Journal of Pure and Applied Algebra 32 (1984) 1–10
North-Holland

PRESENTATIONS FOR GROUPS ACTING ON SIMPLY-CONNECTED COMPLEXES

Kenneth S. BROWN*
Department of Mathematics, White Hall, Cornell University, Ithaca, NY 14853, USA

Communicated by H. Bass
Received 6 January 1983

Suppose a group G acts by homeomorphisms on a simply-connected space X. There are a number of well-known results which show, under suitable hypotheses, how to obtain generators and relations for G by looking at a 'fundamental domain' for the action. See, for instance, [1]–[10]. The purpose of this note is to call attention to a particularly nice such presentation which is available in case X is a CW-complex and the G-action permutes the cells. (We then say that X is a *G-CW-complex*, or simply a *G-complex*.)

The presentation, which is most easily described in terms of the theory of graphs of groups (cf. [7]), has the following form: One builds an abstract group $\tilde{G}$, which is the fundamental group of a certain graph of groups obtained from the action of G on the 1-skeleton of X. The result then is that G is obtained from $\tilde{G}$ by introducing one relation for every 2-cell of X mod G. See Section 3 for more details. See also Theorem 1 in Section 1, where the same result is stated in more concrete language, independent of the theory of graphs of groups.

This theorem, as is probably evident from the rough statement of it given above, is actually a quite trivial corollary of the Bass–Serre structure theorem for groups acting on trees [7, §I.5.4]. Nevertheless, it appears that this corollary has gone unnoticed, or, more to the point, that a number of people using group actions to get presentations have simply overlooked the possibility that [7] (which is about trees) could be fruitfully applied to group actions on spaces of dimension >1. In particular, we show in Section 5 that Theorem 1 leads to the following *finiteness criterion,* which is an improvement of the previously known results of this type (cf. [3], [6]):

Suppose a group G admits a simply-connected G-CW-complex X such that (a) the isotropy group of every vertex is finitely presented; (b) the isotropy group of every edge is finitely generated; and (c) X has a finite 2-skeleton mod G. Then G is finitely presented.

* Partially supported by an NSF grant.

Theorem 1 also leads to an improvement of previously known results in the situation where X has a subcomplex which is a fundamental domain for the action (in the combinatorial sense). This will be discussed in Section 4.

This note grew out of a general investigation of finiteness properties of groups, in which the finiteness criterion above was needed as the starting point (and was initially thought to be well-known). The results of that investigation, which are somewhat more specialized than the results of the present paper, will appear elsewhere.

1. Statement of Theorem 1

Throughout this section G will be an arbitrary group and X a connected non-empty C-CW-complex. We will need to introduce some terminology and notation in order to state Theorem 1.

A 1-cell σ of X is said to be *inverted* under the G-action if there is an element $g \in G$ such that $g\sigma = \sigma$ and g reverses the orientation of σ. By a *tree of representatives* for X mod G we mean a tree $T \subseteq X$ such that the vertex set V of T is a set of representatives for the vertices of X mod G and such that no 1-cell of T is inverted under the G-action. It is easy to see that such a tree T always exists and that the 1-cells of T are inequivalent mod G.

By an *edge* of X we will mean an oriented 1-cell, i.e., a 1-cell σ together with an orientation of σ. Note that G acts on the set of edges of X. It is always possible to choose an orientation for each 1-cell of X which is not inverted under the action, in such a way that these orientations are preserved by G. The corresponding set P of edges, consisting of all the non-inverted 1-cells with the chosen orientations, is then called an orientation of the G-complex X.

Any edge e has a well-defined *initial vertex* (or 'origin') $o(e)$ and a well-defined *final vertex* (or 'terminus') $t(e)$, possibly equal to $o(e)$. Given an edge e, we denote by $\bar{e}$ the *opposite* edge, i.e., the same 1-cell σ with the opposite orientation.

We now make a number of *choices*:

(a) Choose an orientation P.

(b) Choose a tree of representatives T and let V be its set of vertices.

(c) Choose a set E^+ of representatives for P mod G, such that each $e \in E^+$ has $o(e) \in V$ and such that each 1-cell of T (with its chosen orientation) is in E^+. Choose a set E^- of representatives mod G for the edges of X which are inverted under the action, again with $o(e) \in V$ for each $e \in E^-$, and let Σ^- be the corresponding set of 1-cells of X.

(d) For each $e \in E^+$ let $w = w(e)$ be the unique element of V which is equivalent to $t(e)$ mod G, and choose an element $g_e \in G$ such that $t(e) = g_e w$; take $g_e = 1$ if e is in T. Note, then, that the conjugation map c_e given by $g \mapsto g_e^{-1} g g_e$ maps the isotropy group $G_{t(e)}$ onto the isotropy group G_w; in particular, $c_e(G_e) \subseteq G_w$.

(e) Choose a set F of representatives mod G for the 2-cells of X, and choose a

characteristic map $(B, \dot{B}) \to (X, X^{(1)})$ for each $\tau \in F$, where B is the standard 2-cell and $X^{(1)}$ is the 1-skeleton of X. For simplicity, make the choices in such a way that the attaching map $\dot{B} \to X^{(1)}$ is based at a vertex of T for each τ.

Finally, we wish to choose for each $\tau \in F$ a relation which holds in G among the elements g_e and various elements of the isotropy groups G_υ and G_σ of the vertices and 1-cells of X. This requires some preparation.

Note first that any edge of X starting in T has one of the following three forms:

(1) $\quad \upsilon \xrightarrow[he]{} hg_e w \qquad (e \in E^+, \upsilon = o(e), w = w(e), h \in G_\upsilon),$

(2) $\quad w \xrightarrow[hg_e^{-1}\bar{e}]{} hg_e^{-1}\upsilon \qquad (e, \upsilon, w \text{ as in (1)}, h \in G_w),$

(3) $\quad \upsilon \xrightarrow[he]{} ht\upsilon \qquad (e \in E^-, t \text{ inverts } e, h \in G_\upsilon).$

In particular, this gives us a (non-canonical) way of associating to such an edge an element $g \in G$ such that e ends in gT; namely, we set g equal to hg_e in case (1), hg_e^{-1} in case (2), and ht in case (3). [This is non-canonical because the element h in (1)–(3) is not unique, nor is t in (3).]

Now let α be a combinatorial path in X starting at a vertex υ_0 of T. Let e_1 be the first edge of α and let g_1 be associated to e_1 as above. Then e_1 ends in $g_1 T$, so the second edge of α has the form $g_1 e_2$ for some edge e_2 starting in T. Let g_2 be associated to e_2 as above. Then the second edge of α ends in $g_1 g_2 T$. Continuing in this way, we obtain elements $g_1, \ldots, g_n$ of G (where n is the length of α), such that the successive vertices of α lie in $T, g_1 T, g_1 g_2 T, \ldots, g_1 \cdots g_n T$. Set $g = g_1 \cdots g_n$.

In case α is a *closed* path, we have $g\upsilon_0 = \upsilon_0$, i.e., $g \in G_{\upsilon_0}$. Recalling that each g_i was defined in terms of the elements g_e $(e \in E^+)$ and elements of the various groups G_υ $(\upsilon \in V)$ and G_σ $(\sigma \in \Sigma^-)$, we may then view the equation defining g as a *relation* among these elements. To state this more precisely, we introduce an abstract group $\hat{G}$, defined to be the sum (or free product) of the G_υ $(\upsilon \in V)$, the G_σ $(\sigma \in \Sigma^-)$, and a free group generated by elements $\hat{g}_e$ $(e \in E^+)$. Let $\hat{g}_i$ $(i = 1, \ldots, n)$ be obtained by replacing each occurrence of g_e by the corresponding $\hat{g}_e$ in the definition of g_i above, and let $\hat{g} = \hat{g}_1 \cdots \hat{g}_n$. Then, assuming the path is closed, we have a 'relator' $r = \hat{g} g^{-1} \in \hat{G}$, i.e., an element of $\ker\{\hat{G} \to G\}$.

We can now describe the last of our choices:

(f) For each $\tau \in F$, choose a combinatorial path corresponding to the chosen attaching map for τ (cf. (e)), and construct a relator $r_\tau = \hat{g}_\tau g_\tau^{-1}$ from this as above.

Theorem 1. *Suppose X is simply-connected. Then G is generated by the isotropy subgroups G_υ $(\upsilon \in V)$, the isotropy subgroups G_σ $(\sigma \in \Sigma^-)$, and the elements g_e $(e \in E^+)$, subject to the following relations:*

(i) $g_e = 1$ *if e is an edge of T.*

(ii) $g_e^{-1} i_e(g) g_e = c_e(g)$ *for any $e \in E^+$ and $g \in G_e$, where i_e is the inclusion $G_e \hookrightarrow G_{o(e)}$ and $c_e : G_e \to G_{w(e)}$ is as in* (d) *above.* [Thus both sides of the 'relation' are words in

the given generators of G.]

(iii) $i_e(g)=j_e(g)$ *for any* $e\in E^-$ *and* $g\in G_e$, *where* $i_e: G_e \hookrightarrow G_{o(e)}$ *and* $j_e: G_e \hookrightarrow G_\sigma$ *are inclusions,* σ *being the* 1*-cell underlying* e.

(iv) $r_\tau=1$ *(or, equivalently,* $\hat{g}_\tau=g_\tau$*) for any* $\tau\in F$.

We will reformulate Theorem 1 and prove it in Section 3.

Remarks. (1) We have used the somewhat informal language of generators and relations in the statement of Theorem 1, which probably causes no confusion. In case there is any doubt, however, the precise meaning of Theorem 1 is that the map $\hat{G}\to G$ introduced above is surjective, and that its kernel is the normal subgroup of $\hat{G}$ generated by the elements $\hat{g}_e$ (e in T), $\hat{g}_e^{-1}i_e(g)\hat{g}_e c_e(g)^{-1}$ (notation as in (ii)), etc.

(2) If X is simply-connected and the action is *free* (i.e., every isotropy group is trivial), then Theorem 1 reduces to the usual presentation which one obtains from the isomorphism $\pi_1(G\backslash X)\xrightarrow{\sim} G$ and a choice of maximal tree in $G\backslash X$.

2. Extension to the non-simply-connected case

Let Y be a connected but not necessarily simply-connected G-complex. It turns out that one can still obtain a presentation of G in this case, in which the fundamental group of Y plays a role. This result (Theorem 2 below) is actually a consequence of Theorem 1. We are deriving this consequence now, before giving the proof of Theorem 1, because the ideas involved will be needed in that proof.

Let the notation be as in Section 1, but with X replaced by Y, and let $\pi=\pi_1(Y,v)$ for any $v\in V$. [Note that π is well-defined up to canonical isomorphism, since paths in T give a unique change-of-basepoint isomorphism for any two elements of V.] Let $\tilde{G}$ be the group defined by the generators and relations given in the statement of Theorem 1 (with X replaced by Y). More precisely, $\tilde{G}$ is the abstract group generated by the groups G_v ($v\in V$) and G_σ ($\sigma\in\Sigma^-$) and by elements $\tilde{g}_e$ ($e\in E^+$), subject to the relations (i)–(iv), with g_e everywhere replaced by $\tilde{g}_e$). We will show that G is obtained from $\tilde{G}$ by introducing further relations coming from π, analogous to the relations r_τ.

Let α be a closed combinatorial path in Y based at a vertex of T. Recall from the discussion preceding (f) in Section 1 that there is a (non-canonical) procedure for associating to α a relator $\hat{g}g^{-1}$. *Let* $\tilde{g}$ be the image of $\hat{g}$ in $\tilde{G}$ and let $r=\tilde{g}g^{-1}\in\ker\phi$, where $\phi:\tilde{G}\to G$ is the canonical map. Let $[\alpha]$ denote the class in π represented by α, and set $i([\alpha])=r$. Note that i appears to depend on a number of choices, such as the choice of a combinatorial path α representing a given element of π and the choice of group elements g_j associated to α. Using Theorem 1, however, we will prove:

Theorem 2. *The map i is well-defined and is a homomorphism. The resulting sequence*

$$1 \to \pi \xrightarrow{i} \tilde{G} \xrightarrow{\phi} G \to 1$$

is exact.

Proof. Let X be the universal cover of Y. Then the G-action on Y 'lifts' in a well-known way to an action of a group H on X. Namely, H is the set of pairs (g, f) such that g is in G and $f: X \to X$ is a homeomorphism covering the action of g on Y. Note that we have a surjection $H \to G$ whose kernel is the group of deck transformations of the cover and hence is isomorphic to π. [More precisely, a choice of a lifting $\tilde{T}$ of T to X yields a specific such isomorphism.] Note also that the H-orbits of cells of X are in 1–1 correspondence with the G-orbits of cells of Y, and that the isotropy subgroups of H acting on X map isomorphically to the isotropy subgroups of G acting on Y. It is now easy to 'lift' to (H, X) the choices (a)–(f) made for (G, Y), and one deduces from Theorem 1 applied to (H, X) a presentation for H. But this presentation says precisely that $H = \tilde{G}$. Hence ϕ is surjective and its kernel is isomorphic to π.

Let $j: \pi \to \ker \phi$ be the isomorphism just obtained. We will show that $j = i$. Note first that the identification of $\tilde{G}$ with H above yields an action of $\tilde{G}$ on X compatible with the maps $\tilde{G} \to G$ and $X \to Y$. Moreover, the generators G_v, G_σ, and $\tilde{g}_e$ of $\tilde{G}$ have interpretations relative to the $\tilde{G}$-action on X analogous to their interpretations relative to the G-action on Y. For instance the canonical map $G_v \to \tilde{G}$ embeds G_v as the isotropy subgroup of $\tilde{G}$ at the vertex $\tilde{v}$ of $\tilde{T}$ lying over v. Similarly, G_σ can be identified with the isotropy group in $\tilde{G}$ of a 1-cell $\tilde{\sigma}$ of X lying over σ and having a vertex in $\tilde{T}$. Finally, $\tilde{g}_e$ for $e \in E^+$ has the property that $t(\tilde{e})$ is in $\tilde{g}_e \tilde{T}$, where $\tilde{e}$ is the lift of e starting in $\tilde{T}$.

Standard covering space theory, which is what we used to identify π with $\ker \phi$ above, now yields the following description of j. If α is a closed combinatorial path in Y based at $v \in V$ and $\tilde{\alpha}$ is the lift of α to X starting at $\tilde{v}$, then $j([\alpha])$ is the unique element $h \in \ker \phi$ such that $\tilde{\alpha}$ ends at $h\tilde{v}$.

In order to relate this to $r = i([\alpha])$ [which we still do not know to be well-defined], we must refer back to the discussion preceding (f) in Section 1. Let $g_1, \ldots, g_n$ and $e_1, \ldots, e_n$ be the chosen group elements and edges associated to α as in that discussion, and let $\tilde{g}_i$ $(i = 1, \ldots, n)$ be obtained by replacing g_e by $\tilde{g}_e$ in the definition of g_i. Then we have $r = \tilde{g} g^{-1}$, where $g = g_1 \cdots g_n$ and $\tilde{g} = \tilde{g}_1 \cdots \tilde{g}_n$. Let $\tilde{e}_i$ be the lift of e_i to X starting in $\tilde{T}$. Using the remarks above about the action of $\tilde{G}$ on X, it is easy to see that the lift $\tilde{\alpha}$ of α starting in $\tilde{T}$ consists of the edges $\tilde{e}_1, \tilde{g}_1\tilde{e}_2, \tilde{g}_1\tilde{g}_2\tilde{e}_3$, etc., and hence that $\tilde{\alpha}$ ends in $\tilde{g}\tilde{T}$. It follows that $\tilde{\alpha}$ ends at $\tilde{g}\tilde{v}$, which is the same as $r\tilde{v}$ since g is in G_v. But r is in $\ker \phi$, so this says precisely that $j([\alpha]) = r$. Hence r is independent of all the choices, and $j = i$. This completes the proof.

Finally, we need to make a simple observation relating the G-action on paths in

Y to the $\tilde{G}$-action on $\ker \phi$ (by conjugation). Consider a closed combinatorial path α in Y, not necessarily based at a vertex of T. By change of basepoint, α gives rise to a well-defined conjugacy class in π, which we denote by $\{\alpha\}$. We wish to relate the classes $\{\alpha\}$ and $\{g\alpha\}$ for $g \in G$.

Proposition 1. *The classes $i(\{g\alpha\})$ and $i(\{\alpha\})$ in* $\ker \phi$ *are contained in the same $\tilde{G}$-conjugacy class. More precisely, there is an element $\tilde{g} \in \phi^{-1}(g)$ such that $i(\{g\alpha\}) = \tilde{g} i(\{\alpha\}) \tilde{g}^{-1}$.*

Proof. We may assume α is based at $v \in V$. Let $\tilde{\alpha}$ be the lift of α to X starting at $\tilde{v}$. Then $\tilde{\alpha}$ ends at $h\tilde{v}$, where $h = i([\alpha]) = j([\alpha]) \in \ker \phi$. [Here j is as in the proof of Theorem 2.] Now choose a path β from v to gv, lift β to a path $\tilde{\beta}$ in X starting at $\tilde{v}$, and let $\tilde{g} \in \tilde{G}$ be the element covering g such that $\tilde{\beta}$ ends at $\tilde{g}\tilde{v}$. The class $\{g\alpha\}$ is represented by the composite path $\beta(g\alpha)\beta^{-1}$, whose lift to X starting at $\tilde{v}$ is $\tilde{\beta}(\tilde{g}\tilde{\alpha})(\tilde{g}h\tilde{g}^{-1}\tilde{\beta})^{-1}$. Since the latter ends at $\tilde{g}h\tilde{g}^{-1}\tilde{v}$ and $\tilde{g}h\tilde{g}^{-1}$ is in $\ker \phi$, it follows that $i(\{g\alpha\})$ $[=j(\{g\alpha\})]$ is represented by $\tilde{g}h\tilde{g}^{-1}$. Thus $i(\{g\alpha\}) = \tilde{g} i(\{\alpha\}) \tilde{g}^{-1}$.

3. Reformulation and proof of Theorem 1

Throughout this section X will denote a *simply-connected* G-complex as in Theorem 1. We retain all the notation of Section 1. Let K be the abstract graph associated to the 1-skeleton on X. Thus the vertices of K are the vertices of X and the edges of K are the edges of X, as defined in Section 1. (By *abstract graph*, here, we simply mean a graph in the sense of [7]. All unexplained notation and terminology in connection with graphs is that of [7].) Note that G acts on K as a group of graph automorphisms, possibly with inversion.

Let K^+ be the subgraph of K consisting of all the vertices of K and all the edges which are not inverted under the G-action. Let K' be the 'partial barycentric subdivision' of K obtained by subdividing each edge of K not in K^+. Then G acts on K' without inversion and we may apply the results of [7, §I.5.4].

To this end we must make a number of choices, most of which come in a very obvious way from the choices (a)–(d) made in Section 1. The only thing which requires comment is the choice of a tree of representatives for K' mod G. The tree of representatives T for X mod G which we chose in Section 1 can be viewed as a tree of representatives for K^+ mod G. This extends easily to a tree T' of representatives for K' mod G; namely, we adjoin one 'spoke' to T for each $e \in E^-$, this spoke being the 'first half' of the subdivided edge e, going from $o(e)$ to the barycenter of e.

Now let L be the quotient graph $G \backslash K'$. Note that L is obtained from the graph $L^+ = G \backslash K^+$ by attaching a spoke for every $e \in E^-$. As explained in [7, §I.5.4], L supports a graph of groups (G, L), whose vertex and edge groups are the isotropy subgroups of G at the chosen representatives for the vertices and edges of K'. In the present context this takes the following form:

To a vertex of L corresponding to a vertex v of T we assign the isotropy group G_v. To a vertex of L corresponding to the barycenter of an edge $e \in E^-$ we assign the isotropy group at this barycenter, or, equivalently, the isotropy group G_σ (where σ is the 1-cell of X underlying e). To an edge of L corresponding to $e \in E^+$ we assign the isotropy group G_e. To an edge of L corresponding to the first half of a subdivided $e \in E^-$ we assign the isotropy group of that first half, or, equivalently, the isotropy group G_e. Finally, the required injections of an edge group into the two corresponding vertex groups are given by inclusion maps and the conjugation maps c_e introduced in (d) of Section 1. See [7] for more details.

Now let $\tilde{G}$ be the fundamental group of this graph of groups with respect to the chosen maximal tree, i.e., $\tilde{G} = \pi_1(G, L, q(T'))$, where $q: K' \to L$ is the quotient map. For future reference we note that $\tilde{G}$ is an amalgam if $K^+ \neq K$; namely, if we set $\tilde{G}^+ = \pi_1(G, L^+, q(T))$, then $\tilde{G}$ is the sum of $\tilde{G}^+$ and the groups G_σ $(\sigma \in \Sigma^-)$, amalgamated along the subgroups G_e $(e \in E^-)$. This makes sense because G_e is a subgroup of $\tilde{G}^+$ via $G_e \subseteq G_{o(e)} \subseteq \tilde{G}^+$. [It is also worth noting here that G_e is of index 2 in the corresponding G_σ.]

It is now easy to relate $\tilde{G}$ to the presentation given in Theorem 1. In fact, a glance at the definition of the fundamental group of a graph of groups shows that $\tilde{G}$ is generated by the groups G_v $(v \in V)$, the groups G_σ $(\sigma \in \Sigma^-)$, and elements $\tilde{g}_e$ $(e \in E^+)$, subject to the relations (i)–(iii) of Theorem 1 (with g_e everywhere replaced by $\tilde{g}_e$). The content of Theorem 1, then, is that G is obtained from $\tilde{G}$ by introducing the relations r_τ. More precisely, if we now denote by r_τ the image under $\hat{G} \to \tilde{G}$ of the element called r_τ in Section 1, then we have the following reformulation of Theorem 1:

Theorem 1′. *Let X be a simply-connected G-complex as above. Then the canonical map $\tilde{G} \to G$ is surjective and its kernel is the normal subgroup of $\tilde{G}$ generated by the r_τ $(\tau \in F)$.*

Proof. Suppose first that X is 1-dimensional (and hence that there are no r_τ). Then K' is a tree on which G acts without inversion, and Theorem 1′ is simply the result obtained by applying to this action the Bass–Serre structure theorem [7, §I.5.4, Theorem 13] for a group acting without inversion on a tree.

Note that all the results of Section 2 are now available to us for any connected 1-dimensional G-complex Y, since Theorem 1 has been proved in the 1-dimensional case. In particular, if X is of arbitrary dimension, we apply Section 2 to the 1-skeleton Y of X. It is immediate from the definitions that the group $\tilde{G}$ of Section 2 for this Y is precisely the same as the group we have called $\tilde{G}$ in the present section. We therefore have the short exact sequence

$$1 \to \pi \xrightarrow{i} \tilde{G} \xrightarrow{\phi} G \to 1$$

of Theorem 2, where $\pi = \pi_1(Y)$. (Alternatively, instead of using Theorem 2 we

could appeal here to [7, §I.5.4, exercise 3].) Since X is simply-connected, we know that π is normally generated by elements represented by attaching maps for the 2-cells of X. And if we just want a set of generators for $i(\pi)$ as a normal subgroup of $\tilde{G}$, then it suffices to consider attaching maps α_τ for the 2-cells in F. For we can obtain characteristic maps (and hence attaching maps) for the remaining 2-cells by transforming the α_τ by elements of G; in view of Proposition 1, the resulting elements of $i(\pi)$ are conjugate in $\tilde{G}$ to those coming from the α_τ. Recalling now that r_τ, by definition, is equal to $i([\alpha_\tau])$ for some choice of α_τ, we see that $i(\pi) = \ker \phi$ is indeed the normal closure of the r_τ. This completes the proof of Theorems 1 and 1′.

4. Example: Group actions with a subcomplex as fundamental domain

In this section we assume: (a) X is a simply-connected G-CW-complex; (b) W is a subcomplex such that every cell of X is equivalent mod G to a unique cell of W; and (c) G acts without inversion on the 1-cells of X. [*Note.* It follows from (b) that no edge of K with two distinct vertices is inverted by the G-action; so the effect of (c) is simply to rule out the possibility that there are loops which are inverted.] We will denote by V (resp. E) the set of vertices (resp. 1-cells) of W.

This situation was studied in [9], where X was further assumed to be simplicial. The main theorem of [9] says that G is then the sum of the vertex isotropy groups G_v ($v \in V$), amalgamated along their intersections. We will show that Theorem 1 leads to the following more precise result:

Theorem 3. *Under the hypotheses* (a)–(c), *G is the sum of the groups G_v ($v \in V$), amalgamated along the subgroups G_e ($e \in E$).*

[The statement of the theorem means that G is generated by the groups G_v, subject to relations for each $e \in E$ which identify the copy of G_e in G_v with the copy of G_e in G_w, v and w being the vertices of e. Note, however, that this does *not* say that G is an amalgam in the usual sense, except in the special case where the 1-skeleton of W is a tree.]

Proof. We must make a number of choices in order to apply Theorem 1. Note first that W is necessarily connected, since its 1-skeleton is a retract of that of X. [In fact, W is simply-connected, but we do not need to know this.] We can therefore take the tree T in Section 1 to be a maximal tree of W. For E^+ we may take the set E of 1-cells of W, each cell being given some arbitrary orientation, and E^- is of course empty. We take $g_e = 1$ for all e. Finally, let F be the set of 2-cells of W and choose r_τ to be a word in the g_e. [More precisely, if we express an attaching map for τ in terms of the edges in E and their opposites, then we may take r_τ to be the corresponding word in the g_e and their inverses.]

Now look at the presentation given in Theorem 1. This yields generators G_v and g_e, and relations (i), (ii), and (iv). We may replace (i) by the relations $g_e = 1$ for all e, since these relations do in fact hold in G. Hence the generators g_e and relations (i) can be dropped, with g_e being replaced by 1 wherever it occurs in (ii) and (iv). But then (iv) becomes vacuous and (ii) becomes the 'amalgamation' relation which identifies the copy of G_e in $G_{o(e)}$ with that in $G_{t(e)}$ for each $e \in E$. This proves the theorem.

5. Finite presentation

As we stated in the introduction, the following consequence of Theorem 1 is what motivated this paper:

Theorem 4. *Let G be a group which admits a simply-connected G-complex satisfying:*

(a) *Every vertex isotropy group is finitely presented.*

(b) *Every edge isotropy group is finitely generated.*

(c) *X has a finite 2-skeleton* mod G.

Then G is finitely presented.

Proof. In view of (c), it suffices to show that the group $\tilde{G}$ of Theorem 1′ is finitely presented. Now the hypotheses clearly imply that the group $\tilde{G}^+$ defined in Section 3 is finitely presented. And $\tilde{G}$ is obtained from $\tilde{G}^+$ by adjoining finitely many finitely generated groups G_σ, amalgamated along subgroups G_e of index 2. The theorem therefore follows from:

Lemma. *Let H be an amalgam $A *_C B$, where A is finitely presented, B is finitely generated, and C is of finite index in B. Then H is finitely presented.*

Proof. Choose a finite set of generators a_i for A, a finite set of generators c_i for C, and a finite set of coset representatives b_i for B/C. Now let $\tilde{H}$ be the group with generators $\tilde{a}_i$ and $\tilde{b}_i$ (in 1–1 correspondence with the a_i and b_i, respectively), subject to the following three types of relations:

(i) Finitely many relations among the $\tilde{a}_i$ which define the group A.

(ii) Finitely many relations which describe the permutation action of C on B/C; more precisely, for each c_i and b_j there is an equation in B of the form $c_i b_j = b_k c$ for some index k and some $c \in C$. Express c_i and c in terms of the generators a_q of A, let $\tilde{c}_i$ and $\tilde{c}$ be the corresponding expressions in the $\tilde{a}_q$, and introduce the relation $\tilde{c}_i \tilde{b}_j = \tilde{b}_k \tilde{c}$.

(iii) Finitely many relations which describe the permutation action of the b_i on B/C; more precisely, for each pair of representatives b_i, b_j, there is an equation in B of the form $b_i b_j = b_k c$ for some index k and some $c \in C$. Let $\tilde{c}$ be the correspon-

ding word in the $\tilde{a}_q$ as in (ii), and introduce the relation $\tilde{b}_i\tilde{b}_j = \tilde{b}_k\tilde{c}$.

We have a canonical map $\psi : \tilde{H} \to H$, which we will prove to be an isomorphism. Let $\tilde{A}$ be the subgroup of $\tilde{H}$ generated by the $\tilde{a}_q$. Clearly $\tilde{A}$ maps isomorphically to A under ψ. Let $\tilde{C} \subseteq \tilde{A}$ be the subgroup corresponding to C under this isomorphism, and let $\tilde{B}$ be the subgroup of $\tilde{H}$ generated by $\tilde{C}$ and the $\tilde{b}_i$. The relations (ii) and (iii) imply that the cosets $\tilde{b}_i\tilde{C}$ exhaust all the left cosets of $\tilde{C}$ in $\tilde{B}$, and it follows that $\tilde{B}$ maps isomorphically to B under ψ. One can now use the universal mapping property of the amalgam H to define a map $H \to \tilde{H}$ inverse to ψ, so the latter is indeed an isomorphism and H is finitely presented.

References

[1] H. Abels, Generators and relations for groups of homeomorphisms, in: Transformation Groups (Proc. Conf., Univ. Newcastle upon Tyne, 1976), London Math. Soc. Lecture Note Series, No. 26 (Cambridge Univ. Press, Cambridge, 1977) 3–20.

[2] H. Behr, Über die endliche Definierbarkeit von Gruppen, J. Reine Angew. Math. 211 (1962) 116–122.

[3] D. Flöge, Zur Struktur der PSL_2 über einigen imaginär-quadratischen Zahlringen, Dissertation, Frankfurt am Main (1980).

[4] M. Gerstenhaber, On the algebraic structure of discontinuous groups, Proc. Amer. Math. Soc. 4 (1953) 745–750.

[5] A.M. Macbeath, Groups of homeomorphisms of a simply connected space, Ann. of Math. 79 (1964) 473–488.

[6] G. McHardy, Über Struktur und endliche Präsentierbarkeit simplizialer Automorphismengruppen, Diplomarbeit, Frankfurt am Main (1977).

[7] J.-P. Serre, Trees (Springer, Berlin, 1980). (Translation of Arbres, Amalgames, SL_2; Astérisque 46 (1977).)

[8] C.L. Siegel, Discontinuous groups, Ann. of Math. 44 (1943) 674–689.

[9] C. Soulé, Groupes opérant sur un complexe simplicial avec domaine fondamental, C.R. Acad. Sci. Paris (A) 276 (1973) 607–609.

[10] R.G. Swan, Generators and relations for certain special linear groups, Advances in Math. 6 (1971) 1–77.