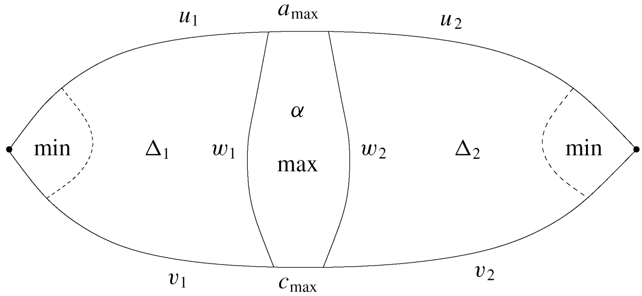

Figure 3.4

Also, α meets both sides of the boundary of Δ: the edges $a_{\max}$ and $c_{\max}$ of α belong to $\partial_+\Delta$ and $\partial_-\Delta$, respectively. Hence, by removing α and the edges and intermediate vertices in $\partial\alpha \cap \partial\Delta$, we disconnect Δ into two subdiagrams Δ_1 and Δ_2. Note that, for $i = 1, 2$, the diagram Δ_i has the following properties.

(1) $\partial\Delta_i$ is the union of three arcs: $\partial\Delta_i \cap \partial_+\Delta$, $\partial\Delta_i \cap \partial_-\Delta$, and $\partial\Delta_i \cap \partial\alpha$. Hence the boundary label of Δ_i has the form $u_i v_i w_i$, where $u_i \in A * B$ is a terminal segment of $u^{\pm 1}$, $v_i^{-1} \in B * C$ is a terminal segment of $v^{\pm 1}$, and w_i is a cyclic subword of $R^{\pm 1}$.

(2) Δ_i contains one of the Gurevich R-regions of Δ, so in particular it has at least one R-region.

The cyclic subwords w_1 and w_2 of R are disjoint. Indeed, there is a cyclic permutation of $R^{\pm 1}$ of the form

$$x_1 a x_2^{-1} w_2^{-1} y_2^{-1} c y_1 w_1,$$

where a is precisely the letter $a_{\max}$ of R and c is precisely the letter $c_{\max}$; $x_1, x_2 \in A * B$; and $y_1, y_2 \in B * C$.

Consider $A * B * C$ as a free product of $A * B$ and $B * C$, amalgamating B:

$$A * B * C = (A * B) \underset{B}{*} (B * C).$$

Without loss of generality, we will assume that the length of the normal form of $z_1 := c y_1 w_1 x_1 a$ in this amalgamated free-product structure is less than or equal to that of $z_2 := c^{-1} y_2 w_2 x_2 a^{-1}$. Note that this length is even, since z_1 begins with c and ends with a. Moreover, it is strictly greater than 2, since otherwise z_1 and hence also w_1 belong to $(B * C) \cdot (A * B)$. But if $w_1 = v'u'$ with $u' \in A * B$ and $v' \in B * C$, then the subdiagram Δ_1 has boundary label

$$v_1 v' u' u_1 \in (B * C) \cdot (A * B),$$

contradicting the minimality of Δ.

If the lengths of the normal forms of z_1 and z_2 are equal then we may assume, also without loss of generality, that the length of z_1 as a word in the free product $A * B * C$ is no greater than that of z_2.

Since Δ_1 has R-regions, it has at least one maximal R-region—say, β.

CLAIM. *$\beta \cap \partial\Delta_1$ is connected.*

Proof. Otherwise, removing β and the interior of $\beta \cap \partial\Delta_1$ from Δ_1 leaves two or more components, each of which contains 2-cells. Taking Δ_0 to be a component that does not contain a Gurevich region of Δ, let γ be a minimal R-region of Δ_0. Then the only regions of $\Delta \setminus \Delta_0$ with which γ can share a- or c-edges are α and β; and $\gamma < \alpha$ and $\gamma < \beta$ by the maximality of α in Δ and of β in Δ_1, respectively. Hence γ is a locally minimal region of Δ but not a Gurevich region, contradicting Lemma 3.1.

Thus we see that Δ_0 contains no R-regions and hence, by Lemma 3.1, no boundary regions of Δ. Hence $\partial\Delta_0 \subset \partial\alpha \cup \partial\beta$ and so the boundary label of Δ_0 has the form $s_1 s_2$, where each s_i is a subword of $R^{\pm 1}$. By Property B (see Section 2.3), $s_1 s_2$ is freely equal to the empty word in the generators of the presentation $\mathcal{P}$; hence Δ_0 can be deleted from Δ, contrary to the minimality assumption. This proves the claim. $\square$

The edge $a_{\max}$ of β lies on $\partial\Delta_1 \setminus \partial_-\Delta = \partial\Delta_1 \cap (\partial_+\Delta \cup \partial\alpha)$. Similarly, the edge $c_{\max}$ of β lies in $\partial\Delta_1 \cap (\partial_-\Delta \cup \partial\alpha)$. There are four possibilities, as follows.

(i) The $a_{\max}$- and $c_{\max}$-edges of β both belong to $\partial\alpha$ (see Figure 3.5).

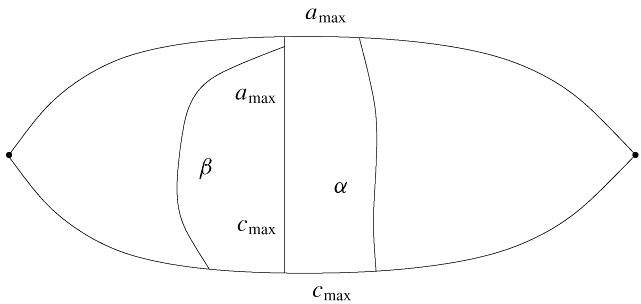

Figure 3.5

Then the $a_{\max}$- and $c_{\max}$-edges of β both belong to a path Q in $\partial\alpha \cap \partial\beta \subseteq \partial\Delta_1 \cap \partial\alpha$. Hence its label is a subword of w_1. Regarding Q as part of $\partial\beta$, we see that it contains a subword $z_1^{\pm 1}$ or $z_2^{\pm 1}$.

Since w_1 is a proper subword of z_1, it cannot contain a subword equal to $z_1^{\pm 1}$, while for w_1 (and hence z_1) to contain a subword equal to $z_2^{\pm 1}$ would contradict the length assumptions on z_1 and z_2.

(ii) The $a_{\max}$- and $c_{\max}$-edges of β both belong to $\partial\Delta$ (Figure 3.6).

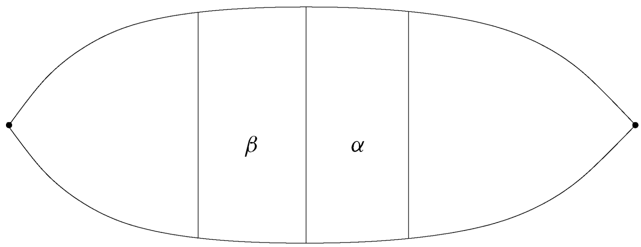

Figure 3.6

Then $\partial\alpha \cap \partial\Delta_1 \subset \partial\beta$. Thus w_1 is identified with a subword of z_2 (if β has the same orientation as α) or z_1 (otherwise). Now w_1 occurs precisely once as a subword of z_1 (e.g., because the first A-letter of z_1 is the first A-letter of w_1). Hence, if α and β have opposite orientations in Δ then they can be canceled to yield a smaller diagram (also representing the equation $u = v$).

On the other hand, if α and β have the same orientation in Δ then we can construct a smaller diagram than Δ by identifying $\beta \in \Delta_1$ with $\alpha \in (\Delta \setminus \Delta_1)$; see Figure 3.7.

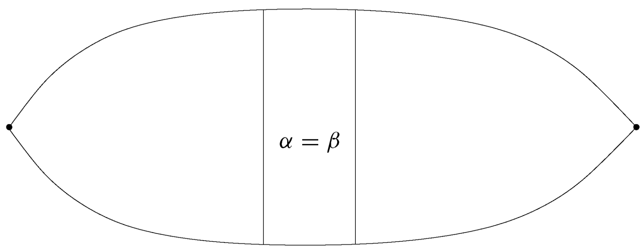

Figure 3.7

(iii) The $a_{\max}$-edge of β belongs to $\partial\alpha$, while the $c_{\max}$-edge belongs to $\partial\Delta$ (Figure 3.8).

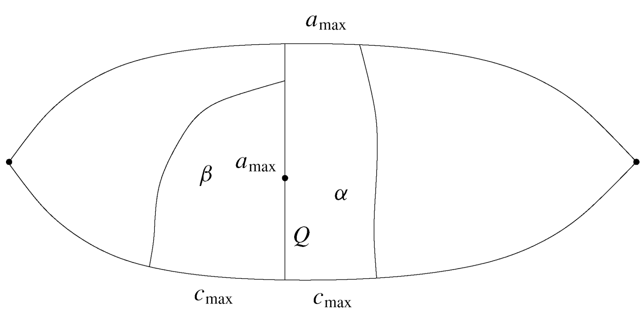

Figure 3.8

Then there is a path Q in $\partial\alpha \cap \partial\beta$ joining $\partial_-\Delta$ to the edge $a_{\max}$ of β. Regarding Q as part of $\partial\Delta_1 \cap \partial\alpha$, we see that its label is an initial segment w', say, of w_1. Regarding Q as part of $\partial\beta$, we see that there exists a word $y \in B * C$ such that $c^{\pm 1} y w' a^{\pm 1} \in \{z_1, z_2\}$.

Write $w_1 = w'w''$. It follows from the length assumptions on z_1 and z_2 that $w'' \in A * B$. Cutting Δ along the arc in $\partial\alpha \cup \partial\Delta_1$ labeled $a^{\pm 1} w''$ creates a new diagram Δ' with the same number of regions as Δ. The boundary label of Δ' is not cyclically reduced, but it is equal in $A * B * C$ to that of Δ. Nevertheless, the regions α and β in Δ' satisfy the conditions of (ii), so we may either cancel or identify α and β (as in (ii)) to obtain a diagram with fewer R-regions than Δ. This diagram may not have cyclically reduced boundary label, but by performing folds along the boundary we can obtain another intersection diagram, again with fewer R-regions than Δ, that does have cyclically reduced boundary label. This contradicts the assumption of minimality.

(iv) The $c_{\max}$-edge of β belongs to $\partial\alpha$, while the $a_{\max}$-edge belongs to $\partial\Delta$. This is the same as case (iii) with A and C interchanged.

In all cases we have obtained a contradiction, and the result follows. □

4. Towers and Magnus Intersections

Recall [7] that a *tower* of 2-complexes is a map $q\colon X \to Y$ between 2-complexes that is a composite of a finite number of maps, each of which is either the inclusion of a subcomplex or a covering projection. One can restrict attention to classes of towers in which restrictions are placed on the subcomplexes and/or on the coverings arising in this decomposition of q. For our purposes, the coverings in the tower will always be *infinite cyclic*; that is, they are connected regular coverings with infinite cyclic covering transformation group. (An infinite cyclic covering of a path-connected space Y is determined by the kernel of an epimorphism $\pi_1(Y) \to \mathbb{Z}$.)

If $f\colon \Delta \to Y$ is a cellular map of 2-complexes (such as, e.g., a van Kampen diagram), then a *tower lift* of f is a cellular map $\bar{f}\colon \Delta \to \bar{Y}$ such that $f = q \circ \bar{f}$ for some tower $q\colon \bar{Y} \to Y$. Since the composite of two towers is a tower, it follows that a tower lift of a tower lift is a tower lift. A tower lift $\bar{f}$ of f with $f = q \circ \bar{f}$ is *proper* if q is not an isomorphism of 2-complexes and is *maximal* if $\bar{f}$ itself has no proper tower lift.

It is shown in [7] that every combinatorial map $f\colon \Delta \to Y$ of 2-complexes, defined on a finite 2-complex Δ, admits a maximal tower lift.

LEMMA 4.1. *Let $L \subset K$ be a pair of 2-complexes such that*:

(i) *$K \setminus L$ consists of two 1-cells a and c and a 2-cell ρ*;
(ii) *each component of L has locally indicable fundamental group*; *and*
(iii) *each 1-cell a, c occurs in the attaching path for the 2-cell ρ.*

Let $\phi\colon \Delta \to K$ be a minimal van Kampen diagram for a minimal equation between a path in $X := L \cup a$ and a path in $Y := L \cup c$. Then there is a tower lift $\bar{\phi}\colon \Delta \to \bar{K}$ of ϕ such that

(1) $\bar{K}$ *has finite* 1*-skeleton,*
(2) $\bar{K} \setminus \bar{\phi}^{-1}(L)$ *has a single* 2*-cell, and*
(3) $\beta_1(\bar{K}) \le 1$.

Proof. Note first that we may assume that ϕ restricts to a surjective map $\Delta^{(1)} \to K^{(1)}$ of 1-skeleta and hence that $K^{(1)}$ and $L^{(1)}$ are finite. To see this, first replace K by the finite subcomplex $K' = \phi(\Delta)$ and L by $L' = L \cap K'$. In general, it will not be true that the components of L' will have locally indicable fundamental groups, since the inclusion-induced maps $\pi_1(L', x) \to \pi_1(L, x)$ (for $x \in L$) are not necessarily injective. However, we may add (possibly infinitely many) 2-cells to L' to form a 2-complex L'', with finite 1-skeleton, such that $\pi_1(L'', x)$ is isomorphic to the image of $\pi_1(L', x) \to \pi_1(L, x)$ for each 0-cell x. Now replace K by $K'' = K' \cup_{L'} L''$.

By minimality of the van Kampen diagram ϕ, it follows that Δ is a topological disc (see Lemma 3.1). Let $R \in \pi_1(K \cup a \cup c)$ be the attaching path for ρ. If Δ contains a single R-region, then any maximal tower lifting of ϕ will satisfy the conclusions of the Lemma (with $\beta_1 = 0$).

Hence we may assume that there exist at least two R-regions in Δ and hence at least two Gurevich R-regions. By Lemma 3.1 and Theorem 3.2, Δ contains precisely one maximal R-region and one minimal R-region, and these are the two Gurevich R-regions. Moreover, this holds for the ordering of 2-cells induced by *any* right ordering of $G = \pi_1(K)$. We argue by induction on $v_1(\phi) := v_0(\Delta) - v_0(K)$, where $v_0(\cdot)$ denotes the number of 0-cells.

If $\beta_1(K) := \dim(H^1(K; \mathbb{Q})) \le 1$ then we may take $\bar{K} = K$ and $\bar{\phi} = \phi$. In particular, this applies if $v_1(\phi) = 0$, for then ϕ has no proper tower lifting and so $\beta_1(K) = 0$. This starts the induction. For the inductive step, assume that $\beta_1(K) \ge 2$. Then $H^1(K; \mathbb{Q}) = \mathrm{Hom}(G, \mathbb{Q})$ has dimension at least 2. Choose a path P in Δ between the base-points of the two Gurevich R-regions. Then we may find a nonzero homomorphism $\psi\colon G \to \mathbb{Q}$ such that $\psi(\phi_*(P)) = 0$. The image of ψ is infinite cyclic (since G is finitely generated). Now choose a right ordering of G dominated by the standard ordering of $\mathbb{Q}$ (via the homomorphism ψ). If α is a 2-cell of Δ and if P_+ and P_- are paths joining the base-point of α to those of the maximal and minimal 2-cells, respectively, then $\psi(\phi_*(P_+)) \ge 0$ and $\psi(\phi_*(P_-)) \le 0$ by definition of maximality and minimality. On the other hand, Δ is simply connected and so $\psi(\phi_*(P_-)) - \psi(\phi_*(P_+)) = \pm\psi(\phi_*(P)) = 0$. Hence $\psi(\phi_*(P_-)) = \psi(\phi_*(P_+)) = 0$.

Let $p\colon \tilde{K} \to K$ denote the (infinite cyclic) regular covering of K corresponding to $\mathrm{Ker}(\psi)$, let $\tilde{\phi}\colon \Delta \to \tilde{K}$ be a lift of ϕ, and define $L_1' = p^{-1}(L)$ and $K_1' = L_1' \cup \tilde{\phi}(\Delta)$. Now take K_1 to be the path component of K_1' containing $\tilde{\phi}(\Delta)$, and define $L_1 = L_1' \cap K_1$. It follows from our previous remarks that $K_1 \setminus L_1$ has only one 2-cell. By standard arguments, the restriction $p\colon K_1 \to K$ is strictly surjective, so that $v_0(K_1) > v_0(K)$ and $v_1(\tilde{\phi}) < v_1(\phi)$.

Now apply the inductive hypothesis to $\tilde{\phi}\colon \Delta \to K_1$ (with respect to the unique 2-cell $\tilde{\rho} \in p^{-1}(\rho)$ and to any choice $\tilde{a} \in p^{-1}(a)$ and $\tilde{c} \in p^{-1}(c)$ of 1-cells that are involved in the attaching map for $\tilde{\rho}$). Since any tower lift of $\tilde{\phi}$ is a tower lift of ϕ, we are done. □

COROLLARY 4.2. *Let $G = (A * B * C)/\langle\langle R\rangle\rangle$ be a one-relator product of locally indicable groups A, B, C, and let uv^{-1} be the boundary label of a minimal exceptional intersection van Kampen diagram, where $u \in (A * B) \setminus B$ and $v \in (B * C) \setminus B$. Then some cyclic conjugate R' of R is contained in a subgroup F_0 of $A * B * C$ such that*

1. *F_0 is finitely generated,*
2. *$\beta_1(F_0) \leq 2$,*
3. *$uv^{-1} \in F_0$, and*
4. *uv^{-1} is a consequence of R' in F_0.*

Proof. We may clearly assume that R is a cyclically reduced word involving each of the free factors A, B, and C. Let L be the disjoint union of 2-complexes L_A, L_B, L_C with fundamental groups A, B, C, respectively. Form a 2-complex K from L by attaching a 1-cell a joining L_A to L_B, a 1-cell c joining L_B to L_C, and a 2-cell ρ with attaching map in the class

$$R \in \pi_1(L_A \cup a \cup L_B \cup c \cup L_C) = A * B * C.$$

Let $\phi\colon \Delta \to K$ be a minimal intersection van Kampen diagram with boundary label uv^{-1}. Then the hypotheses of Lemma 4.1 are satisfied.

Hence we have a tower lift $\bar{\phi}\colon \Delta \to \bar{K}$ of ϕ satisfying the conclusions of Lemma 4.1. Let p_0 be the common initial point of $\partial_+\Delta$ and $\partial_-\Delta$ (which have labels u and v, respectively). By Lemma 3.1, p_0 lies on the boundary of a Gurevich R-region α of Δ. The boundary label of α, read from p_0, is a cyclic conjugate R' of R. Fix $x := \bar{\phi}(p_0)$ as a base-point of $\bar{K}$. Then the single 2-cell $\bar{\rho}$ of $\bar{K} \setminus q^{-1}(L)$ (where $q\colon \bar{K} \to K$ is the tower map) is attached by a closed path P based at x, with $q_*(P) = R'$. Define $F_0 = q_*(\pi_1(\bar{K} \setminus \bar{\rho})) \subseteq A * B * C$ to obtain the result. (Note that $\beta_1(F_0) \leq 2$ since $\beta_1(\bar{K}) \leq 1$ and that F_0 is finitely generated because $\bar{K}$ has finite 1-skeleton.) $\square$

5. Rank-2 Subgroups

Following Corollary 4.2, it is useful to be able to classify the groups F_0 that can arise.

LEMMA 5.1. *Let F be a finitely generated subgroup of the free product $A * B * C$ of locally indicable groups A, B, C such that $\beta_1(F) \leq 2$ and F contains a word of the form uv^{-1}, with $u \in (A * B) \setminus B$ and $v \in (B * C) \setminus B$. Then F has one of the following forms*:

(a) *$\langle uv^{-1}\rangle * D$ for some $D \subset A * B * C$ with $\beta_1(D) \leq 1$;*
(b) *$D * E$, where $u \in D \subset (A * B)$ and $v \in E \subset (B * C)$;*
(c) *$\langle xy\rangle * xDx^{-1}$, where $x \in (A * B) \setminus B$, $y \in (B * C) \setminus B$, $u \in xDx^{-1}$, $v \in y^{-1}Dy$, and $D \subset B$.*

Proof. By the Kuroš subgroup theorem, F is a free product $F = *_{\lambda\in\Lambda} F_\lambda$, with each F_λ either infinite cyclic or a (nontrivial) conjugate of a subgroup of A, B, or C. Moreover, each F_λ is finitely generated and hence indicable. Since $\beta_1(F) \leq 2$, it follows that $|\Lambda| \leq 2$.

Suppose first that $|\Lambda| = 1$. Since $uv^{-1} \in F$ is neither a proper power in $A * B * C$ nor conjugate to an element of A, B, or C, it follows that $F = \langle uv^{-1} \rangle = \langle uv^{-1} \rangle * D$ with $D = \{1\}$. Hence we may assume that $|\Lambda| = 2$, and we write $F = F_1 * F_2$ with $\beta_1(F_1) = \beta_1(F_2) = 1$.

We can express $A * B * C$ as the fundamental group of a graph of groups, where the underlying graph is the tree Γ of Figure 5.1, the vertex groups are A, B, C as indicated, and the edge groups are trivial. By Bass–Serre theory [12], $A * B * C$ acts on a tree T with quotient Γ. The action is free on the edges, and the vertex stabilizers are the conjugates of A, B, and C. We can speak of A-, B-, and C-vertices of T (and of a- and c-edges) according to the image in Γ.

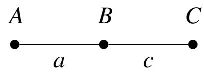

Figure 5.1

Without loss of generality, we may assume that uv^{-1} is cyclically reduced as written. Let t denote the vertex of T whose stabilizer is B. Then the geodesic P from t to $u(t)$ in T consists only of a-edges, while the geodesic Q from t to $v(t)$ consists only of c-edges (since $u \in A * B$ and $v \in B * C$).

By Bass–Serre theory again, F is the fundamental group of a graph of groups whose underlying graph is T/F. Since F is finitely generated with $\beta_1(F) = 2$, it follows that F is actually the fundamental group of a graph of groups whose underlying graph is a finite subgraph X of T/F with $\beta_1(X) \leq 2$. The various possibilities are illustrated in Figures 5.2, where the solid discs represent the images in X of vertices of T whose stabilizers have nontrivial intersection with F.

Let $\bar{P}$ and $\bar{Q}$ denote the images of P and Q in T/F. Since $uv^{-1} \in F$, each of $\bar{P}, \bar{Q}$ is a path joining the image x of t to the common image y of $u(t)$ and $v(t)$, and the path $\bar{P} \cdot \bar{Q}^{-1}$ is contained in X. Moreover, the paths $\bar{P}$ and $\bar{Q}$ have no common edges, since one contains only a-edges and the other contains only c-edges.

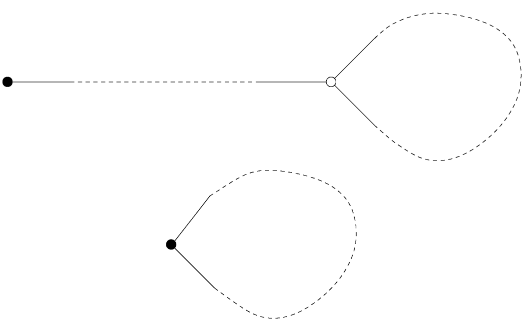

Figure 5.2a

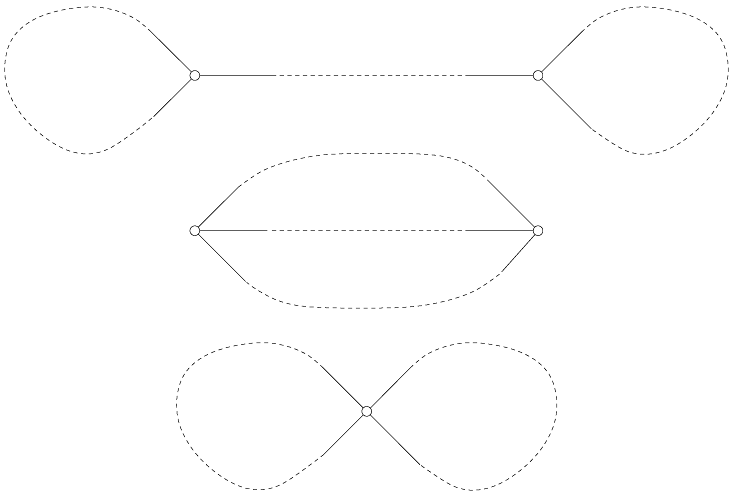

Figure 5.2b

There are essentially three possibilities as follows.

Case 1: $x \neq y$ (see Figures 5.3). Then F can be rewritten as $F_1 * \langle uv^{-1} \rangle$.

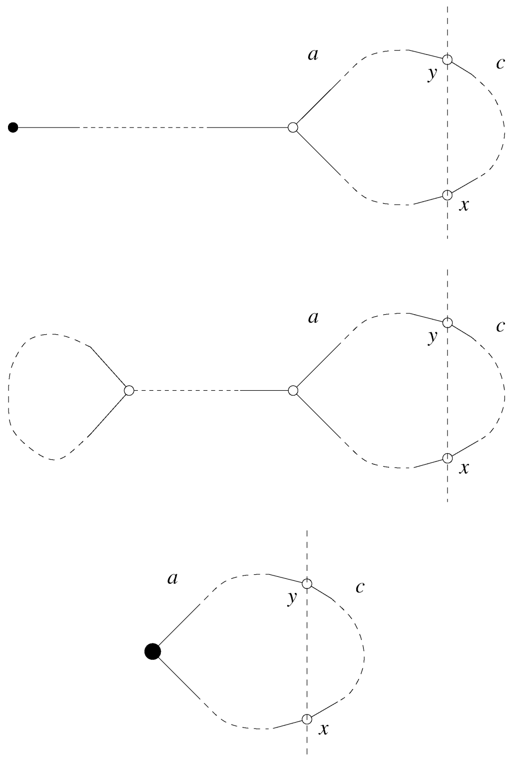

Figure 5.3a

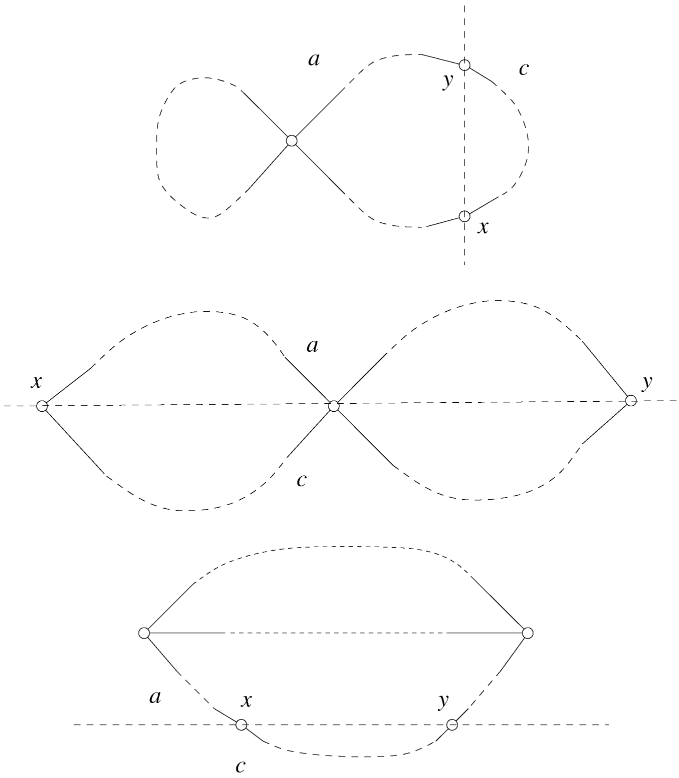

Figure 5.3b

Case 2: $x = y$ is a separating vertex of X (Figure 5.4). In this case we can write $F = F_1 * F_2$, with $u \in F_1 \subset A * B$ and $v \in F_2 \subset B * C$.

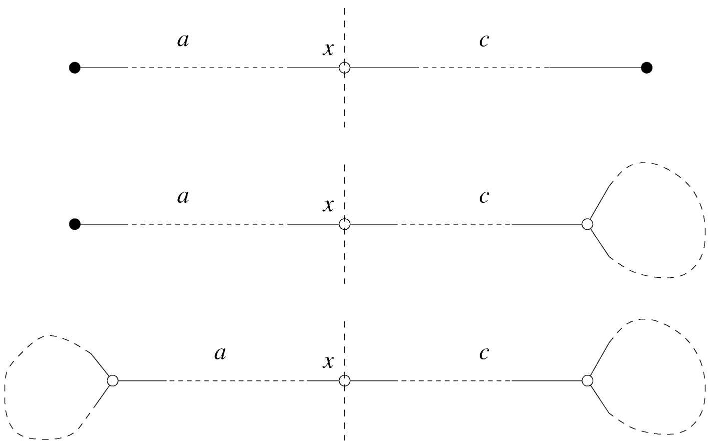

Figure 5.4

Case 3: $x = y$ is a nonseparating vertex of X (Figure 5.5). In this case the solid vertex z must be a B-vertex, and we can write $u \in UDU^{-1}$ and $v \in V^{-1}DV$ for some $U \in A * B$, $V \in B * C$, and $D \subset B$. Moreover, $F = UDU^{-1} * \langle UV \rangle$. $\square$

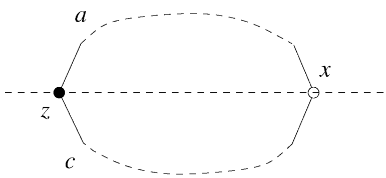

Figure 5.5

We are now ready for the proof of Theorem C. We begin with some lemmas.

LEMMA 5.2. *Let A be a one-relator group, and let R be a cyclically reduced word in $A * \langle t \rangle$ such that $t = 1$ in the one-relator product $G = (A * \langle t \rangle)/\langle\langle R \rangle\rangle$. Then $R = t^{\pm 1}$.*

Proof. By Brodskiĭ's Freiheitssatz for one-relator products of locally indicable groups (Theorem 2.1), since the natural map $\langle t \rangle \to G$ is not injective, R must be conjugate to an element of $\langle t \rangle$. Since also R is cyclically reduced, $R = t^n$ for some n. Then $G = A * \langle t \mid t^{|n|} \rangle$, and the result follows. $\square$

LEMMA 5.3. *Let $G = (A * B * C)/\langle\langle R \rangle\rangle$ be a one-relator product of locally indicable groups A, B, C such that the relator R has the form uv^{-1}, with $u \in A * B$ and $v \in B * C$. Then the intersection of $A * B$ and $B * C$ in G has the form $B * \langle u \rangle = B * \langle v \rangle$.*

Proof. The one-relator product structure expresses G as a free product with amalgamation

$$(A * B) \underset{B*\langle u \rangle = B*\langle v \rangle}{*} (B * C),$$

and the result follows trivially. $\square$

LEMMA 5.4. *Let $G = (A * B * C)/\langle\langle R \rangle\rangle$ be a one-relator product of locally indicable groups A, B, C such that the relator R is contained in a finitely generated subgroup $D * E$ of the free product $A * B * C$, where $D \subset A * B$ and $E \subset B * C$ are subgroups such that $\beta_1(D) = \beta_1(E) = 1$ with $D \not\subset B$ and $E \not\subset B$. Then the intersection of $A * B$ with $B * C$ in G has the form $B * I$, where I is the intersection of D and E in the one-relator product $G_0 = (D * E)/\langle\langle R \rangle\rangle$.*

Proof. Note that D cannot be a free product, since it is a finitely generated and locally indicable group with $\beta_1(D) = 1$. Hence D is either free (of rank 1) or contained in a conjugate of A or B. Since $D \not\subset B$, it follows that the subgroup of $A * B$ generated by $D \cup B$ is a free product $D * B$. Similarly, the subgroup generated by $E \cup B$ is a free product $B * E$.

Hence we may write G as a stem product

$$(A * B) \underset{D*B}{*} (G_0 * B) \underset{E*B}{*} (B * C),$$

and the result follows. $\square$

LEMMA 5.5. *Let $G = (A * B * C)/\langle\langle R\rangle\rangle$ be a one-relator product of locally indicable groups A, B, C such that the relator R is contained in a subgroup $\langle xy\rangle * D$ of the free product $A * B * C$, where $x \in (A * B) \setminus B$, $y \in (B * C) \setminus B$, and $D \subset B$. Then either R is conjugate to $(xyd)^{\pm 1}$ for some $d \in D$, or the intersection of $(A * B)$ with $(B * C)$ in G has the form $B * (x^{-1}Ix)$, where $I \subset D$ is the intersection of D with $(xy)D(xy)^{-1}$ in $(\langle xy\rangle * D)/\langle\langle R\rangle\rangle$.*

Proof. Write R as a word $R_0 = R_0(z, D)$ in the free product $\langle z\rangle * D \subset \langle z\rangle * B$, where $z = xy$. Now form a one-relator product

$$G_0 = (\langle x_0, y_0\rangle * D)/\langle\langle R_0(x_0 y_0, D)\rangle\rangle.$$

Then we can express G as a stem product

$$G = (A * B) \underset{\langle x\rangle * B = \langle x_0\rangle * B}{*} G_1 \underset{\langle y_0\rangle * B = \langle y\rangle * B}{*} (B * C),$$

where G_1 is a free product with amalgamation $G_1 = G_0 *_D B$. Thus the intersection of $A * B$ and $B * C$ in G has the form $E *_D B$, where E is the intersection of $\langle x_0\rangle * D$ and $\langle y_0\rangle * D$ in G_1.

But we can write G_0 as a free product $G_2 * \langle t\rangle$, where $G_2 = (\langle z_0\rangle * D)/\langle\langle R_0(z_0, D)\rangle\rangle$, $z_0 = x_0 y_0$, and $t = x_0$. Suppose $U \in \langle x_0\rangle * D$, $V \in \langle y_0\rangle * D$, and $U = V$ in G_0. We can write U, V in the following form:

$$U = g_0 x_0^{\varepsilon(1)} \cdots x_0^{\varepsilon(k)} g_k,$$

with $g_i \in D$ and $\varepsilon(i) = \pm 1$ for each i; and

$$V = h_0 y_0^{\eta(1)} \cdots y_0^{\eta(l)} h_l,$$

with $h_i \in D$ and $\eta(i) = \pm 1$ for each i.

Comparing U and V in $G_2 * \langle t\rangle$ using $D \subset G_2$ and using $x_0 = t$ and $y_0 = t^{-1}z_0 \in t^{-1}G_1$, we see that $l = k$ and that $\eta(i) = -\varepsilon(i)$ for each i. Moreover, if $\varepsilon(i) = -1 = \varepsilon(i+1)$ for some i, then $U = V$ implies that $g_i = z_0 h_i$ in G_2. But $\langle z_0\rangle * D = \langle z_0 h_i g_i^{-1}\rangle * D$, since $g_i, h_i \in D$. By Lemma 5.2, R_0 must be a cyclic conjugate of $(z_0 h_i g_i^{-1})^{\pm 1}$. Hence R is a cyclic conjugate of xyd with $d = h_i g_i^{-1} \in D$.

A similar argument applies if $\varepsilon(i) = +1 = \varepsilon(i+1)$ for some i, or if $\varepsilon(1) = +1$ or $\varepsilon(k) = -1$.

We may therefore assume that $k = 2m$ is even and that $\varepsilon(i) = (-1)^i$ for all i. In this case, comparing U with V in $G_2 * \langle t\rangle$ shows that $g_i = h_i$ in D for even i, while $g_i = z_0 h_i z_0^{-1}$ in G_2 for odd i.

Hence $U = V \in D * x_0^{-1} I x_0$, where I is the intersection of D with $z_0 D z_0^{-1}$ in G_2. It follows that the intersection of $A * B$ with $B * C$ in G has the form

$$E *_D B = (D * x_0^{-1} I x_0) *_D B = B * x^{-1} I x.$$

This completes the proof. □

Proof of Theorem C. Let $G = A_\Lambda / \langle\langle R \rangle\rangle$ be a one-relator product of locally indicable groups as in the statement of the theorem. Clearly there is no loss of generality in assuming that R is a cyclically reduced word involving all the factors A_λ, $\lambda \in \Lambda$. It also follows from Brodskiĭ's Freiheitssatz (Theorem 2.1) that $M \cup N = \Lambda$ (for otherwise $A_{M \cup N}$ embeds in G and there is no exceptional intersection).

Hence we can express G in the form $(A * B * C)/\langle\langle R \rangle\rangle$, where $A = A_{M \setminus N}$, $B = A_{M \cap N}$, and $C = A_{N \setminus M}$. Let $u = v$ be a minimal intersection equation with $u \in (A * B) \setminus B$ and $v \in (B * C) \setminus B$.

By Corollary 4.2, there is a finitely generated subgroup F_0 of $A * B * C$, with $\beta_1(F_0) \leq 2$, that contains both uv^{-1} and a conjugate R' of R and such that $uv^{-1} = 1$ in $G_0 = F_0 / \langle\langle R' \rangle\rangle$.

By Lemma 5.1, F_0 has one of three possible forms, which we now consider separately.

(a) $F_0 = \langle uv^{-1} \rangle * D$ *with* $\beta_1(D) \leq 1$. Then R' (and hence also R) is conjugate to $(uv^{-1})^{\pm 1}$, by Lemma 5.2. By Lemma 5.3 we see that conclusion (1) of Theorem C holds with $X = \langle u \rangle$ and $Y = \langle v \rangle$.

(b) $F_0 = D * E$, *where* $u \in D \subset A * B$ *and* $v \in E \subset B * C$. In this case, neither D nor E is trivial, since $u \in D$ and $v \in E$. Since F_0 is finitely generated, so are D and E. Since they are also locally indicable, it follows that $\beta_1(D) \geq 1$ and $\beta_1(E) \geq 1$. But then $\beta_1(D) + \beta_1(E) = \beta_1(F_0) \leq 2$ and so we must have $\beta_1(D) = \beta_1(E) = 1$. By Lemma 5.4, we see that conclusion (1) of Theorem C holds with $X = D$ and $Y = E$.

(c) $F_0 = \langle xy \rangle * xDx^{-1}$, *where* $x \in (A * B) \setminus B$, $y \in (B * C) \setminus B$, $u \in xDx^{-1}$, $v \in y^{-1}Dy$, *and* $D \subset B$. By Lemma 5.5, either R' (and hence R) is conjugate to $(xyd)^{\pm 1}$ for some $d \in D$, or the intersection of $(A * B)$ with $(B * C)$ in G has the form $B * (x^{-1} I x)$, where $I \subset D$ is the intersection of D with $(xy)D(xy)^{-1}$ in $(\langle xy \rangle * D)/\langle\langle R' \rangle\rangle$. The first of these possibilities reduces to case (a), with (say) $u = x^{-1}$ and $v = yd$. The second gives conclusion (2) of Theorem C. □

Proof of Theorem D. A theorem of Brodskiĭ [2] states that, in a one-relator product $G = (X * Y)/\langle\langle R \rangle\rangle$ of locally indicable groups X and Y, the intersection $X \cap Y$ is cyclic—as is the intersection $X \cap g^{-1} X g$ for any $g \in G \setminus X$. The result now follows directly from Theorem C. □

6. Algorithms

We now restrict attention to one-relator groups in order to consider the problem of algorithmically recognizing and identifying any exceptional intersection of two Magnus subgroups. It is clear from Lemma 5.4 that a solution of this problem will include the ability to recognize the intersection of two cyclic Magnus subgroups $\langle x \rangle \cap \langle y \rangle$ in a two-generator, one-relator group $G = \langle x, y \mid R \rangle$, which will in particular be central. There is an algorithm due to Baumslag and Taylor [1] for computing the center of a one-relator group, but technically this is not exactly what we require. The Baumslag–Taylor algorithm will provide a generator for $Z(G)$ (which is always cyclic except in the case where G is free abelian of rank 2 [10]) as a word in x and y, but we need further to decide which powers of x and of y belong to the center. A slight variation of the Baumslag–Taylor algorithm does just that; we include it here for completeness.

LEMMA 6.1. *Let $G = \langle x, y \mid R(x, y)\rangle$ be a two-generator, one-relator group, and let M be the Magnus subgroup $M = \langle x\rangle$ of G. Then the intersection $M \cap Z(G)$ is algorithmically computable.*

Proof. By a theorem of Murasugi [10], $M \cap Z(G)$ is trivial unless there is a homomorphism $\varepsilon\colon G \to \mathbb{Z}$ that is injective on M. Put $t = |\varepsilon(x)| > 0$, and let $G_1 = \langle z, y \mid R(z^t, y)\rangle$ be the one-relator group obtained from G by adjoining a tth root to x:

$$G_1 = \langle z\rangle \underset{z^t=x}{*} G.$$

Then ε extends to an epimorphism $G_1 \to \mathbb{Z}$ (with $\varepsilon(z) = \pm 1$), and $\langle z\rangle \cap Z(G_1) = \langle x\rangle \cap Z(G)$. Hence, without loss of generality we may assume that our original homomorphism $\varepsilon\colon G \to \mathbb{Z}$ restricts to an isomorphism on M, so that $G = \mathrm{Ker}(\varepsilon) \rtimes M$. We may also assume without loss of generality that $\varepsilon(y) = 0$. (If not, use Tietze transformations to replace y by yx^s for $s = -\varepsilon(y) \in \mathbb{Z}$.)

Murasugi's theorem [10] tells us further that $Z(G)$ is trivial unless $F = \mathrm{Ker}(\varepsilon)$ is free and of finite rank. But standard rewriting techniques show that G is an HNN extension of a one-relator group $G_2 = \langle y_0, y_1, \ldots, y_k \mid R_2\rangle$, where the associated subgroups are the Magnus subgroups $M_1 = \langle y_0, y_1, \ldots, y_{k-1}\rangle$ and $M_2 = \langle y_1, \ldots, y_k\rangle$. We thus require that the inclusions $M_1, M_2 \to G_2$ be isomorphisms or (equivalently) that R_2 involve each of the letters y_0 and y_k exactly once.

The rewriting R_2 is algorithmically computable from R, so it is visibly checkable whether or not R_2 has the necessary form. If so, then we can rewrite the relation $R_2 = 1$ as $y_k = W = W(y_0, y_1, \ldots, y_{k-1})$. Then the automorphism θ of $F = M_1$ arising from conjugation by x in G is

$$y_0 \mapsto y_1 \mapsto \cdots \mapsto y_{k-1} \mapsto W. \tag{$*$}$$

Then $\langle x\rangle \cap Z(G)$ is nontrivial if and only if θ has finite order (m, say) in $\mathrm{Aut}(F)$, in which case $\langle x\rangle \cap Z(G) = \langle x^m\rangle$. But there is a computable upper bound for the order of a torsion element of $\mathrm{Aut}(F)$ [1; 14]: indeed, any finite subgroup of $\mathrm{Aut}(F)$ maps isomorphically (since F is torsion-free) to a subgroup of the same order in $\mathrm{Out}(F)$, and the order of such a subgroup is at most $2^n \cdot n!$.

The rule $(*)$ allows us to compute θ^m for all positive m up to this upper bound and hence to determine $\langle x\rangle \cap Z(G)$, as required. □

COROLLARY 6.2. *Let $G = \langle x, y \mid R\rangle$ be a two-generator, one-relator group. Then the intersection of the Magnus subgroups $\langle x\rangle$ and $\langle y\rangle$ in G is algorithmically computable.*

Proof. Clearly, we may ignore the case where $G/[G,G]$ is free abelian of rank 2, so $Z(G)$ is cyclic by Murasugi's theorem [10]. As remarked previously, the intersection of $\langle x\rangle$ and $\langle y\rangle$ in G is central. By Lemma 6.1 we can compute natural numbers m,n such that $\langle x\rangle \cap Z(G) = \langle x^m\rangle$ and $\langle y\rangle \cap Z(G) = \langle y^n\rangle$. If $mn = 0$ then $\langle x\rangle \cap \langle y\rangle$ is trivial. Otherwise, by [10] there is an epimorphism $\varepsilon\colon G \to \mathbb{Z}$ that is injective on $Z(G)$ and hence also on $\langle x\rangle$ and on $\langle y\rangle$.

This epimorphism is unique up to multiplication by ± 1 in $\mathbb{Z}$, since $G/[G,G]$ has torsion-free rank 1, so it can be computed. In particular we can compute $\varepsilon(x) = a$ and $\varepsilon(y) = b$. Note that, for $s,t \in \mathbb{Z}$, $x^s = y^t$ if and only if $s \in m\mathbb{Z}$, $t \in n\mathbb{Z}$, and $as = bt$. Thus $\langle x\rangle \cap \langle y\rangle = \langle x^{mk}\rangle$, where k is the least positive integer for which amk is divisible by bn (which is clearly computable). □

We are now ready to prove the main algorithmic theorem.

Proof of Theorem E. By Corollary 4.2 and Lemma 5.1, the intersection can contain exceptional elements only if R is contained in a rank-2 subgroup $F_0 = \langle x,y\rangle$ of

$$F = \langle a_1,\dots,a_k,b_1,\dots,b_l,c_1,\dots,c_m\rangle$$

such that either:

(a) $x \in A * B$ and $y \in B * C$; or
(b) $x \in B$ and $y \in (A * B)\cdot(B * C)$.

Let X be the single-vertex graph with an edge for each of the generators $a_1,\dots,a_k$, $b_1,\dots,b_l,c_1,\dots,c_m$ of F (so that $F = \pi_1 X$). Then each two-generator subgroup F_0 arises as $\pi_1(X_0,v)$ for some finite graph X_0 (with $\beta_1(X_0) \le 2$) that admits an immersion $\eta\colon X_0 \to X$ and for some vertex v of X_0 [13].

Without loss of generality, we may assume that $R = \eta_*(R')$ for some closed path R' in X_0, based at v, that involves *all* the edges of X_0. This provides an upper bound on the number of edges of X_0 and hence a finite list of candidates for X_0, v, and η.

For each (X_0,v,η) in our finite list, we can check for closed paths $x',y' \in \pi_1(X_0,v)$ such that x',y' generate $\pi_1(X_0,v)$ and such that $x = \eta(x')$ and $y = \eta(y')$ satisfy conditions (a) or (b). Note that, since $R = \eta(R')$ is a cyclically reduced word involving all the generators of F, the path x' must represent a primitive element of $\pi_1(X)$ and must omit at least one edge of X_0 that occurs in R'. In case (a), a similar argument applies to y'; in case (b), we may deduce the same conclusion by choosing y' of minimal length in the double coset $\langle x'\rangle\cdot y'\cdot\langle x'\rangle$. This gives an upper bound (of twice the number of edges in X_0) for the lengths of x' and y', so again we are reduced to checking a finite list of potential candidates.

Given a choice of x,y satisfying (a), Lemma 5.3 tells us that $(A*B)\cap(B*C) = B * I$, where I is the intersection of $\langle x\rangle$ and $\langle y\rangle$ in the one-relator group $G_0 := \langle x,y \mid R(x,y)\rangle$, which is computable by Corollary 6.2.

Hence, for the rest of this proof we may assume that case (a) does not occur for any of our finite list of potential candidate triples (X_0,v,η). We may assume that, for some (X_0,v,η), there is a pair x',y' of paths giving rise to $x,y \in F$ satisfying the conditions of case (b). Then, by Lemma 5.5, the intersection $(A*B)\cap(B*C)$ has the form $B*I$, where I is conjugate to the intersection of $\langle x\rangle$ and $\langle yxy^{-1}\rangle$ in G_0.

Suppose that y occurs in G_0 with exponent sum zero. Then the standard Magnus rewriting method allows us to express some conjugate of R as a cyclically reduced word R_1 in $\{x_i;\ i=0,\dots,k\}$ for some $k\geq 1$, where $x_i=y^ixy^{-i}$. By choosing k as small as possible, we can also assume that both x_0 and x_k occur in R_1, in which case G_0 is an HNN extension of the one-relator group $G_1=\langle x_0,\dots,x_k \mid R_1\rangle$. Note that $k>1$, for otherwise we have $R_1\in\langle x_0,x_1\rangle$. Since $y=U\cdot V$ with $U\in(A*B)$ and $V\in(B*C)$, it follows that R is conjugate to a word in $\langle U^{-1}xU\rangle*\langle VxV^{-1}\rangle$ with $U^{-1}xU\in A*B$ and $VxV^{-1}\in B*C$. This falls into case (a), contrary to assumption.

Moreover, the intersection of $\langle x\rangle$ with $\langle yxy^{-1}\rangle$ in G_0 is equal to the intersection of $\langle x_0\rangle$ and $\langle x_1\rangle$ in G_1. This intersection is trivial, since $R_1\notin\langle x_0,x_1\rangle$.

Hence we may assume that the exponent sum of y in R is nonzero. There is a unique epimorphism $\varepsilon\colon G_0\to\mathbb{Z}$ (up to composition with $\pm\mathrm{id}\colon\mathbb{Z}\to\mathbb{Z}$), and $\varepsilon(x)\neq 0$. Now note that, if $x^m=yx^ny^{-1}$ for some m,n, then

$$m\varepsilon(x)=\varepsilon(x^m)=\varepsilon(yx^ny^{-1})=n\varepsilon(x)$$

and hence we have $m=n$ and $x^m\in Z(G_0)$.

Conversely, if $x^m\in Z(G_0)$ then $x^m=yx^my^{-1}$, so $x^m\in\langle x\rangle\cap\langle yxy^{-1}\rangle$. Thus $\langle x\rangle\cap\langle yxy^{-1}\rangle=\langle x\rangle\cap Z(G_0)$, which is computable by Lemma 6.1. $\square$

References

[1] G. Baumslag and T. Taylor, *The centre of groups with one defining relator,* Math. Ann. 175 (1968), 315–319.

[2] S. D. Brodskiĭ, *Anomalous products of locally indicable groups* (Russian), Algebraic systems, pp. 51–77, Ivanov. Gos. Univ., Ivanovo, 1981.

[3] ———, *Equations over groups, and groups with one defining relation* (Russian), Sibirsk. Mat. Zh. 25 (1984), 84–103.

[4] R. G. Burns and V. W. D. Hale, *A note on group rings of certain torsion-free groups,* Canad. Math. Bull. 15 (1972), 441–445.

[5] D. J. Collins, *Intersections of Magnus subgroups of one-relator groups,* Groups: Topological, combinatorial and arithmetic aspects (T. W. Müller, ed.), London Math. Soc. Lecture Note Ser., 311, pp. 255–296, Cambridge Univ. Press, Cambridge, 2004.

[6] J. Howie, *On locally indicable groups,* Math. Z. 180 (1982), 445–461.

[7] ———, *How to generalize one-relator group theory,* Combinatorial group theory and topology (Alta, 1984), Ann. of Math. Stud., 111, pp. 53–78, Princeton Univ. Press, Princeton, NJ, 1987.

[8] J. Howie and S. J. Pride, *A spelling theorem for staggered generalized 2-complexes, with applications,* Invent. Math. 76 (1984), 55–74.

[9] R. C. Lyndon and P. E. Schupp, *Combinatorial group theory,* Springer-Verlag, Berlin, 1977.

[10] K. Murasugi, *The center of a group with a single defining relation,* Math. Ann. 155 (1964), 246–251.

[11] B. B. Newman, *Some results on one-relator groups,* Bull. Amer. Math. Soc. 74 (1968), 568–571.
[12] J.-P. Serre, *Trees,* Springer-Verlag, Berlin, 1980.
[13] J. R. Stallings, *Topology of finite graphs,* Invent. Math. 71 (1983), 551–565.
[14] S. C. Wang and B. Zimmermann, *The maximum order of finite groups of outer automorphisms of free groups,* Math. Z. 216 (1994), 83–87.

Department of Mathematics
Heriot-Watt University
Edinburgh EH14 4AS
Scotland

J.Howie@hw.ac.uk