Computing Group Resolutions

Graham Ellis

Mathematics Department, National University of Ireland Galway, Ireland

Abstract

We describe an algorithm for constructing a reasonably small CW–structure on the classifying space of a finite or automatic group G. The algorithm inputs a set of generators for G, and its output can be used to compute the integral cohomology of G. A prototype GAP implementation suggests that the algorithm is a practical method for studying the cohomology of finite groups in low dimensions. We also explain how the method can be used to compute the low–dimensional cohomology of finite crossed modules. The paper begins with a review of the notion of *syzygy* between defining relators for groups. This topological notion is then used in the design of the algorithm.

Key words:
Eilenberg–Mac Lane space, finite group, automatic group, free resolution, cohomology, GAP implementation
1991 MSC: 20J06

1 Introduction

For any group G one can construct a reduced CW–space X whose fundamental group is isomorphic to G, and whose universal cover $\tilde{X}$ is contractible. We

Email address: graham.ellis@nuigalway.ie (Graham Ellis).
URL: www.maths.nuigalway.ie (Graham Ellis).

call the n–skeleton of such a space X an *n–presentation* of the group G. A 2–presentation in this sense is equivalent to the conventional algebraic notion of a presentation for G in terms of generators and relators. A 3–presentation is equivalent to a collection of generators and relators for G together with a set of *identities* among the relators [5,34]. More geometrically, one can regard 3–presentations as collections of homotopy classes of 1–dimensional and 2–dimensional *syzygies* [26,30]. We begin this paper with some explicit examples and precise definitions of the notions of *identity* and *syzygy*. Using these notions we then describe an algorithm for computing a reasonably small n–presentation of a finite group G, starting from a set of generators for G. The output can be used to compute the $(n-1)$–dimensional cohomology of G.

We first describe a basic form of the algorithm which requires a listing of all of the elements of G. We then outline a refinement which works relative to a subgroup $H \leqslant G$. A prototype GAP implementation can be downloaded from [15]. Output from this implementation is presented in Section 5 and suggests that the algorithm is a practical method for studying cohomology of finite groups in low dimensions.

It is well–known [18] that automatic groups, although generally infinite, admit classifying spaces with only finitely many cells in each dimension. We explain how our algorithm can in principle be adapted to obtain such classifying spaces. We have not yet implemented the algorithm in this generality, and realistically such an implementation could only be applied to certain fairly special automatic groups. General methods for calculating cohomology of finitely generated infinite groups tend to involve Gröbner bases or Knuth–Bendix type rewrite methods. Since these general methods can fail on automatic groups, an implementation of our algorithm for automatic groups might be of some interest. (For instance, both the Magnus [32] and Bergman [2] packages have difficulty calculating the homology of the 4–string braid group from the standard presentation $B_4 = \langle x, y, z : xyx = yxy,\, yzy = zyz,\, xz = zx\rangle$ because the relevant Gröbner basis is infinite. Our method would yield a resolution for this particular example, though Squier's specialized method for Artin groups [37] is more efficient. See [17] for a generalisation and implementation of Squier's method.)

Our algorithm can also be adpated to the problem of calculating the low–dimensional cohomology of the classifying space of a finite crossed module (*cf.* [14,11]). A crossed module is an algebraic object that captures the homotopy type of a connected CW–space Y with $\pi_i Y = 0, i \geqslant 3$. We explain how the algorithm could be used to calculate the cohomology of Y in dimensions $\leqslant 3$. Again, we have not implemented the algorithm in this context.

Procedures for calculating cohomology of groups are available in a number of computer algebra packages. Both GAP [19] and MAGMA [31] contain functions for computing the first and second cohomology $H^2(G, A), (n = 1, 2)$ of a finite group with coefficients in a finite G–vector space A. These are based on group theoretic techniques due to Derek Holt and work well even when G has fairly large order. In MAGNUS [32] there is a function for computing the integral homology of a finitely presented group G. The method is due to J. Groves and begins by trying to construct a complete rewrite system for the group. The function works well in low dimensions on an impressive variety of groups, both finite and infinite. (Interestingly though, it has difficulty computing the 6–dimensional homology of the dihedral group of order 12 from the presentation $D_6 = \langle x, y : x^2 = y^6 = (xy)^2 = 1\rangle$. It seems that in this case it tries to construct a resolution with an extremely large number of generators.) Procedures for calculating the cohomology rings $H^*(G, \mathbb{Z}_2)$ of finite 2–groups have been developed by Jon Carlson [8] and implemented in MAGMA. They use linear algebra over finite fields to construct resolutions, and have been used to compute most of the rings for the 2–groups of order up to 64. Integral resolutions for finite p–groups can be computed using the methods of Grabmeier and Lambe [20]. Previous work of Lambe [28] provides an algorithm for constructing integral resolutions of finitely generated torsion free groups. Squier [37] and Salvetti [35] have independently described methods for computing small resolutions of another special class of groups, namely the spherical Artin groups. Their methods also yield resolutions for finite Coxeter groups.

Our interest in the problem of computing n–presentations of groups was sparked by the articles of J. Groves [21] and R. Brown and A. Razak Salleh [6]. These both give procedures for converting a group rewrite system into, respectively, a free $\mathbb{Z}G$–resolution of $\mathbb{Z}$ and a free crossed G–resolution. As mentioned above, Groves' algorithm is implemented in MAGNUS. The procedure of Brown and Razak Salleh has been implemented by A. Heyworth and C.D. Wensley [25] using a logged Knuth–Bendix algorithm and Gröbner basis techniques. Both implementations apply to fairly general groups and, consequently, might be less efficient than implementations directed at a specific class of groups. It was with this point of view in mind that the author and I. Kholodna described in [16] a method for constructing n–presentations of *finite* groups G based on integer arithmetic calculations in the finitely generated free abelian group underlying the integral group ring $\mathbb{Z}G$. The algorithm given in the present paper is again aimed primarily at finite groups, but has a more homotopy–theoretic flavour in keeping with the approaches taken in [21] and [6], and for large groups is considerably faster than the arithmetic method of [16]. In common with [21] and [6] it involves the Cayley graph of G viewed as the 1–skeleton of the universal cover of a $K(G, 1)$, and in common with Heyworth and Wensley's implementation of [6] it involves a 'logging procedure'. Some differences between our approaches are that: we have opted for a set of generators of a finite (or automatic) group as starting data, rather than an arbitrary complete rewrite system of a group; we use geometric arguments involving elementary homotopy collapses in place of the algebraic theory of crossed complexes underlying [6]; we make no use of Gröbner bases or the theory of string rewriting; we incorporate an algorithmic procedure for eliminating redundancies from the constructed n–presentation; (Heyworth and Reinert [24] are currently developing such a procedure based on Gröbner bases); in order to handle large finite groups G we outline how one could work with the cosets of a subgroup H for which an n–presentation is already known;

We should mention that in dimension two our algorithm essentially coincides with J. Cannon's method [7,33] for constructing a group presentation from a set of generators of a finite group.

The paper is organised as follows. Section 2, which is very much based on the articles of Brown and Huebschmann [5], Kapranov and Saito [26] and Loday [30], contains examples and details on homotopical syzygies and identities among relations, and describes a procedure for passing from syzygies to identities. In Section 3 we explain how our basic algorithm works on finite groups by applying it, in low dimensions, to the symmetric group S_3. In Section 4 we recall how to compute the cohomology of a group. Section 5 gives details of a prototype GAP implementation. In Section 6 we outline the modifications needed to run the algorithm relative to a subgroup. In Section 7 we consider automatic groups and describe the modifications needed to run the algorithm on these groups. Section 8 explains how the methods can be applied to the cohomology of crossed modules.

2 Syzygies and identities

Let G be a group. We wish to construct a *classifying space* for G, by which we mean a reduced (*i.e.* has just one 0–cell) CW–space X with fundamental group $\pi_1 X = G$ and with homotopy groups $\pi_i X = 0$ in dimensions $i \geqslant 2$. We want the construction to show explicitly how each $(n+1)$–dimensional cell is attached to the n–skeleton $X(n)$ of X. To get a feel for the issues involved let us consider some specific examples.

Example 1. [5,30] For every group G we take the 0–skeleton of X to be a point, and the 1–skeleton to have one 1–cell for each generator in some generating set for G. Consider for instance $G = \mathbb{Z}^3$, the free abelian group on the set $\underline{x} = \{x, y, z\}$ of three generators. In this case the 1–skeleton $X(1)$ is the one–point union of three circles. We orientate each circle by means of an arrow head.

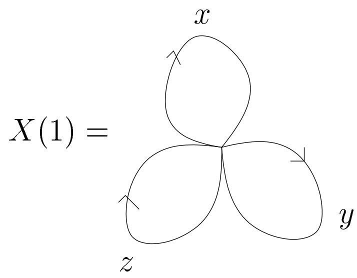

In order for X to have the appropriate fundamental group we must attach 2–cells to $X(1)$, one for each relator in some set $\underline{r}$ of defining relators for G. (A *relator* is a word in the free group on $\underline{x}$ with trivial image in G.) We can take $\underline{r}$ to be the set of three relators $\underline{r} = \{R_{[z,y]} := zyz^{-1}y^{-1}, R_{[x,z]} := xzx^{-1}z^{-1}, R_{[y,x]} := yxy^{-1}x^{-1}\}$. The space $X(2)$ is then obtained by attaching three 2–cells to $X(1)$ so that their boundaries spell these relators. The 2–cells can be represented by square discs with oriented edges labelled by generators.

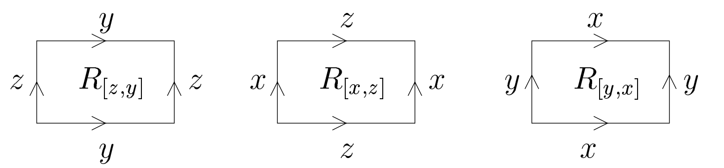

Van Kampen's theorem implies that $\pi_1 X(2) = G$ as required. But $\pi_2 X(2) \neq 0$, as the following cubical picture suggests.

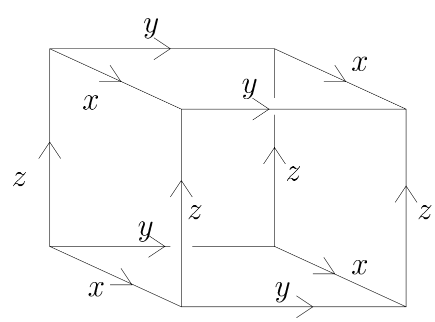

Figure 1

This can be viewed as a CW–decomposition of a 2–sphere S^2 involving eight 0–cells, twelve 1–cells, and six 2–cells. The labelling describes a map $f_\sigma \colon S^2 \to X(2)$ which sends each corner vertex to the unique 0–cell in $X(2)$, each edge to a 1–cell, and each face to a 2–cell. We denote by σ the data consisting of the CW–decomposition of S^2 together with the map f_σ, and we refer to σ as a *syzygy*.

The term 'syzygy' derives from the Greek word for 'yoke', and σ can be thought of as yoking together the relators $R_{[z,y]}$, $R_{[x,z]}$ and $R_{[y,x]}$. The terms *homotopical* syzygy [30] or *nonabelian* syzygy [26] are used to emphasise that σ is not a syzygy in the usual commutative algebraic sense; the term 2–*syzygy* is used to emphasise its 2–dimensional nature.

By considering the universal cover of $X(2)$ we see that $\pi_2 X(2) \cong \pi_2 \tilde{X}(2)$ is generated as a $\mathbb{Z}G$–module by the homotopy class of f_σ. So we construct $X(3)$ by attaching one 3–cell to $X(2)$ via the boundary map f_σ. Then $X(3)$ is homeomorphic to the direct product of three circles, and so its homotopy groups are trivial in all but the first dimension. We can thus set $X = X(3)$.

The syzygy σ is a geometric means of explaining a dependency between the relators $R_{[z,y]}$, $R_{[x,z]}$ and $R_{[y,x]}$. The dependency is explained algebraically by the formal product

$$S := (R_{[z,y]}\, {}^{y}x\, R^{-1}_{[z,y]}\, {}^{y}x^{-1})(R_{[y,x]}\, {}^{x}z\, R^{-1}_{[y,x]}\, {}^{x}z^{-1})(R_{[x,z]}\, {}^{z}y\, R^{-1}_{[x,z]}\, {}^{z}y^{-1})$$

in which ${}^{u}v$ denotes the conjugate uvu^{-1}. This product is an *identity between relators* in the sense that it is an expression, involving relators and their conjugates, which represents the identity element in the free group on x, y and z. Using the commutator convention $[u, v] := uvu^{-1}v^{-1}$, the identity $S = 1$ captures the well–known Jacobi–Hall–Witt identity between commutators:

$$[[z, y], {}^{y}x]\, [[y, x], {}^{x}z]\, [[x, z], {}^{z}y] = 1\,.$$

Precise details of the correspondence between S and σ are given below. For the moment the reader might note that: each of the three relators and their inverses occur in S; each pair of opposite faces in σ represents the two orientations of a relator disc.

(The syzygy σ could also be represented as a hexagon

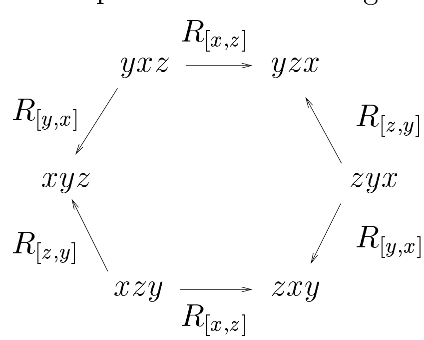

whose vertices are certain elements in the free group on x, y, z, and whose edges $\xrightarrow{R_{[u,v]}}$ correspond to rewrite rules $uv \mapsto vu$.)

Example 2. The preceding example can be extended to free abelian groups of arbitrary finite rank. Consider for instance $G = \mathbb{Z}^4$, the free abelian group on the set $\underline{x} = \{w, x, y, z\}$ of four generators. Continuing with previous notation, we have a set

$$\underline{r} = \{R_{[z,y]}, R_{[x,z]}, R_{[y,x]}, R_{[w,x]}, R_{[w,y]}, R_{[w,z]}\}$$

of six relators for G. The corresponding 2–skeleton $X(2)$ has four 1–cells and six 2–cells.

Figure 1 again yields a 2–syzygy which we now denote by $\sigma_{x,y,z}$. There are corresponding 2–syzygies $\sigma_{w,y,z}$, $\sigma_{x,y,w}$ and $\sigma_{x,w,z}$. A suitable space $X(3)$ is obtained from $X(2)$ by attaching one 3–cell for each of these four syzygies. Then $\pi_1 X(3) = G$, $\pi_2 X(3) = 0$ and $\pi_3 X(3)$ is generated, as a $\mathbb{Z}G$–module, by a single element. The generator is represented by a 3–syzygy τ (*i.e.* a labelled CW–decomposition of the 3–sphere S^3) consisting of two copies of each of the four 2–syzygies. The following picture of τ omits the 'outer' 2–syzygy.

$\tau =$ 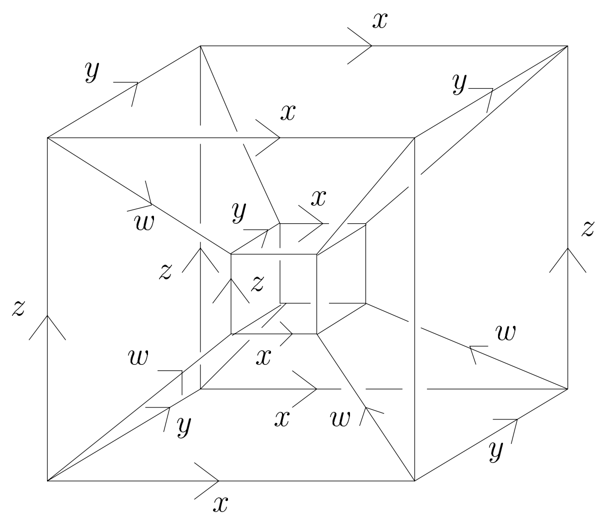

We construct $X(4)$ from $X(3)$ by attaching one 4–cell via the attaching map represented by τ. Then $X(4)$ is homeomorphic to a direct product of four circles and we can set $X = X(4)$.

The 3–syzygy τ represents a dependency between the 2–syzygies $\sigma_{x,y,z}$, $\sigma_{w,y,z}$, $\sigma_{x,y,w}$ and $\sigma_{x,w,z}$. Letting $S_{x,y,z}$, $S_{w,y,z}$, $S_{x,y,w}$ and $S_{x,w,z}$ denote the corresponding identities among relators, the dependency is represented algebraically by the formal sum

$$\begin{aligned} T := {}^{w}S_{x,y,z} \; - \; S_{x,y,z} \; + \; S_{w,y,z} \; - \; {}^{x}S_{w,y,z} \\ + S_{x,y,w} \; - \; {}^{z}S_{x,y,w} \; + \; S_{x,w,z} \; - \; {}^{y}S_{x,w,z} \, . \end{aligned}$$

This sum is an *identity among identities among the relators* in the sense that it is equal to zero when regarded, in the obvious way, as an element in the free $\mathbb{Z}G$–module on the set of symbols $\underline{r}$. (The sum can alternatively be expressed as a product of elements in the free crossed F–module associated to the presentation $\langle \underline{x} \mid \underline{r} \rangle$ with F the free group on $\underline{x}$, *cf.* [5,34]. Using the primary identity property described in Section 8 of [5], one can show that the product represents the identity element in the crossed module. Expanding the product using the above formula for $S_{x,y,z}$, the front face of the inner cube appears as ${}^{zwz^{-1}}R_{[x,z]}^{-1}$ in $S_{x,w,z}$ and as ${}^{w}R_{[x,z]}$ in ${}^{w}S_{x,y,z}$. The two terms cancel by the primary identity property. Eighteen pairs cancel in this way.)

Example 3. [30] Let G be an arbitrary group, and consider the uneconomical presentation with generating set $\underline{x} = \{x_g \ : \ 1 \neq g \in G\}$ and relator set $\underline{r} = \{R_{g,h} := x_g x_h x_{gh}^{-1} \ : \ 1 \neq g, h \in G\}$. In the relator $R_{g,g^{-1}}$ we read x_1^{-1} as the empty word. The associated 1–skeleton is a one–point union of circles $X(1) = \vee_{\underline{x}} S^1$. The 2–skeleton is obtained by attaching one 2–cell for each relator in $\underline{r}$, the attached 2–cell being represented by a triangular disc.

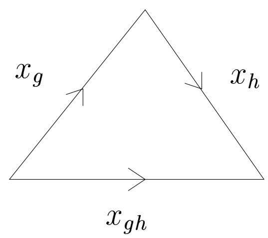

The fundamental group of $X(2)$ is isomorphic to G, but $\pi_2 X(2)$ is in general far from trivial. In fact, for each triple of non–trivial elements $g, h, k \in G$ the triangular discs corresponding to the relations $R_{h,k}, R_{gh,k}, R_{g,hk}, R_{g,h}$ fit together to form a tetrahedral syzygy which we denote by $\sigma_{g,h,k}$.

$\sigma_{g,h,k} =$

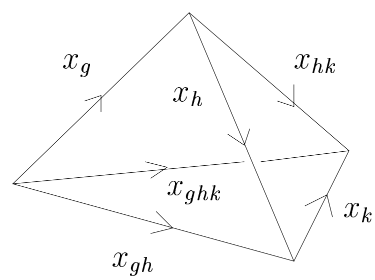

The corresponding identity between relators is

$$S_{g,h,k} := {}^{x_g}(R_{h,k})\, R_{g,hk}\, R_{gh,k}^{-1}\, R_{g,h}^{-1}.$$

A suitable 3–skeleton $X(3)$ is obtained from $X(2)$ by attaching a 3–cell, via the map $f_\sigma \colon S^2 \to X(2)$ represented by $\sigma = \sigma_{g,h,k}$, for all $1 \neq g, h, k \in G$.

The third homotopy group of $X(3)$ is in general non–trivial since for $1 \neq g, h, k, l \in G$ there is a 3–syzygy represented by the formal sum

$$T_{g,h,k,l} := {}^{g}S_{h,k,l} - S_{gh,k,l} + S_{g,hk,l} - S_{g,h,kl} + S_{g,h,k}.$$

One 4–cell must be attached to $X(3)$ for each such 3–syzygy. Continuing on to higher dimensions in this fashion one obtains a classifying space $X = \cup_{n\geqslant 0} X(n)$ whose n–cells are indexed by n–tuples of non–trivial elements in G. The corresponding $(n-1)$–syzygies can be expressed in terms of the boundary of an n–simplex. The space X is nothing but the geometric realisation of the categorical nerve of G (where G is thought of as a category with one object).

Our aim in this paper is to develop a computer method for constructing classifying spaces of finite (and automatic) groups involving relatively few cells in each dimension. The next examples illustrate the type of classifying spaces we are aiming at.

Example 4. [27] Let $G = D_{2k+1}$ be the dihedral group of order $m = 4k + 2, k \geqslant 1$. The construction in Example 3, when applied to this group, yields a space with $(m-1)^n$ n–dimensional cells. One can of course construct a classifying space for G with far fewer n–cells. To this end note that G is the group generated by $\underline{x} = \{x, y\}$ subject to the relators $\underline{r} = \{R_1 := x^2, R_2 := xy^k x^{-1} y^{-k-1}\}$. The associated 2–skeleton $X(2)$ has just two 1–cells and two 2–cells. The 2–cells (four copies of each) fit together to form the following 2–syzygy σ.

$\sigma =$

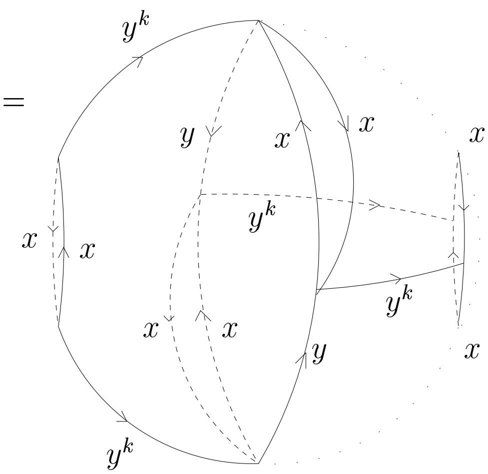

The computer program in [16] has been applied to this example for various values of k. For all of these values the program shows that the map $f_\sigma \colon S^2 \to X(2)$ determined by σ represents a homotopy class that generates $\pi_2 X(2)$ as a $\mathbb{Z}G$–module. So we can glue one 3–cell e^3 to $X(2)$ via f_σ to produce a suitable 3–skeleton $X(3)$. Furthermore, the computer program shows that m copies of e^3 can be glued together (or more precisely, the m images of e^3 under the action of D_{2k+1} can be added together in the relative homotopy group $\pi_3(X(3), X(2))$) to form a 3–syzygy that generates $\pi_3 X(3)$ as a $\mathbb{Z}G$–module. In fact, the program can be used to produce a classifying space X with two cells in dimensions congruent to 1 or 2 modulo 4, and a single cell in dimensions congruent to 0 or 3 modulo 4. A precise description and theoretical justification of this classifying space, valid for all $k \geqslant 1$, is given by Irina Kholodna in [27]. (The cellular chain complex of the universal cover of X is a periodic $\mathbb{Z}D_{2k+1}$–resolution of $\mathbb{Z}$ with minimal period. This resolution for $k = 1$ is essentially the periodic $\mathbb{Z}D_3$–resolution given in an appendix to Swan's paper [38].)

Example 5. Let $G = S_{m+1}$ be the symmetric group of degree $m + 1$. This group admits a Coxeter presentation with generators $\underline{x} = \{x_1, \cdots, x_m\}$ and relators

$$\begin{aligned}\underline{r} = \{R_{i,i} &:= x_i^2 && (1 \leqslant i \leqslant m),\\ R_{i,j} &:= x_i x_j x_i^{-1} x_j^{-1} && (1 \leqslant i < j \leqslant m, j - i > 1),\\ R_{i,j} &:= x_i x_j x_i x_j^{-1} x_i^{-1} x_j^{-1} && (1 \leqslant i < j \leqslant m, j - i = 1)\}.\end{aligned}$$

The associated 2–skeleton $X(2)$ has m 1–dimensional cells and $(m^2 + m)/2$ 2–dimensional cells. There is a 2–syzygy $\sigma_{i,j,k}$ for each triple of integers $1 \leqslant i \leqslant j \leqslant k \leqslant m$ which we now describe. For $i < j < k$ with $|k - i| = 2$ we set $x = x_i, y = x_j, z = x_k$ and picture the 2–syzygy $\sigma_{i,j,k}$ as follows.

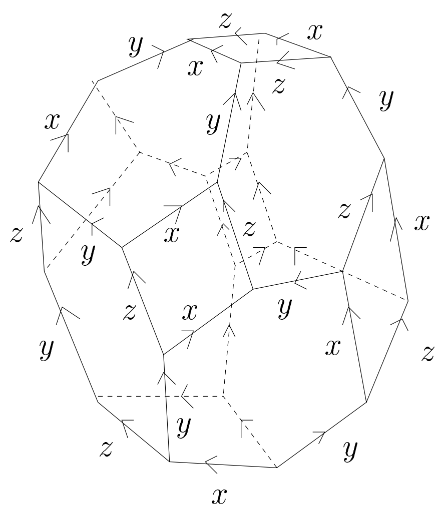

Figure 2

For $|k - i| > 2$ with $j = i + 1$ we set $x = x_i, y = x_j, z = x_k$, and for $|k - i| > 2$ with $k = j + 1$ we set $x = x_j, y = x_k, z = x_i$. In both cases we picture the 2–syzygy $\sigma_{i,j,k}$ as follows.

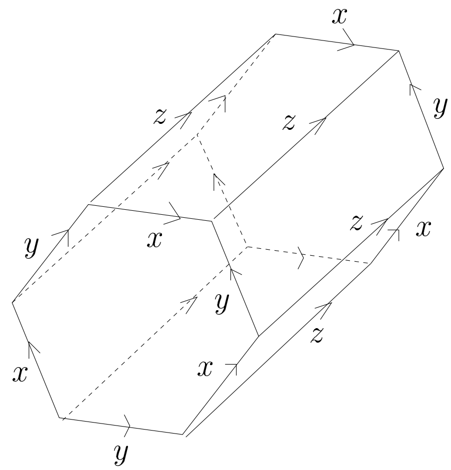

Figure 3

When $|j-i| > 1$ and $|k-j| > 1$ we set $x = x_i, y = x_j, z = x_k$ and picture the 2–syzygy $\sigma_{i,j,k}$ as in Figure 1 above.

For each $1 \leqslant i \leqslant m$ there is a 2–syzygy $\sigma_{i,i,i}$ obtained by gluing together two copies of the relator disc $R_{i,i}$ (with one copy rotated). For $1 \leqslant i < j \leqslant m$ there are 2–syzygies $\sigma_{i,i,j}$ and $\sigma_{i,j,j}$ obtained by gluing together two copies of the relator disc $R_{i,j}$ and copies of the relator discs $R_{i,i}$ or $R_{j,j}$. We set $x = x_i, y = x_j$ and, for $|i-j| > 2$, picture these syzygies as follows.

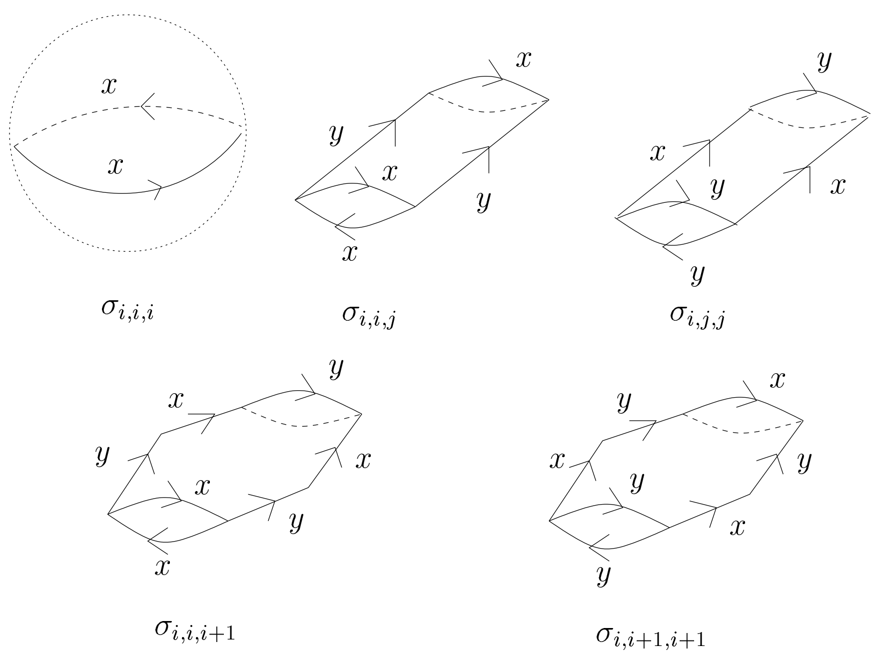

Figure 4

In total there are $m(m-1)(m-2)/6$ syzygies $\sigma_{i,j,k}$. Attaching one 3–cell to $X(2)$ for each of these syzygies yields a space $X(3)$ with the correct fundamental group and, we claim, with $\pi_2 X(3) = 0$. To prove the claim we consider the

Cayley graph $\Gamma(\underline{x})$. This is the directed graph whose vertices are the elements of S_{m+1}, with an edge from vertex g to vertex gx for each $g \in S_{m+1}, x \in \underline{x}$. There is a maximal acyclic subgraph (i.e. one without directed loops) which is isomorphic to the 1–skeleton of the m–dimensional convex polytope P_m known as the Permutahedron (*cf.* [30]). This polytope is the convex hull of the points $(g(1), \cdots, g(m+1))$ in $\mathbb{R}^{m+1}$, where g ranges over all permutations in S_{m+1}. The Cayley graph can be viewed as the 1–skeleton of $\tilde{X}(2)$. Moreover, the CW–space $\tilde{X}(2)$ has a CW–subspace P^2 that is isomorphic to the 2–skeleton of the Permutahedron. Let Y be the CW–space obtained from $X(2)$ by attaching one 3–cell for each of the syzygies shown in Figure 4. Then clearly any element in $\pi_2(Y) = \pi_2(\tilde{Y})$ can be represented by a map $S^2 \to P^2 \subset \tilde{Y}$. The 3–cells of the Permutahedron have boundaries corresponding to the syzygies $\sigma_{ijk}, (i < j < k)$. Since the Permutahedron is contractible, it follows that $\pi_2 X(3) = \pi_2(\tilde{X}(3)) = \pi_2(P_m) = 0$.

This example can be extended to higher dimensions, and generalized to all Coxeter groups. Details are given in [23] where they are used to obtain explicit formulas for the low–dimensional integral homology of arbitrary Coxeter groups. In the case of a finite Coxeter group with generators $\underline{x}$, the associated classifying space has one k–cell for each monomial in $\underline{x}$ of degree k. This classifying space for finite Coxeter groups has also been obtained by De Concini and Salvetti [12] using arguments based on hyper-plane arrangements.

A complete set of syzygies for general Artin groups is implicit in the work of Squier [37] (see also [35,13,17]). Examples of syzygies for the parametrized braid group, Steinberg group, Stasheff group, and Saneblidze group can be found in [26] and [30]. However, the above five examples should suffice to motivate the following more precise definitions of 'syzygy' and 'identity'.

Let G be an arbitrary group, $\underline{x}$ a set, and $\theta\colon F(\underline{x}) \twoheadrightarrow G$ a surjective homomorphism from the free group on $\underline{x}$ onto G. An element in the kernel of θ is called a *relator*, and a set $\underline{r}$ of relators is said to be *complete* if its normal closure in $F(\underline{x})$ equals the kernel of θ. A generating set together with a complete set of relators is called a *presentation* for G.

A presentation $\langle \underline{x} \mid \underline{r} \rangle$ can be viewed topologically in terms of a 2–dimensional CW–space $X(2)$. This space has one 0–cell and its 1–cells are in bijection with the generators $\underline{x}$. Its 2–cells are in bijection with the relators $\underline{r}$. Each 1–cell is oriented and labelled by the corresponding generator. The 2–cell corresponding to a relator $x_{i_1}^{\epsilon_1} \cdots x_{i_m}^{\epsilon_m}$ $(x_{i_j} \in \underline{x}, \epsilon_j = \pm 1)$ can be thought of as an oriented m–sided polygonal disc with attaching map that spells the relator (where inverses are represented by the direction of the attaching map). Recall that a CW–space is *regular* if the attaching map of each cell is a homeomorphism. Following Loday [30] we define a 2–*syzygy* σ to be a regular CW–decomposition of the

2–sphere S^2 together with a cellular map $f_\sigma: S^2 \to X(2)$ which sends each vertex to the unique 0–cell, each edge interior homeomorphically onto a 1–cell, and each face interior either homeomorphically onto a 2–cell or degenerately into the 1–skeleton. (Loday's definition does not allow for such degeneracies.) So, a 2–syzygy is a polytope decomposition of the sphere S^2 in which each face is oriented, labelled by some relator in $\underline{r}$ or labelled as degenerate, and has a preferred vertex (base point). The labelling and orientation are such that the edges can be oriented and labelled by generators in such a way that the boundary of a face, when read from the preferred vertex in the positive direction, spells the relator labelling it. The map f_σ represents a homotopy class in $\pi_2 X(2)$, and we refer to this class as a *homotopical* 2*–syzygy*. A set $\underline{s}$ of homotopical 2–syzygies is said to be *complete* if they generate $\pi_2 X(2)$ as a $\mathbb{Z}\pi_1$–module. In other words, $\underline{s}$ is complete if $\pi_2 X(3) = 0$ where $X(3)$ is the space obtained from $X(2)$ by attaching one 3–ball for each syzygy $\sigma \in \underline{s}$ via the attaching map f_σ.

By a 3*–syzygy* τ we mean a regular CW–decomposition of the 3–sphere S^3 together with a cellular map $f_\tau: S^3 \to X(3)$ which sends each 3–cell either homeomorphically onto a 3–cell or degenerately into the 2–skeleton, and such that the restriction of f_τ to the boundary of any 3–cell of S^3 is a 2–syzygy. This definition extends to higher dimensions.

The definition does not fully capture the notion of syzygy suggested by the above examples. We shall say that an n–syzygy σ is *regular* if all i–cells of S^n are mapped homeomorphically onto i–cells of $X(n)$ for $0 \leqslant i \leqslant n$. The syzygies in the above examples are regular in this sense.

A complete set of homotopical n–syzygies for $n = 1, 2, 3, \cdots$ gives rise to a chain of CW–spaces $X(1) \subset X(2) \subset X(3) \subset \cdots$ in which $X(n)$ is the n–dimensional CW–space obtained by gluing n–cells to $X(n-1)$ via the complete $(n-1)$–syzygies. The union $X = \cup_{1 \leqslant n < \infty} X(n)$ is a classifying space for G with n–skeleton $X(n)$.

To handle homotopical syzygies on a computer it can be useful to have an algebraic representation of them. Homotopical 1–syzygies are conveniently represented by relators. Homotopical 2–syzygies can be represented by *identities among relators* of the presentation $\langle \underline{x} \mid \underline{r} \rangle$. Such an identity is a formal expression

$$S := {}^{f_1}R_1^{\epsilon_1}\ {}^{f_2}R_2^{\epsilon_2}\ \cdots\ {}^{f_m}R_m^{\epsilon_m}$$

where: each f_i lies in the free group $F = F(\underline{x})$; each R_i lies in the set $\underline{r}$; each ϵ_i has the value +1 or -1; and the expression $\partial(S) := f_1 R_1^{\epsilon_1} f_1^{-1} \cdots f_m R_m^{\epsilon_m} f_m^{-1}$, considered as a word in F, is the identity element. One could think of S as living in the free group generated by the set $F \times \underline{r}$ (though this is not its ideal habitat) and then ∂ defines a group homomorphism. The term *identity among relators* is a translation (*cf.* [5]) of the expression *Identitäten zwischen*

Relationen due to R. Peiffer and K. Reidemeister. The correspondence between identities among relators and homotopical 2–syzygies is described explicitly by Loday [30] via the useful intermediate notion of an *Igusa picture* (see also [5,34]). The correspondence can also be described directly, as we now explain with the help of a specific example.

Consider the identity $S :=$

$$({}^{x^{-1}}R_2)(R_1^{-1})({}^{xy^{-1}}R_2^{-1})({}^{xy^{-1}}R_1)({}^{xy^{-1}x^{-1}}R_2^{-1})({}^{y^{-1}}R_1^{-1})({}^{y^{-1}}R_2)({}^{y}R_1)$$

among the relators of the presentation $\langle x, y \mid R_1 := x^2,\ R_2 := xyx^{-1}y^{-2}\rangle$ of the dihedral group $G = D_3$. To verify that S is indeed an identity it suffices to replace R_1 by x^2, replace R_2 by $xyx^{-1}y^{-2}$, and then evaluate the expression as a word in the free group F on x and y with actions ${}^{f}R_i^{\epsilon}$ interpreted as conjugation $fR_i^{\epsilon}f^{-1}$. The expression evaluates to the identity element of F.

To each relator R_i let us associate an oriented polygonal disc D_i with preferred vertex whose boundary is the corresponding 1–syzygy. Thus the boundary edges are oriented and labelled so as to spell R_i when read from preferred vertex in the direction of the disc's orientation. Let D_i^{-1} denote a copy of D_i with reversed orientation (whose boundary thus spells R_i^{-1}). We represent the conjugated relator ${}^{f}R_i^{\epsilon}$ by attaching to the preferred vertex of D_i^{ϵ} a piecewise linear tail whose edges, when read from the unattached end, spell the word f. For example, the conjugated relator ${}^{xy^{-1}x^{-1}}R_2^{-1}$ is represented as:

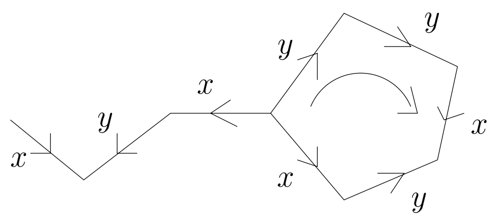

As explained in Brown's exposition [4] of a theorem of J.H.C. Whitehead, the identity among relators S can then be thought of as a bouquet of discs with tails all emanating from a single base point. See Figure 5. We denote this bouquet by B and consider it as a subspace of the 2–dimensional closed disc E^2.

The bouquet B has a natural CW–structure, and the labelling on B determines (up to homotopy) a cellular map $\phi\colon B \to X(2)$ into the 2–dimensional CW–complex associated to the presentation. The image of the 1–skeleton of B is contractible to a point, the contraction taking place entirely in the 1–skeleton $X(1)$, since the image spells a word representing the identity element in the free group $F(\underline{x})$. Hence ϕ can be extended to a map $\phi\colon E^2 \to X(2)$ that sends $E^2 \setminus B$ into the 1–skeleton of $X(2)$ and sends the boundary of E^2 into the

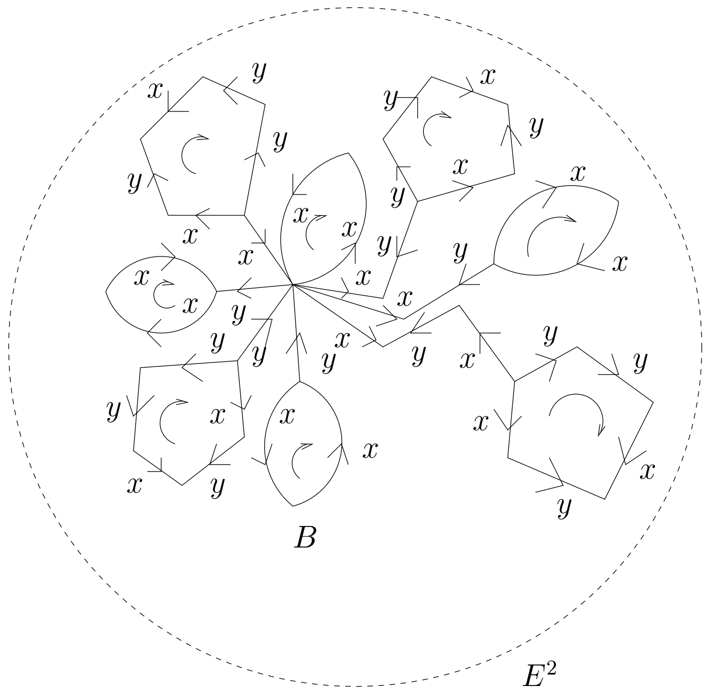

Figure 5

base point. We can thus view ϕ as a 2–syzygy $\phi\colon S^2 \to X(2)$ in which the CW–structure on S^2 is such that precisely one 2–cell maps degenerately into the 1–skeleton $X(1)$.

Working up to homotopy equivalence of maps, we can try to simplify the syzygy ϕ to a regular one. The simplification procedure can be viewed in terms of identifications and/or removals of certain edges in the bouquet B. In this particular example B simplifies to the following union of four pentagonal dics and four two–sided discs. The base point of the bouquet corresponds to the vertex marked by an asterix.

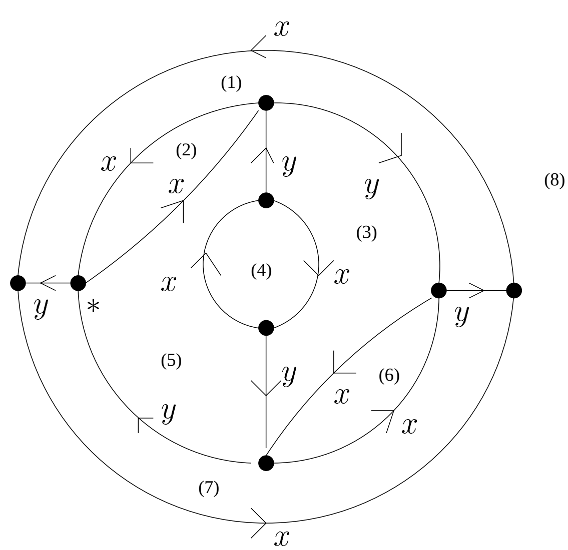

On identifying the top and bottom edges we see that this is precisely the regular 2–syzygy σ pictured in Example 4 (with $k = 1$). The regions of the syzygy have been numbered such that the first region corresponds to the first term in the above identity S, the second region to the second term, and so on.

The process of converting an identity among relators into a 2–syzygy is turned into a bijection by working up to suitable equivalence classes, namely Peiffer equivalence of identities [5,36] and homotopy equivalence of 2–syzygies:

$$\text{identities among relators}/\simeq \quad \xrightarrow{\cong} \quad \text{homotopical 2–syzygies}/\simeq\ .$$

This bijection follows from J.H.C. Whitehead's result [40] that the boundary homomorphism $\partial\colon \pi_2(X(2), X(1)) \to \pi_1(X(1))$ from the second relative homotopy group to the fundamental group has the structure of a free crossed module. (The second homotopy group $\pi_2(X(2))$ is isomorphic to $\ker \partial$ and every homotopy class in $\pi_2(X(2))$ is represented by a homotopical 2–syzygy.) Further details can be found in [5,36] though we shall not need them in what follows.

A 3–presentation for a group G is thus represented algebraically by the data $\langle \underline{x} \mid \underline{r} \mid \underline{s} \rangle$ where $\underline{x}$ is a set of generators for G, $\underline{r}$ is a complete set of relators for G, and $\underline{s}$ is a complete set of identities among relators.

Homotopical 3–syzygies can be represented algebraically by formal expressions

$$T := {\epsilon_1}^{f_1} S_1 \,+\, {\epsilon_2}^{f_2} S_2 \,+\, \cdots \,+\, {\epsilon_t}^{f_t} S_t$$

where: each f_i lies in the free group $F = F(\underline{x})$; each S_i is an identity among relators; each ϵ_i is an integer; the expression T represents (in the obvious manner) the zero element in the free $\mathbb{Z}G$–module $\oplus_{\underline{r}} \mathbb{Z}G$ freely generated by the elements of $\underline{r}$. We call T a 3–*identity* (in preference to an *identity among identities among relators*) and regard T as an element of $\oplus_{\underline{s}} \mathbb{Z}G$ where $\underline{s}$ is the set of identities.

As an illustration consider the identity

$$S := R_1\, {}^x R_1^{-1}$$

among the relators of the above presentation of the dihedral group $G = D_3$. Then

$$T := S \,+\, {}^x S$$

is a 3–identity because, evaluated as an element of the free $\mathbb{Z}G$–module on

$\{R_1, R_2\}$, we have

$$\begin{aligned} T &= S \, + \, {}^{x}S \\ &= R_1 - x \cdot R_1 + x \cdot (R_1 - x \cdot R_1) \\ &= R_1 - x \cdot R_1 + x \cdot R_1 - x^2 \cdot R_1 \\ &= 0 \, . \end{aligned}$$

An arbitrary 3–identity T corresponds to a homotopical 3–syzygy in the following manner. The cellular integral homology group $H_3(\tilde{X}(3), \tilde{X}(2))$ can be regarded as the free $\mathbb{Z}G$–module generated by the 3–cells of the universal cover $\tilde{X}(3)$. So T can be naturally regarded as an element of this homology group. Indeed, an element of this homology group is a 3–identity precisely when it lies in the kernel of the boundary homomorphism

$$d_3 \colon H_3(\tilde{X}(3), \tilde{X}(2)) \to H_2(\tilde{X}(2), \tilde{X}(1)).$$

Via the isomorphisms

$$\ker d_3 \cong H_3(\tilde{X}(3)) \cong \pi_3(\tilde{X}(3)) \cong \pi_3(X(3))$$

the 3–identity T corresponds to an element of $\pi_3(X(3))$, and thus to a homotopical 3–syzygy (since each element in $\pi_3(X(3))$ is represented by a 3–syzygy as defined above).

This notion of 3–identity extends to higher n–identities. For $n \geqslant 2$ we define an *n–identity* for the group G to be an element of the $\mathbb{Z}G$–module $\ker d_n$ where d_n is the boundary homomorphism

$$d_n \colon H_n(\tilde{X}(n), \tilde{X}(n-1))) \to H_{n-1}(\tilde{X}(n-1), \tilde{X}(n-2))).$$

There is a one–to–one correspondence between such n–identities and homotopy classes of n–syzygies.

The main aim of this paper is to describe an algorithm for computing, from a set $\underline{x}$ of generators of a finite or automatic group G, a sequence $\underline{x}, \underline{r}, \underline{s}, \underline{s}_3, \underline{s}_4, \cdots$ in which $\underline{s}_n$ is a reasonably small complete set of n–identities among generators. The computed n–identities will correspond to regular n–syzygies.

3 The basic algorithm for finite groups

In this section we consider an arbitrary finite group G.

The algorithm inputs a set $\underline{x}$ of generators for G, and we assume that some method is available for deciding when two words in these generators are equal in G. For example, $\underline{x}$ could be a set of permutations or matrices.

The algorithm inductively computes the skeleta of a classifying space X for G.

The 1–skeleton $X(1)$ is taken to be a wedge of circles, one circle for each generator in $\underline{x}$. We denote by $\tilde{X}(1)$ the topological space underlying the Cayley graph $\Gamma(\underline{x})$ of G. Recall that $\Gamma(\underline{x})$ is the oriented graph whose vertices are the elements of G, with an edge from vertex g to vertex gz for each $g \in G, z \in \underline{x}$. There is a fixed–point free cellular action of G on the space $\tilde{X}(1)$. On vertices the action is just pre–multiplication by elements of G.

Given $X(n), \tilde{X}(n)$ and a suitable action of G on $\tilde{X}(n)$ the algorithm computes $X(n+1), \tilde{X}(n+1)$ and an action of G on $\tilde{X}(n+1)$ roughly as follows:

(1) It first constructs a maximal contractible CW–subspace $Y(n)$ in $\tilde{X}(n)$.
(2) Each n–cell in $\tilde{X}(n)\backslash Y(n)$ determines a homotopical n–syzygy. Using the elimination procedure explained below, redundant syzygies are removed, leaving a complete set $\underline{s}_n$ of n–syzygies. Details are recorded of precisely why each redundant syzygy is redundant (and are subsequently used to express $(n+1)$–syzygies algebraically).
(3) The space $X(n+1)$ is obtained by attaching one $(n+1)$–ball to $X(n)$ for each syzygy in $\underline{s}_n$. The space $\tilde{X}(n+1)$ is the universal cover of $X(n+1)$. The fundamental group $G = \pi_1 \tilde{X}(n+1)$ acts in the usual way on the universal cover.

Tasks 2 and 3 are performed simultaneously using the following.

Elimination Procedure.

We initially set $U(n) := \tilde{X}(n)$, $W(n) := Y(n)$ and then repeatedly perform the following two steps until $W(n) = \tilde{X}(n)$.

(a) Choose some n–cell $\tilde{e}^n \in \tilde{X}(n) \setminus W(n)$ with corresponding n–syzygy μ. The action of G yields in an obvious way an n–syzygy $g \cdot \mu$ for all $g \in G$. For each $g \in G$ attach an $(n+1)$–cell $g \cdot \tilde{e}^{n+1}_\mu$ to $U(n)$ via $g \cdot \mu$. That is, set $U(n) := U(n) \cup \{g \cdot \tilde{e}^{n+1}_\mu : g \in G\}$. Also set $W(n) := W(n) \cup \{\tilde{e}^n\}$.
(b) Define an n–cell $\tilde{f}^n_0 \in \tilde{X}(n) \setminus W(n)$ to be *contractible into* $W(n)$ if and only if there exists an $(n+1)$–cell $\tilde{f}^{n+1} \in U(n)$ such that the boundary of $\tilde{f}^{n+1}$ in $\tilde{X}(n)$ involves only the n–cell $\tilde{f}^n_0$ and n–cells in $W(n)$.

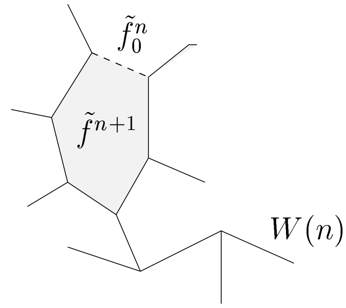

While possible, repeat the following:

Choose some $\tilde{f}_0^n \in \tilde{X}(n) \setminus W(n)$ which is contractible into $W(n)$ and set $W' := W(n) \cup \{\tilde{f}_0^n\}$; record precisely why $\tilde{f}_0^n$ is contractible into $W(n)$; then set $W(n) := W'$.

□

Let us illustrate the basic algorithm on the symmetric group $G = S_3$. As input data we must supply a set of generators for G. We take the set of permutations

$$\underline{x} = \{(1,2), (1,2,3)\}$$

and for convenience write $x = (1,2)$ and $y = (1,2,3)$.

Step (i). The algorithm begins by constructing a maximal tree $Y(1)$ in the Cayley graph $\Gamma = \Gamma(\underline{x})$, and setting $\tilde{X}(1)$ equal to the underlying space of Γ. A particular choice of maximal tree $Y(1)$ is represented by the solid arrows in the following picture of Γ.

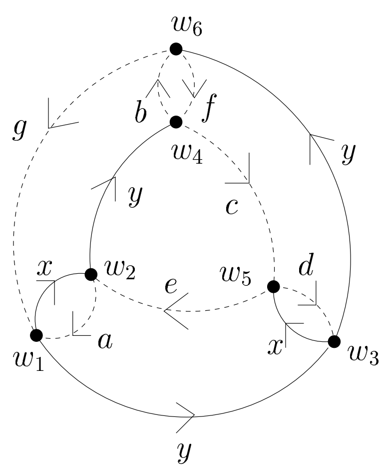

The tree $Y(1)$ has been constructed in a breadth–first manner (this is probably a good strategy in general), and during the construction those edges in $\tilde{X}(1) \setminus Y(1)$ have been labelled $a, b, c, \cdots$ in the order that they were discovered. During the construction the elements of G are enumerated in some order, and

$Y(1)$ associates to the ith element $w_i \in G$ a particular representation as a product of generators. We always assume that w_1 is the identity in G.

Note that to construct the Cayley graph we have to be able to decide when two words in the generators are equal. Once constructed, the Cayley graph provides an efficient method for multiplying elements of G. For example, one can read directly from Γ that the product of $w_4 = xy$ with $w_5 = yx$ is $w_4 w_5 = y = w_3$.

The space $\tilde{X}(1) = \Gamma$ is a 1–dimensional CW–space whose 0–cells are the vertices of the graph, and whose 1–cells are the edges. For each non–trivial $g \in G$ the permutation $G \to G, w \mapsto gw$ extends to a continuous fixed–point free map $\alpha_g: \tilde{X}(1) \to \tilde{X}(1)$, thus giving a fixed–point free cellular action of G on $\tilde{X}(1)$. We set $X(1) := \tilde{X}(1)/G$, the space obtained by killing the action of G. The orbit space $X(1)$ is homeomorphic to a wedge of circles, one circle for each generator, and the quotient map $p: \tilde{X}(1) \to X(1)$ is a $|G|$–fold covering map.

Step (ii) A 2–presentation $X(2)$ for G is constructed by first attaching 2–cells to $\tilde{X}(1)$ so as to form a simply connected 2–dimensional CW–space $\tilde{X}(2)$ with fixed–point free cellular G–action, and then taking the orbit space $X(2) := \tilde{X}(2)/G$. Thus $\tilde{X}(2)$ is to be the universal cover of $X(2)$. We construct $\tilde{X}(2)$ by attaching one 2–cell $g \cdot \tilde{e}^2_\lambda$ to $\tilde{X}(1)$ for each 1–cell $\lambda \in \tilde{X}(1) \setminus Y(1)$ and $g \in G$. The 2–cell $\tilde{e}^2_\lambda = 1 \cdot \tilde{e}^2_\lambda$ is attached so that its boundary $\partial \tilde{e}^2_\lambda$ is the unique simple closed circuit in $\tilde{X}(1)$ involving only edge λ and edges in $Y(1)$. For convenience we orient $\tilde{e}^2_\lambda$ by choosing a positive direction for its boundary, say the direction of the generator on edge λ. (Then, starting at the identity in G the boundary $\partial \tilde{e}^2_\lambda$ spells a word R_λ in the generators $\underline{x}$ and their inverses. Each R_λ is a word in $F = F(\underline{x})$ representing the identity in G.) The action of G on $\tilde{X}(1)$ induces an action of G on the simple closed circuits in $\tilde{X}(1)$. For $g \in G$ the 2–cell $g \cdot \tilde{e}^2_\lambda$ is attached in such a way that its boundary is the image of $\partial \tilde{e}^2_\lambda$ under the action of g. The action of G on $\tilde{X}(1)$ extends to a fixed–point free cellular action on $\tilde{X}(2)$.

The 2–presentation $X(2)$ has one 1–cell e^1_z for each generator $z \in \underline{x}$ and one 2–cell e^2_λ for each edge $\lambda \in \tilde{X}(1) \setminus Y(1)$. It corresponds to the generator–relator presentation $G = \langle \underline{x} \mid R_\lambda \ (\lambda \in \tilde{X}(1) \setminus Y(1)) \rangle$. Explicitly, for S_3 the seven relators are

$$R_a := x^2$$

$$R_b := xyxy^{-2}$$

$$R_c := xy^2x^{-1}y^{-1}$$

$$R_d := x^2$$

$$R_e := yxyx^{-1}$$

$$R_f := y^2xy^{-1}x^{-1}$$
$$R_g := y^3$$

(If, as is the case in this example, the maximal tree $Y(1)$ represents each element of G as a positive word in the generators, then the relators R_λ in fact present G as a monoid.)

Step (iii). The 2–presentation $X(2)$ has a number of redundant 2–cells which can be eliminated to form a smaller 2–presentation $X'(2)$. A version of the elimination procedure described above yields a 2–presentation $X'(2)$ for S_3 with just two 2–cells $\overline{e}^2_a$ and $\overline{e}^2_b$. The corresponding generator–relator presentation is

$$S_3 = \langle x, y \mid x^2, xyxy^{-2} \rangle .$$

Step (iv). While removing redundant 2–cells in Step (iii) a record can be kept of precisely why $\tilde{X}'(2)$ is simply connected. This space is a subcomplex of the original $\tilde{X}(2)$ with 1–skeleton $\tilde{X}(1)$. It is simply connected because each edge $\mu \in \tilde{X}(1) \setminus Y(1)$ can be deformed into the (varying) 1–dimensional space $W(1)$ through some 2–cell $g \cdot \tilde{e}^2_\lambda$ where $g \in G$ and e^2_λ is a 2–cell in $X'(2)$. For each μ a record of the corresponding deformation 2–cell $g \cdot \tilde{e}^2_\lambda$ is kept. Furthermore, the boundary of $g \cdot \tilde{e}^2_\lambda$ corresponds to the relator $R_\lambda \in F$, and in the same fashion this boundary also spells a word W_μ in the letters labelling the edges of Γ. The formal equation $W_\mu = R_\lambda$ is recorded, and an expression for μ is derived from it. These records for $G = S_3$ are tabulated as follows.

Edge μ	2–cell $g \cdot \tilde{e}^2_\lambda$	Corresponding equation	Resulting expression for μ
a	$\tilde{e}^2_a$	$xa = R_a$	$a = x^{-1}R_a$
b	$\tilde{e}^2_b$	$xyby^{-2} = R_b$	$b = y^{-1}x^{-1}R_by^2$
d	$y \cdot \tilde{e}^2_a$	$xd = R_a$	$d = x^{-1}R_a$
f	$xy \cdot \tilde{e}^2_a$	$bf = R_a$	$f = y^{-2}R_b^{-1}xyR_a$
c	$x \cdot \tilde{e}^2_b$	$ayxc^{-1}y^{-1} = R_b$	$c = y^{-1}R_b^{-1}x^{-1}R_ayx$
e	$yx \cdot \tilde{e}^2_b$	$dyfy^{-1}e^{-1} = R_b$	$e = R_b^{-1}x^{-1}R_ay^{-1}R_b^{-1}xyR_ay^{-1}$
g	$y^2 \cdot \tilde{e}^2_b$	$fcdy^{-1}g^{-1} = R_b$	$g = R_b^{-1}y^{-2}R_b^{-1}xyR_ay^{-1}R_b^{-1}x^{-1}R_a yR_ay^{-1}$

Table 1

Now set $X(2) := X'(2)$. Also, let $\Lambda = \{a, b\}$ be the set indexing the 2–cells of $X(2)$.

Step (v). A 3–presentation $X(3)$ for G is obtained by attaching 3–cells to $X(2)$ so as to kill the second homotopy group. The construction is performed by finding a minimal set of generators for the abelian group $\pi_2 X(2)$ and then attaching one 3–cell for each generator. It is convenient to use the isomorphism $\pi_2 X(2) \cong \pi_2 \tilde{X}(2)$ and work in the universal cover. The minimal generating set is derived from a maximal contractible subcomplex $Y(2)$ in $\tilde{X}(2)$.

In our 2–presentation $X(2)$ for $G = S_3$ the number of 1–cells is $d = 2$, the number of 2–cells is $k = 2$, and S_3 is of order $n = 6$. Hence the number of 2–cells in $\tilde{X}(2)$ is $nk = 12$, precisely $n(d-1)+1 = 7$ of which appear in Table 1. Let $Y(2)$ be the subcomplex of $\tilde{X}(2)$ comprising the 1–skeleton $\tilde{X}(1)$ and those 2–cells recorded in the table. Since $Y(2)$ is clearly contractible, $\tilde{X}(2)$ is homotopy equivalent to a wedge of $nk - n(d-1) - 1$ 2–spheres. Thus, as an abelian group, $\pi_2 \tilde{X}(2)$ is free abelian of rank $n(k-d+1) - 1 = 5$.

Each 2–cell $g \cdot \tilde{e}^2_\lambda$ of $\tilde{X}(2)$ not in $Y(2)$ gives rise to a homotopical 2–syzygy $\phi\colon S^2 \to \tilde{X}(2)$ in the following way. Let E^+, E^- denote the closed northern and southern hemispheres of S^2, and choose the 0–cell $g \in \Gamma = \tilde{X}(1)$ as the base point of $\tilde{X}(2)$. Then ϕ maps E^+ homeomorphically onto the closure of $g \cdot \tilde{e}^2_\lambda$, and E^- into $Y(2)$. The boundary of $g \cdot \tilde{e}^2_\lambda$ spells the relator R_λ, and the image of E^- is an oriented polytope decomposition of a disc with edges labelled by generators, with faces labelled by relators, and whose boundary spells R_λ. This polytope decomposition of the disc is represented algebraically by a formal expression

$$w_{(g,\lambda)} := {}^{f_1}R^{\epsilon_1}_{\lambda_1}\, {}^{f_2}R^{\epsilon_2}_{\lambda_2} \cdots {}^{f_t}R^{\epsilon_t}_{\lambda_t}$$

with $\lambda_i \in \Lambda$, $f_i \in F$ and $\epsilon_i = \pm 1$. Let $\tilde{g} \in F$ be the representative for $g \in G$ specified by the maximal tree $Y(1)$ and set

$${}^{\tilde{g}}w_{(g,\lambda)} := {}^{\tilde{g}f_1}R^{\epsilon_1}_{\lambda_1}\, {}^{\tilde{g}f_2}R^{\epsilon_2}_{\lambda_2} \cdots {}^{\tilde{g}f_t}R^{\epsilon_t}_{\lambda_t}\,.$$

When considered as a word in F, ${}^{\tilde{g}}w_{(g,\lambda)}$ equals $\tilde{g}R_\lambda\tilde{g}^{-1}$. The generator of $\pi_2 X(2) \cong \pi_2 \tilde{X}(2)$ corresponding to the 2–cell $g \cdot \tilde{e}^2_\lambda$ is represented by the identity between relators

$${}^{\tilde{g}}w_{(g,\lambda)}\, {}^{\tilde{g}}R^{-1}_\lambda\,.$$

To illustrate all of this, consider the 2–cell $y^2 \cdot \tilde{e}^2_a$ in the space $\tilde{X}(2)$. This 2–cell is not in the contractible subcomplex $Y(2)$ and so gives rise to a generating homotopical 2–syzygy whose restriction to E^- can be viewed as the following oriented polytope decomposition of a disc.

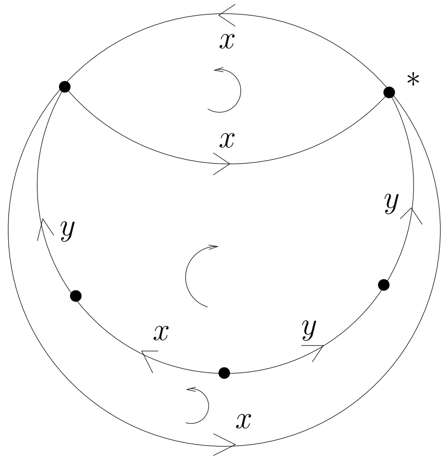

The boundary of this diagram spells the relator $R_a = x^2$. Furthermore, the diagram is obtained by piecing together three discs with boundary labels R_b, R_b, R_a and (with $w_6 = y^2$ as the base point) corresponds to the product of conjugates $w_{(y^2,a)} = {}^{y^{-2}}R_b^{-1}\,{}^{y^{-2}xy}R_a\,{}^{y^{-2}}R_b$. The identity between relators

$$R_b^{-1}\,{}^{xy}R_a\,R_b\,{}^{y^2}R_a^{-1}$$

thus represents the generator of $\pi_2\tilde{X}(2) \cong \pi_2 X(2)$ corresponding to the 2–cell $y^2 \cdot \tilde{e}_a^2$.

The identity between relators ${}^{\tilde{g}}w_{(g,\lambda)}\,{}^{\tilde{g}}R_\lambda^{-1}$ corresponding to a 2–cell $g \cdot \tilde{e}_\lambda^2$ of $\tilde{X}(2) \setminus Y(2)$ can in fact be derived from Table 1 by a direct algebraic calculation. Noting that $w_{(g,\lambda)}$ is an expression involving conjugates of the R_{λ_i} which equals R_λ when evaluated in F, we can determine $w_{(g,\lambda)}$ as follows. Let $L = \{a, \cdots, g\}$ be the the set of letters labelling the edges of the complement of $Y(1)$ in the Cayley graph. Take the word u in the letters $\underline{x} \cup L$ and their inverses spelled by the boundary of $g \cdot \tilde{e}_\lambda^2$, and let u' be obtained from u by substituting for each letter of L the corresponding expression given in Table 1. It is not difficult to see that the word u', considered as an element in the free group on $\underline{x} \cup \{R_{\lambda_i} \ : \ \lambda_i \in \Lambda\}$, is equal to a product of conjugates of the R_{λ_i}, namely $w_{(g,\lambda)}$.

For example, the boundary of the 2–cell $y^2 \cdot \tilde{e}_a^2$ spells fb and Table 1 yields

$$fb = y^{-2}R_b^{-1}xyR_ay^{-1}x^{-1}R_by^2$$

$$= {}^{y^{-2}}R_b^{-1}\,{}^{y^{-2}xy}R_a\,{}^{y^{-2}}R_b\,.$$

So $w_{(y^2,a)} = {}^{y^{-2}}R_b^{-1}\,{}^{y^{-2}xy}R_a\,{}^{y^{-2}}R_b$ and again we see that the corresponding identity between relators is $R_b^{-1}\,{}^{xy}R_a\,R_b\,{}^{y^2}R_a^{-1}$. The identity representing each generator of $\pi_2\tilde{X}(2) \cong \pi_2 X(2)$ is calculated by this algebraic method and recorded in Table 2.

2–cells in $\tilde{X}(2)\setminus Y(2)$	Boundary equations	Generating identities among relators
$x\cdot\tilde{e}^2_a$	$ax=R_a$	$R_a\,{}^xR_a^{-1}$
$yx\cdot\tilde{e}^2_a$	$dx=R_a$	${}^yR_a\,{}^{yx}R_a^{-1}$
$y^2\cdot\tilde{e}^2_a$	$fb=R_a$	$R_b^{-1}\,{}^{xy}R_a\,R_b\,{}^{y^2}R_a^{-1}$
$y\cdot\tilde{e}^2_b$	$xeag^{-1}y^{-1}=R_b$	${}^{yx}R_b^{-1}\,{}^yR_a\,R_b^{-1}\,{}^{xy}R_a\,R_a\,{}^yR_a^{-1}\,R_a^{-1}$ ${}^xR_b\,{}^{xy}R_a^{-1}\,R_b\,{}^{y^2}R_b\,{}^yR_b^{-1}$
$xy\cdot\tilde{e}^2_b$	$bgxe^{-1}c^{-1}=R_b$	$R_b\,{}^{y^2}R_b^{-1}\,R_b^{-1}\,{}^{xy}R_a\,{}^xR_b^{-1}$ $R_a\,{}^yR_a\,{}^{xy}R_a^{-1}\,R_b\,{}^yR_a^{-1}$ ${}^{yx}R_b\,R_a^{-1}\,{}^xR_b\,{}^{xy}R_b^{-1}$

Table 2

The space $X(3)$ is obtained by attaching one 3–cell to $X(2)$ for each generating identity between relators so that the corresponding generator of $\pi_2X(2)$ is killed. Let $\tilde{X}(3)$ denote the universal cover of $X(3)$.

Step (vi). The 3–presentation $X(3)$ has a number of redundant 3–cells which can be eliminated to form a smaller 3–presentation $X'(3)$. The redundancy arises because $\pi_2\tilde{X}(2)$ is naturally a $\mathbb{Z}G$–module and the generating set in Step (v) could be taken to be a set of module generators rather than group generators. Removing these redundancies using a version of the elimination procedure described above, we obtain a 3–presentation $X'(3)$ for $G=S_3$ with just the two 3–cells (which, for future reference, we denote by $e^3_{(x,a)}$ and $e^3_{(y,b)}$). In terms of generators, relators, and identities among relators this 3–presentation is

$$S_3 := \langle x,y \mid R_a:=x^2, R_b:=xyxy^{-2} \mid S_{(x,a)}:=R_a\,{}^xR_a^{-1},\ S_{(y,b)}:= {}^{yx}R_b^{-1}\,{}^yR_a\,R_b^{-1}\,{}^{xy}R_a\,R_a\,{}^yR_a^{-1}\,R_a^{-1}\,{}^xR_b\,{}^{xy}R_a^{-1}\,R_b\,{}^{y^2}R_b\,{}^yR_b^{-1}\rangle\,.$$

The 2–syzygy corresponding to the identity $S_{(y,b)}$ has the following stereographic projection (with $*$ marking the base point).

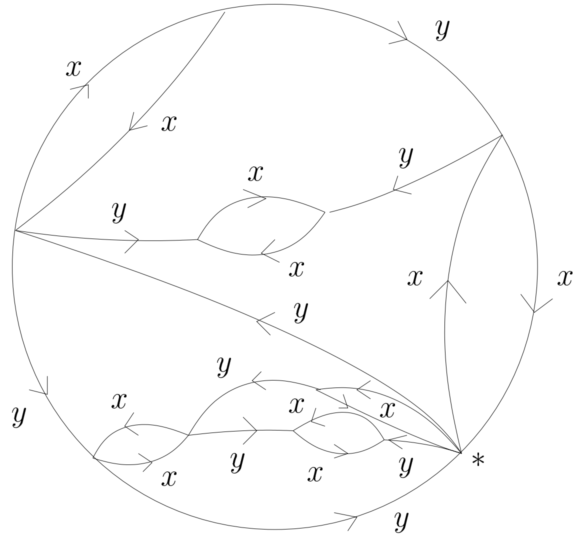

In order to record details of how redundant 2–syzygies were eliminated we denote by $S_{(g,\lambda)}$ the generating identity ${}^{\tilde{g}}w_{(g,\lambda)}\,{}^{\tilde{g}}R_\lambda^{-1}$ associated to a 2–cell $g\cdot\tilde{e}^2_\lambda \in \tilde{X}(2)\setminus Y(2)$. We denote by $e^3_{(g,\lambda)}$ the corresponding 3–cell of $X(3)$, and denote by $h\cdot\tilde{e}^3_{(g,\lambda)}$ ($h\in G$) the corresponding 3–cells of $\tilde{X}(3)$. For $h\in G$ we set ${}^hS_{(g,\lambda)} := {}^{\tilde{h}\tilde{g}}w_{(g,\lambda)}\,{}^{\tilde{h}\tilde{g}}R_\lambda^{-1}$. The identity ${}^hS_{(g,\lambda)}$ can be thought of as a polytope decomposition of the sphere S^2 with each non–degenerate face labelled by a 2–cell of $\tilde{X}(2)$. This polytope decomposition describes the attaching map of the 3–cell $h\cdot\tilde{e}^3_{(g,\lambda)}$. For example, $S_{(y^2,a)}$ is the identity $R_b^{-1}\,{}^{xy}R_a\,R_b\,{}^{y^2}R_a^{-1}$ and, as ilustrated in Step (v), can be pictured a polytope with four faces labelled $\tilde{e}^2_b, xy\cdot\tilde{e}^2_a, \tilde{e}^2_b, y^2\cdot\tilde{e}^2_a$. The identity ${}^xS_{(y^2,a)}$ can be pictured as a polytope with four faces labelled $x\cdot\tilde{e}^2_b, y\cdot\tilde{e}^2_a, x\cdot\tilde{e}^2_b, xy^2\cdot\tilde{e}^2_a$.

At this stage we set $X(3) := X'(3)$ and take $\tilde{X}(3)$ to be the universal cover.

Further steps. A 4–presentation for G is constructed by attaching 4–cells to $X(3)$, one 4–cell for each generator of the group $\pi_3X(3)$. Using the isomorphism $\pi_3X(3)\cong\pi_3\tilde{X}(3)$ a suitable set of generators is derived from a maximal contractible subcomplex $Y(3)$ in $\tilde{X}(3)$. As it in fact suffices to take generators for $\pi_3X(3)$ considered as $\mathbb{Z}\pi_1$–module, redundant cells can be eliminated from the 4–presentation $X(4)$ to produce a smaller presentation $X'(4)$.

Note that $\tilde{X}(3)$ has twelve 3–cells, precisely five of which correspond to the generating identities in Table 2. Let $Y(3)$ be the subcomplex of $\tilde{X}(3)$ comprising the 2–skeleton $\tilde{X}(2)$ and the five 3–cells corresponding to Table 2. Since $Y(3)$ is contractible, $\tilde{X}(3)$ is homotopy equivalent to a wedge of seven 3–spheres. Thus, as an abelian group, $\pi_3\tilde{X}(3)$ is free abelian of rank 7 with generators in one–one correspondence with the 3–cells in $\tilde{X}(3)\setminus Y(3)$. The 4–presentation $X(4)$ will therefore have seven 4–cells, some of which are redundant.

For the explicit construction of $X'(4)$ we represent the generators of $\pi_3\tilde{X}(3)$ algebraically as 3–identities. The 3–identity corresponding to a 3–cell in $\tilde{X}(3)\setminus Y(3)$ can be derived as follows. First note that the oriented 2–cells in $Y(2)$ are naturally labelled by conjugates of relators. Label the oriented 2–cells in $\tilde{X}(2)\setminus Y(2)$ by letters $A,\cdots,E$ in some fashion, express the corresponding identities between relators as a product of these letters and conjugates of relators, and then deduce an expression for each letter as a product of identities between relators and conjugates of relators. For example:

Face α	2−cells in $\tilde{X}\setminus Y(2)$	Corresponding equation	Resulting expression for α
A	$x\cdot\tilde{e}^2_a$	$S_{(x,a)}=R_aA^{-1}$	$A=S^{-1}_{(x,a)}R_a$
B	$yx\cdot\tilde{e}^2_a$	$S_{(yx,a)}={}^yR_aB^{-1}$	$B=S^{-1}_{(yx,a)}{}^yR_a$
C	$y^2\cdot\tilde{e}^2_a$	$S_{(y^2,a)}=R_b^{-1xy}R_aR_bC^{-1}$	$C=S^{-1}_{(y^2,a)}R_b^{-1xy}R_aR_b$
D	$y\cdot\tilde{e}^2_b$	etc.	
E	$xy\cdot\tilde{e}^2_b$	etc.	

Table 3

Suppose that $\tilde{e}^3_{(g,\lambda)}$ is the 3–cell in $Y(3)$ attached so as to kill the identity between relators $S_{(g,\lambda)}$. Suppose that for some $h\in G$ the 3–cell $h\cdot\tilde{e}^3_{(g,\lambda)}$ is not in $Y(3)$. Then the 3–cell $h\cdot\tilde{e}^3_{(g,\lambda)}$ corresponds to some 3–identity which we denote by $T_{(h,g,\lambda)}$. The boundary of $h\cdot\tilde{e}^3_{(g,\lambda)}$ is represented by ${}^hS_{(g,\lambda)}$. But the boundary of $h\cdot\tilde{e}^3_{(g,\lambda)}$ is a union of 2–cells and can be naturally expressed as a word w in the letters $A,\cdots,E$ and conjugates of relators. The final column in Table 3 can be used to rewrite w as a word $u_{(h,g,\lambda)}v$ where $u_{(h,g,\lambda)}$ is a word in conjugates of identities among relators and where v is a word in conjugates of relators; moreover, v will be an identity between relators representing the trivial element in $\pi_2\tilde{X}(2)$. In multiplicative (rather than the usual additive) notation the 3–identity $T_{(h,g,\lambda)}$ is

$$T_{(h,g,\lambda)}:=u_{(h,g,\lambda)}\,{}^hS^{-1}_{(g,\lambda)}.$$

For example, consider the 3–cell $\tilde{e}^3_{(x,a)}$ in $\tilde{X}(3)$ attached so as to kill the generating identity between relators $S_{(x,a)}$. We include $\tilde{e}^3_{(x,a)}$ in the subcomplex $Y(3)$. The 3–cell $x\cdot\tilde{e}^3_{(x,a)}$ is then not in $Y(3)$ and so gives rise to a 3–identity $T_{(x,x,a)}$. To determine this 3–identity note that the boundary of $x\cdot\tilde{e}^3_{(x,a)}$ is represented by ${}^xS_{(x,a)}$. Using Table 3 we see that

$$ {}^xS_{(x,a)} = {}^x(R_a\, {}^xR_a^{-1}) = A\, {}^{x^2}R_a^{-1} = S_{(x,a)}^{-1}\, R_a\, {}^{x^2}R_a^{-1} $$

from which we deduce that

$$ T_{(x,x,a)} := S_{(x,a)}^{-1}\, {}^xS_{(x,a)}^{-1}. $$

In additive notation this becomes

$$ T_{(x,x,a)} := -S_{(x,a)} \; - \; {}^xS_{(x,a)}. $$

A resolution for the group S_3 was also constructed by Brown and Salleh in [6]. At this point we should further elaborate on the comparison of their method and our method. Both employ the same basic idea, namely the recursive construction of an $(n+1)$–skeleton $X(n+1)$ from an n–skeleton $X(n)$ and data on how to contract the n–skeleton. The algebraic language of crossed complexes is used in [6], whereas we have opted for the language of CW–spaces. Although there is a well established equivalence between the homotopy categories of free crossed complexes and CW–spaces, the languages do seem to yield different insights and suggest slightly different strategies for implementing the basic idea. In [6] the data on how to contract the n–skeleton is stored algebraically as a contracting homotopy. We store the data as a maximal contractible subcomplex in $\tilde{X}(n)$. In constructing the subcomplex we have used a generalisation of Prim's Algorithm for finding spanning trees. It would be interesting to try generalisations of other spanning tree algorithms such as Kruskal's algorithm. There is scope for futher experimentation here.

The language of crossed complexes has the definite advantage that it leads to precise statements involving explicit algebraic formulae. On the other hand, the notion of a finite regular contractible CW–space is easily encoded and manipulated on a computer. In subsequent sections we will see that it is also a convenient notion when working relative to subgroups and when dealing with automatic groups.

There is much interest in constructing resolutions which, in low dimensions, correspond to a specified generator/relator presentation $G = \langle \underline{x} \mid \underline{r} \rangle$. The method in [6] produces such a resolution. Indeed, it's input is a specific presentation (or rewrite system). In its present form the GAP implementation of our method can also be used to produce such a resolution under certain mild resitrictions on the presentation.

4 Computing (co)homology

A classifying space $X = K(G,1)$ gives rise to a free $\mathbb{Z}G$–resolution of $\mathbb{Z}$, namely the cellular chain complex $C_*(\tilde{X})$:

$$\to H_n(\tilde{X}^n, \tilde{X}^{n-1}) \stackrel{\partial_n}{\to} H_{n-1}(\tilde{X}^{n-1}, \tilde{X}^{n-2}) \to \cdots \to H_1(\tilde{X}^1, \tilde{X}^0) \to \mathbb{Z}G$$

of the universal cover. In terms of the corresponding n–presentation $\langle \underline{x} \mid \underline{r} \mid \underline{s} \mid \underline{s}_3 \mid \cdots \mid \underline{s}_{n-1} \rangle$ the lower dimensions of the resolution have the form

$$\oplus_{\underline{s}_{n-1}} \mathbb{Z}G \xrightarrow{\partial_n} \oplus_{\underline{s}_{n-2}} \mathbb{Z}G \to \cdots \to \oplus_{\underline{s}} \mathbb{Z}G \xrightarrow{\partial_3} \oplus_{\underline{r}} \mathbb{Z}G \xrightarrow{\partial_2} \oplus_{\underline{x}} \mathbb{Z}G \xrightarrow{\partial_1} \mathbb{Z}G.$$

For systematic notation we set $s_{-1} := \{*\}$, $\underline{s}_0 := \underline{x}$, $\underline{s}_1 := \underline{r}$, $\underline{s}_2 := \underline{s}$. Then, for $i \geqslant 0$, $C_i(\tilde{X}) = \oplus_{\underline{s}_{i-1}} \mathbb{Z}G$ is the free $\mathbb{Z}G$–module generated by the symbols e^i_S where $S \in \underline{s}_{i-1}$.

For $i \geqslant 4$ the boundary homomorphism ∂_i maps each generator e^i_S as follows. The $(i-1)$–identity S is by definition a formal sum $S := \epsilon_1{}^{f_1} S_1 + \epsilon_2{}^{f_2} S_2 + \cdots + \epsilon_t{}^{f_t} S_t$ of $(i-2)$–identities. The homomorphism ∂_i is defined by

$$\partial_i(e^i_S) = \epsilon_1 f_1 . e^{i-1}_{S_1} + \epsilon_2 f_2 \cdot e^{i-1}_{S_2} + \cdots + \epsilon_t f_t \cdot e^{i-1}_{S_t}$$

where $f_i \in F = F(\underline{x})$ acts via the quotient homomorphism $\phi\colon F \twoheadrightarrow G$. The boundary homomorphism ∂_3 is also defined in this way, except for the notational difference that in this case S is a formal product rather than a formal sum.

The homomorphism $\partial_1\colon \oplus_{\underline{x}} \mathbb{Z}G \to \mathbb{Z}G$ is defined by setting

$$\partial_1(e^1_x) = \phi(x) - 1$$

for $x \in \underline{x}$.

To describe the homomorphism ∂_2 we need to recall details on the Fox derivative. Let W be a $\mathbb{Z}G$–module, and let an element $f \in F$ act on an element $w \in W$ by $f.w = \phi(f)w$. A function $\chi\colon F \to W$ is said to be a *derivative* if it satisfies the rule $\chi(ff') = \chi f \ + \ f \cdot \chi f'$. A consequence of this rule is that $\chi(f^{-1}) = -f^{-1}(\chi f)$. It is readily seen that for each generator $x \in \underline{x}$ there is a unique derivative $\frac{\partial}{\partial x}\colon F \to \mathbb{Z}G$ that satisfies $\frac{\partial}{\partial x}(x) = 1$ and $\frac{\partial}{\partial x}(y) = 0$ for $x \neq y \in \underline{x}$.

The homomorphism $\partial_2\colon \oplus_{\underline{r}} \mathbb{Z}G \to \oplus_{\underline{x}} \mathbb{Z}G$ is defined on generators by

$$\partial_2(e^2_R) = \sum_{x \in \underline{x}} (\frac{\partial R}{\partial x})\, e^1_x.$$

More details on this description of the resolution $C_*(\tilde{X})$ can be found, for instance, in [5,36].

The resolution $C_*(\tilde{X})$ can be used to compute the (co)homology of G with coefficients in a $\mathbb{Z}G$–module A. In particular, to compute the homology of G with coefficients in the trivial module $\mathbb{Z}$ we use the chain complex $C_*(\tilde{X}) \otimes_{\mathbb{Z}G} \mathbb{Z}$:

$$\longrightarrow \oplus_{\underline{s}_{n+1}} \mathbb{Z} \xrightarrow{\overline{\partial}_{n+1}} \oplus_{\underline{s}_n} \mathbb{Z} \xrightarrow{\overline{\partial}_n} \cdots \longrightarrow \oplus_{\underline{s}} \mathbb{Z} \xrightarrow{\overline{\partial}_3} \oplus_{\underline{r}} \mathbb{Z} \xrightarrow{\overline{\partial}_2} \oplus_{\underline{x}} \mathbb{Z} \xrightarrow{\overline{\partial}_1} \mathbb{Z} \xrightarrow{\overline{\partial}_0} 0.$$

The integral homology group $H_n(G,\mathbb{Z}) = \ker \overline{\partial}_n / \mathrm{image}\, \overline{\partial}_{n+1}$ is calculated by using the Smith Normal Form algorithm [10] to determine a basis for $\ker \overline{\partial}_n$ and a basis for $\mathrm{image}\, \overline{\partial}_{n+1}$.

In calculating $\mathrm{image}\, \overline{\partial}_{n+1}$ there is the option of using an unreduced set of n–identities $\underline{s}_n$. The process of eliminating redundant n–identities is then, in effect, carried out by the Smith Normal Form algorithm. For example, to calculate the second homology $H_2(D_{400},\mathbb{Z})$ of the dihedral group of order 800, we could use the algorithm in Section 3 to compute a 3–presentation with 2 generators, 3 relators, and 1599 identities between relators, and then apply the Smith Normal Form algorithm.

Homology and cohomology are functors. Given a group homomorphism $\phi: G' \to G$ we would like to compute the associated momorphisms in homology and cohomology. The crux of this computation is the construction of a chain map $\phi_*: C_*(\tilde{X}') \to C_*(\tilde{X})$ realizing ϕ, where $X = K(G,1)$, $X' = K(G',1)$. The construction of ϕ_* in turn reduces to the problem of expressing an arbitrary n–syzygy of G as a combination of generating n–syzygies for G. This expression can be found algorithmically from a maximal contractible $(n+1)$–subcomplex of X. The maximal contractible subcomplexes $Y(1), Y(2), \cdots$ in effect constitute a contracting homotopy for the space $\tilde{X}$.

Consider for example the homomorphism $\phi: C_2 \to S_3$ that sends the generator of the cyclic group $C_2 = \langle z \mid z^2 \rangle$ to the element xy in $S_3 = \langle x, y \mid x^2, xyxy^{-2} \rangle$. The group C_2 admits a resolution $C_*(\tilde{X}')$ with just one generator e^n in each dimension n. The presentation for S_3 yields the low dimensions of a resolutiuon $C_*(\tilde{X})$ for S_3. An induced chain map ϕ_* is given in dimensions $0,1$ by setting $\phi_0(e^0) = e^0$ and $\phi_1(e^1) = e^1_x + e^1_y$. The map ϕ_1 corresponds to the lifted homomorphism of free groups $\tilde{\phi}: F(z) \to F(x,y), z \mapsto xy$. In order to define ϕ_2 we are faced with the problem of expressing the element $\tilde{\phi}(z^2) = xyxy$ as a product of conjugates of the relators $R_a := x^2$ and $R_b := xyxy^{-2}$. Using the labelled graph Γ from Step (i) in Section 3 we see that, as a word in Γ, the element $\phi(z^2)$ spells $xybg$. Using Table 1 we can re–express $xybg$ as

$$R_b y^2 R_b^{-1} y^{-2} R_b^{-1} xy R_a y^{-1} x^{-1} x R_b^{-1} x^{-1} R_a y R_a y^{-1}.$$

We thus define

$$\phi_2(e^2) = (1 + y + xy)e^2_{R_a} - (x + y^2)e^2_{R_b}.$$

The method can be continued to obtain definitions of ϕ_n for $n \geqslant 3$.

5 GAP implementation

In order to obtain some idea of how well the above algorithm performs, we have produced a prototype implementation using the GAP computer language [19]. This can be downloaded from [15] and loaded into GAP using the command `Read("downloaded-file-name")`. The algorithm is then run using the command `Resolution(Gens,n)` where `Gens` is a set of generators for a finite group G and n is a positive integer. The output is a pair $[C, A]$ where C is a list of length n representing the first n terms of a free $\mathbb{Z}G$–resolution R_* of $\mathbb{Z}$, and A is a function representing the action of G on R_*. The i–th term of C is a pair $[c_i, d_i]$ where c_i is an integer and d_i is a function. The integer c_i is the number of free generators for the $\mathbb{Z}G$–module R_i, and the function d_i represents the boundary homomorphism.

We have represented the group G by the set of integers $[1, \cdots, |G|]$. Let $g \in G$ be represented by the integer $\overline{g}$ say, and let e^i_j denote the j–th free generator of R_i. The element $g \cdot e^i_j \in R_i$ is represented by the integer $(j-1)|G| + \overline{g}$, and for each $1 \leqslant j \leqslant c_i$ the function d_i returns a list of integers $d_i(j) = [a_1, \cdots, a_t]$ such that, using our integer representation, the boundary $\partial_i(e^i_j) \in R_{i-1}$ is

$$\partial_i(e^i_j) = a_1 + \cdots + a_t.$$

(Let X be the classifying space correpsonding to R_*. The integers in the list $d_i(j)$ in fact correspond to a set of oriented $(i-1)$–cells in $\tilde{X}$ which form a homotopical $(i-1)$–syzygy. The function d_i can thus be used to construct syzygies.)

Given an integer b representing $g \cdot e^i_j \in R_i$ and an integer $\overline{h}$ representing $h \in G$ the function A returns the integer $A(b, \overline{h})$ representing $(hg) \cdot e^i_j \in R_i$.

For example, the following GAP session

```
gap>R:=Resolution([(1,2),(1,2,3)],3);;
gap>c1:=R[1][1][1]; c2:=R[1][2][1]; c3:=R[1][3][1];
```

```
2
3
4
gap>d1:=R[1][1][2];; d2:=R[1][2][2];; d3:=R[1][3][2];;
gap>d3(1);
[-1,3]
gap>Elts:=Elements(Group([(1,2),(1,2,3)]));;
gap>Elts[3];
(1,2)
```

shows that S_3 admits a resolution R_* with two generators in dimension 1, three generators in dimension 2, and four generators in dimension 3. Furthermore, the boundary of the first generator of R_3 is $\partial_3(e_1^3) = -e_1^2 + (1,2)\cdot e_1^2$.

Continuing this session with

```
gap>Elts;
[ (), (2,3), (1,2), (1,2,3), (1,3,2), (1,3) ]
gap>d1(1); d1(2);
[-1, 3]
[-1, 4]
gap>d2(1); d2(2); d2(3);
[1, 3]
[10, 7, 11]
[-8, 1, -7, 2]
```

shows that S_3 is generated by two elements x and y subject to the three relations $x^2 = 1$, $y^3 = 1$, $xy^{-1}xy^{-1}$. At the chain complex level this gives $\partial_2(e_1^2) = (1+x)\cdot e_1^1$, $\partial_2(e_2^2) = (1+y+y^2)\cdot e_2^1$ and $\partial_2(e_3^2) = (1+xy^{-1})\cdot(e_1^1 - e_2^1)$. Then

```
gap>d3(4);
[8, -2, 13, -5, 15, 7, -1, 16]
```

shows that the fourth generator of R_3 corresponds to the following syzygy.

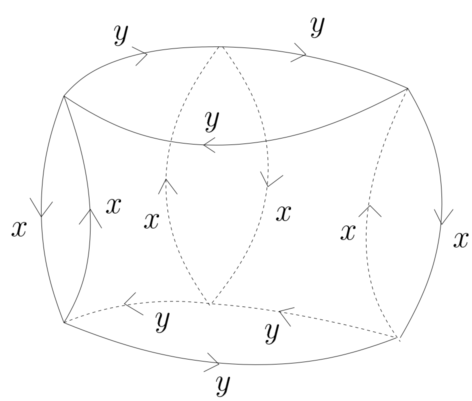

The elements in `Gens` do not have to be permutations. For instance,

```
gap>R:=Resolution(GeneratorsOfGroup(CyclicGroup(120)),100);;
gap>c100:=R[1][100][1];
1
```

establishes the well–known fact that the cyclic group of order 120 admits a resolution with precisely one generator in dimension 100. Here the GAP function `CyclicGroup(120)` creates the cyclic group of order 120 as a polycyclic group on five generators. Our implementation uses only the first of these generators.

The function `Resolution(Gens,n)` has been incorporated into a function `IntegralHomology(Gens,n)` for calculating the integral homology group $H_n(G,\mathbb{Z})$ of a group. For example

```
gap>IntegralHomology([(1,2),(1,2,3)],101);
[ 2 ]
gap>IntegralHomology([(1,2),(1,2,3)],102);
[  ]
gap>IntegralHomology([(1,2),(1,2,3)],103);
[ 6 ]
```

gives the well–known results $H_{101}(S_3,\mathbb{Z}) = \mathbb{Z}_2$, $H_{102}(S_3,\mathbb{Z}) = 0$ and $H_{103}(S_3,\mathbb{Z}) = \mathbb{Z}_6$.

The performance of our implementation for various groups G is summarised in Table 4. In each case a 5–presentation $\langle \underline{x} \,|\, \underline{r} \,|\, \underline{s} \,|\, \underline{s}_3 \,|\, \underline{s}_4 \rangle$ was computed and the number of generators, relators, and so on was recorded. The explicit generators are available in [15]. (We should remark that the table contains no interesting new resolutions since, for each group in the table, a smallish resolution could be obtained using standard theoretical methods.)

Name	Order	$\lvert\underline{x}\rvert$	$\lvert\underline{r}\rvert$	$\lvert\underline{s}\rvert$	$\lvert\underline{s}_3\rvert$	$\lvert\underline{s}_4\rvert$
S_3	6	2	3	4	5	6
A_4	12	2	3	4	5	6
S_4	24	2	3	4	6	9
S_4	24	3	6	10	15	20
$(C_5)^2$	25	2	3	4	5	6
D_{25}	50	2	3	5	6	7
A_5	60	2	3	4	6	9
D_{50}	100	2	3	5	6	7
S_5	120	2	4	7	12	19
S_5	120	4	10	20	35	56
$(C_5)^3$	125	3	6	10	15	21
D_{100}	200	2	3	4	5	6
A_6	360	3	7	15	51	?
D_{200}	400	2	3	4	5	6
D_{300}	600	2	3	4	5	7
$(C_5)^4$	625	4	10	20	35	56
S_6	720	2	5	13	102	?
S_6	720	5	15	35	78	161
D_{400}	800	2	3	4	5	6

Table 4

The symbol ? means that the computation did not complete in a reasonable time or else failed due to lack of space on a 1.8GHz Linux PC with 512Mb of RAM. The requirement for such a large workspace when handling these modest examples is a feature of our particular implementation rather than of the algorithm itself. The computer algebra language GAP provides an extremely friendly environment for programming the algorithm, and the inevitable price for this friendliness is a reduction in running speed. In an attempt to increase speed the implementation stores a large number of intermediate computations in an array so that they can be accessed quickly if needed at a later stage. Experimentation suggests that a C version of the implementation would certainly perform more than an order of magnitude faster than the prototype

GAP implementation.

The current implementation is extremely unsophisticated in the way it searches for a maximal subcomplex and in the way it searches for redundancies. More sophisticated searches should drastically increase speed. As an indication of the speed of the current implementation, we note that it took 50 seconds of CPU to compute the 5–presentation of the largest group D_{400} in Table 4. It took 55 seconds to produce the 5–presentation of S_5 with four (Coxeter) generators. By contrast, it took 1 second to calculate the first four dimensions of the two–generator presentation of S_5 and over 100 minutes to calculate the fifth dimension in this case!

Of course, for large groups the speed of the algorithm will be a problem even for a C implementation with sophisticated searches. One approach to this problem is to work relative to a subgroup $H < G$. This is discussed in the Section 6.

A second problem with the algorithm, which costs workspace and time, is that even for certain small groups the "size" of the boundary map can become problematically large. Let us define the *area* of an n–syzygy to be the number of faces in the corresponding polytope decomposition of the n–sphere. In the resolution R_* we can then define the *size* of the boundary map $\partial_n\colon R_n \to R_{n-1}$ to be the maximum of the areas of the $(n-1)$–syzygies corresponding to the free generators of R_n.

The dihedral group D_6 of order 12 provides a good example of how the size of the boundary map can suddenly explode even for an extremely small group. The permutations $x = (1,7)(2,9)(3,8)(4,12)(5,11)(6,10)$ and $y = (1,4,2,5,3,6)(7,10,8,11,9,12)$ generate this group. For the resolution R_* of D_6, produced by applying our implementation to these generators, the following table lists the number $|\underline{s}_{n-1}|$ of free generators of R_n and the size $|\partial_n|$ of the boundary map.

n	1	2	$\cdots$	12	$\cdots$	16	17	18	19	20	21
$\lvert\underline{s}_{n-1}\rvert$	2	3	$\cdots$	23	$\cdots$	48	58	71	83	96	111
$\lvert\partial_n\rvert$	2	8	$\cdots$	10	$\cdots$	106	141	217	348	700	1576

21–presentation for D_6

The explosion in $|\partial_n|$ is very much dependent on the choice of generators. For instance, when our implementation is applied to the permutation generators $x = (1,2,3,4,5,6)$, $y = (2,6)(3,5)$ of D_6 the boundary map has size $|\partial_{17}| = 3282$ in dimension 17. In general, as the size of ∂_n grows the algorithm requires more workspace and CPU time.

A possible method (not yet implemented) for keeping the boundary size in check would be to apply higher–dimensional versions of the well–known Tietze operations for 2–presentations of groups. Any n–syzygy σ can be decomposed (in many ways) into two n–dimensional discs D_1, D_2 whose boundaries lie in the $(n-1)$–skeleton. Suppose that σ and σ' are two n–syzygies admitting decompositions D_1, D_2 and D_1', D_2' with $D_2 = D_2'$. Then D_1 and D_1' combine to form an n–syzygy σ''.

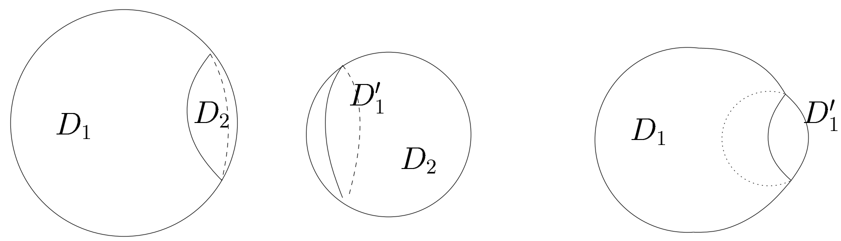

Suppose that σ and σ' belong to some complete set of n–syzygies $\underline{s}_n$. The basic higher–dimensional operation deletes σ from $\underline{s}_n$ and adds σ''. Under some obvious restrictions the resulting set is complete. If D_1' has smaller area than D_2 then clearly the area of σ'' is less than that of σ.

6 A refinement of the algorithm

The algorithm described in Section 3 is only practical for small groups. One approach to handling larger groups G is to work relative to a subgroup H. For normal subgroups we can directly use a result of C.T.C. Wall [39]: if H is *normal* in G and if C_* (resp. D_*) is a free $\mathbb{Z}H$ (resp. $\mathbb{Z}(G/H)$) resolution of $\mathbb{Z}$, then there exists a differential d and G–action on the graded abelian group $C_* \otimes_{\mathbb{Z}} D_*$ making it a free $\mathbb{Z}G$–resolution of $\mathbb{Z}$. From this result it follows that an n–presentation of the normal subgroup

$$H = \langle \underline{x}^H, \underline{r}^H, \underline{s}^H, \underline{s}_3^H, \underline{s}_4^H \cdots \rangle,$$

together with an n–presentation of the quotient group

$$Q = G/H = \langle \underline{x}^Q, \underline{r}^Q, \underline{s}^Q, \underline{s}_3^Q, \underline{s}_4^Q \cdots \rangle,$$

give rise to an n–presentation

$$G = \langle \underline{x}^G, \underline{r}^G, \underline{s}^G, \underline{s}_3^G, \underline{s}_4^G \cdots \rangle$$

in which $\underline{s}_i^G$ consists of one identity $s \otimes t$ for each $s \in \underline{s}_p^H, t \in \underline{s}_q^Q$ with $p+q=i$. (Here $\underline{s}_0^Q = \underline{x}^Q$, $\underline{s}_1^Q = \underline{r}^Q$ etc. .)

Example 6. Consider the central extensions $\mathbb{Z} \to G \to \mathbb{Z}\times\mathbb{Z}$ of the free abelian group on two generators by the infinite cyclic group. Since $H^2(\mathbb{Z}\times\mathbb{Z},\mathbb{Z}) \cong \mathbb{Z}$ there are countably many such extensions up to Yoneda equivalence. The group G occuring in these extensions is parametrized by non–negative integers n and can be presented as $G = \langle x, y, z : [x,y]z^n = 1, [x,z] = 1, [y,z] = 1\rangle$. When $n = 1$ it is the Heisenberg group. The group $\mathbb{Z}\times\mathbb{Z}$ has a 2–dimensional classifying space, namely the torus, with two 1–cells and one 2–cell. The group $\mathbb{Z}$ has a 1–dimensional classifying space, namely the circle, with one 1–cell. Wall's result thus implies that the extension G admits a 3–dimensional classifying space with three 1–cells, three 2–cells and one 3–cell. It can be shown [17] that the 3–cell is attached according to the following picture.

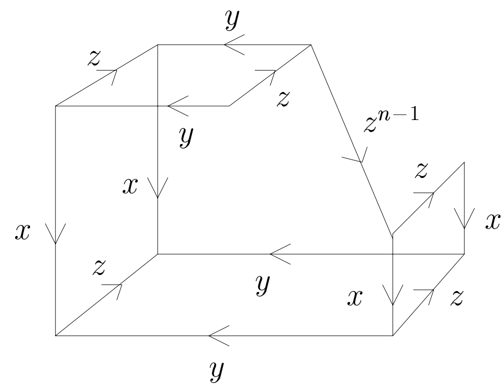

To calculate with Wall's result one needs a general method for describing the boundary of the identities $s \otimes t$. The Perturbation Lemma of R. Brown [3] and V. Gugenheim [22] can be used here, as discussed in [1]. One can also use a naive computer search which we now explain. We first attempt to give a pictorial explanation, up to dimesnion 3, in which Wall's assumption on the normality of H is dropped. Then, for the case when H is normal, we give a brief algebraic explanation of Wall's method that applies to all dimensions.

So let G be a finite group with subgroup H. As input data for the modified algorithm we supply two sets $\underline{x}^H$, $\underline{x}^Q$ of elements in G where $\underline{x}^H$ generates H and where $\underline{x}^Q$ generates a subgroup of G containing a right transversal of H in G. We assume that $Hgz \neq Hg$ for all $z \in \underline{x}^Q, g \in G$.

As a specific example let us consider the symmetric group $G = S_4$ with

$$\underline{x}^H = \{(1,2),(1,2,3)\}$$

$$\underline{x}^Q = \{(1,2,3,4)\}$$

and for convenience label these permutations $x = (1,2)$, $y = (1,2,3)$, $z = (1,2,3,4)$. The subgroup H is S_3.

Step (i). The algorithm begins by constructing a maximal tree Y^H in the Cayley graph $\Gamma^H = \Gamma(\underline{x}^H)$ of the group H, rooted at 1. It also constructs a maximal tree Y_*^Q in the Schreier graph $\Gamma^Q = \Gamma(\underline{x}^Q, G/H)$ of right cosets of H in G, rooted at H. Here Γ^Q is an oriented graph whose vertices are the right cosets of H in G, with an edge from coset Hg to coset Hgz for all $z \in \underline{x}^Q$ and $g \in G$. The tree Y_*^Q provides a choice of representative t_g for each right coset Hg. These representatives form a transversal $T = \{t_g\}$.

We let Y^Q denote a tree in the Cayley graph $\Gamma^G = \Gamma(\underline{x}^H \cup \underline{x}^Q)$ of G obtained by lifting Y_*^Q. Explicitly, the vertices of Y^Q are the elements of the transversal T; there is an edge from vertex t_g to vertex $t_{g'}$ if and only if $t_{g'} = t_g z$ for some $z \in \underline{x}^Q$.

For our particular choice of G and H the graph Γ^Q is a square cycle with four vertices and four edges. The tree Y^Q thus consists of four vertices and three edges. A choice of tree Y^H was given in Section 3 above.

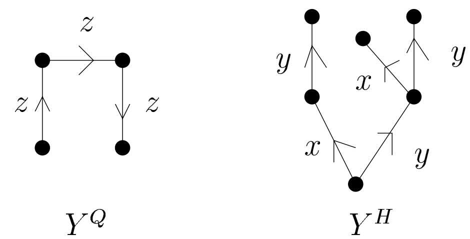

A maximal tree Y^G in the Cayley graph Γ^G can be constructed by "attaching" one copy of Y^Q to each vertex in Y^H. We denote by vY^Q the copy of Y^Q attached to vertex $v \in Y^H$.

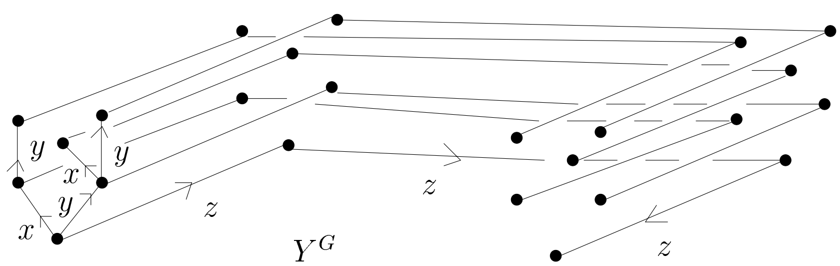

Steps (ii) and (iii). Those edges in Γ^G which lie in the complement of Y^G can be divided into three types, represented by dotted edges in the following pictures. The solid lines in the pictures represent paths in the tree Y^G (or in subtrees Y^H or vY^Q), v and v' are vertices in Y^H, and $x \in \underline{x}^H$, $z \in \underline{x}^Q$.

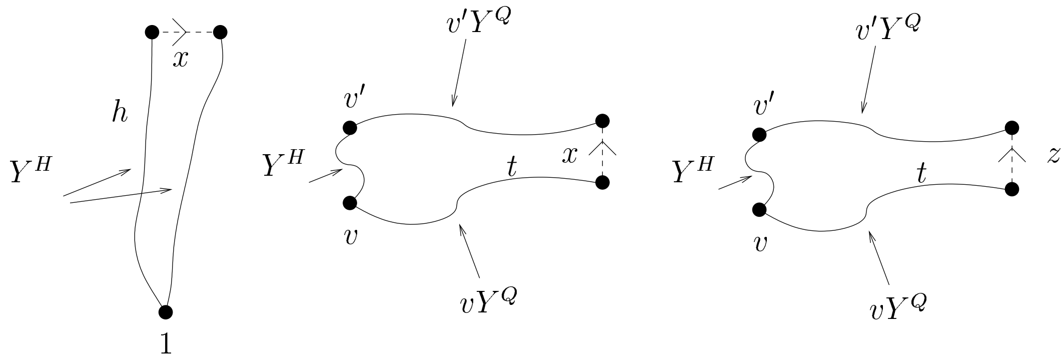

Type H　　　Type HQ　　　Type Q

As in Section 3 we take $\tilde{X}(1)$ to be the underlying space of Γ^G, and note that there is a fixed–point free action of G on this space. We set $X(1) := \tilde{X}(1)/G$. We need to attach 2–cells to $\tilde{X}(1)$ so as to form a simply connected space $\tilde{X}(2)$ admitting a fixed–point free cellular action of G. We will then have a 2–presentation $X(2) := \tilde{X}(2)/G$. We could of course simply attach one 2–cell $\tilde{e}^2_\lambda$ (and all its translates $g \cdot \tilde{e}^2_\lambda$) for each edge $\lambda \in \Gamma^G \setminus Y^G$, where the boundary of $\tilde{e}^2_\lambda$ spells the unique reduced relator determined by the edge λ and tree Y^G. But that would be uneconomical. We can get away with attaching just a smallish subset of these 2–cells $\tilde{e}^2_\lambda$ together with their translates, as we now explain.

Firstly, note that in general we need only attach 2–cells $\tilde{e}^2_\lambda$ of type HQ or Q with $v = 1$ since the translates of such cells will ensure the contractibility of all loops of type HQ or Q for arbitrary v. We denote by $\tilde{e}^2_{(x,t)}$ and $\tilde{e}^2_{(z,t)}$ (where $x \in \underline{x}^H, z \in \underline{x}^Q, t \in T$) the 2–cells of type HQ and Q with $v = 1$. A version of the reduction procedure described in Section 3 can be applied to try to eliminate other redundant 2–cells of type HQ and Q. (The reduction can be performed by iterating over T if H is normal or if T is a subgroup.)

It is often sufficient to attach only those 2–cells $\tilde{e}^2_{(x,z)}$ of type HQ with $x \in \underline{x}^H, z \in \underline{x}^Q$. This is certainly the case when H is normal in G because then the loops of type HQ have the form

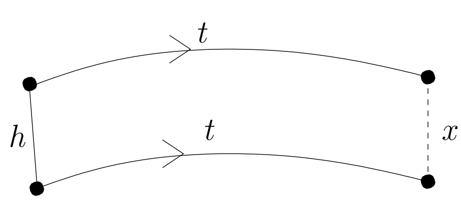

Type HQ (Normal subgroup H)

with $h \in H, t \in T, x \in \underline{x}^H$, and such a loop can be decomposed into smaller such loops with $t \in \underline{x}^Q$.

The reduction procedure of Section 3 can be applied (to the subgroup H) to eliminate redundant 2–cells of type H.

For the groups $H = S_3, G = S_4$ in question we obtain in this manner the 2–presentation

$$S_4 = \langle x, y, z \mid R_a := x^2, R_b := xyxy^{-2}, R_c := z^4, \\ R_{x\otimes z} := xzy^{-2}x^{-1}z^{-1}, R_{y\otimes z} := yzx^{-1}z^{-2}\rangle$$

where the relators R_a, R_b are of type H, the relators $R_{x\otimes z}, R_{z\otimes z}$ are of type HQ, and relator R_c is of type Q.

Steps (iv) - (vi). Let $X(2)$ be the 2–presentation that we have constructed for G. Its 1–cells and 2–cells correspond to generators and relators for G. A maximal contractible subcomplex $Y^G(2)$ in $\tilde{X}(2)$ can be constructed by attaching an appropriate 2–cell to $\tilde{X}(1)$ for each 1–cell in $\tilde{X}(1) \setminus Y^G$. (We assume that some translate of a 2–cell of type H is attached for each edge of type H, and that some translate of a 2–cell of type HQ is attached for each edge of type HQ.) The 2–cells in $\tilde{X}(2) \setminus Y^G(2)$ correspond to a generating set for the free abelian group $\pi_2(X(2))$. Let us define a generator $\sigma \in \pi_2(X(2))$ to be of type H, HQ or Q if it corresponds to a 2–cell $g \cdot \tilde{e}^2_\lambda$ where e^2_λ is of type H, HQ or Q respectively. Attaching a 3–cell to $\tilde{X}(2)$ for each generator of $\pi_2(X(2))$ we obtain a 3–presentation

$$G = \langle \underline{x}^H, \underline{x}^Q \mid \underline{r}^H, \underline{r}^{HQ}, \underline{r}^Q \mid \underline{s}^H, \underline{s}^{HQ}, \underline{s}^Q\rangle$$

in which there are three types of relator and three types of 2–syzygy. There are of course many redundant 2–syzygies. The geometry of the situation leads to the following two observations that can be used to remove at least some of these redundancies.

(1) If $\underline{r}^{HQ}$ just consists of one relator for each $(x, z) \in \underline{x}^H \times \underline{x}^Q$ then all 2–cells of type HQ can be included in $Y^G(2)$ in which case there are no syzygies of type HQ.

(2) Corresponding to each $1 \neq t \in T$ and $r \in \underline{r}^H$ there is a syzygy which we denote by $t \otimes r$ (though t and r don't fully determine the syzygy). This syzygy can be pictured as a cylinder whose right–hand end is sealed by a single disc corresponding to the relator r, and whose left–hand end is sealed by a patchwork of discs corresponding to relators of type H. The central part of the cylinder is a patchwork of discs corresponding to relators of type HQ.

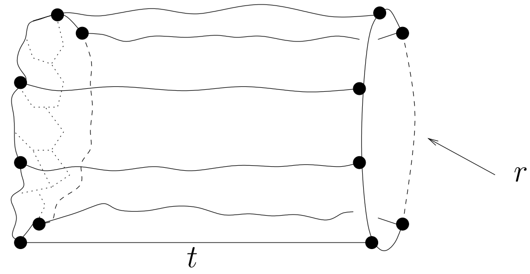

By applying the elimination procedure of Section 3 to the subgroup H we can probably find some smallish subset $\underline{s}'^H \subset \underline{s}^H$ that yields a 3–presentation

$$H = \langle \underline{x}^H \mid \underline{r}^H \mid \underline{s}'^H \rangle.$$

In our 3–presentation for G the set $\underline{s}^H$ can be replaced by the set

$$\underline{s}'^H \cup \{t \otimes r \mid 1 \neq t \in T, r \in \underline{r}^H\}.$$

(If H is normal in G then it suffices to include only those syzygies $t \otimes r$ with $t \in \underline{x}^Q$.)

Brute force searches could be used to try to remove further redundant 2–syzygies.

It seems that the above pictorial approach will run into difficulties in higher–dimensions, at least when H is not normal in G. So we now restrict attention to the case when H is normal and describe a more algebraic approach. As this case is covered in [39] our treatment is brief.

Let H now be a normal subgroup of the finite group G and set $Q = G/H$. The basic algorithm can be used to construct classifying spaces X^H and X^Q for the groups H and Q (up to some dimension). The chain complexes $C_*(\tilde{X}^H)$ and $C_*(\tilde{X}^Q)$ are then free $\mathbb{Z}H$– and $\mathbb{Z}Q$–resolutions of $\mathbb{Z}$. As explained in Section 4 the family of maximal contractible subcomplexes used to construct X^H can be used to algorithmically lift any group homomorphism $K \to H$ to a chain map $C_* \to C_*(\tilde{X}^H)$ from any free $\mathbb{Z}K$–chain complex C_*.

Let $\underline{b}_n$ denote the set of free $\mathbb{Z}Q$–generators for $C_n(\tilde{X}^Q)$. Now we can regard $C_*(\tilde{X}^Q)$ as a $\mathbb{Z}G$–chain complex. The stabilizer in G of each generator $b \in \underline{b}_n$ is precisely the subgroup H. The direct sum of chain complexes

$$S_{n*} := \bigoplus_{b \in \underline{b}_n} C_*(\tilde{X}^H) \otimes_{\mathbb{Z}H} \mathbb{Z}G$$

is thus a free $\mathbb{Z}G$–resolution of $C_n(\tilde{X}^Q)$. The $\mathbb{Z}G$–morphism $\partial_n \colon C_n(\tilde{X}^Q) \to C_{n-1}(\tilde{X}^Q)$ can be lifted to a chain map $\partial^h \colon S_{n*} \to S_{(n-1)*}$. This lifting can be constructed algorithmically using the family of maximal contractible subcomplexes in X^H. The differential in $C_*(\tilde{X}^H)$ induces a graded map $\partial^v \colon S_{*n} \to$

$S_{*(n-1)}$. Clearly $\partial^v\partial^v = 0$.

The collection S_{**} of free $\mathbb{Z}G$–modules is not in general a bicomplex because $\partial^h\partial^h \neq 0$. Nevertheless, we can construct a free $\mathbb{Z}G$–chain complex TS_* with $TS_n = \oplus_{p+q=n} S_{pq}$ and differential given by $\partial|_{S_{pq}} = \partial^h + (-1)^p(\partial^v + \epsilon)$ where $\epsilon: TS_n \to TS_{n-1}$ is a 'perturbation'. The definition of this perturbation homomorphism can be found in [39]. A spectral sequence argument involving the exactness of each column in S_{**} can be used to show that the total complex is a free $\mathbb{Z}G$–resolution of $\mathbb{Z}$. The main point is that both ∂^v and ϵ can be algorithmically computed from the family of maximal contractible subcomplexes in X^H.

7 Automatic groups

Our approach to computing classifying spaces for finite groups can, in principle, be applied to certain infinite groups. To see this let G be a group of *type* FP_∞. This means that there exists a reduced CW–classifying space X for G which has only finitely many cells in each dimension. As usual, the 1–skeleton of the universal cover $\tilde{X}$ can be identified with the Cayley graph $\Gamma = \Gamma(G, \underline{x})$ for some set $\underline{x}$ of generators for the group G. Furthermore, the 1–skeleton can be viewed as a metric space in which the distance between two vertices equals the least number of edges in any path between the two vertices. Let us define the *diameter* of an arbitrary n–cell $g \cdot \tilde{e}^n$ in $\tilde{X}$ to be the maximum distance between any pair of vertices in the boundary of the n–cell. Let us define the *diameter* of an n–cell e^n in X to be the diameter of its pre–image $\tilde{e}^n$ under the covering map $\tilde{X} \to X$.

Suppose that we do not explicitly know the CW–structure of a classifying space X for G but that we have to hand an n–presentation $X(n)$ for G and an integer r_{n+1} such that an $(n+1)$–presentation can be constructed by attaching to $X(n)$ a finite number of $(n+1)$–cells each of diameter at most $2r_{n+1}$. Suppose also that we have some means of computing a unique normal form for words in the generators of G and their inverses. Then we can compute an $(n+1)$–presentation for G as follows.

Let B denote the CW–subspace of $\tilde{X}(n)$ consisting of the $(n-1)$–skeleton $\tilde{X}(n-1)$ together with those n–cells whose vertices lie a distance at most r_{n+1} from the identity vertex. Let Y be a subcomplex of B which contains $\tilde{X}(n-1)$ and is maximal with respect to the property that Y is homotopy equivalent to $\tilde{X}(n)$. (Equivalently, Y is maximal with respect to the property that $H_n(Y, \mathbb{Z}) = 0$. This equivalent condition allows one to algorithmically construct Y since $H_n(Y, \mathbb{Z})$ is completely determined by the boundary of the

n–cells in Y.) There are only finitely many n-cells in B, some of which lie in Y and some of which lie in the complement $B \setminus Y$. Each n-cell of B lying in the complement of Y determines an n-syzygy. By the definition of r_{n+1} the finite set $\underline{s}_n$ of such syzygies generates the $\mathbb{Z}G$–module $\pi_n(X(n))$. A version of the elimination procedure given in Section 3 could be used to eliminate redundant generating syzygies, thus producing a smaller generating set $\underline{s}'_n$. The space $X(n+1)$ is obtained from $X(n)$ by attaching one $(n+1)$–cell for each syzygy in s'_n.

The above construction can be applied recursively to form a classifying space X, starting with $X(1)$ equal to a wedge of circles, provided there is some means of calculating the numbers r_{n+1}. This is clearly the case if G is finite. More generally, we now explain why it is the case when G is *automatic*.

Recall that a *finite state automaton* W consists of an alphabet $\underline{x}$, a finite set of states S and a function $\underline{x} \times S \to S$. One state $s_0 \in S$ is deemed to be the *start state*, and a (possibly empty) subset of states are deemed to be *accept states*. The idea is that the automaton starts in state s_0 and reads in a *word*, *i.e.* a finite (possibly empty) string of letters. After each letter is read the machine changes its state. If the final state is an accept state then the input word is deemed to have been accepted. The collection $L(M)$ of accepted words is referred to as the *language* of M.

An automaton can be represented as a directed graph, with a node for each state and an arrow for each transition between states. For example, the following directed graph represents an automaton over the alphabet $\underline{x} = \{x, y, z\}$ with eight states, all but one of which are accept states. The non–accept state has been omitted from the graph, as have all transition arrows to this state. The language of the automaton is finite and consists of 24 words representing the elements of the symmetric group S_4 generated by $x = (1,2), y = (2,3), z = (3,4)$.

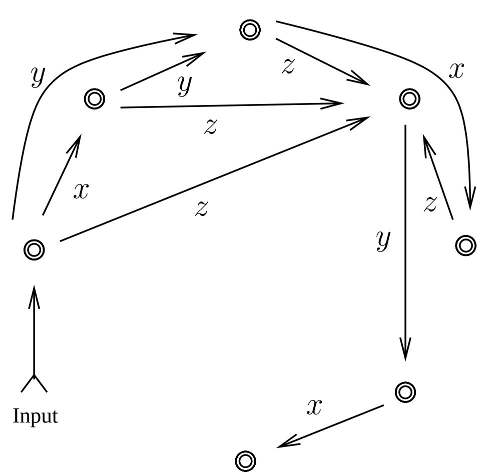

We recall from [18] that an *automatic structure* on a group G consists of a set $\underline{x}$ of monoid generators of G, a finite state automaton W over $\underline{x}$, and finite state automata M_x over $(\underline{x},\underline{x})$, for $x \in \underline{x} \cup \{\epsilon\}$, satisfying the following conditions. (The alphabet $(\underline{x},\underline{x})$ is by definition the set $(\underline{x} \cup \{\epsilon\}) \times (\underline{x} \cup \{\epsilon\}) \setminus \{(\epsilon,\epsilon)\}$ where ϵ represents the identity in G.)

(1) The canonical map $L(W) \to G$ is surjective.
(2) For $x \in \underline{x} \cup \{\epsilon\}$, we have $(w_1, w_2) \in L(M_x)$ if and only if both w_1 and w_2 are words in $L(W)$ and the two words $w_1 x$, w_2 represent the same element in G. (Words in $L(M_x)$ are strings of pairs, but here we are viewing them as pairs of strings.)

A group G is said to be *automatic* if it admits an automatic structure. The automatic structure is a succinct form for representing G in a computer, and yields an effective solution to the word problem for G. Clearly every finite group admits an automatic structure. Many infinite groups do too. For example all euclidean groups (characterised by having a free abelian subgroup of finite index) and braid groups are automatic [18].

Let G be an automatic group. For each $x \in \underline{x} \cup \{\epsilon\}$ let c_x denote the most transitions needed to get from any non–accept state of M_x to any given accept state. Set

$$k = 2(\max_{x \in \underline{x} \cup \{\epsilon\}} c_x) + 1.$$

This integer k is called a *lipschitz constant* for the automatic structure.

To explain the terminology, note that a word $w = x_{i_1} x_{i_2} \cdots x_{i_n}$ in the generators $\underline{x}$ corresponds to a path in the Cayley graph $\Gamma = \Gamma(G, \underline{x})$ starting at the identity vertex and passing through the vertices $w(t)$ $(0 \leqslant t \leqslant n)$ where $w(t)$ is the element of G represented by $x_{i_1} x_{i_2} \cdots x_{i_t}$. It is convenient to think of this path as a map $w\colon [0,\infty] \to \Gamma$ and so we define $w(t) = w(n)$ for $t \geqslant n$. If two words $w_1, w_2 \in L(W)$ and $x \in \underline{x} \cup \{\epsilon\}$ are such that $w_1 x$ and w_2 both represent the same element in G, then

$$d(w_1(t), w_2(t)) \leqslant k \qquad \text{for all } t \geqslant 0$$

where d is our metric on the Cayley graph Γ. This inequality is an important and well–known property of automatic groups [18]. To obtain the inequality note that, for any given $t \geqslant 0$, there exist words u_1, u_2 such that $w_1(t) u_1 x$ and $w_2(t) u_2$ represent the same element in G; the lengths of u_1 and u_2 need be no more than c_x.

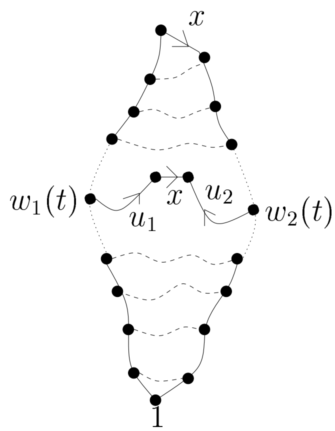

Since the word $u_1 x u_2^{-1}$ has length at most k, we get $d(w_1(t), w_2(t)) \leqslant k$.

We now consider the construction of a classifying space X for an automatic group G with monoid generators $\underline{x}$ and lipschitz constant k. In the construction we shall denote by $B(\tilde{X}(n), r)$ the subcomplex of the CW-space $\tilde{X}(n)$ consisting of those cells whose vertices lie a distance at most r from the identity vertex.

Let $\underline{x}'$ be a maximal subset of $\underline{x}$ with the property that $x \in \underline{x}'$ only if $x^{-1} \notin \underline{x}'$. We take the 1–skeleton $X(1)$ to be a wedge of circles, one circle for each element in $\underline{x}'$. We set r_2 equal to the least integer $\geqslant (k+1)/2$. We use the automatic structure to compute the subgraph $B(\tilde{X}(1), r_2)$ of the Cayley graph $\tilde{X}(1) = \Gamma(G, \underline{x}')$. Starting from a maximal tree in $B(\tilde{X}(1), r_2)$ we determine a smallish set of 1–syzygies $\{\sigma^1_\lambda : S^1 \to B(\tilde{X}(1), r_2)\}_{\lambda \in \Lambda_1}$ with the following property: let $\tilde{X}(2)$ be obtained from $\tilde{X}(1)$ by attaching one 2–cell $g \cdot e^2_\lambda$ via $g \cdot \sigma^1_\lambda$ for each $g \in G, \lambda \in \Lambda_1$; the fundamental group of the finite space $B(\tilde{X}(2), r_2)$ is to be trivial. To see that this property implies $\pi_1 \tilde{X}(2) = 0$ consider an arbitrary 1–syzygy $\sigma: S^1 \to \tilde{X}(1)$. From each vertex $\sigma(t)$ of this 1–syzygy there is at least one path w_t in $\tilde{X}(1)$ connecting $\sigma(t)$ to the identity vertex, and such that the word corresponding to w_t lies in the language $L(W)$.

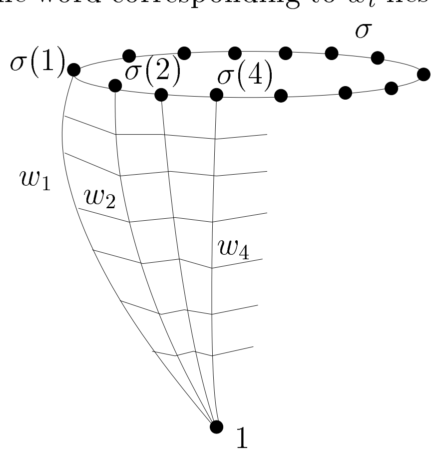

The 1–syzygy σ can thus be decomposed into smaller 1–syzygies with the following maximum dimensions (with respect to our path metric on $\tilde{X}(1)$).

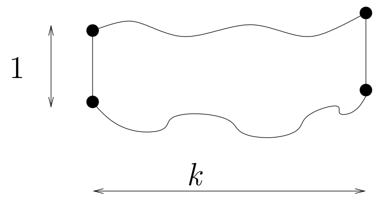

Each of these smaller syzygies can be isometrically translated to a 1–syzygy in $B(\tilde{X}(2), r_2)$ and can thus be further decomposed into syzygies of the form $g \cdot \sigma^1_\lambda$ ($g \in G, \lambda \in \Lambda_1$). This means that our arbitrary 1–syzygy σ can be homotopied to a point in $\tilde{X}(2)$, proving that $\pi_1 \tilde{X}(2) = 0$.

For the construction of $X(3)$ we choose the least integer d_1 such that no two vertices in any 1–syzygy σ^1_λ ($\lambda \in \Lambda_1$) lie more than a distance d_1 apart (with respect to our metric on G). In other words, d_1 is the maximum diameter of any 1–syzygy σ^1_λ. For any $n \geqslant 0$ we let $\overline{n}$ denote an integer with the property that, if $g \in G$ lies a distance $d(1, g) \leqslant n$ from the identity, then g is represented by a word $w \in L(W)$ of length at most $\overline{n}$. We set $r_3 = \overline{kd_1 + 1}$.

Starting from a maximal contractible CW–subspace in $B(\tilde{X}(2), r_3)$ we determine a smallish set of 2–syzygies $\{\sigma^2_\lambda : S^2 \to B(\tilde{X}(2), r_3)\}_{\lambda \in \Lambda_2}$ with the following property: let $\tilde{X}(3)$ be obtained from $\tilde{X}(2)$ by attaching one 3–cell $g \cdot e^3_\lambda$ via $g \cdot \sigma^2_\lambda$ for each $g \in G, \lambda \in \Lambda_2$; the second homotopy group of the finite space $B(\tilde{X}(3), r_3)$ must be trivial. To see that this implies $\pi_2 \tilde{X}(3) = 0$ consider an arbitrary 2–syzygy $\tau : S^2 \to \tilde{X}(2)$. From each vertex of this 2–syzygy there is at least one path in $\tilde{X}(1)$ connecting the given vertex to the identity vertex, and such that the word corresponding to this path lies in the language $L(W)$.

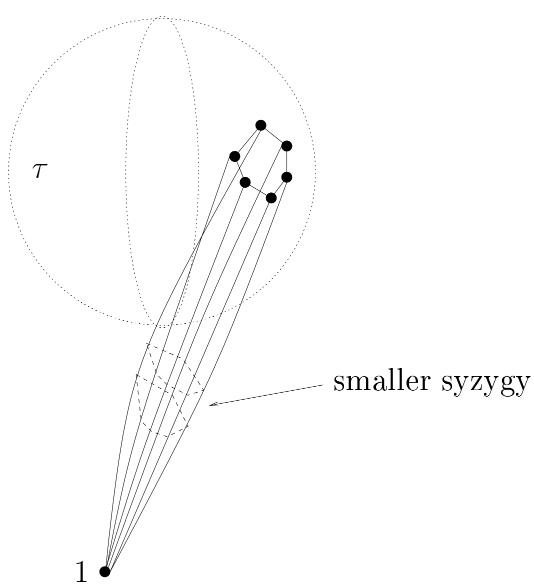

The 2–syzygy τ can thus be decomposed into smaller 2–syzygies with the following maximum dimensions.

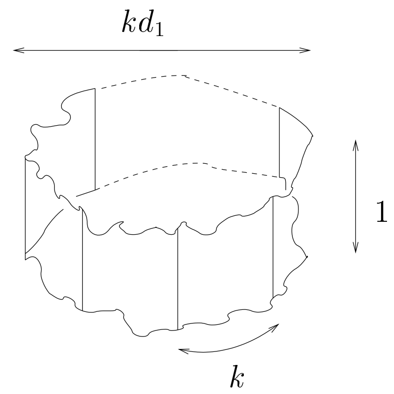

Each of these smaller syzygies can be isometrically translated to a 2–syzygy in $B(\tilde{X}(3), r_3)$ and can thus be further decomposed into syzygies of the form $g \cdot \sigma^2_\lambda$ $(g \in G, \lambda \in \Lambda_2)$. Thus τ can be homotopied to a point in $\tilde{X}(3)$.

The construction can be repeated in higher dimensions, using $r_{n+2} = \overline{kd_n + 1}$ with d_n is the maximum diameter of an $(n+1)$–cell.

8 Crossed modules

The preceding sections describe a method for calculating the cohomology of a space with given fundamental group and trivial homotopy groups in dimensions $i \neq 1$. We now briefly turn our attention to the problem of computing the cohomology of a path–connected space X with $\pi_i X = 0$ for $i \geqslant 3$. The homotopy type of such a space is determined by its fundamental group π_1, its second homotopy group π_2 regarded as a π_1–module, and a cohomology class $k \in H^3(\pi_1, \pi_2)$. For current purposes the homotopy type of X can be more conveniently encoded as a cat^1–group (see [29] for details).

We recall that a *cat^1–group* consists of a group G and two homomorphisms $s, t \colon G \to G$ satisfying: $ss = s, st = t, tt = t, ts = s$ and $[\ker t, \ker s] = 1$. Given a cat^1–group we define the groups $\pi_1 = s(G)/t(\ker s)$, $\pi_2 = \ker s \cap \ker t$, $M = \ker s$, $P = \text{image}\, s$ and take $\partial \colon M \to P$ to be the restriction of the homomorphism t. It is an easy exercise to show that π_2 is abelian and that congugation in G determines an action of π_1 on π_2 and an action of P on M. The *crossed extension*

$$1 \to \pi_2 \to M \xrightarrow{\partial} P \to \pi_1 \to 1$$

represents a cohomology class in $H^3(\pi_1, \pi_2)$. The homotopy groups and cohomology class thus defined agree with those of the space X represented by the cat^1–group.

For any π_1–module A we can define the (co)homology

$$H_n((G,s,t),A) = H_n(X,A), \qquad H^n((G,s,t),A) = H^n(X,A).$$

Some results on this (co)homology can be found in [14,11,9]. When the group G is finite the methods of the previous sections can be used to compute this (co)homology in dimensions $n \leqslant 3$.

As input data for our computation we take a set $\underline{y}$ of generators for the finite group G together with the sets $s\underline{y} = \{s(y) : y \in \underline{y}\}$ and $t\underline{y} = \{t(y) : y \in \underline{y}\}$. These three sets completely determine the cat^1–group G. The output will be the 4–skeleton $X(4)$ of a CW–space X whose (co)homology equals the (co)homology of the cat^1–group (G,s,t) in dimensions $\leqslant 3$. The associated chain complex $C_*(\tilde{X}(4))$ can be used to compute the cohomology of X in dimensions $\leqslant 3$.

We begin our computation of $X(4)$ by testing to see whether there is some small quotient of G that represents the same homotopy type. This is done by searching for a large subgroup N in the image of s such that: i) N is normal in G, and ii) $N \cap t(\ker s) = 1$. If N exists then the quotient group $\overline{G} = G/N$ inherits the structure of a cat^1–group and the quotient homomorphism $G \twoheadrightarrow \overline{G}$ induces isomorphisms on π_1 and π_2. Thus $\overline{G}$ represents the same homotopy type as G and so we set $G := \overline{G}$. As above, set $P = \operatorname{image} s$ and $M = \ker s$.

We next compute:

(1) a set $\underline{x}$ of generators for the group P.
(2) a subset $\underline{z}$ of M that: (i) generates M as a P–group, and (ii) contains a subset $\underline{z}_0 \subset \underline{z}$ of module generators for the π_1–module $\pi_2 = \ker(\partial\colon M \to P)$.
(3) a set $\underline{r}_1 = \{r_z \in F(\underline{x}) : z \in \underline{z}\}$ where r_z is an element in the free group on $\underline{x}$ whose image in P is equal to $\partial(z)$. We take r_z to be the identity element if $\partial(z) = 1$.
(4) a set $\underline{r}_2$ of relators in $F(\underline{x})$ that present the group P. The set r_2 can be computed using the method described in Section 3.

We take $X(2)$ to be the 2–dimensional CW–space corresponding to the presentation $\langle \underline{x} \,|\, \underline{r}_1 \cup \underline{r}_2 \rangle$ for π_1. There is a surjective homomorphism $\pi_2 X(2) \twoheadrightarrow \pi_2$ because $\underline{z}$ contains a generating set for π_2.

We next compute a maximal contractible subspace $Y(2)$ in $\tilde{X}(2)$. The 2–cells in the complement of $Y(2)$ correspond to free generators of the free abelian group $\pi_2 X(2)$. We construct $X(3)$ by attaching one 3–cell to $X(2)$ for each 2–cell in the complement $\tilde{X}(2) \setminus Y(2)$. The attaching maps are determined from calculations in the finite module π_2 and are such that $\pi_2 X(3) = \pi_2$.

There will usually be many redundant 3–cells (in the sense that many could be omitted without changing $\pi_2 X(3)$. Using a version of the elimination procedure described in Section 3 we remove redundant 3–cells to obtain a smaller space $X(3)$.

Let $X'(4)$ denote a space obtained by attaching 4–cells to $X(3)$ in such a way that $\pi_3 X'(4) = 0$. Since the Hurewicz homomorphism $\pi_3 \tilde{X}'(4) \twoheadrightarrow H_3(\tilde{X}'(4), \mathbb{Z})$ is surjective we have $H_3(\tilde{X}'(4)) = 0$. By contrast, let $X(4)$ denote any space obtained by attaching 4–cells to $X(3)$ in such a way that $H_3(\tilde{X}(4)) = 0$. As explained in Proposition 5 of [14] there are isomorphisms $H_n(X'(4), A) \cong H_n(X(4), A)$ and $H^n(X'(4), A) \cong H^n(X(4), A)$ for $n \leqslant 3$ and any π_1–module A.

Let $Y(3)$ be a subspace of $\tilde{X}(3)$ containing $\tilde{X}(2)$ and and maximal with respect to the property that $H_3(Y, \mathbb{Z}) = 0$. The 3–cells in the complement $\tilde{X}(3) \setminus Y(3)$ correspond to free generators of the free abelian group $H_3(\tilde{X}(3), \mathbb{Z})$. Using the method of Section 3 we can attach a smallish number of 4–cells to $X(3)$ to obtain a space $X(4)$ with $H_3(\tilde{X}(4), \mathbb{Z}) = 0$.

Acknowledgments

The author would like to thank the referee for many detailed and helpful comments on an earlier draft of the paper. The author would also like to thank Robert Morse for help with GAP.

References

[1] D.W. Barnes and L.A. Lambe, 'A fixed point approach to homological perturbation theory', *Proc. Amer. Math. Soc.* 112 (1991), 881-892.

[2] The Bergman gröbner basis package, University of Stokholm, http://servus.matematik.su.se/bergman .

[3] R. Brown, 'The twisted Eilenberg–Zilber theorem', *Simposio Topologia (Messina, 1964)* (1965), 33-37.

[4] R. Brown, 'On the second relative homotopy group of an adjunction space: an exposition of a theorem of J. H. C. Whitehead', *J. London Math. Soc. (2)*, no. 1, 22 (1980), 146-152.

[5] R. Brown and J. Huebschmann, 'Identities among relations', in *Low dimensional topology (ed. R. Brown and T.L. Thickstun)*, LMS Lecture Note Series 48 (1982), 153-202.

[6] R. Brown and A. Razak Salleh, 'Free crossed resolutions of groups and presentations of modules of identities among relations', *LMS J. Comput. Math.*, (1999), 28-61.

[7] J. Cannon, 'Construction of defining relators for finite groups', *Discrete Math.* 5 (1973), 105-129.

[8] J. Carlson, 'Computing group cohomology: tests for completion', Computational Algebra and Number Theory (Milwaukee, 1996), *J. Symbolic Comput.* 31 (2001), 229-242.

[9] J.M. Casas, G. Ellis, M. Ladra and T. Pirashvili, 'Derived functors and the homology of n–types', *J. Algebra* 256 no. 2 (2002), 583-598.

[10] H. Cohen, *A course in computational algebraic number theory*, Graduate Texts in Math., 138 (Springer–Verlag, 1993).

[11] T. Datuashvili and T. Pirashvili, 'On (co)homology of 2–types and crossed modules', *J. Algebra* 244, no. 1 (2001), 352-365.

[12] C. De Concini and M. Salvetti, 'Cohomology of Coxeter groups and Artin groups', *Math. Res. Lett.* 7 (2000), 213-232.

[13] P. Dehornoy and Y. Lafont, 'Homology of Gaussian groups', *Annales Inst. Fourier (Grenoble)*, 53 (2003), 489-540.

[14] G. Ellis, 'Homology of 2–types', *J. London Math. Soc. (2)* 46 (1992), 1-27.

[15] G. Ellis, GAP code for computing homotopical syzygies, http://www.maths.nuigalway.ie

[16] G. Ellis and I. Kholodna, 'Three–dimensional presentations for the groups of order at most 30', *LMS J. Math. Comp.* Vol. 2 (1999), 93-117.

[17] G. Ellis and E. Sköldberg, 'Cohomology rings of some non-spherical Artin groups', preprint (http://www.maths.nuigalway.ie).

[18] D.B.A. Epstein, J. Cannon, D. Holt, S. Levy, M. Paterson, W. Thurston, *Word processing in groups*, (Jones and Bartlett, Boston, 1992).

[19] The GAP Group, 'GAP – Groups, Algorithms and Programming', School of Mathematical and Computational Sciences, University of St. Andrews, Scotland (1998).

[20] J. Grabmeier and L. Lambe, 'Computing resolutions over finite p–groups', *Algebraic combinatorics and applications (Gössweinstein, 1999)*, 157-195 (Springer 2001).

[21] J.R.J. Groves, 'An algorithm for computing homology groups', *J. Algebra* 194 (1997), 331-361.

[22] V. Gugenheim, 'On a theorem of E.H. Brown', *Illinois J. Math.* 4 (1960), 292-311.

[23] J. Harris, PhD thesis, National University of Ireland, Galway (in preparation).

[24] A. Heyworth and B. Reinert, 'Reduction in $\mathbb{Z}G$–modules with applications to identities among relators', preprint (http://www.mcs.le.ac.uk/~ah83).

[25] A. Heyworth and C.D. Wensley, 'Logged rewriting and identities among relators' in *Groups–St.Andrews 2001 in Oxford*, London Math. Soc. Lecture Note Ser. 304, eds. Colin Campbell, Edmund Robertson & Geoff Smith, Cambridge Univ. Press (2003).

[26] M. Kapranov and S. Saito, 'Hidden Stasheff polytpoes in algebraic K–theory and the space of Morse functions', *Contemp. Math. 227* (1999), 191-225.

[27] I. Kholodna, *Low–dimensional homotopical syzygies*, PhD thesis, National University of Ireland, Galway, 2001.

[28] L. Lambe, 'Homological perturbation theory, Hochschild homology, and formal groups', *Contemp. Math.* 134, (1992), 183-218.

[29] J.–L. Loday, 'Spaces with finitely many non–trivial homotopy groups', *J. Pure Applied Algebra* 24 (1982), 179-202.

[30] J.–L. Loday, 'Homotopical sysygies', *Contemp. Math.* 265 (2000), 99-127.

[31] *The Magma computational algebra system*, (http://magma.maths.usyd.edu.au)

[32] *MAGNUS software*, (http://zebra.sci.ccny.cuny.edu/web).

[33] J. Neubüser, 'An elementary introduction to coset table methods in computational group theory' in *Groups–St. Andrews 1981*, London Math. Soc. Lecture Note Ser., 71 (1982), 1-45.

[34] S.J. Pride, 'Identities among relations of group presentations' in: *Group theory from a geometrical viewpoint, Trieste 1990* (E. Ghys, A. Haefliger, A. Verjofsky editors), World Scientific Publishing (1991), 687-717.

[35] M. Salvetti, 'The homotopy type of Artin groups', *Math. Res. Lett.* 1, no. 5 (1994), 565-577.

[36] A.J. Sieradski, 'Algebraic topology for two dimensional complexes' in *Two–dimensional homotopy and combinatorial group theory, (ed. C. Hog–Angeloni, W. Metzler and A.J. Sieradski)*, LMS Lecture Note Series 197 (1993), 51-96.

[37] C.C. Squier, 'The homological algebra of Artin groups', *Math. Scand.* 75, no. 1 (1994), 5-43.

[38] R.G. Swan, 'Periodic resolutions for finite groups', *Ann. of Math. (2)* 72 (1960), 267-291.

[39] C.T.C. Wall 'Resolutions for extensions of groups', *Proc. Cambridge Philos. Soc.* 57 (1961), 251-255.

[40] J.H.C. Whitehead, 'Combinatorial homotopy II', *Bull. Amer. Math. Soc.* 55 (1949), 453-496.