Computing subgroup presentations, using the coherence arguments of McCammond and Wise

Oliver Payne and Sarah Rees,
School of Mathematics and Statistics, University of Newcastle,
Newcastle NE1 7RU, UK, Sarah.Rees@ncl.ac.uk

January 27th 2006

Abstract

We describe an algorithm which may be used to compute a finite presentation of a finitely generated subgroup of a finitely presented group G, provided that G satisfies appropriate hypotheses. The algorithm is based on an algorithm of McCammond and Wise, but is extended to cover a wider class of groups, including all those satisfying the path reduction or weak 2-cell reduction hypotheses of McCammond and Wise. The proofs of correctness of our algorithm emerge from McCammond and Wise' own proofs that their hypotheses imply coherence of the groups satisfying them. We also demonstrate that the algorithm may be extended further to cover groups satisfying appropriate conditions on fans (strings of 2-cells) within disc diagrams. The core of this work originally appeared in the PhD thesis of the first author.

Dedicated to Charles Leedham-Green on the occasion of his 65th birthday

1 Introduction

Suppose that G is a finitely presented group and H a finitely generated subgroup of G. If H has finite index in G then H is finitely presented [19]. A finite presentation can be found either by the Reidemeister-Schreier process [21] or by a modification (see, for example [23]) of the Todd-Coxeter procedure; the latter process gives a presentation for H on the given generating set.

Since the Reidemeister-Schreier and modified Todd-Coxeter processes rely on enumeration of the cosets of H in G, as algorithms for computing finite presentations they are essentially restricted to the finite index case. But H may have a finite presentation even when its index is not finite. In particular, if G is free, a surface group, polycyclic or [22] the fundamental group of a 3-manifold, then all of its finitely generated subgroups are finitely presented and hence (by definition) G is coherent. Recently, coherence of other classes of groups has been investigated by various authors [5, 14, 15].

This article has been motivated by the results of McCammond and Wise in [14], where coherence is proved for groups for which an associated 2-complex satisfies various combinatorial conditions. Under these conditions it is proved that any finitely generated subgroup of the group under consideration can be found as the fundamental group of a finite 2-dimensional cell complex (or '2-complex'). Under the more restrictive of these conditions a clear algorithm to construct that 2-complex and hence to compute a subgroup presentation emerges.

We report on a generalisation of the algorithm of [14], so far recorded in the Ph.D. thesis [17]; the essence of the proofs of correctness of our algorithm under certain hypotheses emerges from proofs of coherence under the same hypotheses in [14]. Hence both the thesis [17] and this article draw very heavily on that work. A secondary aim of this article is to present the topological ideas of [14] and the earlier work on which that is based in a language which is more algebraic, combinatorial and algorithmic than topological, suitable for an audience of computational group theorists.

Section 2 of this article, which follows this introductory section, sets up the basic terminology and notation for the article, introducing the language of 2-complexes and maps between them, of homotopy and fundamental groups, and standard and Schreier complexes for a group. Section 3 introduces the notions of perimeter and of perimeter reducing moves which are central to the work of [14], and describes the basic algorithm A1 which McCammond and Wise prove terminates correctly for groups satisfying their 2-cell reduction hypothesis (stated in this paper as definition 3.1). Section 4 describes a more general algorithm A2, and proves it terminates correctly under the path reduction hypothesis and, more generally, the weak 2-cell reduction hypothesis of [14] (stated in this paper as definitions 4.1 and 4.2). We list a number of groups for which it is shown (in [14]) that the relevant hypotheses hold, and so to which the algorithm A2 can be applied. Section 5 is devoted to the description of the execution of the algorithm on a pair of examples. Section 6 contains a description of a further generalisation A3 of the algorithm A1, for which the basic moves are attachments of fans (particular sets of 2-cells) rather than of 2-cells as in A1 and A2. It also contains an outline of a proof of the correct termination of this algorithm under certain hypotheses which again rise from [14], where they were proved to imply coherence.

We observe that, although in this article we have restricted attention to the computation of subgroup presentations, there seems to be scope for further development, along similar lines. McCammond and Wise have certainly taken their techniques further, to prove the related property of local quasiconvexity (quasiconvexity of all finitely generated subgroups), and to investigate Howson's property (that the intersections of finitely generated subgroups are finitely generated, known, for example, to hold for free groups) and the generalised word problem. Other authors have developed similar techniques, to solve a range of algorithmic problems, particularly to deal with free groups, free products and related classes of groups, and have developed an automaton interpretation of the subgroup complexes; see for example [1, 3, 4, 7].

We should also comment that two further and quite distinct approaches to the computation of finite presentations of finitely generated subgroups of finitely groups have been developed in [9, 6, 11]; presentations are calculated in [9] given the existence of an automatic coset system for a group G relative to its subgroup H and in [6, 11] given the existence of a confluent rewrite system of G relative to H.

We are especially grateful to Jon McCammond for very useful discussion of the work in [14] on a number of occasions. We would also like to thank Keith Goda, Armando Martino, Andrew Duncan, Derek Holt and the referee for helpful comments.

2 Graphs and 2-complexes

In this article we take a rather combinatorial view of 2-complexes, slightly different from that of [14], but of course equivalent in the context in which we are working. This forces some changes in notation. Our notation also varies slightly from that of the thesis [17], in order to maintain consistency with the most recent version of [14], which uses in some place different notation from previous preprint versions.

Basically, we shall view a 2-complex as a directed, edge-labelled graph with 2-cells attached. Most (but not all) of the 2-complexes in this article will be finite. In this section we set up our notation for 2-complexes, and define the standard concepts of homotopy, fundamental groups, morphisms, immersions, covering maps and covers within that framework.

2.1 The basics

Let A be a finite set (which will always be the generating set of G). By a *directed graph over A* we mean a system of vertices, forming a set $\mathcal{V}$, and directed edges, forming a set $\mathcal{E}$. Each directed edge e has a *source* $s(e)$, a *terminus* $t(e)$, and *label* $l(e)$ taken from A. We allow an edge to be a loop, that is, to have its source and terminus equal. An edge is said to be *incident* with its source and terminus.

Let Γ be such a directed graph. For each $e \in \mathcal{E}$, we define its (formal) *reverse* e^{-1} (not itself an element of E) to have source $t(e)$, terminus $s(e)$, and label $l(e)^{-1}$. We define a *path* p in Γ to be a sequence $e_1, \ldots e_n$ of edges or reverses of edges with $t(e_i) = s(e_{i+1})$ for each i. The path p has source $s(e_1)$, terminus $t(e_n)$ and label the concatenation $l(p) = l(e_1) \ldots l(e_n)$. We say that it passes through the vertices $s(e_1), \ldots s(e_n), t(e_n)$. It is a closed path if $s(e_1) = t(e_n)$. Its *reverse* p^{-1} has source e_n, terminus e_1 and the formal inverse $l(e_n)^{-1} \ldots l(e_1)^{-1}$ of $l(p)$ as its label.

If Γ_1, Γ_2 are directed graphs over the same set A, with vertex sets $\mathcal{V}_1, \mathcal{V}_2$ and directed edge sets $\mathcal{E}_1, \mathcal{E}_2$, we define a *morphism* from Γ_1 to Γ_2 to be a map

$$\phi : (\mathcal{V}_1, \mathcal{E}_1) \to (\mathcal{V}_2, \mathcal{E}_2),$$

such that for all $e \in \mathcal{E}_1$,

$$s(\phi(e)) = \phi(s(e)),\ t(\phi(e)) = \phi(t(e)),\ l(\phi(e)) = l(e).$$

Given any word $w = x_1 \ldots x_n$ over A (that is a concatenation of symbols x_i from $A \cup A^{-1}$), we shall denote by *cycmin*(w) the lexicographically least word amongst all cyclic conjugates $x_i \ldots x_n x_1 \ldots x_{i-1}$ of w and all cyclic conjugates of w^{-1}. Then we define $C(w)$ to be the directed graph over A with n vertices $1, \ldots n$ and n edges, such that a single closed path labelled by *cycmin*(w) joins the n vertices in the order $1, 2, \ldots n, 1$. Note that $C(w) = C(w')$, for any cyclic permutation w' of w or of its inverse.

Now, given a directed graph Γ over A, we define a *2-cell* γ over Γ to be a morphism to Γ from some $C(w)$, where w is a cyclically reduced word over A (that is, a word $x_1 \ldots x_n$ for which $x_i \neq x_{i+1}^{-1}$ for any $i < n$, and $x_n \neq x_1^{-1}$). The image of the morphism is the subgraph induced by a closed path p in Γ and labelled by w; we say that γ is *attached* on p, and that p is its *boundary*. We also say that γ is incident with each of the edges and vertices in that path (we use the symbol I to denote incidence), and that it is labelled by w. Note that γ is also attached on the cyclic permutations and reverses of p, and further that if $w = u^k$ is a proper power, and u itself is not a power, then up to k distinct 2-cells may be attached on the same path p labelled by

w in Γ (less than k if the path is itself a proper power of a path). For each of the distinct vertices v in p from which a subpath labelled by u commences we can define a distinct morphism from $C(w)$ to Γ mapping the vertex of $C(w)$ labelled 1 to the vertex v, and $C(w)$ to the subgraph of Γ spanned by p.

Now we define a *2-dimensional cell complex* X to be a pair $(\Gamma, \mathcal{C}(X))$, where Γ is a directed edge-labelled graph, with vertex set $\mathcal{V}(X)$ and edge set $\mathcal{E}(X)$, and $\mathcal{C}(X)$ is a set of 2-cells over Γ. We call Γ the *1-skeleton* of X.

In fact, it seems that the definition of a 2-complex in [14] is slightly wider than the one we have just defined, and would allow any number of identical 2-cells to be attached on the same closed path in Γ. (In [14], where two identical copies of a 2-cell are attached on the same path the second one is called *redundant*.) So in theory $\mathcal{C}(X)$ could be not so much as a set as a multiset; in practice we shall not see redundant 2-cells in any of the 2-complexes we meet in our algorithms.

We also need the concept of a *side* of a 2-cell. Where γ is a 2-cell, we define a side of γ to be a pair $s = (\gamma, f)$ for which f is an edge of $C(w)$. Where e is an edge of X equal to $\gamma(f)$ or $\gamma(f^{-1})$, we say that e is *on* the side s, and that s is *through* e. We shall denote by $\mathcal{S}(\gamma)$ the set of sides of γ, by $\mathcal{S}(X)$ the set of all sides of 2-cells in a 2-complex X.

2.2 Weights

We may be given a specific *weight function* on X, that is a map $\mathrm{wt} : \mathcal{S}(X) \to \mathbb{N} \cup \{0\}$, with the property that (a) the sum of the weights of all sides through any one edge is finite and (b) the sum of the weights of all sides of any one 2-cell is strictly positive. If no weight function is given, for any complex relevant to this article (which will normally be finite) we can always define such a function, by specifying that every side has weight 1.

We extend the definition of wt as follows. Where e is an edge incident with γ, we define

$$\mathrm{wt}(\gamma, e) = \sum_{s=(f,\gamma), s\in\mathcal{S}(\gamma), \gamma(f)=e} \mathrm{wt}(s),$$

and in general, for a 2-cell γ of X,

$$\mathrm{wt}(\gamma) = \sum_{e\in\mathcal{E}(X), eI\gamma} \mathrm{wt}(\gamma, e),$$

for an edge e of X,

$$\mathrm{wt}(e) = \mathrm{wt}(e^{-1}) = \sum_{\gamma\in\mathcal{C}(X), eI\gamma} \mathrm{wt}(\gamma, e),$$

and for a path $p = e_1 \dots e_k$,

$$\mathrm{wt}(p) = \sum_{i=1}^{k} \mathrm{wt}(e_i).$$

2.3 Homotopy and fundamental groups

Let p be a path in a 2-complex X. We call a section of p consisting of an edge followed by or preceded by its reverse a *spur*. Given two paths p, q in a 2-complex X, such that $s(p) = s(q)$ and $t(p) = t(q)$ we say that p and q are elementary homotopic if either p is formed from q by the insertion or deletion of a spur, or p has the form $p_1p_2p_3$ and q the form $p_1p_4p_3$ where $p_1, \dots p_4$ are subpaths (some of which could be trivial), and some 2-cell is attached on $p_2p_4^{-1}$. We define p and q to be homotopic if there is a finite sequence of paths $q_0 = p, q_1, \dots q_r = q$, with q_i and q_{i+1} elementary homotopic for each i. Then we define the fundamental group $\pi_1(X, v)$ of X based at a vertex v to be the group of homotopy classes of closed paths with source v, composed by concatenation; if X is connected, and hence all such fundamental groups are isomorphic, we call this simply the fundamental group of X, written $\pi_1(X)$.

2.4 Standard complex and Schreier complex

Let $G = \langle A \mid R \rangle$ be a finitely presented group, with finite generating set A, and finite set of relators R. We assume that each element of R is freely and cyclically reduced, and that if $r \in R$ then R does not contain any cyclic conjugate of r or its inverse which is distinct from r. We define the *standard complex* $X(G)$ for G as follows. The 1-skeleton has a single vertex, v_0, and $|A|$ directed edges, each labelled by a distinct element a of A, and with v_0 as both source and terminus. Then $X(G)$ is formed from this by attaching one 2-cell labelled by r, for each $r \in R$. It has fundamental group G.

Now suppose that H is a finitely generated subgroup of G, generated (throughout this article) by a finite set $B = \{b_1, \dots b_m\}$ of words over A. We define the *Schreier complex* $X(G, H)$ as follows. Its 1-skeleton $\Gamma(G, H)$ is just the Schreier graph, or coset graph for H in G; that is, its vertices are in correspondence with the right cosets of H in G, with the vertex corresponding to a coset Hg joined by a directed edge labelled by a to the vertex corresponding to Hga, for each $a \in A$. Then, for each vertex v, for each relator r, a distinct 2-cell is attached along the unique path starting at v and labelled by r. The complex $X(G, H)$ has fundamental group H.

2.5 Morphisms, covers and immersions

Given two 2-dimensional cell complexes X_1, X_2 over the same alphabet, we define a *morphism* from X_1 to X_2 to be a map

$$\phi : (\Gamma(X_1), \mathcal{C}(X_1)) \to (\Gamma(X_2), \mathcal{C}(X_2))$$

which induces a graph morphism from $\Gamma(X_1)$ to $\Gamma(X_2)$, maps 2-cells of X_1 to 2-cells of X_2 (that is, whenever $\gamma : C(w) \to \Gamma(X_1)$ is a 2-cell of X_1, then the composite $\phi \circ \gamma : C(w) \to \Gamma(X_2)$ is a 2-cell of X_2), and preserves incidences between 2-cells, edges and vertices.

If v is a vertex in a 2-complex X, we define the *star* of v, $\mathrm{st}(v)$, to be the subset of $\mathcal{E}(X) \cup \mathcal{C}(X)$ consisting of all edges with source or terminus v and all 2-cells attached on a path through v.

Given 2-complexes X_1, X_2 over A, we define a morphism $\phi : X_1 \to X_2$ to be *locally injective* or an *immersion* if for any vertex v of X_1 the restriction of ϕ to $\mathrm{st}(v)$ is injective. We shall sometimes need the more general concepts of 1-immersion and near-immersion. We define a morphism $\phi : X_1 \to X_2$ to be a *1-immersion* if its restriction to the 1-skeleton of X_1 is locally injective as a graph morphism; in fact, a 1-immersion between 2-complexes without redundant 2-cells is necessarily an immersion [14, lemma 4.13]. We define a morphism $\phi : X_1 \to X_2$ to be a *near-immersion* if its restriction to the subsystem $(\mathcal{E}(X_1), \mathcal{C}(X_1))$ is locally injective, that is, for any edge y of Y the restriction of ϕ to the set of sides through y is injective [14, remark 2.9].

We define a morphism $\phi : X_1 \to X_2$ to be a *covering map* if the restriction of ϕ to $\mathrm{st}(v)$ maps $st(v)$ bijectively to $\mathrm{st}(\phi(v))$. If $\phi : X_1 \to X_2$ is a covering map then we say that the 2-complex X_1 is a *cover* for X_2.

It is well known that the Schreier complex $X(G, H)$ covers the standard complex $X(G)$, and that the Schreier complex for the trivial subgroup, which is known as the Cayley complex, and whose 1-skeleton is the Cayley graph, is a universal cover.

If H has finite index in G, then the Schreier complex is finite, and it is elementary to compute a presentation for H as its fundamental group. This article addresses the situation where H does not have finite index, and so the Schreier complex is infinite. Hence we search for another, finite, related complex with the same fundamental group. This might be a finite subcomplex of the Schreier complex [14, theorem 3.7], in which case it is called a *finite core* of that complex. But it might only map into the Schreier complex; for the fact that H is finitely presented does not force the Schreier complex to have a finite core [2, theorem 1], and even if it does, it may not be natural for an algorithm to construct precisely that [14, section 6].

3 Finding subgroup presentations

3.1 The basic process

The idea of finding a subgroup presentation by finding the subgroup as the fundamental group of a finite complex is not new. Schreier's original article [21] (and see also [27]) already interprets his construction topologically. More recently, the article [26] presents a graphical version of the Todd-Coxeter algorithm via an iterative construction of the Schreier graph $\Gamma(G, H)$. It is elementary to attach 2-cells to that complex and then compute a presentation for H as the fundamental group. But since the Schreier graph and associated 2-complex can only be finite when H has finite index in G, this method has to be modified in order to extend it beyond the finite index case.

The foundations for such an extension are laid down by Stallings in the article [25], which stresses the use of 1-immersions between graphs as maps which are not much weaker than covering maps, and sets up much of the machinery used in [14]. The basic method of [14] is to construct (or at least prove the existence of) a finite complex with fundamental group H by an iterative procedure which constructs a sequence of complexes $Y_0, Y_1, \ldots$ each equipped with a morphism $\phi_i : Y_i \to X(G)$, such that the image of $\pi_1(Y_i)$ in $\pi_1(G)$ induced by the morphism ϕ_i is the subgroup H. If the sequence $Y_0, Y_1, \ldots$ terminates, then it will do so in a complex $\hat{Y}$ equipped with a morphism $\hat{\phi} : \hat{Y} \to X(G)$ for which the map from $\pi_1(\hat{Y})$ to $\pi_1(X)$ induced by $\hat{\phi}$ is injective (that is, ϕ is π_1-injective). Hence H is found as the fundamental group of $\hat{Y}$.

The natural starting complex Y_0 for the iterative procedure described in the last paragraph is the following 2-complex, which we shall call $R(H, B)$ (where R stands for 'rose'). This complex has no 2-cells, so it is basically a graph. We construct it in 2 phases.

The first phase defines a central vertex v_0 and adjoins a closed path with label b, and source and terminus v_0, for each word b in the set B. There is then a natural morphism from this complex to $X(G)$ which maps every vertex to the single vertex of $X(G)$ and each edge to the unique edge of $X(G)$ with the same label. We shall call this complex $R(H, B)'$. Its fundamental group is free on B.

The second phase identifies certain pairs of edges of the graph just defined, by an operation called *folding*, which we now define.

Following [25], we define a pair of edges e, e' in a 2-complex X to be *admis-*

sible if e and e' have the same label and either the same source or the same terminus (possibly both), and say that X admits a *folding*. We define a folding of a complex X to be the natural morphism from X to the complex X' formed from X by identifying the edges e, e' in an admissible pair as well as $s(e)$ and $s(e')$ or $t(e)$ and $t(e')$, if these are distinct in X. Then X' has one fewer edge than X, and either one fewer vertex or the same number of vertices. Note that X and X' need not have the same fundamental group, for example, if e, e' have both the same source and the same terminus.

The 2-complex $R(H,B)$ is formed from $R(H,B)'$ as the image of a sequence of foldings, which successively identify each pair of admissible edges until there are no more admissible pairs. In general, the fundamental group of $R(H,B)$ is the same as that of $R(H,B)'$ (though it would not be if B contained some redundant elements such as words freely reducing to the trivial word, or words which could be expressed as products of powers of other elements) and still maps onto H within $\pi_1(X)$. But the map from $R(H,B)$ to $X(G)$ is now an immersion.

The basic step in the iterative procedure of [14] is, given a morphism $\phi : Y \to X$, to define a complex Y' and a morphism $\phi' : Y' \to X$, which is in some sense better than $\phi : Y \to X$. This is done by a combination of two basic operations to Y, folding and the addition of (sets of) 2-cells, basically the attachment of 2-cells to paths in Y for which more 2-cells are currently attached to $\phi(p)$ in X than are attached to p within Y. The value of the *perimeter* of ϕ is used to monitor the progress of this procedure towards a complex with fundamental group H.

3.2 Perimeter

We need to define perimeter and related concepts which arise from [14]. Here, partly because of our decision to use combinatorial rather than topological terminology, our notation is slightly different from that of [14]; in particular we have sometimes chosen to use the word 'weight' where 'perimeter' was used in [14]. Of course, where we quote results from [14], we have translated them to the notation of this article.

Suppose that X, Y are 2-complexes and that $\phi : Y \to X$ is a morphism from Y to X. Let wt be a weight function on X.

Now suppose that f is an edge of Y, and that $\phi(f) = e$. Suppose that s is a side through e. We that s is *present* at f if there is a side s' through f which maps under ϕ to s, and that s is *missing* at f is there is no such side s'. Then we define the *perimeter* of f, $\mathrm{per}(f)$, to be the sum of the weights of those sides through e which are missing at f; this is a non-negative integer. Further we define the perimeter of ϕ, $\mathrm{per}(\phi)$, to be the sum of all the perimeters of edges of Y; this is also a non-negative integer. It is sometimes also convenient to call this the perimeter of Y.

Provided that ϕ is a near-immersion we have an alternative but equivalent definition of perimeter, which may be more useful in practice. For we can extend wt to Y by defining, for any side s of Y, $\mathrm{wt}(s) = \mathrm{wt}(\phi(s))$. Then we have the equation

$$\mathrm{per}(f) = \mathrm{wt}(\phi(f)) - \mathrm{wt}(f)$$

In fact this equation holds precisely when ϕ is a near-immersion [14, remark 2.9]. Provided that Y is also finite, we also have the equation

$$\mathrm{per}(\phi) = \sum_{f \in \mathcal{E}(Y)} \mathrm{wt}(\phi(f)) - \sum_{\gamma \in \mathcal{C}(Y)} \mathrm{wt}(\phi(\gamma)) \tag{1}$$

(see [14, Lemma 2.18]).

We hope to improve our morphism $\phi : Y \to X$ by replacing it by a morphism $\phi' : Y' \to X$ with lower perimeter than ϕ. By [25, Section 3], any morphism $\phi : Y \to X$ can be written as a product of a sequence of foldings and a 1-immersion. It is easily seen that folding does not affect perimeter. Hence we need only consider how to improve $\phi : Y \to X$ in the case where ϕ is a 1-immersion.

The *perimeter reduction hypothesis* of [14, definition 3.6] is satisfied by a 2-complex X if whenever $\phi : Y \to X$ is a 1-immersion, with Y finite and connected, and ϕ not π_1-injective, then Y can be extended to Y^+ and ϕ to $\phi^+ : Y^+ \to X$, which induces the same image of $\pi_1(Y^+)$ in $\pi_1(X)$ as ϕ does of $\pi_1(Y)$ in $\pi_1(X)$, and with $\mathrm{per}(\phi^+) < \mathrm{per}(\phi)$.

For a 2-complex $X(G)$ satisfying this hypothesis, there must exist a sequence $Y_0 = R(H, B), Y_1, \ldots Y_i, \ldots$ (with Y_{i+1} constructed from Y_i using a combination of foldings and perimeter reducing moves) terminating in a finite 2-complex $\hat{Y}$ with H as its fundamental group. But it is not necessarily clear how to find this sequence. Hence G is proved coherent, but there need not be a clear algorithm to find a presentation for a subgroup. A number of groups are proved coherent in [14] by verification of the perimeter reduction hypothesis.

In order to find an actual algorithm, we need perimeter reducing moves to be visible. Stronger hypotheses on X will force this, as we shall see below.

3.3 Attaching 2-cells

The basic perimeter reducing move of most of [14] is *2-cell reduction*. (Later on in the article, the more general *fan reduction* is introduced.)

We need first to define some notions relevant to 2-cells.

If X is a 2-complex and γ is a 2-cell in it, then the *exponent* of γ is defined to be the largest integer n such that the boundary of γ is a concatenation q^n of n copies of a closed path q with itself.

Now suppose that $\phi : Y \to X$ is a morphism, and that γ is a 2-cell of X which is attached on a path p, of exponent n and labelled by w^n. Let p' be a path in Y such that $\phi(p') = p$. Let $\phi^{-1}(\gamma, p')$ be the set of 2-cells of Y which are attached on p' and map to γ under ϕ; call this the *fibre* of γ on p'. Note that every 2-cell in that fibre has the same exponent n'. We call the fibre a *packet* if it is non-empty and contains n/n' 2-cells. We say that the map ϕ is *packed* if every non-empty fibre of a 2-cell of X is a packet.

Note that, even when the fibre of γ on p' is non-empty and every 2-cell in it has the same weight as γ, for any edge f of p', the inequality

$$\sum_{\gamma' \in \phi^{-1}(\gamma, p')} \mathrm{wt}(\gamma', f) \leq \mathrm{wt}(\gamma, \phi(f))$$

holds, with equality if and only if the fibre is a packet. So if ϕ is not packed, Y must contain edges with strictly positive perimeter.

Now suppose that $\phi : Y \to X$ is a packed 1-immersion.

Suppose that p is a closed path in X along which a 2-cell γ is attached, that q is a path in Y, mapping under ϕ to a subpath of p, and maximal in the sense that no path properly containing q maps to a subpath of p. If $\phi(q) = p$, then suppose also that $\phi^{-1}(\gamma, q)$ is empty. Then q is called a *2-cell attachment site*.

There is now a unique way to extend Y to a 2-complex Y^+ and ϕ to a packed 1-immersion $\phi^+ : Y^+ \to X$ such that the following conditions are satisfied:

(a) Y^+ contains a packet of 2-cells attached on a path q^+, which contains q as a subpath and maps under ϕ^+ to p. The packet maps under ϕ^+ to γ.

(b) the only edges in Y^+ not in Y are those in the path q^+ not also in q.

(c) the only 2-cells in Y^+ not in Y are those in $(\phi^+)^{-1}(\gamma, q^+)$.

(d) the restriction of ϕ^+ to Y is ϕ.

We shall call an extension of this type a *packet attachment on q over γ*. The attachment is called *complete* if $\phi(q) = p$, *incomplete* otherwise.

Note that in the case where $\phi(q)$ is the whole of p, it is possible that q might not be closed in Y (although its image under ϕ is closed). In that case the source and terminus of q, though distinct in Y, are identified in Y^+, and hence Y does not actually embed as a subcomplex of Y^+.

Where a packet attachment is complete, the weight of the packet which is attached is $\frac{n}{n'}\mathrm{wt}(\gamma)$, and hence the perimeter of ϕ^+ is equal to

$$\mathrm{per}(\phi) - \frac{n}{n'}\mathrm{wt}(\gamma).$$

This is certainly less than $\mathrm{per}(\phi)$ since every 2-cell has positive weight.

Given an incomplete packet attachment, let s be the complement of q in q^+ (which we shall call the *extending path*). Then each 2-cell in the packet has exponent 1 (that is, $n' = 1$), and hence the total weight of the packet is $n\mathrm{wt}(\gamma)$. Then, since Y^+ is formed from Y by adding the edges of the path s together with the 2-cells in the packet $\phi^{-1}(\gamma, q^+)$, the perimeter of ϕ^+ is equal to

$$\mathrm{per}(\phi) + \mathrm{wt}(\phi^+(s)) - \exp(\gamma)\mathrm{wt}(\gamma)$$

(see equation (1) in section 3.2).

Hence a packet attachment extends ϕ to ϕ^+ with perimeter strictly less than that of ϕ provided that either the attachment is complete or, where s is the extending path,

$$\mathrm{wt}(\phi^+(s)) < \exp(\gamma)\mathrm{wt}(\gamma).$$

A packet attachment which strictly reduces perimeter is called a *2-cell perimeter reduction*, or simply *2-cell reduction*.

An incomplete packet attachment with extending path s for which

$$\mathrm{wt}(\phi^+(s)) \leq \exp(\gamma)\mathrm{wt}(\gamma)$$

is called a *weak 2-cell (perimeter) reduction*

3.4 The 2-cell reduction hypothesis and an algorithm

Definition 3.1 (2-cell reduction hypothesis [14]) *A 2-complex X satisfies the* 2-cell reduction hypothesis *if for any 2-complex Y and any map $\phi : Y \to X$ such that ϕ is a packed 1-immersion but not a π_1-injection, there* *is some 2-cell γ in X, and some path q in Y such that the attachment of a packet over γ on q is a 2-cell perimeter reduction.*

Given the 2-cell reduction hypothesis, the following algorithm, described in [14, theorem 6.1], terminates and computes a presentation for H.

Algorithm A1

Input : a finite presentation for G, a finite set of generators for H.

Initially :

$X = X(G), Y = R(H, B)$, $\phi : Y \to X$ (this is uniquely defined as a morphism of 2-complexes), halt:= false;

repeat :

if Y admits a 2-cell reduction, first select one (any one) and apply it, then apply any foldings admitted by the resulting complex, until no more are admitted. Reset Y to be the resultant complex, and modify ϕ accordingly.

else halt:=true;

while halt=false;

Output : a finite presentation for $\pi_1(Y)$.

Note that since the initial morphism $\phi : R(H, B) \to X(G)$ is certainly packed, and since when $\phi : Y \to X$ is packed so also is $\phi' : Y' \to X$, where Y' is the result of either applying a 2-cell reduction or a folding to Y, then the map $\phi : Y \to X$ remains packed at the end of each step of the algorithm.

Termination is guaranteed by the fact that each step of the procedure reduces the value of the pair of non-negative integers $(\text{per}(Y), |\mathcal{E}(Y)|)$ with respect to the (lexicographical) ordering on $(\mathbb{N} \cup \{0\})^2$ which sets (i, j) less than (i', j') if either $i < i'$ or $i = i'$ and $j < j'$.

However the 2-cell reduction hypothesis is hard to satisfy.

On the other hand, a range of groups satisfy the weaker *path reduction hypothesis* (definition 4.1 below), which also implies coherence [14, theorem 7.6]. Section 4 is devoted to the description of a generalisation of the algorithm A1 to an algorithm A2 which accommodates weak 2-cell reductions. It is proved that A2 terminates correctly when either the path reduction hypothesis or the *weak 2-cell reduction hypothesis* (definition 4.2 below) holds.

The algorithm A2 is described in the PhD thesis [17], and implemented in the computer algebra system GAP [20].

A further generalisation to accommodate, under certain hypotheses, more general *fan reductions* is described briefly in section 6.

4 Modifying the algorithm to accommodate weak 2-cell reductions

4.1 The hypotheses under which we shall work

We shall consider complexes satisfying either the path reduction hypothesis or the weak 2-cell reduction hypothesis.

Definition 4.1 (Path reduction hypothesis [14]) *A 2-complex X satisfies the* path reduction hypothesis *if, given any non-trivial closed null-homotopic path p in X, there is a sequence of closed paths $p_1, \dots p_t$, such that $p_1 = p$, the path p_t is trivial, and, for each i, the path p_{i+1} is obtained from p_i either by replacing a subpath q of p_i by the reverse s^{-1} of its complement s within the boundary of a 2-cell γ, where*

$$\mathrm{wt}(s) \leq \exp(\gamma)\mathrm{wt}(\gamma),$$

or by the removal of a backtrack.

Certain small cancellation conditions (see [12] for a definition of these) together with a condition on the weight function imply satisfaction of the path reduction hypothesis. In particular, we have [14, theorem 9.6]:

let X be a 2-complex with a weight function, satisfying C(6)-T(3) (respectively C(4)-T(4)). Suppose that whenever γ is a 2-cell, and s a path within the boundary of γ which is the concatenation of at most 3 (respectively 2) consecutive pieces, then $\mathrm{wt}(s) \leq \exp(\gamma)\mathrm{wt}(\gamma)$. Then X satisfies the path reduction hypothesis.

In particular [17], if the weight condition in the hypotheses of [14, theorem 9.6] is strict, then G satisfies a strict form of the path reduction hypothesis. In this case the 2-cell reduction hypothesis (definition 3.1) is satisfied and the algorithm A1 can be applied.

The weak 2-cell reduction hypothesis is a weakening of the 2-cell reduction hypothesis (definition 3.1) to allow weak 2-cell reductions.

Definition 4.2 (Weak 2-cell reduction hypothesis [14]) *A 2-complex X satisfies the* weak 2-cell reduction hypothesis *if for any 2-complex Y and any map $\phi : Y \to X$ such that ϕ is a packed 1-immersion but not a π_1-injection, there is some 2-cell γ in X, and some path q in Y such that the attachment of a packet over γ on q is a weak 2-cell perimeter reduction.*

In fact the path reduction hypothesis itself implies the weak 2-cell reduction hypothesis. This implication follows from proposition 4.9 below, and a modification of the proof gives the strict version of that result referred to above (both versions appeared in an early preprint of [14]). So our theorem 4.10 below is in fact a special case of theorem 4.12, and it is not necessary to prove both. However we choose to do so, partly because the essential part of the proof of theorem 4.10, namely proposition 4.9, is constructive, partly because it is this proposition which needs to be generalised in order to extend the algorithm to include fan reductions as in A3.

We note also that the weak 2-cell reduction hypothesis is not clearly implied by easy to check criteria, other than those which imply the path reduction hypothesis.

4.2 Some groups known to satisfy the path reduction hypothesis.

For convenience, we reproduce from [14] the following list of examples, which are proved to satisfy the criteria of theorem 9.6 of [14] and hence the path reduction hypothesis.

$\langle a_1, \dots a_r \mid W^n \rangle$, where W is cyclically reduced and $|W| \leq n + 1$. [14, theorem 8.3]

$\langle a_1, \dots a_r \mid W^n \rangle$, where W is cyclically reduced, and a_1 occurs k times in W, where $n \geq k > 0$, alternatively where each a_i occurs in W and $n \geq |W|/r$. [14, theorem 8.4]

$\langle a, b, c, d, e \mid (abcde)a(abcde)b(abcde)c(abcde)d(abcde)e \rangle$. [14, example 9.7]

$< a_1, \dots a_g \mid a_1^2 a_2^2 \dots a_g^2 \rangle$. [14, example 13.2]

Any $\langle a_1, \dots \mid R_1, \dots \rangle$ satisfying the small cancellation condition $C'(1/n)$ and such that each a_i occurs at most $n/3$ times among the R_j. [14, theorem 13.3]

Any $\langle a_1, \dots \mid W_1^{n_1}, \dots \rangle$ for which each W_i is cyclically reduced and not a proper power, W_i is not freely conjugate to W_j for $i \neq j$, and for each

i, $n_i \geq 6\frac{|W_{max}|}{|W_{min}|}\Sigma|W_i|$, [14, theorem 13.7],
in particular $\langle a, b \mid a^N, (aaba^{-1}b^{-1})^{N+1}, b^N, (bbab^{-1}a^{-1})^{N+1}\rangle$, [14, corollary 13.8].

4.3 Constructing and verifying the algorithm

From now on we assume that every side of $X(G)$ has a strictly positive weight.

As in A1, the basic step in the procedure will be the attachment of a single packet of 2-cells followed by as many foldings as are admitted by the resulting complex. We shall always demand that an attachment does not increase perimeter, and then call it a *weak 2-cell reduction*. If an attachment strictly reduces perimeter, we call it a *strict 2-cell reduction*, and if an attachment is complete, we call it a *complete 2-cell reduction*. A complete 2-cell reduction is always strict.

Note that if $\phi : Y \to X(G)$ is a 1-immersion, then the application of a complete 2-cell reduction to Y may produce a complex Y' which admits foldings (in the case where a 2-cell is attached along a path which is not closed in Y although its image in $X(G)$ is). However if an incomplete 2-cell reduction is applied, the result is a complex Y' which admits no foldings, and in which Y embeds as a subcomplex.

In our algorithm we need to address the question of how to select the right 2-cell reduction to apply at each stage, and how to recognise termination. Failure to do this correctly could lead us into an infinite chain of reductions. (an example of this is given as [14, example (7.1)]).

The algorithm will essentially take the form

Algorithm A2: outline

Input : a finite presentation for G, a finite set of generators for H.

Initially : halt:= false, $X = X(G), Y = R(H, B), \phi : Y \to X$;

repeat :

> **If** Y admits a finite sequence of weak 2-cell reductions the last of which is complete, select one such sequence and apply it. Fold the resulting complex until no more foldings are admitted, and set Y to be the result.
>
> **else** halt:=true;

while halt=false;

Output : a finite presentation for $\pi_1(Y)$.

In order to identify an appropriate finite sequence of reductions, we examine the *2-cell reduction tree* $T(Y)$ for Y.

We define $T(Y)$ to be a tree whose vertices are labelled by 2-complexes and whose edges are labelled by weak 2-cell reductions, as follows. The root is labelled by Y itself. Where a vertex v is labelled by a 2-complex Y_v, the directed edges out of v are labelled by the various weak 2-cell reductions which can be applied to Y_v, and the terminus of an edge e out of v is labelled by the 2-complex which is obtained from Y_v by applying the reduction by which e is labelled (followed by all possible foldings). Labels are attached to vertices and edges, but do not define them; clearly the same 2-complex may label more than one vertex.

We note and we shall use the following, which is straightforward to prove.

Lemma 4.3 *The 2-cell reduction tree associated with a finite complex Y has finite degree.*

Nonetheless the tree $T(Y)$ will normally have infinite depth and so be infinite. But it turns out that we need only search a finite portion of it. We are able to restrict our search by examining the *active graphs* of the 2-complexes which label its vertices.

Given a 2-complex Y' and a 1-immersion $\phi' : Y' \to X$, we call an edge of the 1-skeleton of Y' *active* if it has strictly positive perimeter, *inactive* otherwise. We define the *active graph* of Y' to be the subgraph of the 1-skeleton of Y' which consists of all active edges and all vertices incident with those. Each edge is marked by its label as an edge of the 1-skeleton of Y' and also by the sides of the 2-cells which are already attached there. We define the perimeter of the active graph to be the sum of perimeters of its edges; this is then equal to the perimeter of the 2-complex Y'. We define the active graphs of two distinct complexes to be isomorphic as active graphs if they are isomorphic as graphs by an isomorphism which respects the edge markings. Of course, isomorphic active graphs must have the same perimeter.

We shall often (for convenience) abuse notation and refer to an active graph rather than the 2-complex to which it is associated as the label of a vertex of a reduction tree.

Note that a reduction of a 2-complex Y can be completely described by its effect on the active graph of Y. Hence we have the following lemma

Lemma 4.4 *If Y_1, Y_2 are 2-complexes whose active graphs are isomorphic then the same sequences of reductions may be applied to both.*

Now suppose that $\phi : Y_1 \to X$ is a 1-immersion and that r is a 2-cell reduction taking Y_1 to Y_2. Then certainly the active graph of Y_2 has perimeter no higher than the active graph of Y_1. Hence we can bound the number of distinct active graphs which can be associated with 2-complexes in the 2-cell reduction tree $T(Y)$ by a (computable) integer $N(Y)$.

Then we can deduce

Lemma 4.5 *Any path of length greater than $N(Y)$ from the root down the 2-cell reduction tree $T(Y)$ must contain two vertices labelled by complexes with isomorphic active graphs.*

We define a non-root vertex of $T(Y)$ to be *repeating* if it is the second of two vertices on a path from the root with isomorphic active graphs.

For any positive integer k, we define $T(Y)_k$ to be the full subtree of $T(Y)$ whose vertices are all those vertices within distance k of the root. We define the core $\overline{T(Y)}$ of $T(Y)$ to be the full subtree of $T(Y)_{N(Y)}$ which excludes any repeating vertex as well as any vertex for which the path from the root to that vertex includes a repeating vertex.

The following result is crucial for our algorithm.

Proposition 4.6 **(a)** *Any active graph which labels a vertex of $T(Y)$ must label a vertex in the core of $T(Y)$.*

(b) *Any weak 2-cell reduction which labels an edge in $T(Y)$ must also label an edge in its core. In particular*

- **(i)** *if $T(Y)$ contains an edge labelled by a strict 2-cell reduction then so does its core.*
- **(ii)** *if $T(Y)$ contains an edge labelled by a complete 2-cell reduction, then so does its core.*

Proof: To prove part (a), suppose that Γ is an active graph which labels a vertex in $T(Y)$ but no vertex within the core. Let v be a vertex at minimal distance from the root of $T(Y)$ which is labelled by Γ. Then there must be two vertices v_1, v_2 on the path from the root to v which are labelled by the same active graph. If v_1 is the higher of those two vertices, and s is the length of the path from v_2 to v, then a chain of r reductions transforms the complex labelling v_2 to the complex labelling v, and hence the same r reductions transform the complex labelling v_1 to a complex with the same active graph as the complex labelling v. This is closer to the root than v, and hence we have a contradiction. The first sentence of part (b) now follows immediately from lemma 4.4. Parts (i) and (ii) are direct consequences of this (but we state them separately, because they are useful). $\square$

Now in more detail, our algorithm can be defined by

Algorithm A2

Input : a finite presentation for G, a finite set of generators for H.

Initially : halt:= false, $X = X(G), Y = R(H, B), \phi : Y \to X$;

repeat : Simultaneously construct and traverse $\overline{T(Y)}$ in a depth-first fashion.

> **If** an edge is found labelled by a complete 2-cell reduction , apply to Y the sequence of 2-cell reductions corresponding to the path from the root up to and including that edge. Then fold the resulting complex until no more foldings are admitted, and replace Y by the folded complex.
>
> **else** halt := true;

while halt=false;

Output : a finite presentation for $\pi_1(Y)$.

Of course we halt the traversal of the tree $\overline{T(Y)}$ as soon as we have found an appropriate sequence of reductions, and so only traverse the whole tree if no appropriate sequence can be found.

We see easily that

Proposition 4.7 *The algorithm A2 terminates in finite time.*

Proof: The tree $\overline{T(Y)}$ is finite, of finite degree (by lemma 4.3) and finite depth at most $N(Y)$, and so can be traversed in finite time. If a reduction is found it reduces the perimeter of Y, and hence the algorithm terminates after at most $\mathrm{per}(R(H, B))$ steps. $\square$

Correctness of the algorithm is slightly more complicated to prove, and it is for this that we need some hypotheses on $X(G)$.

Where $X(G)$ satisfies the path reduction hypothesis (definition 4.1), the meat of our proof of correctness, namely lemma 4.8 and proposition 4.9, is extracted from the proof of coherence of G under that hypothesis, given in [14, theorem 7.6].

We need the following basic fact.

Lemma 4.8 *Suppose that a weak 2-cell reduction reduces Y to Y'. If Y contains a loop which is not null-homotopic but whose image in Y' is null-homotopic, then the reduction must be complete.*

PROOF: An incomplete 2-cell reduction adds a single loop to Y together with a 2-cell which makes that loop null-homotopic. Hence it does not change the fundamental group. $\square$

For complexes satisfying the path reduction hypothesis we have the following.

Proposition 4.9 *Suppose that X satisfies the path reduction hypothesis, and that $\phi : Y \to X$ is a packed 1-immersion, but is not π_1-injective. Then there exists a finite sequence of weak 2-cell reductions and foldings which may be applied to Y, such that one of the weak 2-cell reductions is complete.*

PROOF: Since ϕ is not π_1-injective, Y must contain a loop p which is not null-homotopic but such that $\phi(p)$ is null-homotopic. The path reduction hypothesis ensures the existence of a sequence of closed loops $p'_1 = \phi(p), p'_2, \ldots p'_k$ in X, such that for each i, p'_{i+1} is obtained from p'_i either by the removal of a backtrack or by replacing a section q'_i of p'_i by the reverse of a path s'_i such that $q'_i s'_i$ is the boundary of a 2-cell γ_i and $\mathrm{wt}(s'_i) \leq \exp(\gamma_i)\mathrm{wt}(\gamma_i)$.

Now we define successively 2-complexes $Y_1 = Y, Y_2, \ldots Y_k$, and morphisms $\phi_1 = \phi, \phi_2, \ldots \phi_k$ from these to X, where each Y_i is obtained from the previous one either by a fold or by a weak 2-cell reduction, as follows.

Where p'_{i+1} is formed from p'_i by dragging across a 2-cell γ_i, and p_i is a loop in Y satisfying $\phi_i(p_i) = p'_i$, we form Y_{i+1} and $\phi_{i+1} : Y_{i+1} \to X$ by attaching a 2-cell over γ_i on a loop which intersects p_i in a path q_i such that $\phi_i(q_i) = q'_i$. The weight condition ensures that this is a weak 2-cell reduction. We then define p_{i+1} to be the loop obtained by replacing the subpath q_i in p_i by the reverse of its complement in the boundary of the 2-cells in the attached packet.

Otherwise, where p'_{i+1} is formed from p'_i by the removal of a backtrack, we fold Y_i to get Y_{i+1}.

Since the image of p in Y_k is null-homotopic, but by lemma 4.8 a sequence of incomplete 2-cell reductions alone cannot change the fundamental group (and such reductions do not provoke foldings), one of the 2-cell reductions must be complete. $\square$

We are now ready to prove

Theorem 4.10 *Suppose that $X(G)$ satisfies the path reduction hypothesis (definition 4.1), and that all of its sides have strictly positive weights. Then A2 terminates with a finite presentation for H.*

PROOF: Since termination was proved in proposition 4.7, we merely need to prove correctness. So suppose that at some stage the algorithm has reached a complex Y and morphism $\phi : Y \to X$ which is not π_1-injective. It is clear that ϕ is a 1-immersion since foldings are always applied whenever possible. Then proposition 4.9 ensures that some edge in $T(Y)$ is labelled by a complete reduction. Proposition 4.6 ensures that the same reduction labels an edge in the core of $T(Y)$, and hence that the algorithm will find a reduction and continue on to the next step. $\square$

In order to prove the analogous theorem for a complex satisfying the weak 2-cell reduction hypothesis (definition 4.2), we need the following result, the basis of whose proof was explained to us recently by Jon McCammond.

Proposition 4.11 *[13] Suppose that X satisfies the weak 2-cell reduction hypothesis, and that $\phi : Y \to X$ is a packed 1-immersion, but is not π_1-injective. Then there exists a finite sequence of weak 2-cell reductions and foldings which may be applied to Y, such that one of the 2-cell reductions is complete.*

PROOF: Consider the 2-cell reduction tree $T(Y)$. Suppose that it contains no complete reductions. Then applying all possible weak 2-cell reductions in sequence gives an infinite sequence $Y = Y_1, Y_2, Y_3, \dots$ of 2-complexes, each embedding in its successor, and each mapping via a packed 1-immersion to X. We can form the limit of this sequence and hence obtain an infinite 2-complex $\hat{Y}$ which admits no 2-cell reduction, but maps via a packed 1-immersion to X. But the existence of the 2-complex violates the weak 2-cell reduction hypothesis, and hence we have a contradiction. $\square$

Hence we are also able to prove the following theorem.

Theorem 4.12 *Suppose that $X(G)$ satisfies the weak 2-cell reduction hypothesis (definition 4.2), and that all of its sides have strictly positive weights. Then A2 terminates with a finite presentation for H.*

The proof is just as for theorem 4.10, but with the use of proposition 4.9 replaced by the use of proposition 4.11.

Note that in our algorithm we have chosen to search for complete reductions. The algorithm might however find an incomplete strict reduction first. It is conceivable that it might run more quickly if it at any stage it applied the first sequence it could find containing a strict reduction, rather than necessarily wait for a complete reduction. Since every complete reduction is strict, the proof above remains valid if strict rather than complete reductions are used.

5 Some examples which display the operation of the algorithm

These examples appear in [17]. That thesis also describes the implementation of the algorithm A2 within GAP [20]. The implementation uses a version of the GRAPE package [24] which incorporates a subpackage dealing with homotopy invariants [18].

Our first example appears as [14, example 7.1], where it demonstrates the failure of the algorithm A1 (which loops indefinitely for this example), in the absence of the 2-cell reduction hypothesis. We set

$$G = \langle a, b \mid aba^{-1}b^{-1} \rangle$$

and $H = \langle ab^{-1} \rangle$. Of course H is infinite cyclic.

The standard complex $X(G)$ has one vertex, two edges, labelled a, b and a single 2-cell attached on $aba^{-1}b^{-1}$.

G satisfies the C(4)-T(4) small cancellation conditions. Now if every side of $X(G)$ is given unit weight, we see that each 2-cell of $X(G)$ has weight 4 and exponent 1, and each edge has weight 2. Hence the weight condition (for length 2 paths) of [14, theorem 9.6] is satisfied with equality, and hence G satisfies the path reduction hypothesis. Hence the algorithm A2 must terminate with a presentation for H.

The initial subgroup complex Y_0 has two vertices, joined by two edges labelled a, b, to give a simple circuit labelled ab^{-1}, as shown in figure 1, with no 2-cells attached.

The reduction tree $T(Y_0)$ is simply a binary tree.

The complex Y_0 (shown on the left of the top row of the figure) admits two possible reductions, since there are two ways to attach a single 2-cell. Both reductions preserve perimeter and the resulting complexes are isomorphic. One reduction extends Y_0 to the right, and the other to the left. Either reduction yields a 2-complex with 3 vertices, 4 edges and a single 2-cell, as shown on the left of the second row of the figure. Each of these complexes also admits two reductions, by attaching a 2-cell either to the left or the right, of which the resulting complexes are isomorphic. And so it continues. Whichever of the two possible reductions is chosen at each stage, we find a sequence of complexes with 1-skeletons as shown on the left of the figure.

The active graphs are shown on the right of the figure. The integer against each edge indicates which of the four sides of a 2-cell labelled by $aba^{-1}b^{-1}$ is attached there.

We see that after any three reductions to Y_0 the active graph is the same as it was after just two reductions. Hence we can deduce that any vertex at depth three in the reduction tree is repeating, and, since $T(Y_0)_3$ contains no complete reduction, there is no useful sequence of reductions to Y_0, and hence Y_0 itself has H as its fundamental group.

For our second example we consider the group

$$G = \langle a, b, c, d \mid (ab)^3, (bc)^3, (cd)^3, (ad)^3, (bd)^3, (ac)^3 \rangle$$

For this example too, G satisfies the small cancellation conditions $C(4)$ and $T(4)$. If each side of each 2-cell of the standard complex $X(G)$ is assigned weight 1, then each 2-cell has total weight 6 and exponent 3. Each of the edges has weight 9. Hence each path of length 2 has weight 18. Hence for this example too, $X(G)$ satisfies the appropriate weight condition of [14, theorem 9.6] with equality, and hence X satisfies the path reduction hypothesis and algorithm A2 can be used.

Now we let H be the subgroup

$$\begin{aligned} H = \langle dcb^{-1}d^{-1}, d^2, da^2d^{-1}, &daba^{-1}d^{-1}, \\ &daca^{-1}d^{-1}, db^2a^{-1}b^{-1}d^{-1}, dba^2b^{-1}d^{-1}, dbadb^{-1}d^{-1} \rangle, \end{aligned}$$

which has infinite index in G, as can be verified using, for instance, the KBMAG [10, 8] package within GAP.

Now let Y_0 be the initial subgroup complex $R(H, B)$. Its one skeleton is the top graph in Figure 2.

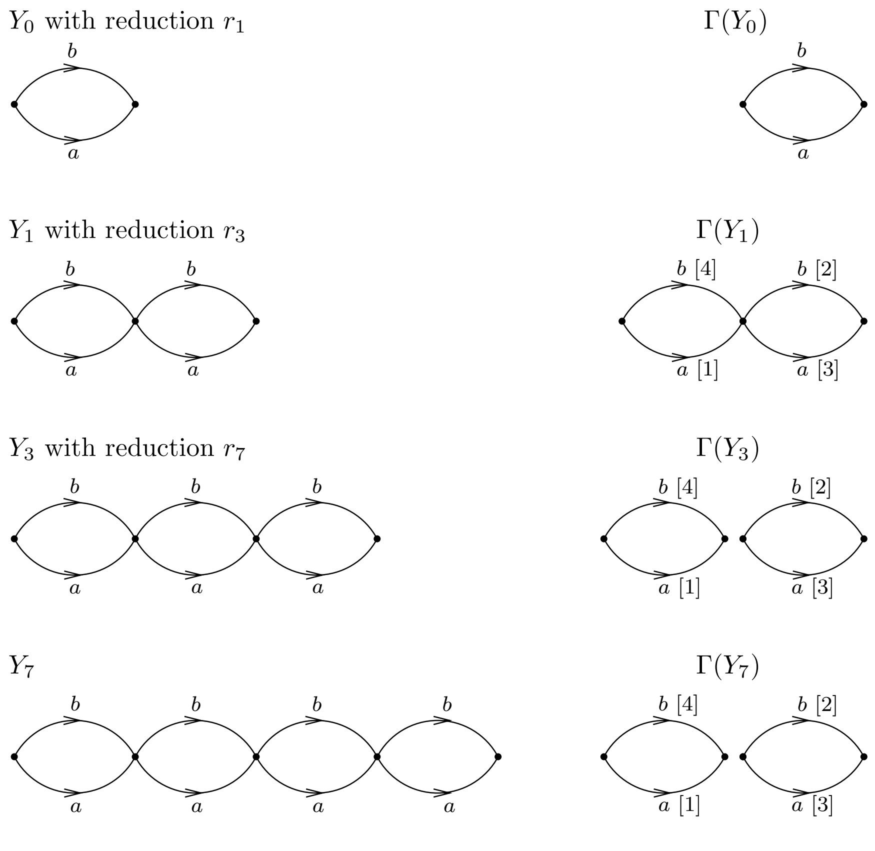

Figure 1: Subgroup Complexes with Reductions and Active Graphs

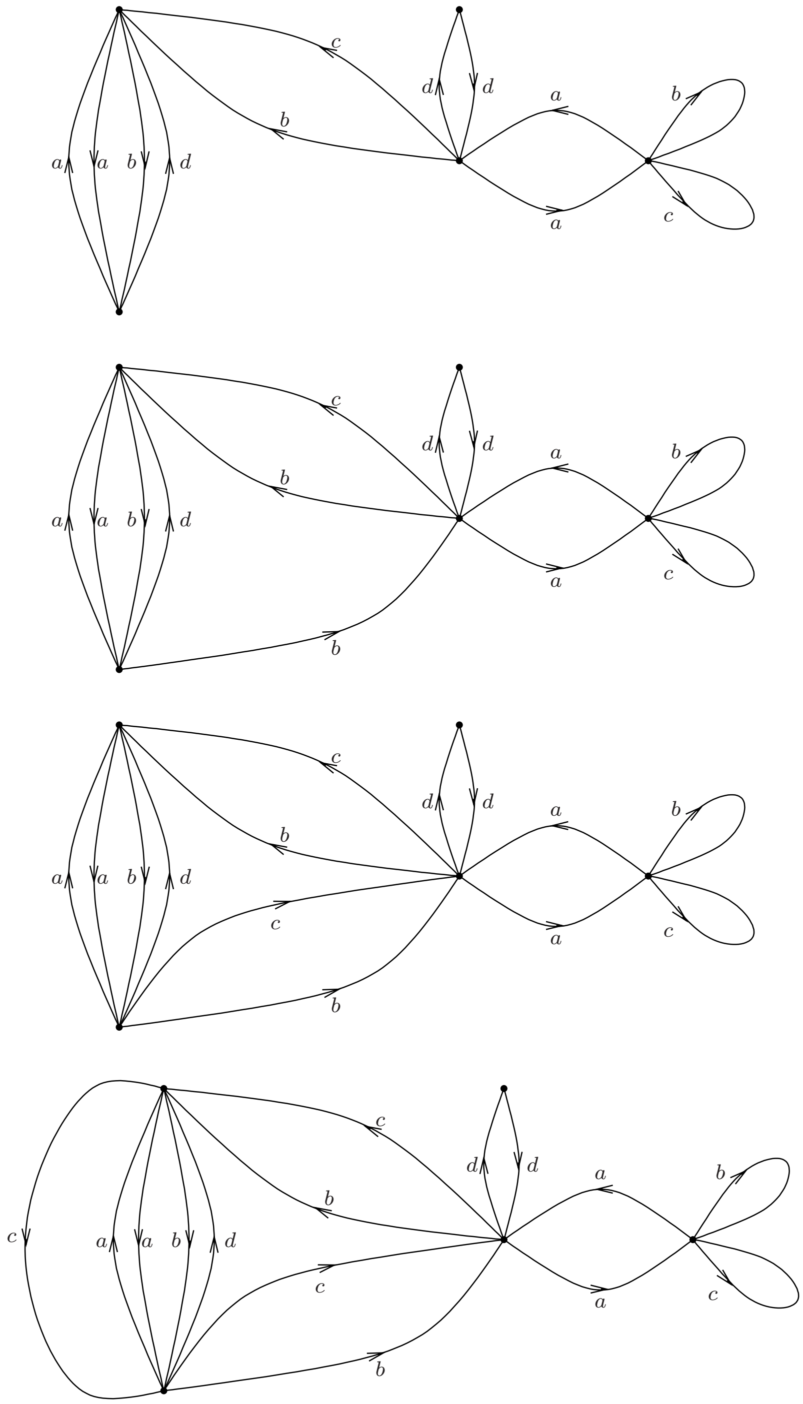

Figure 2: Sequence of one skeletons of subgroup complexes.

The complex Y_0 admits four complete reductions. In the first of these (ordered as in the implementation of [17]) packets can be attached with boundary $(bc)^3$ along the two loops on the right of the figure, and packets with boundaries $(ab)^3$, $(ad)^3$ and $(bd)^3$ along the four edges on the left of the figure to give a complex Y_1 (whose 1-skeleton is the same as that of Y_0, so is not redrawn).

The complex Y_1 admits twelve incomplete reductions, but no complete reductions; hence the algorithm A2 needs to examine the reduction tree. In the implementation of [17], the first to be examined is a reduction r_1, which attaches the packet with boundary path $(ab)^3$ on the path labelled $(ab)^2a$ which starts at the central vertex of the graph. This transforms Y_1 to a complex Y_2, whose 1-skeleton is the second graph in Figure 2.

The active graph of Y_2 is distinct from that of Y_1. The complex Y_2 now admits six incomplete reductions. The first examined by the algorithm of [17] is a reduction r_2, which attaches a packet with boundary path $(ac)^3$ on the path labelled $(ac)^2a$ which starts at the central vertex of the graph. This transforms Y_2 to a complex Y_3, whose 1-skeleton is the third graph in Figure 2.

The active graph of Y_3 is not isomorphic to either of the active graphs of Y_1 or Y_2, so the algorithm needs to look further down the tree. Now Y_3 has six incomplete reductions. The first examined by the algorithm of [17] is a reduction r_4, which attaches a packet with boundary path $(bc)^3$ on the path labelled $(bc)^2b$ which starts at the bottom left vertex of the graph. This leads to a complex Y_4, whose 1-skeleton is the fourth graph in Figure 2.

Finally we have arrived at a complex with a complete reduction; the packet with boundary $(ac)^3$ can be attached along the two leftmost edges in the fourth graph. Hence we have found a useful sequence of reductions. The 2-complex Y_5 resulting from this complete reduction has no further reductions and so the algorithm A2 terminates. We see that a presentation for the fundamental group of this complex is

$$\begin{array}{lll}\langle x_1,\ldots,x_8 & \mid & (x_3x_8^{-1})^3,(x_1x_2^{-1}x_7^{-1})^3,(x_1x_2^{-1}x_8^{-1})^3,\\ & & (x_2x_4^{-1}x_3x_7^{-1})^3,(x_2x_4^{-1}x_3x_8^{-1})^3,(x_3^{-1}x_1x_5^{-1}x_4^{-1}x_5^{-1})^3\rangle.\end{array}$$

6 Dealing with fan reductions

We finish by explaining briefly how further ideas of McCammond and Wise may be used to generalise the algorithm further, under certain small cancellation like hypotheses. We avoid detail, since there seems little point in recording what would in essence simply be a rewrite of the material in [14], For just as our theorems 4.10 and 4.12 were derived from McCammond and Wise's proofs of coherence for groups satisfying the path reduction and weak 2-cell hypotheses respectively, the proof of theorem 6.4 is directly derived from McCammond and Wise's proof of coherence under hypotheses which generalise the path reduction hypothesis.

The algorithm A2 has only applied 2-cell reductions. But a more general notion of *fan reduction* is developed in the article [14]. In essence a fan reduction is a finite sequence of 2-cell attachments, which together reduce perimeter, although the individual 2-cell attachments might not, and such that the 2-cells fit together well within X (basically in such a way that they fit together in a line.). We shall show in this section that we can define an analogue of algorithm A2 incorporating fan reductions which terminates correctly when G satisfies certain conditions of small cancellation type (concerning fans), introduced in [14, 16]; basically these conditions form the fan analogue of the path reduction hypothesis for $X(G)$.

Following in essence the definitions of [14], we define a *(stand alone) fan* F with outer path q and inner path s to be a finite 2-complex as follows. $\mathcal{C}(F)$ is a finite set $\{\gamma_1, \dots \gamma_n\}$, and the 1-skeleton of F is just the union of the images $\mathrm{im}(\gamma_i)$. For each i, $\mathrm{im}(\gamma_i) \cap \mathrm{im}(\gamma_{i+1})$ is a path of length at least 1, with distinct endpoints, but for $j > i$, $j \neq i+1$, $\mathrm{im}(\gamma_i) \cap \mathrm{im}(\gamma_j)$ is at most a single vertex, and is contained within $\cap_{k=i,i+1,\dots j} \mathrm{im}(\gamma_k)$. The paths q, s are the two distinct paths which connect a vertex in $\mathrm{im}(\gamma_1) \setminus \mathrm{im}(\gamma_2)$ to a vertex in $\mathrm{im}(\gamma_n) \setminus \mathrm{im}(\gamma_{n-1})$.

Equivalently, each γ_i is attached to F along a closed circuit r_i, where for $i > 1$ r_i is a concatenation $a_{i-1} q_i a_i^{-1} s_i^{-1}$, while $r_1 = q_1 a_1^{-1} s_1^{-1}$, with $q = q_1 \dots q_n$, $s = s_1 \dots s_n$, and with each a_i having length at least 1.

We call the path qs^{-1} the boundary of F, denoted ∂F.

We define a *fan in* X to be the image F of a stand alone fan F' under a 2-complex morphism to X. The outer path of F is then defined to be the image under that morphism of the outer path of F' and the inner path of F to be the image of the inner path of F'.

We need also the concept of a *disc diagram over* X *for a word* w. We define this to be a 2-complex D as follows. D maps into X via a 2-complex morphism. The 1-skeleton of D is a finite planar graph for which one planar embedding has been selected. The outer boundary of D, with respect to the selected embedding is labelled by the word w, and is null homotopic. We define the area of D to be the number of 2-cells it contains, and say that D has *minimal area* if its area is minimal over the areas of all disc diagrams with the same boundary label w.

Suppose that F is a fan in X, and that D is a disc diagram over X. We say that D contains F if there exists a stand alone fan F' and morphisms $\theta_1 : F' \to D$, $\theta_2 : D \to X$ such that $\theta_2 \circ \theta_1$ maps F' to X and $\theta_1(F')$ maps the outer path of F' to the boundary of D and the inner path of F' to a path in the interior of D.

Where F is a fan in X we define its packet $\tilde{F}$ to be the union of the packets of its 2-cells. We say that a fan F is *perimeter reducing* if the weight of its packet $\tilde{F}$ is strictly less than the weight of its outer path q, and that it is *weakly perimeter reducing* if the weight of its packet is less than or equal to the weight of q.

We say that a fan F in X is *spread out* when no side of a 2-cell occurs more than once as a side through an edge in the path q.

Now if $\phi : Y \to X$ is a packed 1-immersion, F is a fan in X with outer path q, and q' is a path in Y which maps under ϕ to q, then we say that F is *missing* in Y along q' if no 2-cell incident with an edge of q' maps under ϕ to a 2-cell in F.

Suppose that $\phi : Y \to X$ is a packed 1-immersion, and that F is a spread out fan in X, containing 2-cells $\gamma_1, \dots \gamma_n$ as described above, and with outer path q which is missing along a path q' in Y which maps under ϕ to q. Then we can perform a sequence of 2-cell attachments over $\gamma_1, \dots \gamma_n$. We call this a *fan attachment*. It extends Y to a complex Y^+, and ϕ to a packed 1-immersion $\phi^+ : Y^+ \to X$. It is shown in [14] that this attachment reduces perimeter strictly (i.e. $\mathrm{per}(\phi^+) < \mathrm{per}(\phi)$) precisely when F is a perimeter reducing fan, and reduces perimeter weakly (i.e. $\mathrm{per}(\phi^+) \leq \mathrm{per}(\phi)$) precisely when F is a weakly perimeter reducing fan. We call an attachment of this type which strictly reduces perimeter a *fan reduction*, and one which only weakly reduces perimeter a *weak fan reduction*. We shall call a reduction followed by all foldings admitted by the resulting complex a *folded reduction*.

A weighted 2-complex satisfies the *fan reduction hypothesis* provided that any packed 1-immersion which is not π_1-injective admits a fan perimeter reduction. For a 2-complex satisfying the fan reduction hypothesis, a generalisation of the algorithm A1, which replaces each 2-cell reduction by a fan reduction, terminates after finitely many steps in a complex $\hat{Y}$ and a packed 1-immersion $\hat{\phi} : \hat{Y} \to X$ which is π_1-injective [14].

Under slightly weaker conditions we can define an algorithm A3 which uses fan reductions to generalise A2.

We define a *fan reduction* tree $T'(Y)$ much as we defined the 2-cell reduction tree, except that now each edge is labelled by a fan reduction. We note that since for a given finite complex Y, there is a limit on the length of a path in Y along which a spread out fan may be missing, we have

Lemma 6.1 *$T'(Y)$ has finite degree.*

We define the core $\overline{T'(Y)}$ just as before. As before this is a finite tree.

Hence we can define A3 as follows.

Algorithm A3

Input : a finite presentation for G, a finite set of generators for H.

Initially : halt:= false, $X = X(G), Y = R(H, B), \phi : Y \to X$;

repeat : Simultaneously construct and traverse the fan reduction tree $\overline{T'(Y)}$ in a depth-first fashion.

> **If** an edge is found labelled by a strict fan reduction , apply to Y the sequence of folded fan reductions corresponding to the path from the root up to and including that edge. Replace Y by its image under that sequence of reductions.
>
> **else** halt := true;

while halt=false;

Output : a finite presentation for $\pi_1(Y)$.

The analogue to proposition 4.6 is immediate to prove.

Proposition 6.2 **(a)** *Any active graph which labels a vertex of $T'(Y)$ must label a vertex in the core of $T'(Y)$.*

(b) *Any weak fan reduction which labels an edge in $T'(Y)$ must also label an edge in its core.*

As we have already done under the two hypotheses considered in section 4, we can extract from a proof of coherence under certain conditions on fans, namely the proof of theorem 10.11 of [14], a proof that our algorithm works correctly under the same hypotheses.

Specifically we consider groups G for which $X(G)$ satisfies the following hypothesis, which we shall call the *diagram fan reduction hypothesis*:

There exists a collection $\mathcal{T}$ of spread out weakly perimeter reducing fans in $X(G)$ such that for each fan F in $\mathcal{T}$, $\mathrm{wt}(\tilde{F}) < \mathrm{wt}(\partial F)$. Every non-trivial minimal area diagram for G contains a spur or a fan in $\mathcal{T}$.

Proposition 6.3 *Suppose that G satisfies the diagram fan reduction hypothesis, and that $\mathcal{T}$ is the associated set of fans. Let $\phi : Y \to X(G)$ be π_1-surjective but not π_1-injective. Then Y admits a sequence of attachments of weakly perimeter reducing spread-out fans, terminating in an attachment which is strictly perimeter reducing.*

We shall not give a proof here. The proof is contained within the proof of [14, theorem 10.11]. That is analogous to the proof of [14, theorem 7.6] on which the proof of proposition 4.9 was based.

Now we can deduce

Theorem 6.4 *Suppose that G satisfies the diagram fan reduction hypothesis. Then the algorithm A3 terminates in finite time to give a finite presentation for H.*

The proof is as for theorem 4.10, except that propositions 6.3, 6.2 and lemma 6.1 are used in place of propositions 4.9, proposition 4.6 and lemma 4.3.

References

[1] J.-C. Birget, S. Margolis, J. Meakin, P. Weil, PSPACE-complete problems for subgroups of free groups and inverse finite automata, Theor. Comp. Sci. 242 (2000) 247-281.

[2] A. Clifford, Compact cores of coverings, Topology and its applications 97 (1999) 267–271.

[3] M. Delgado, S. Margolis and B. Steinberg, Combinatorial Group Theory, Inverse Monoids, Automata and Global Semigroup Theory, Int. J. Alg. Comput. 12 (2002) 179–211.

[4] L. Markus Epstein, PhD thesis, Bar Ilan University, 2005.

[5] M. Feighn and M. Handel, Mapping tori of free group automorphisms are coherent, Ann. Math (2) , 149 (1999) 1061-1077.

[6] N. Ghani and A. Heyworth, A rewriting alternative to Reidemeister-Schreier, preprint 2003.

[7] R. Gitik, S. Margolis and B. Steinberg, On the Kurosh theorem and separability properties, J. Pure Appl. Algebra 179 (2003) 87–97.

[8] D.F. Holt, The Warwick automatic groups software, in *Geometrical and Computational Perspectives on Infinite Groups*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 25 (1996) 69–82, G. Baumslag et al.

[9] D. F. Holt and D. F. Hurt, Computing automatic coset systems and subgroup presentations, J. Symbolic Comput. 27 (1999) 1–19.

[10] Derek F. Holt, KBMAG—Knuth-Bendix in Monoids and Automatic Groups, software package, 1995, `www.maths.warwick.ac.uk/ dfh/download/kbmag2/`.

[11] D. F. Hurt, The use of Knuth-Bendix methods amd automatic coset systems for solving the generalized word problem and finding subgroups presentations, Ph. D. thesis, University of Warwick, 1996.

[12] R. C. Lyndon, P. E. Schupp, Combinatorial Group Theory, Springer-Verlag, Berlin, Heidelberg, New York, 1977.

[13] J. P. McCammond, Private communication, 2005.

[14] J. P. McCammond and D. T. Wise, Coherence, local quasiconvexity and the perimeter of 2-complexes, to appear in Geom. Funct. Anal.

[15] J. P. McCammond and D. T. Wise, Windmills and extreme 2-cells, preprint.

[16] J. P. McCammond and D. T. Wise, Fans and ladders in small cancellation theory, Proc. London Math. Soc. (3) 84 (2002) 599–644.

[17] O. Payne, Computing presentations of subgroups, Ph.D. thesis, University of Newcastle upon Tyne, 2002.

[18] S. Rees and L.H. Soicher, An algorithmic approach to fundamental groups and covers of combinatorial cell complexes, J. Symb. Comp. 29 (2000) 59–77.

[19] K. Reidemeister, Knoten und Gruppen, Abh. Math. Sem. Univ. Hamburg 5 (1927) 7–23.

[20] M. Schönert et al. GAP – *Groups, Algorithms and Programming, fourth edition*, RWTH Aachen: Lehrstuhl D für Mathematik, 1994.

[21] O. Schreier, Die Untergruppen der freien Gruppen, Abh. Math. Sem. Univ. Hamburg 5 (1927) 161–183.

[22] G P Scott, Finitely generated 3-manifold groups are finitely presented, J. London Math. Soc. 6 (1973) 437–440.

[23] C. C. Sims, Computation with finitely presented groups, Encyclopaedia of Mathematics 48, Cambridge University Press, 1994.

[24] L. H. Soicher, GRAPE: a system for computing with graphs and groups, DIMACS series in discrete mathematics and theoretical computer science 11 (1993) 287–291.

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

[26] J. R. Stallings and A. R. Wolf, The Todd-Coxeter process, using graphs, Combinatorial group theory and topology (Alta, Utah, 1984) 157–161, Princeton Univ. Press, Princeton N.J., 1987.

[27] J. Stillwell, Classical Topology and Combinatorial Group Theory, 2nd edition, Springer, New York, Berlin, Heidelberg, 1993.