J. Symbolic Computation (1999) **11**, 1–000

An Algorithmic Approach to Fundamental Groups and Covers of Combinatorial Cell Complexes

SARAH REES and LEONARD H. SOICHER

Department of Mathematics, University of Newcastle, Newcastle NE1 7RU, UK,
e-mail: Sarah.Rees@ncl.ac.uk
School of Mathematical Sciences, Queen Mary and Westfield College,
University of London, Mile End Road, London E1 4NS, UK,
e-mail:L.H.Soicher@qmw.ac.uk

(Received)

We first develop a construction, originally due to Reidemeister, of the fundamental group and covers of a 2-dimensional combinatorial cell complex. Then, we describe a practical algorithmic approach to the computation of fundamental groups and first homology groups (as finitely presented groups), of first homology groups mod p (as vector spaces), of deck groups (as permutation groups), and of covers of finite simple such complexes. In the case of clique complexes of finite simple graphs, the algorithms described have been implemented in GAP, making use of the GRAPE package.

Mathematics Subject Classifications (1991): 57M05, 57M10, 57M15, 57M20.
Keywords: cell complex, algorithms, homotopy, fundamental group, covers.

Acknowledgments

We thank Klaus Lux, Armando Martino and Martin Roller for very helpful discussions during the preparation of this paper, and the EU for its support of our collaboration via HCM grant no. ERBCHRXCT930418, 'Computational Group Theory'. The first author thanks the RWTH-Aachen and the Universität Bielefeld for their warm hospitality while much of this work was done, and the DAAD (and also the EU) for financial support in Germany.

1. Introduction

Finite 2-complexes and their covers arise naturally in the studies of finite geometries, groups and graphs. In particular, we are interested in problems such as determining the fundamental group of a finite 2-complex and classifying the r-fold covers of that complex for a given r. This has led us to a combinatorial and algorithmic approach to the study of finite 2-complexes, which is the main subject of this article. Furthermore, we have been interested in complexes with an associated group of automorphisms, and the relationship of this group action to the covers of that complex.

We shall describe a practical algorithmic approach to the computation of fundamental

0747–7171/90/000000 + 00 $03.00/0

groups, first homology groups, deck groups, and covers, in the general context of finite simple 2-dimensional combinatorial cell complexes. For the case of clique complexes of finite simple graphs, the algorithms described have been implemented in GAP (Schönert *et al.*, 1994), making use of the GRAPE package (Soicher, 1993). This implementation should eventually be included in GRAPE.

Our definition of a 2-dimensional combinatorial cell complex basically follows the definition of a 'Flächenkomplex' in (Reidemeister, 1951, chapter 5), chapter 5; we are slightly more restrictive than Reidemeister. The definitions and the basic results of section 3, and the basis of the graph-theoretic imagery of section 4 are found in (Reidemeister, 1951, chapters 4, 5, and 6), although our notation is slightly different from Reidemeister's. One of our aims has been to consolidate and extend Reidemeister's point of view, and we do this in the later part of section 4. General topological background is provided by (Maunder, 1980).

The term 'combinatorial cell complex' is also used in (Aschbacher, 1996) to denote a partially ordered set with some additional structure. Although we prefer to think of our combinatorial 2-complexes as graphs with specified sets of subgraphs, any one of them can certainly be interpreted in a natural way as one of Aschbacher's. In fact each of our combinatorial 2-complexes gives a 'restricted combinatorial cell complex' in the sense of Aschbacher when the set of all cells in the complex is equipped with the partial order which relates pairs of distinct incident cells in an order determined by dimension.

The remainder of this paper is divided into four parts. Section 2 contains a very brief history of the use of combinatorial topology which motivated our own study, section 3 contains the basic definitions and results for the algebraic topology of combinatorial cell complexes necessary for this paper, section 4 describes the view of the fundamental group and covering spaces of such objects which is used by the algorithms, and section 5 describes the algorithms.

2. Some History

The standard algebraic topological notions of homotopy, coverings, fundamental groups, and homology have been studied in various, essentially equivalent, combinatorial settings.

Reidemeister developed a theory for graphs, and then for 'Flächenkomplexe' (and hence simplicial complexes) in (Reidemeister, 1951). His approach (which is also described in (Seifert and Threlfall, 1934), and in its English translation (Seifert and Threlfall, 1980), where Flächenkomplexe become 'surface complexes') has descended into mathematical folklore, and is the basis of our own treatment.

In (Tits, 1981), Tits developed a theory of various types of covers for chamber systems, and in particular for locally finite incidence geometries, and derived a local characterisation of buildings. Further, he proved that certain local properties of an incidence geometry (that is, the fact that certain subdiagrams do not occur in its diagram) ensure that its universal (2-) cover is a building, and hence that it can be found as a quotient of that building. Geometries of this type for various sporadic and exceptional groups are described in (Kantor, 1981; Kantor, 1987; Kantor, 1985; Ronan, 1984). Tits' theory of covers for chamber systems was further developed in (Ronan, 1980; Surowski, 1984; Pasini, 1994); in particular Ronan proved that the systems of covers of a chamber system correspond to the systems of topological covers of a related simplicial complex.

Tits developed in (Tits, 1986) a theory of homotopy and covers for partially ordered sets, and in particular for the system of flags of a finite geometry; he extended results of

Serre for groups acting on trees, and proved that a flag-transitive group of automorphisms of a simply connected finite geometry (that is, one with trivial fundamental group) is a free product with amalgamation (over flag stabilisers) of vertex stabilisers, for the vertices in a fixed maximal flag. (This result was also proved independently by Pasini (see (Pasini, 1994)) and by Shpectorov.) Subsequently, various sporadic simple groups were described as free amalgams, through their relationship to simply connected finite geometries, as, for example in (Ivanov and Shpectorov, 1990; Aschbacher and Segev, 1991). In (Aschbacher and Segev, 1992a) a general theory of uniqueness systems is developed which allows proofs of the uniqueness of many of the sporadic simple groups, through their representation as free amalgams, associated with a simply connected finite simplicial complex.

The homotopy invariants of the partially ordered sets $\mathcal{S}_p(G)$ of p-subgroups, and $\mathcal{A}_p(G)$ of elementary abelian p-subgroups of a finite group G were investigated by Quillen in (Quillen, 1978). Quillen's conjecture for various classes of groups was investigated in (Aschbacher and Kleidman, 1990; Aschbacher and Smith, 1993); appropriate topological machinery in the context of simplicial complexes is developed in (Aschbacher and Segev, 1992b).

We also remark that covers of finite clique complexes arise naturally in the study of graphs that are locally a given graph and in the study of distance-regular antipodal covers of distance-regular graphs (see (Brouwer, Cohen and Neumaier, 1989; Godsil, Jurišić and Schade, 1990)).

3. Definitions

3.1. Definition of a combinatorial cell complex

We define a *2-dimensional combinatorial cell complex* Γ to be a non-empty undirected graph, possibly with multiple edges but without loops, in which certain subgraphs, each consisting of the vertices and edges of a simple circuit, are specified. A specified subgraph is called a *face* or *2-cell* of Γ, an edge a *1-cell* and a vertex a *0-cell*. An i-cell is defined to have *dimension* i. A symmetric *incidence* relation, inherited from the graph, relates some pairs of cells of different dimensions; each edge is incident with the two vertices at its ends, and a 2-cell with the vertices and edges which make it up. We define the *star* of a cell to be the set consisting of that cell together with all cells incident with it. For brevity, we shall often use the term *2-complex* to mean 2-dimensional combinatorial cell complex.

In our definition we differ from Reidemeister's definition of Flächenkomplexe in (Reidemeister, 1951) by outlawing loops and requiring that the circuits which define 2-cells are simple (i.e. traverse no vertex or edge twice). Reidemeister also equips each of his faces ('Flächenstücke') with an orientation, that is, he sees it as a directed subgraph, defined by a directed circuit, and defines a second face with the opposite orientation. For us (see below) a circuit is a directed path, with a beginning and an end, but a 2-cell does not inherit that direction. Since the subgraph defined by a circuit does not depend on which vertex is seen as the beginning and end, it follows that $2n$ different circuits of length n can define the same 2-cell.

Throughout this paper we shall study Γ through its underlying graph (or *1-skeleton*), and so we shall use the language and notation of graph theory, alongside more topological language. We shall use the same label Γ for the graph as for the 2-complex. The vertex-set of Γ is denoted $V(\Gamma)$, and the edge-set $E(\Gamma)$.

If Γ has no multiple edges, then it is *simple.* In sections 4 and 5 we shall assume that this restriction holds, and also that Γ is *connected*, which for our purposes means that Γ is connected as a graph. If Γ is simple and all 2-cells are triangles, then Γ is a simplicial complex, but we shall not require this in general; however, we observe that at most two barycentric subdivisions will transform any 2-complex Γ, simple or otherwise, into a simplicial complex. (Similarly, we observe that since barycentric subdivision of a graph with loops eliminates the loops, it is not difficult to modify our treatment to allow Γ to have loops.) The complex Γ might be the 2-skeleton of a higher dimensional complex, but higher dimensional cells are irrelevant to us, since we are only interested in this paper in the low dimensional topology of Γ.

3.2. Homotopy, fundamental group and covers

We define a *path* p in Γ to be a sequence, $v_0, e_1, v_1, e_2, \ldots, e_n, v_n$ of vertices v_i and edges e_i, such that each e_i is incident with the vertices v_{i-1} and v_i; of course, if Γ is simple, then p is completely determined by the sequence of vertices. We call v_0 the *initial vertex* and v_n the *terminal vertex* of p. The *reverse* of p, p^{-1}, is defined to be the path $v_n, e_n, \ldots, e_2, v_1, e_1, v_0$, and for any path q whose initial vertex is the terminal vertex of p, pq is defined to be the concatenation of p and q (with the terminal vertex of p identified with the initial vertex of q). If $v_0 = v_n$, then p is called a *circuit based at* v_0.

Let b be a specified vertex of Γ, which we shall call the *base point* of Γ. Suppose that c and c' are two circuits based at b. We say that c and c' are *elementary homotopy equivalent* if there are paths p, q, q', r of Γ (possibly of trivial length) such that $c = pqr$, $c' = pq'r$ and either qq'^{-1} defines a 2-cell of Γ or one of q, q' is trivial and the other consists of a path along an edge and back again. More generally, c and c' are *homotopy equivalent* if there is a sequence $c_0, c_1, \ldots, c_k$ of circuits in Γ, each based at b, such that $c_0 = c$, $c_k = c'$ and for each $i = 0, \ldots, k-1$, c_i and c_{i+1} are elementary homotopy equivalent.

The *fundamental group of* Γ *based at* b, $\pi_1(\Gamma, b)$, is defined to be the group of all homotopy equivalence classes of circuits based at b, composed by concatenation of equivalence class representatives. Provided that Γ is connected (as a graph), the isomorphism type of the group is independent of the choice of b. This is our situation, and we shall often simply denote the fundamental group of the 2-complex Γ by G. Γ is said to be *simply connected* if it is connected as a graph and has trivial fundamental group.

For our purposes, the *first homology group of* Γ *based at* b is defined to be the abelianized quotient $G/[G, G]$ of the fundamental group G, and, for p a prime, the *first homology group mod* p is defined to be the quotient of the first homology group by the subgroup consisting of the p-th powers of all the elements.

A *morphism* from a 2-complex Γ' to Γ is defined to be a map from Γ' to Γ which maps i-cells of Γ' to i-cells of Γ, for $i = 0, 1, 2$, and preserves incidence between cells. A bijective morphism is called an *isomorphism*, and an isomorphism from Γ to itself is called an *automorphism* of Γ.

A morphism θ from a 2-complex $\hat{\Gamma}$ to Γ is called a *covering map* if θ maps $\hat{\Gamma}$ onto Γ and is locally an isomorphism, in the sense that the restriction of θ to the star of any cell $\hat{\sigma}$ of $\hat{\Gamma}$ is injective and maps onto the star of $\hat{\sigma}\theta$. In particular, this means that, for any 0,1 or 2-cell σ of Γ, $\sigma\theta^{-1}$ is a set of cells no two of which are incident with any common cell of $\hat{\Gamma}$. The pair $(\hat{\Gamma}, \theta)$ (or sometimes simply the 2-complex $\hat{\Gamma}$) is then called a *cover* of Γ. For any cell σ of Γ the set $\sigma\theta^{-1}$ of cells of $\hat{\Gamma}$ which map under θ to σ is called the *fibre* of σ; provided that Γ is connected, the cardinality of the fibre of σ is independent of the choice of σ. Where m is an integer, θ is called an *m-fold* covering map and $(\hat{\Gamma}, \theta)$ an m-fold cover if for each σ, $\sigma\theta^{-1}$ has cardinality m.

If $(\hat{\Gamma}, \theta)$ is a cover of Γ, then, where $\hat{b}\theta = b$, the morphism θ maps circuits of $\hat{\Gamma}$ based at $\hat{b}$ to circuits of Γ based at b. If c, c' are distinct circuits based at $\hat{b}$, then $c\theta$ and $c'\theta$ are distinct, and c, c' are homotopy equivalent if and only if $c\theta$ and $c'\theta$ are. Hence clearly θ induces an isomorphism from the fundamental group of $\hat{\Gamma}$ to a subgroup of the fundamental group of Γ.

More generally, suppose that (Γ_1, θ_1) and (Γ_2, θ_2) are covers of Γ, and that b_1, b_2, b are vertices in Γ_1, Γ_2, Γ respectively such that $b_1\theta_1 = b_2\theta_2 = b$. Let H_1 and H_2 be the subgroups of $\pi_1(\Gamma, b)$ which are the images under θ_1 and θ_2 of $\pi_1(\Gamma_1, b_1)$ and $\pi_1(\Gamma_2, b_2)$. Then there is a covering map ψ from Γ_1 to Γ_2 which maps b_1 to b_2 and satisfies $\theta_1 = \psi\theta_2$ if and only if H_1 is a subgroup of H_2.

A cover $(\tilde{\Gamma}, \theta)$ of Γ is said to be a *universal cover* of Γ if $\tilde{\Gamma}$ is simply connected. By the above, $\tilde{\Gamma}$ is determined uniquely up to isomorphism, and covers all other covers of Γ. Further, if θ_1 and θ_2 are two covering maps from $\tilde{\Gamma}$ to Γ then there is an isomorphism ψ of $\tilde{\Gamma}$ such that $\theta_1 = \psi\theta_2$. In particular this implies that any automorphism α of Γ lifts to an automorphism $\tilde{\alpha}$ of $\tilde{\Gamma}$ with the property that $\tilde{\alpha}\theta = \theta\alpha$, since $\theta\alpha$ must be a covering map. Note also that if $(\hat{\Gamma}, \theta)$ is any cover of Γ, then for some ϕ, $(\tilde{\Gamma}, \phi)$ is a universal cover of $\hat{\Gamma}$.

An automorphism α of $\hat{\Gamma}$ with the property that $\alpha\theta = \theta$ is called a *deck transformation* of the cover $(\hat{\Gamma}, \theta)$ of Γ. The deck transformations of $(\hat{\Gamma}, \theta)$ form a group. The group of deck transformations of $(\tilde{\Gamma}, \theta)$ is isomorphic to the fundamental group of Γ, provided that Γ is connected.

A *fundamental region* for Γ is defined to be a set F of cells of $\tilde{\Gamma}$ which contains a unique element of the fibre of σ, for each cell σ of Γ. The images of F under the group of deck transformations of Γ partition $\tilde{\Gamma}$; Γ can be described as the quotient of $\tilde{\Gamma}$ under the action of that group.

4. An Edge-Labelled Graph Viewpoint

From now on, we assume that our 2-complex Γ is simple (no multiple edges) and connected. The basis of the description of the fundamental group and covers of Γ given in subsections 4.1, 4.2 and 4.3 below is found in (Reidemeister, 1951); we believe that the further development of this viewpoint in the remainder of section 4 is new.

4.1. The Fundamental Group

Let T be a spanning tree of the connected graph Γ. To each edge $e = \{v, w\}$ of Γ not in T, associate a pair of directed edges (v, w) and (w, v), and a circuit of Γ based at b, formed by joining b by a simple path in T to v and then w and then back to b by a simple path in T. It is well-known (see for example (Reidemeister, 1951)) that the homotopy equivalence classes of these circuits form a generating set for the fundamental group G. Let $g_{v,w}$ be the generator of G corresponding to the circuit through (v, w). Now where each (directed) edge (v, w) is assigned the label $g_{v,w}$ (and its reverse the label $g_{v,w}^{-1}$), and each edge in T is labelled with the trivial word, any circuit in Γ is naturally labelled by a word in the $g_{v,w}$'s and their inverses which is formed by composing the labels on the edges of the directed circuit (from left to right); this gives an expression for the corresponding element of the fundamental group in terms of the generators $g_{v,w}$. The set of all labels of circuits of the form pcp^{-1}, for which p is any path with initial vertex b, and c a circuit defining a 2-cell, forms a full set of relators for the group.

Where b' is a vertex distinct from b, and q is a path in T from b to b', there is a natural correspondence between circuits based at b and circuits based at b', which matches a circuit c based at b with the circuit $q^{-1}cq$ based at b'. Hence it is clear not only that the fundamental groups based at b and b' are isomorphic but that the set of generators of the form $g_{\overline{e}}$ described above is natural for both groups, in the sense that it depends only on T and not on b or b'.

4.2. The Covers of Γ

Suppose that $(\hat{\Gamma}, \theta)$ is a cover of Γ. Let b be the base point of Γ, and label the elements of the fibres of b in $\hat{\Gamma}$ by pairs (b, i) for i in some indexing set I. Let T be a spanning tree of Γ. Then the vertices of $\hat{\Gamma}$ can be partitioned by trees T_i such that $(b, i) \in T_i$, $T_i\theta = T$, and θ restricted to T_i is an isomorphism. For each vertex v of Γ, let (v, i) label the single vertex in $T_i \cap v\theta^{-1}$.

Now if $\{v, w\}$ is an edge of Γ, let $\rho_{v,w}$ be the function from I to I defined by $i\rho_{v,w} = j$ if $\{(v, i), (w, j)\}$ is an edge of $\hat{\Gamma}$. It follows from the fact that θ is a covering map that $\rho_{v,w}$ is a permutation of I. If $v_0, v_1, \ldots, v_k$, where $v_k = v_0$, is a circuit c defining a 2-cell of Γ, then $c\theta^{-1}$ is a disjoint union of circuits of length k in $\hat{\Gamma}$ (defining 2-cells of $\hat{\Gamma}$), and so $\rho_{v_0,v_1}\rho_{v_1,v_2} \cdots \rho_{v_{k-1},v_k}$ is the identity permutation. From this it is clear that the map $\rho\colon G \to Sym(I)$ defined by $(g_{v,w})^\rho = \rho_{v,w}$ extends to a homomorphism from G to $Sym(I)$. The cover $\hat{\Gamma}$ is completely defined by the homomorphism ρ; every homomorphism from G to $Sym(I)$ defines a cover. Connected covers correspond to transitive permutation representations of G, and for intransitive permutation representations, the number of orbits is equal to the number of connected components of the corresponding cover. From now on, given a specified homomorphism $\rho\colon G \to Sym(I)$ we shall label the associated cover of Γ by Γ_ρ.

Where H is the fundamental group $\pi_1(\Gamma_\rho, (b, i))$ of a connected cover Γ_ρ based at the vertex (b, i), H is isomorphic to the preimage in G of the stabiliser in G^ρ of i. Indeed, Γ_ρ can be defined in terms of the right cosets of H in G. The vertices of Γ_ρ correspond to pairs (v, Hx), for v a vertex of Γ and $x \in G$. Whenever v and w are joined in Γ, (v, Hx) and $(w, Hxg_{v,w})$ are joined in Γ_ρ. The 2-cells of Γ_ρ are the connected components of the inverse images of the 2-cells of Γ.

4.3. The Universal Cover

The right regular representation of G defines the universal cover of Γ, that is, $\tilde{\Gamma}$ is isomorphic to the cover Γ_{reg} whose vertices correspond to pairs (v, x), where v is a vertex of Γ and $x \in G$. Edges of Γ_{reg} join pairs (v, x) and $(w, xg_{v,w})$ where v and w are joined in Γ and $g_{v,w}$ is the label on that edge of Γ. For any $g \in G$, the product of the labels of a path from (v, x) to (v', xg) must be a word equal to g in G.

To verify that this construction really does yield the universal cover, we need only to see that Γ_{reg} is simply connected. Let θ be the covering map from Γ_{reg} to Γ, defined by $(v, x)\theta = v$. Now suppose that c is a circuit in Γ_{reg} based at $(b, 1)$. Then, as a path from $(b, 1)$ to $(b, 1)$, c can only be labelled by a word equal to the identity element. The same is true of $c\theta$. Since G is the fundamental group of Γ relative to b, this implies that $c\theta$ is homotopically equivalent to the trivial circuit, Since homotopies of Γ lift to homotopies of Γ_{reg}, the same is true of c.

4.4. Lifting an Automorphism to the Universal Cover

Suppose that α is an automorphism of Γ. Then for each element of G a lift of α to $Aut(\tilde{\Gamma})$ is defined. More precisely, the following is true.

Proposition 4.4.1. *Each automorphism α of Γ lifts to an automorphism of $\tilde{\Gamma}$. For each automorphism $\alpha \in Aut(\Gamma)$ and for each element $g \in G$, there is a unique automorphism $\tilde{\alpha}$ of $\tilde{\Gamma}$ such that*

$$(b,1)\tilde{\alpha} = (b\alpha, g)$$

More precisely, $\tilde{\alpha}$ is defined by the rule

$$(v,x)\tilde{\alpha} = (v\alpha, gx^{\alpha}t_{v,\alpha}) \tag{4.1}$$

for each $(v,x) \in \tilde{\Gamma}$, where x^{α} is the label on the image under α of any circuit based at b with label x, and $t_{v,\alpha}$ is the label on the image under α of the path in T from b to v.

Proof. The fact that α lifts to an automorphism of $\tilde{\Gamma}$ is simply a consequence of the universality of $\tilde{\Gamma}$.

Uniqueness is determined once the rule (4.1) has been verified.

First we observe that α induces an automorphism of the fundamental group G, viewed as an abstract group with the sets of generators and relators described, rather than as the fundamental group based at the specific base point b.

For α induces an action on the set of circuits of Γ, mapping circuits based at b to circuits based at $b\alpha$. Since α is an automorphism of Γ, two circuits through b are homotopic if and only if their images (through $b\alpha$) are. Since two circuits through the same vertex are homotopic if and only if they are labelled by the same element of G, α induces an injective map from the set of labels on circuits through b to the set of labels on circuits through $b\alpha$, that is, from G to G. Every element of G labels some circuit through $b\alpha$ (or indeed any point), and so since α maps the set of circuits through b onto the set of circuits through $b\alpha$, the map induced by α on G is also surjective. Since the concatenation of two circuits through b is mapped by α to the concatenation of their images, the map induced by α on G is a group homomorphism. Hence as a bijective homomorphism it is an automorphism of G. x^{α} is then the image of x under this automorphism.

Note that it is not in general true that α^{-1} induces the inverse of the group automorphism induced by α.

Now suppose that p and p' are two paths from b to some vertex v, both labelled by the same element x of G. Then pp'^{-1} is a circuit based at b and labelled with the identity element 1 of G, and hence is mapped under α to a circuit based at $b\alpha$ also labelled by the identity element, since $1^{\alpha} = 1$. So $p\alpha$ and $p'\alpha$ carry the same label, which we shall call x^{α_v}.

Let q be the path in T from b to v. Then $t_{v,\alpha}$ is defined to be the label of $q\alpha$. (Note that $t_{v,\alpha} = 1^{\alpha_v}$.) So the circuit $p\alpha(q\alpha)^{-1}$ has label $x^{\alpha_v}t_{v,\alpha}^{-1}$. However $p\alpha(q\alpha)^{-1}$ is the image of the circuit pq^{-1} under α, and so must have label x^{α}. Hence

$$x^{\alpha_v} = x^{\alpha}t_{v,\alpha}.$$

Let $\tilde{\alpha}$ be a lift of α defined by $(b,1)\tilde{\alpha} = (b\alpha, g)$.

Now any edge or path in $\tilde{\Gamma}$ has the same label as its projection in Γ. Hence if v and w are adjacent vertices of Γ, then $(v,x)\tilde{\alpha}$ is joined to $(w, xg_{v,w})\tilde{\alpha}$ by an edge with the label $g_{v\alpha,w\alpha}$. So

$$(v,x)\tilde{\alpha} = (v\alpha, y) \quad \text{implies that} \quad (w, xg_{v,w})\tilde{\alpha} = (w\alpha, yg_{v\alpha,w\alpha}). \tag{4.2}$$

Any path c from $(b,1)$ to (v,x) has label x, and x^{α_v} is the label of its image under $\tilde{\alpha}$. Hence it follows from (4.2) that

$$(v,x)\tilde{\alpha} = (v\alpha, gx^{\alpha_v}) = (v\alpha, gx^{\alpha}t_{v,\alpha}). \tag{4.3}$$

□

As an immediate corollary of Proposition 4.4.1 we have the following:-

COROLLARY 4.4.2. *For any $g \in G$, the rule*

$$(v,x)\tau_g = (v, gx),$$

for all v, x, defines an automorphism τ_g of $\tilde{\Gamma}$ with the property that $\tau_g\theta = \theta$. Any automorphism τ of $\tilde{\Gamma}$ with the property that $\tau\theta = \theta$ must be of this form for some $g \in G$.

PROOF. Any automorphism τ of $\tilde{\Gamma}$ with the property that $\tau\theta = \theta$ is a lift of the trivial automorphism ι of Γ. By Proposition 4.4.1, for any $g \in G$, there is a unique lift τ_g of ι with the property that

$$(b,1)\tau_g = (b,g),$$

and by (4.3), for all v, x,

$$(v,x)\tau_g = (v, gx^{\iota}1) = (v, gx).$$

□

This verifies that the fundamental group G of Γ is isomorphic to the full group of deck transformations of the universal cover.

Clearly we have shown that the lifts of α are all composites with deck transformations of a lift $\tilde{\alpha}_1$ which maps $(b,1)$ to $(b\alpha, 1)$, and hence satisfies

$$(v,x)\tilde{\alpha}_1 = (v\alpha, x^{\alpha}t_{v,\alpha}).$$

We shall call this lift the *principal* lift of α to $\tilde{\Gamma}$.

The principal lifts of automorphisms of Γ do not necessarily form a subgroup of $Aut(\tilde{\Gamma})$. For, if a lift $\tilde{\alpha}$ of α and its inverse $\tilde{\alpha}^{-1}$ are defined by the rules

$$(v,x)\tilde{\alpha} = (v\alpha, g_1x^{\alpha}t_{v,\alpha}), \quad (w,y)\tilde{\alpha}^{-1} = (w\alpha^{-1}, g_2y^{\alpha^{-1}}t_{w,\alpha^{-1}})$$

then g_1 and g_2 are related by the equation

$$g_1g_2^{\alpha}t_{b\alpha^{-1},\alpha} = 1. \tag{4.4}$$

For any deck transformation τ_g of $\tilde{\Gamma}$, the conjugate of τ_g by any lift $\tilde{\alpha}$ of α is an automorphism of $\tilde{\Gamma}$ which induces the identity automorphism on Γ, and so is itself a deck transformation. So $\tilde{\alpha}$ is an element of the subgroup of automorphisms of $\tilde{\Gamma}$ which normalise the subgroup $T_G = \{\tau_g\colon g \in G\}$ of $Aut(\tilde{\Gamma})$.

In fact, where $\tilde{\alpha}$ and its inverse are defined as above, using (4.4), we see that

$$\tilde{\alpha}^{-1}\tau_g\tilde{\alpha} = \tau_{g_1 g^\alpha g_1^{-1}}$$

and so in particular, when $\tilde{\alpha}$ is the principal lift of α,

$$\tilde{\alpha}^{-1}\tau_g\tilde{\alpha} = \tau_{g^\alpha},$$

that is, the action by conjugation of the principal lift of α on G as a group of deck transformations of $\tilde{\Gamma}$ is identical to the action of α itself on G as the fundamental group of Γ based at b. The action of any other lift of α is conjugate to the action of the principal lift.

4.5. Mapping Down an Automorphism of the Universal Cover.

The preceding analysis gives a precise formula for the lift to $\tilde{\Gamma}$ of any automorphism of Γ, and hence gives a precise condition for an automorphism of $\tilde{\Gamma}$ to induce an automorphism of Γ.

An automorphism β of $\tilde{\Gamma}$ induces an automorphism of Γ if and only if it permutes the orbits of the deck group T_G of $\tilde{\Gamma}$; in that case β is a lift of the automorphism α defined by the rule

$$v\alpha = (v,x)\beta\theta$$

and the action of β is completely defined once one of its images is defined, using the above analysis.

We have already shown that such an automorphism β must normalise the subgroup T_G of deck transformations. Conversely, if γ is an automorphism of $\tilde{\Gamma}$ which normalises T_G, then γ must permute the orbits of T_G. Thus the automorphisms of $\tilde{\Gamma}$ which induce automorphisms of Γ are precisely those which normalise the subgroup T_G of deck transformations.

Suppose that H is a subgroup of G, and that H^β is the image of H under the action of β described above. Then β induces an isomorphism β' between the two covers $\Gamma_1 = \tilde{\Gamma}/H$ and $\Gamma_2 = \tilde{\Gamma}/H^\beta$ of Γ, mapping, for all v, x, (v, Hx) to $(v\alpha, H^\beta x^\beta)$. Where θ_1 is the covering map from Γ_1 to Γ defined by $(v, Hx)\theta_1 = v$ and θ_2 the covering map from Γ_2 to Γ defined by $(v, H^\beta x^\beta)\theta_2 = v$, then $\theta_2 = \beta'\theta_1$. If $H^\beta = H$, then β induces an automorphism of $\tilde{\Gamma}/H$.

4.6. Lifting an Automorphism to an Intermediate Cover

Let $\hat{\Gamma} = \Gamma_\rho$ be an intermediate cover of Γ, with fundamental group H; then the vertices of $\hat{\Gamma}$ can be described as pairs (v, Hx) for v a vertex of Γ, and $x \in G$, as described above. Thus the vertices of $\hat{\Gamma}$ are seen naturally as the orbits of H as a group of deck transformations on the vertices of $\tilde{\Gamma}$.

Proposition 4.6.1. *Suppose that α is an automorphism of Γ, and that b is a vertex of Γ and g an element of G. For any $x \in G$, define x^α to be the label of the image under α of any circuit based at b with label x. Let $\hat{\Gamma}$ be a connected cover of Γ with fundamental group $H \subseteq G$.*

Then α lifts to an automorphism $\hat{\alpha}$ of $\hat{\Gamma}$ which maps (b, H) to $(b\alpha, Hg)$ if and only if $H^\alpha \subseteq H^g$.

PROOF. The automorphism α lifts to an automorphism of $\hat{\Gamma}$ precisely when one of its lifts $\tilde{\alpha}$ in $Aut(\tilde{\Gamma})$ maps down to an automorphism of $\hat{\Gamma}$.

Clearly $\tilde{\alpha}$ as above maps down to an automorphism of $\hat{\Gamma}$ provided it permutes the orbits of H acting as a group of deck transformations. More precisely, $\tilde{\alpha}$ maps down if and only if, given $(v,x) \in \tilde{\Gamma}$, for each $h \in H$, there is $h' \in H$ such that

$$(v, hx)\tilde{\alpha} = (v, x)\tilde{\alpha}\tau_{h'}$$

By Proposition 4.4.1 above, we see that, where $(b,1)\tilde{\alpha} = (b\alpha, g)$, this condition translates as

$$(v\alpha, g(hx)^{\alpha}t_{v,\alpha}) = (v\alpha, h'gx^{\alpha}t_{v,\alpha}),$$

or equivalently

$$gh^{\alpha} = h'g,$$

Hence $\tilde{\alpha}$ maps down to $\hat{\Gamma}$ if and only if $H^{\alpha} \subseteq H^{g}$.

Thus α has a lift which is an automorphism mapping (b,H) to (b,Hg) for each $g \in G$ such that $H^{\alpha} \subseteq H^{g}$. When $\hat{\alpha}$ is such a lift, then

$$(v, Hx)\hat{\alpha} = (v\alpha, Hgx^{\alpha}t_{v,\alpha})$$

□

4.7. Equivalence of Covers

Suppose that (Γ_1, θ_1) and (Γ_2, θ_2) are two connected covers of Γ, associated with subgroups H_1 and H_2 of the fundamental group G. We need to decide under what circumstances we should consider (Γ_1, θ_1) and (Γ_2, θ_2) to be equivalent. Let ϕ_1 be the covering map from $\tilde{\Gamma}$ to Γ_1 and let ϕ_2 be the covering map from $\tilde{\Gamma}$ to Γ_2.

First observe that if Γ_1 and Γ_2 are isomorphic as 2-complexes, related by an isomorphism $\psi : \Gamma_1 \to \Gamma_2$, then since $\phi_1\psi$ and ϕ_2 are then both covering maps from $\tilde{\Gamma}$ to Γ_2, and $\tilde{\Gamma}$ is a universal cover for Γ_2, ψ lifts to an automorphism $\tilde{\psi}$ of $\tilde{\Gamma}$.

It is standard to consider (Γ_1, θ_1) and (Γ_2, θ_2) to be equivalent as covers when Γ_1 and Γ_2 are related by an isomorphism ψ with the property that $\theta_1 = \psi\theta_2$. From a computational point of view it is relevant also to consider other, broader, notions of equivalence.

We find isomorphisms ψ with $\theta_1 = \psi\theta_2$ using deck transformations. For suppose that H_1 and H_2 are subgroups of G such that $H_2 = gH_1g^{-1}$. Let (Γ_1, θ_1) and (Γ_2, θ_2) be the covers of Γ associated with H_1 and H_2. Then the deck transformation τ_g induces an isomorphism between Γ_1 and Γ_2 such that (v, H_1x) maps to (v, H_2gx). Clearly $\theta_1 = \tau_g\theta_2$. We shall say that Γ_1 and Γ_2 are equivalent under the action of the group of deck transformations.

In fact, any ψ: $\Gamma_1 \to \Gamma_2$ with $\theta_1 = \psi\theta_2$, must arise in this way. For such a ψ induces the trivial automorphism on Γ. Hence the lift $\tilde{\psi}$ of ψ to $\tilde{\Gamma}$ acts as a deck transformation by Corollary 4.4.2.

We can find a more general notion of equivalence, by considering also any isomorphism $\psi' : \Gamma_1 \to \Gamma_2$ which induces a (possibly non-trivial) automorphism α on Γ. In this case $\theta_1\alpha = \psi'\theta_2$.

An isomorphism ψ' of this type is induced by the action of a lift $\beta = \tilde{\alpha}$ of α to $\tilde{\Gamma}$. For suppose that β is the (principal) lift of α defined by the rule

$$(v, x)\beta = (v\alpha, x^{\alpha}t_{v,\alpha})$$

(Any other lift is a composite of this with a deck transformation.) Now let H_1 be any subgroup of G and H_2 its image under the action of α. Then the set of vertices of $\tilde{\Gamma}$ of the form (v, hx) for $h \in H_1$ maps under β to the set of vertices of the form $(v\alpha, h^\alpha x^\alpha t_{v,\alpha})$. So β induces an isomorphism ψ' between Γ_1 and Γ_2 with the property that $\theta_1\alpha = \psi'\theta_2$.

Conversely, suppose that ψ' is an isomorphism between covers Γ_1 and Γ_2 with the property that $\theta_1\alpha = \psi'\theta_2$ for some automorphism α of Γ. Suppose that

$$(b, H_1)\psi' = (b\alpha, H_2)$$

(This can be achieved if necessary by replacing Γ_2 by a cover related to it by an isomorphism induced by a deck transformation, replacing ψ' by its composite with that isomorphism, and replacing H_2 by a conjugate.) Then ψ' arises exactly as described above.

We see that we can classify covers of Γ up to the standard notion of equivalence ($\theta_1 = \psi\theta_2$) by enumerating the conjugacy classes of subgroups of G. If the r-fold covers of Γ are required, for r at most some fixed positive integer m, then we may apply the low-index subgroup algorithm with input m and a finite presentation for G, to determine up to conjugacy the subgroups of G of index of at most m. However, this can be an extremely time-consuming process. The low-index subgroup algorithm is described in (Sims, 1994, section 5.6) and in (Neubüser, 1982). A method for the case when G is polycyclic is described in (Lo, 1997).

From a computational point of view, it might make more sense to consider the more general definition of equivalence suggested above. In this case we enumerate the equivalence classes of subgroups of G for which H_1 and H_2 are equivalent if related by a composite of conjugation and an automorphism of Γ (in its action on G as a fundamental group), or (equivalently) if related by an automorphism of $\tilde{\Gamma}$ which normalises G (in its action on G by conjugation as a group of deck transformations). We could require that automorphism to be in a specified subgroup of $Aut(\tilde{\Gamma})$. From a computational point of view, such a definition of equivalence could well be very useful.

This discussion suggests that it would be worthwhile to attempt to design a version of the low-index subgroup algorithm which would classify the subgroups of index at most m in a finitely presented group H, up to action by a composite of conjugation and an element of an explicit group A of automorphisms of H, where A is given by specifying the images of the generators of H under the generators of A.

5. The Algorithms

We describe our algorithms for an arbitrary finite, simple, connected 2-complex Γ. However, the current GAP/GRAPE implementation of our algorithms is only for the special case of a clique complex of a finite simple graph (in which the 2-cells are precisely the triangles). Work remains to be done on efficient implementation for more general classes of complexes. We have used our current implementation successfully on complexes with over 1,000 vertices and over 100,000 edges, but the range of applicability depends heavily on the nature of the fundamental group.

5.1. Building a Fundamental Record for Γ

Let Γ be a finite, simple, connected 2-complex, and to each edge $\{v, w\}$ of Γ, associate the pair of directed edges (v, w), (w, v). A *fundamental record* (P, f) of Γ, produced by the algorithm of this section, consists of a finite presentation $P = (X; R)$ (with generators X and relators R) for the fundamental group G of Γ, and a labelling (mapping) f from the directed edges of Γ to the free group on X, such that the G-image of the label $f_{v,w}$ of (v, w) is the edge-label $g_{v,w}$ encountered in section 4. The labelling f will be used to construct covers of Γ as described in section 4.2.

The presentation $(X; R)$ for the fundamental group is constructed in such a way as to attempt to minimize the number of generators. This approach is appropriate if the fundamental group is cyclic or if we are abelianizing the fundamental group as we proceed to compute the first homology group (as is often the case for complexes in which we are interested). Otherwise, the length of relators in R may explode.

Roughly speaking, we build a copy Δ of Γ, one edge at a time, labelling each directed edge of Δ (if necessary with a new generator) as we do so. At any stage, if a label can be found for a directed edge which is a word in existing generators, then this label is used. If two distinct such labels can be found, then a relation has been discovered.

More precisely, the algorithm runs as follows. The input is a finite, simple, connected 2-complex Γ, and the output is a fundamental record for Γ.

We initialize the complex Δ to be a spanning tree T of Γ, set P to be the presentation with no generators and no relators, and set f to map each directed edge of Δ to the trivial word. If at this stage $\Delta = \Gamma$, we are done, and output (P, f). Otherwise, we shall extend Δ edge by edge until it is equal to Γ. At each stage that we modify Δ we shall update (P, f) to remain a fundamental record for Δ.

Suppose that $\Delta \neq \Gamma$. Our basic step is to search for an edge $\{v, w\} \in E(\Gamma) \setminus E(\Delta)$, such that, for some 2-cell c of Γ incident with $\{v, w\}$, say c defined by the circuit

$$v, v_1, \ldots, v_n, w, v,$$

we have that $v, v_1, \ldots, v_n, w$ is a path in Δ.

Suppose that this search is successful. Then we add the edge $\{v, w\}$ to Δ, together with all 2-cells of Γ incident with $\{v, w\}$, such that these 2-cells are circuits in Δ. We then define

$$f_{v,w} := f_{v,v_1} f_{v_1,v_2} \cdots f_{v_n,w},$$

and

$$f_{w,v} := (f_{v,w})^{-1}.$$

If P is not obviously a presentation for the trivial group, then we must check if we need to add more relators to P. To do this, for each 2-cell $c' \neq c$ in Δ incident with $\{v, w\}$, say c' defined by the circuit $v, w_1, \ldots, w_m, w, v$, we add to P the relator

$$f_{v,w_1} f_{w_1,w_2} \cdots f_{w_m,w} f_{w,v}.$$

After we have done this, (P, f) is a fundamental record for Δ. We may choose to try to simplify the relators of P at this stage, and perhaps to change the edge-labels to shorter words (if possible) representing the same elements of the fundamental group of Δ.

If our search of $E(\Gamma) \setminus E(\Delta)$ is not successful (i.e. there is no edge with the desired properties), a new generator g is added to P. In this case, we choose an edge $\{v, w\} \in E(\Gamma) \setminus E(\Delta)$, add this edge to $E(\Delta)$, and define

$$f_{v,w} := g \quad \text{and} \quad f_{w,v} := g^{-1}.$$

After we have done this, (P, f) is still a fundamental record for Δ.

If at this stage we still have $\Delta \neq \Gamma$, we go back and repeat our basic search step. Otherwise we are done, and output (P, f).

5.1.1. Implementing the "basic step"

In our implementation for clique complexes we do no actual searching in the "basic step". Whenever an edge $\{x, y\}$ is added to Δ, we determine the edges e in triangles of Γ containing $\{x, y\}$, such that $e \notin E(\Delta)$ and adding e to Δ would complete a 2-claw containing $\{x, y\}$ to a triangle, and each such edge e is added to a queue Q (if not already in Q) for future addition to Δ. We use bit-arrays (boolean lists in GAP) for fast membership testing in this queue. Thus, the "basic step" simply consists of taking an edge off Q, or determining that Q is empty (in which case a generator must be added to P).

Let k be the maximum valency of a vertex of Γ. The data structures used in our implementation allow us to add the appropriate edges to Q in $O(k)$ time (assuming that an array access takes constant time), for each edge $\{x, y\}$ added to Δ.

5.2. Computing the First Homology Group of Γ

A presentation for the first homology group of Γ can of course be found by first computing a presentation for the fundamental group of Γ, and then abelianizing this presentation. In order to keep the size of our computation down from the start, however, we choose instead to follow the algorithm described above for the fundamental record, except that the presentation P is always abelianized, and the labelling f is a mapping from the directed edges of Δ to the free abelian group on the generators X of P. Doing this, it is often possible to compute quickly the first homology group when computing the fundamental group with the algorithm above would not be practical. We call the output (P, f) of this abelianized form of the algorithm the *abelianized fundamental record* of Γ.

Algorithms to determine the structure of a finitely presented abelian group are discussed in (Sims, 1994, Chapter 8) and (Havas and Majewski, 1997). We also mention that a different approach from ours to computing first homology groups has been developed by Steve Linton (unpublished).

5.3. Computing the First Homology Group mod p

Suppose that p is a prime and that we wish to compute the first homology group mod p of Γ, that is, $H_1(\Gamma, \mathbb{F}_p)$. Then we follow an algorithm similar to that above for the first homology group, except that instead of a presentation, we (implicitly) maintain the homology group mod p of the subcomplex Δ as a vector space V of all d-tuples over $\mathbb{F}_p$, where $d = \dim(H_1(\Delta, \mathbb{F}_p))$. The labelling f is then a mapping from the directed edges of Δ to V. When a non-trivial relator is found, the dimension d of V is decreased by one, and the edge-labels are appropriately rewritten to lie in V.

5.3.1. Some Complexity Analysis

Suppose now that our 2-complex Γ is a clique complex. We discuss the complexity of our implementation for computing the first homology group mod p, H say, of Γ, for a fixed prime p (a similar analysis holds for the more general case of a simplicial complex, which is the next case we shall implement).

Suppose Γ has n vertices and maximum valency k. Then Γ has $O(nk)$ directed edges and $O(nk^2)$ 2-cells. As described in section 5.1.1, we do not explicitly search Γ for edges to add to Δ, but maintain a queue of edges to add. The work to maintain this queue takes $O(nk^2)$ time in total ($O(k)$ time for each edge of Γ). Most of the work in computing H usually comes from checking relators and rewriting edge-labels, but we aim to perform the latter task as few times as possible.

Let m be the total number of generators introduced by the algorithm in the computation of H; that is, m is the number of times we find the queue empty when trying to add an edge to Δ. We have that $m \leq |E(\Gamma)| - |V(\Gamma)| + 1$ (the number of edges of Γ not on a spanning tree), but in practice we expect often to have m much smaller than this upper bound.

For each triangle of Γ, $O(m)$ time is taken to check the relator corresponding to that triangle (we add three row-vectors of dimension at most m over $\mathbb{F}_p$, and check whether this sum is the zero vector). The time taken to rewrite the edge-labels (when necessary) is $O(nkm)$ ($O(m)$ to rewrite each edge-label), and such a rewriting takes place at most m times. Thus, checking relators and rewriting edge-labels take

$$O(nk^2m + nkm^2)$$

time in total. The other steps in the implementation, such as setting up and initializing the data structures, take no more than $O(n^2 + nk^2)$ time in total.

5.4. Computing Covers

Suppose that Γ is a finite, simple, connected 2-complex with fundamental group G. Then, as described in section 4.2, every cover Γ is of the form Γ_ρ, where ρ is a permutation representation of G.

Suppose that we are given a fundamental record (P, f) of Γ, and a finite-degree permutation representation ρ: $G \to Sym(I)$. For example, ρ may have been obtained from the presentation P by using coset enumeration or the low-index subgroup algorithm (see (Neubüser, 1982) or (Sims, 1994)). We can construct the cover Γ_ρ of Γ as described in section 4.2, after defining, for each directed edge (v, w) of Γ,

$$\rho_{v,w} := (g_{v,w})^\rho,$$

where $g_{v,w}$ is the natural image of $f_{v,w}$ in G. We remark that if G^ρ is abelian then we may use an abelianized fundamental record instead of a fundamental record.

5.5. Computing the Deck Group of a Cover

Suppose that we are given a permutation representation ρ: $G \to Sym(I)$ of the fundamental group G of the finite, simple, connected 2-complex Γ. Then a typical fibre of the cover (Γ_ρ, θ) is

$$v\theta^{-1} = \{(v, i) : i \in I\},$$

where $v \in V(\Gamma)$. The deck group $D \leq Aut(\Gamma_\rho)$ of Γ_ρ acts on each such fibre.

Let $\tau \in D$, $v \in V(\Gamma)$, $i \in I$, and suppose

$$(v, i)\tau = (v, j).$$

Then if (v, w) is a directed edge of Γ, we must have

$$(w, i\rho_{v,w})\tau = (w, j\rho_{v,w}). \tag{5.1}$$

If (v, w) is a directed edge of the fixed spanning tree T of Γ on which the edge-labels are trivial then $\rho_{v,w} = 1$ and we have that $(w, i)\tau = (w, j)$. Indeed, if w is any vertex of Γ, then by induction on the length of a path in T from v to w, we still must have $(w, i)\tau = (w, j)$. This gives us a natural faithful representation

$$\sigma\colon D \to Sym(I)$$

defined by $i\tau = j$ if and only if $(v, i)\tau = (v, j)$, and this definition does not depend on the choice of $v \in V(\Gamma)$. Equation (5.1) further tells us that D^σ must centralize G^ρ. Conversely, each element in the $Sym(I)$-centralizer of G^ρ defines a deck transformation of Γ_ρ.

Note that the above gives an alternate way of seeing that the deck group D of the universal cover $\tilde{\Gamma}$ of Γ is isomorphic to the fundamental group G of Γ, since the $Sym(G)$-centralizer of the right-regular representation of G is the left-regular representation of G.

We can compute the deck group of an r-fold cover Γ_ρ of Γ by computing

$$D := C_{Sym(r)}(G^\rho).$$

If G^ρ is transitive, which is equivalent to saying that Γ_ρ is connected, then this computation is very easy. In this case, for $i = 1, \ldots, r$, there is at most one element of D mapping 1 to i. For each i, it is easy to determine if there is a permutation mapping 1 to i and centralizing G^ρ, and if so, to determine this permutation. (Suppose that $\tau \in D$, $1\tau = i$, and $j \in \{1, \ldots, r\}$. Then $j = 1g^\rho$ for some $g \in G$, and so $j\tau = 1g^\rho\tau = 1\tau g^\rho = ig^\rho$.)

5.6. Lifting an Automorphism to a Cover

We now suppose that we have a fundamental record (P, f) for our finite, simple, connected 2-complex Γ, with fundamental group G, and that $\hat{\Gamma}$ is a connected r-fold cover of Γ, defined by a transitive permutation representation

$$\rho\colon G \to Sym(r).$$

Hence, for any directed edge (v, w) of Γ we can determine the associated permutation $\rho_{v,w} \in G^\rho$.

Suppose that α is an automorphism of Γ. The theory of section 4 tells us exactly when α should lift to an automorphism of $\hat{\Gamma}$. If there is a lift, its effect is completely defined by the image of one vertex. Hence we attempt to lift α by defining an image for a chosen vertex of $\hat{\Gamma}$ and then attempting to extend the image outwards from that vertex. If the image extends, then α has such a lift and we have found it; if it does not, then α has no such lift.

The vertices of $\hat{\Gamma}$ are the pairs (v, i), where v is a vertex of Γ and $i \in \{1, \ldots, r\}$.

We want to see if α can lift to an automorphism $\hat{\alpha}$ of $\hat{\Gamma}$ which maps $(v_1, 1)$ to $(v_1\alpha, i)$.

We form a queue, which initially consists just of $(v_1, 1)$, containing vertices of $\hat{\Gamma}$ whose images under $\hat{\alpha}$ have been defined. Then, as long as this queue is non-empty, we remove a vertex a from (the head of) the queue, and do the following for each vertex b adjacent to a:

1 If no image of b under $\hat{\alpha}$ has been defined then we define an image: if $a = (v, x)$, $a\hat{\alpha} = (v\alpha, y)$, and $b = (w, x\rho_{v,w})$, then $b\hat{\alpha}$ must be (and is) defined as the vertex $b' = (w\alpha, y\rho_{v\alpha,w\alpha})$. If the image of a vertex different from b has already been defined to have image b' under $\hat{\alpha}$, then $\hat{\alpha}$ cannot define a permutation of $V(\hat{\Gamma})$ and so the required $\hat{\alpha}$ does not exist. Otherwise we add b to the queue and continue.
2 If, on the other hand, an image of b under $\hat{\alpha}$ has already been defined, then we check that the image of b is adjacent to the image of a, and if not, then the required $\hat{\alpha}$ does not exist.

Hence, we shall eventually construct a lift $\hat{\alpha}$ of α, such that $\hat{\alpha}$ maps $(v_1, 1)$ to $(v_1\alpha, i)$, or show that no such lift exists.

Suppose that D is the deck group for the cover $\hat{\Gamma}$ of Γ, and that S is a set of orbit representatives of D on the fibre $\{(v_1\alpha, j) : 1 \leq j \leq r\}$. Then $\alpha \in Aut(\Gamma)$ has a lift to an element of $Aut(\hat{\Gamma})$ if and only if α has such a lift taking $(v_1, 1)$ to an element in S.

5.6.1. Another Approach

Let Γ and $\hat{\Gamma} = \Gamma_\rho$ be as in the previous section, such that G is the fundamental group of Γ with respect to a base point b, and $H \leq G$ is the fundamental group of $\hat{\Gamma}$. Here we describe another approach (which we have not yet implemented) to determining whether an automorphism α of Γ lifts to the cover $\hat{\Gamma}$, and if so, determining one such lift.

We proceed as follows:

1 Define the action of α on G by finding, for each element g of the (finite) generating set X of G, a circuit of Γ based at b and labelled by g, and tracing out the label of the image circuit. In order to define the $t_{v,\alpha}$ (which we need in order to construct the lifts of α, but not actually to prove they exist), for each vertex v of Γ, compute the label of the image under α of the path from b to v in our fixed spanning tree T of Γ.
2 Compute generators h_j for H as words in the generators of G. (We may already know such generators, or we can compute Schreier generators.)
3 For each such h_j compute h_j^α, and hence the permutation in G^ρ corresponding to that element. Provided there are points in $\{1, 2, \ldots, r\}$ fixed by all such permutations, α has lifts to $\hat{\Gamma}$. In fact there is one lift for each point k which is fixed by all such permutations; where $g \in G$ satisfies $1g = k$, then there is a lift $\hat{\alpha}$ with

$$(v, Hx)\hat{\alpha} = (v\alpha, Hgx^\alpha t_{v,\alpha}).$$

This approach could be particularly useful if we are interested in studying several covers of Γ. The information computed in the first step can be used for any one of them, as it is independent of H.

The subgroup of $Aut(\Gamma)$ consisting of all automorphisms which lift to $\hat{\Gamma}$ is just the set of all such α which satisfy $H^\alpha \subseteq H^g$ for some $g \in G$.

5.7. A Sample Calculation

We now give an example of the use of our implementation of the algorithms above. We use GAP (version 3.4.4) and its share package GRAPE (version 2.31) together with this implementation to construct explicitly two distance-regular (in fact, distance-transitive) graphs discovered by Thomas Meixner; see (Meixner, 1991, Proposition 4.4). These graphs are respectively 4-fold and 2-fold covers of a graph having 672 vertices and valency 176, and automorphism group containing $U_6(2) = PSU(6, \mathbb{F}_4)$.

This calculation is given in the form of a GAP-logfile, and was originally performed for Aleksandar Jurišić, who wished to study the Meixner graphs. The computer used was a 233 MHz Pentium PC running Linux. The calculation of the fundamental record took about 56 seconds, and the total CPU-time used was about 103 seconds.

```
gap> RequirePackage("grape");

Loading  GRAPE 2.31  (GRaph Algorithms using PErmutation groups),
by L.H.Soicher@qmw.ac.uk.

gap> GRAPE_RANDOM:=true;;
gap> # We will use certain randomized methods in GRAPE
gap> # (which do not affect the correctness of results).
gap> SU:=SpecialUnitaryGroup(6,2);;
gap> orb:=Orbit(SU,GF(4).one*[1,1,Z(4),0,Z(4),0],OnLines);;
gap> Length(orb);
672
gap> #
gap> #  orb  is an orbit of non-isotropic projective points.
gap> #
gap> gamma:=EdgeOrbitsGraph(Operation(SU,orb,OnLines),[1,2]);;
gap> if VertexDegree(gamma,1) >= OrderGraph(gamma)/2 then
>    gamma:=ComplementGraph(gamma);
> fi;
gap> GlobalParameters(gamma);
[ [ 0, 0, 176 ], [ 1, 40, 135 ], [ 48, 128, 0 ] ]
gap> #
gap> # gamma  is the primitive quotient of the Meixner graphs.
gap> # We have associated the group  U_6(2) ( <= Aut(gamma) )
gap> # with  gamma.
gap> #
gap> # We now build the universal cover of the clique complex of
gap> # gamma,  using the algorithms described in this paper.
gap> #
gap> Read("/home/alg1/leonard/gapprogs/complex.g");
gap> Runtime(); # runtime in milliseconds
5170
gap> F:=FundamentalRecCliqueComplex(gamma);;
gap> Runtime();
61400
gap> G:=F.group;
Group( _x1, _x2 )
gap> Size(G);
4
gap> IsElementaryAbelian(G);
true
gap> #
gap> # G  is the fundamental group of the clique complex of
gap> # gamma,  and is isomorphic to  C2 x C2.
gap> #
```

```
gap> H:=TrivialSubgroup(G);
Subgroup( Group( _x1, _x2 ), [  ] )
gap> delta:=CoveringGraph(gamma,G,F.edgeLabels,H);;
gap> #
gap> # delta  is the (1-skeleton of the) universal cover of
gap> # the clique complex of  gamma.
gap> #
gap> GlobalParameters(delta);
[ [ 0, 0, 176 ], [ 1, 40, 135 ], [ 12, 128, 36 ], [ 135, 40, 1 ],
  [ 176, 0, 0 ] ]
gap> #
gap> # delta  is the Meixner 4-fold cover.
gap> #
gap> H:=Subgroup(G,[G.generators[1]]);
Subgroup( Group( _x1, _x2 ), [ _x1 ] )
gap> Size(H);
2
gap> epsilon:=CoveringGraph(gamma,G,F.edgeLabels,H);;
gap> GlobalParameters(epsilon);
[ [ 0, 0, 176 ], [ 1, 40, 135 ], [ 24, 128, 24 ], [ 135, 40, 1 ],
  [ 176, 0, 0 ] ]
gap> #
gap> # epsilon  is the Meixner 2-fold cover.
gap> #
gap> Runtime();
102550
gap>
```

References

Aschbacher, M., (1996). Combinatorial cell complexes, *Advanced Studies in Pure Mathematics* **24**, 1–80.

Aschbacher, M., Kleidman, P. B. (1990). On a conjecture of Quillen and a lemma of Robinson, *Archiv der Math* **55**, 209–217.

Aschbacher, M., Segev, Y. (1991). The uniqueness of groups of type J_4, *Inv. Math.* **105**, 589–607.

Aschbacher, M., Segev, Y. (1992). Extending morphisms of groups and graphs, *Annals of Math.* **135**, 297–323.

Aschbacher, M., Segev, Y. (1992). Locally connected simplicial maps, *Israel J. Math.* **77**, 285–303.

Aschbacher, M., Smith, S.D. (1993). Michael Aschbacher and Stephen D. Smith, On Quillen's conjecture for the p-group complex, *Annals of Math.* **137**, 473–529.

Brouwer, A.E., Cohen, A.M., Neumaier, A. (1989). *Distance-Regular Graphs*. Berlin and New York: Springer.

Godsil, C.D., Jurišić, A., Schade, T. (1990). Distance regular antipodal covers of strongly regular graphs, *Combinatorics and Optimization Research Report, University of Waterloo* **CORR91-02**.

Havas, G., Majewski, B.S. (1997). Integer matrix diagonalization, *J. Symbolic Computation* **24**, 399–408.

Ivanov, A.A., Shpectorov, S.V. (1990), The P-geometry for M_{23} has no non-trivial 2-coverings, *Europ. J. Comb.* **11**, 373–379.

Kantor, W.M. (1981). Some geometries that are almost buildings, *Europ. J. Comb.* **2**, 239–247.

Kantor W.M. (1987). Some locally finite flag transitive buildings, *Europ. J. Comb.* **8**, 429–436.

Kantor, W.M. (1985). Some exceptional 2-adic buildings, *J. Alg.* **92**, 208–233.

Lo, E.H. (1997). Enumerating finite index subgroups of polycyclic groups, preprint.

Maunder, C.R.F. (1980). *Algebraic Topology*, Cambridge University Press.

Meixner, Th. (1991). Some polar towers, *Europ. J. Comb.* **12**, 397–415.

Neubüser, J. (1982). An elementary introduction to coset table methods in computational group theory, in *Groups – St Andrews 1981* (C.M. Campbell and E.F. Robertson, eds), LMS Lecture Notes **71**, Cambridge University Press, 1–45.

Pasini, A. (1994). *Diagram Geometries*, Oxford University Press.

Quillen, D. (1978). Homotopy properties of the poset of nontrivial p-subgroups of a group, *Advances in Math.* **28**, 101–128.

Reidemeister, K. (1951). *Einführung in die kombinatorische Topologie*. Braunschweig: Friedr. Vieweg & Sohn.

Ronan, M.A. (1980). Coverings of automorphisms of chamber systems, *Europ. J. Comb.* **1**, 259–269.

Ronan, M.A. (1984). On triangle geometries, *J. Comb. Thy. Ser. A* **37**, 294–319.

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

Seifert, H., Threlfall, W. (1934). *Lehrbuch der Topologie.* New York: Chelsea.

Seifert, H., Threlfall, W. (1980). *A Textbook of Topology.* London: Academic Press.

Sims, C.C. (1994). *Computation with Finitely Presented Groups.* Cambridge University Press.

Soicher, L.H. (1993). GRAPE: a system for computing with graphs and groups, *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* **11**, 287–291.

Surowski, D.B. (1984). Covers of simplicial complexes and applications to geometry, *Geometriae Dedicata* **16**, 35–62.

Tits, J., (1981). A local approach to buildings, *The Geometric Vein, The Coxeter Festschrift (Davis et al., eds).* Berlin and New York: Springer. 519–547.

Tits, J., (1986). Ensembles ordonnés, immeubles et sommes amalgamées, *Bulletin de la Société Mathématique de Belgique* **38**, 367–387.