

ON THE FP_3 -CONJECTURE FOR METABELIAN GROUPS

ROBERT BIERI AND JENS HARLANDER

1. Introduction

A metabelian group G fits into an exact sequence

$$1 \rightarrow M \rightarrow G \rightarrow Q \rightarrow 1, \quad (1)$$

where both M and Q are abelian. G is uniquely given, up to isomorphism, by the abelian group Q , the Q -module M (by conjugation in G), and the cohomology class of the extension in $H^2(Q, M)$.

The group G is of type F_n , $n \geq 0$, if it admits a $K(G, 1)$ -complex with finite n -skeleton. Every group is of type F_0 . The group G is of type F_1 if and only if it is finitely generated, and of type F_2 if and only if it is finitely presented. The metabelian groups constitute a class of groups G for which the property of being of type F_n relates in a very subtle way to the internal structure of G . The F_n -conjecture proposes a complete description of this relationship in terms of the Q -module M alone. We state its assertion in detail below. It requires the notion of the *geometric invariant* $\Sigma(M)$ of a Q -module M was introduced in [5] to settle the case $n = 2$. The precise wording and first evidence of the F_n -conjecture for $n > 2$ came up only in [2]. The full conjecture is still an open question. The purpose of this paper is to prove it in the case when $n = 3$ and extension (1) splits.

The extension G of M by Q is finitely generated if and only if M is finitely generated as a Q -module and Q is finitely generated. This obvious fact is all there is to the F_1 -conjecture, and we shall henceforth assume that G is finitely generated and $n \geq 2$. Moreover, the property that G be of type F_n is inherited and detected by subgroups of finite index. Hence we can pass to a subgroup of finite index and henceforth assume that Q is free abelian of finite rank m .

We write Q as a multiplicative group, but we use an embedding of Q into \mathbb{R}^m that identifies Q with the integral lattice of \mathbb{R}^m and thus converts the multiplicative notion into an additive one.

The F_n -conjecture requires some elementary geometry that is introduced via the *support of group ring elements* $\lambda \in \mathbb{Z}Q$. If $\lambda = \sum_{q \in Q} n_q q$ is the unique expansion of λ in the free abelian group $\mathbb{Z}Q$, then we put

$$\text{supp}(\lambda) = \{q \in Q \mid n_q \neq 0\}.$$

This is a finite subset of Q with a geometric flavor arising from the embedding of Q into \mathbb{R}^m . We endow the latter with the standard inner product and consider the unit sphere $S^{m-1} \subseteq \mathbb{R}^m$. For each unit vector $u \in S^{m-1}$ and each non-zero $\lambda \in \mathbb{Z}Q$, we consider the *u -shift* of λ , defined by

$$\chi_u(\lambda) = \min \langle u, \text{supp}(\lambda) \rangle.$$

We shall also need the *diameter* of λ , which is given by

$$\text{diam}(\lambda) = \text{diam}(\text{supp}(\lambda)).$$

The centralizer $C(M)$ of the Q -module M is the set of all group ring elements $\lambda \in \mathbb{Z}Q$ with $\lambda m = m$ for all $m \in M$. One way to define the *geometric invariant* $\Sigma(M)$ in [5] is to identify it with a subset of S^{m-1} by putting

$$\Sigma(M) = \{u \in S^{m-1} \mid \text{sup}(\chi_u(C(M))) > 0\}.$$

It is a remarkable fact (although it is not needed in this paper) that the set $\Sigma(M)$ is polyhedral, that is, it is a finite union of finite intersections of open hemispheres. This was established in [3]. It can be shown that the above equation holds with $C(M)$ replaced by a finite subset of $C(M)$; see [4].

Following [2], we call the Q -module M *n-tame*, $n \in \mathbb{N}$, if every n -point subset of S^{m-1} not contained in an open hemisphere includes at least one point from $\Sigma(M)$, in other words, if an equation $\sum_{i=1}^n r_i u_i = 0$, with $0 < r_i \in \mathbb{R}$ and $u_i \in S^{m-1}$, can only be satisfied when one of the u_j is in $\Sigma(M)$.

F_n -CONJECTURE. The finitely generated metabelian group G is of type F_n if and only if the Q -module M is n -tame. (By [5], we know that for metabelian groups, the homological condition that G be of type FP_n coincides with the homotopical condition that G be of type F_n . Therefore the F_n -conjecture is often referred to as the FP_n -conjecture.)

The charm of the F_n -conjecture lies in the fact that both of its directions are difficult and still unsettled, in general. The F_2 -conjecture was established in [5]. In [2], it was shown that if G is of type F_n , then $M \otimes K$ is n -tame as a KQ -module for every field K . In a remarkable paper [1], Hans Åberg established the full F_n -conjecture for the case when M , as an abelian group, is virtually torsion-free of finite rank. Noskov [13] and Kochloukova [12] extended Åberg's 'only-if' method considerably. By [12], it is known that the 'only-if' part of the F_n -conjecture holds true if either the additive group of M is a torsion group or if the short exact sequence (1) splits. Rather less is known about the 'if' part of the F_n -conjecture. Bux [8] and Kochloukova extended Åberg's 'if' method to situations where the additive group of M is a torsion group. The best result here so far is the full F_n -conjecture for the case when M is torsion and of Krull dimension 1 as a $\mathbb{Z}Q$ -module [11, 12].

In this paper, we contribute to the 'if' part of the F_3 -conjecture by proving the following theorem.

MAIN THEOREM. *If the exact sequence (1) splits, then the F_3 -conjecture holds.*

To prove this theorem, we lift the 'if' method of [5], which used generators and defining relations, by one dimension, using methods from combinatorial homotopy theory.

2. A reminder about spherical diagrams

Maps from the 2-sphere into a piecewise-linear CW-complex can be described conveniently (up to homotopy) by spherical diagrams. As we will be using diagrams at various places in this paper, we briefly recall some of their features in this section.




FIGURE 2.1. Bouquet of spheres.




FIGURE 2.2. Replace move.

We rely on diagrams only in an intuitively obvious way. More details about planar and spherical diagrams can be found in papers by Lyndon and Schupp [10], Collins and Huebschmann [9] and Bogley and Pride [6].

Let K be a piecewise linear CW-complex. For every cell σ , fix a convex polyhedron P_σ and a characteristic map $\phi_\sigma : P_\sigma \rightarrow K$ such that the restriction of ϕ_σ to a face of P_σ is, up to PL-homeomorphism, of the form ϕ_τ for some cell τ . By $-\phi_\sigma$, we denote ϕ_σ precomposed with an orientation-reversing homeomorphism $s_\sigma : P_\sigma \rightarrow P_\sigma$. Suppose that $f : S \rightarrow K$ is a map from an oriented tessellated 2-sphere into K such that the restriction of f to an oriented cell ξ agrees with (up to orientation-preserving PL-homeomorphism) $\pm\phi_\sigma$ for some cell σ of K with $\dim(\sigma) = \dim(\tau)$. Label the face ξ by $\pm\sigma$. The map f is (up to homotopy) encoded by the labeled oriented tessellated 2-sphere S . The sphere S , together with the labeling, is called a spherical diagram over K . One can replace the 2-sphere by other 2-complexes and obtain more general diagrams.

In the following we will be concerned with a PL-3-complex X (with single 0-cell), where the attaching maps of the 3-cells are given by spherical diagrams, and its second homotopy module $\pi_2(X)$. Although not every element from $\pi_2(X)$ can be represented by a spherical diagram, the module is generated by such elements. If α is an element of $\pi_1(X)$ represented by an edge loop w , and β is an element of $\pi_2(X)$ represented by a spherical diagram S , then the element $\alpha\beta$ can be represented by S ‘with a tail added on’. The tail is a subdivided interval that is mapped combinatorially onto the loop w . From this, it is clear that every element of $\pi_2(X)$ can be represented by a bouquet of spherical diagrams (see Figure 2.1).

There are certain ‘moves’ one can make in diagrams that do not change the homotopy class of the map represented. The one that is important to us is the ‘replace’ move: if S is a spherical diagram over X that contains a subdiagram A that is also part of the boundary of a 3-cell (oppositely oriented), $\partial(e^3) = -A \cup B$, then A

can be replaced in S by B to obtain a new spherical diagram S' that represents the same homotopy class as S . The situation (in dimension 1) is shown in Figure 2.2.

The homotopy simply ‘pushes’ part of S through the 3-cell e^3 .

3. Constructing part of a classifying complex for M

3.1.

To prove the main result, it suffices to construct a connected 3-dimensional Q -CW-complex X with the following properties:

- (a) X has a single 0-cell $*$ (fixed by Q), the fundamental group of X is isomorphic to M as Q -modules, and the second homotopy group of X is trivial.
- (b) The quotient complex X/Q is finite.

The proof of this is fairly straightforward. Let \tilde{X} be the universal covering of X , and let $\tilde{*}$ be a vertex that maps to $*$ under the covering projection. The Q -action on X lifts uniquely to a Q -action on \tilde{X} , fixing the vertex $\tilde{*}$, and one uses the M -action on \tilde{X} by covering transformations to construct a G -action, where G is the split extension of M by Q . If \tilde{x} is a point in \tilde{X} and $(m, q) \in G$, then simply define

$$(m, q)\tilde{x} = m(q\tilde{x}).$$

One can check that this makes \tilde{X} into a G -complex with cell stabilizers all isomorphic to subgroups of Q and hence finitely generated abelian. It now follows from standard arguments (see [7]) that G is F_3 .

3.2.

We will first construct a 3-complex X satisfying condition (a) in Section 3.1. Choose a resolution of the Q -module M by finitely generated Q -modules

$$\mathbb{Z}B \xrightarrow{\partial} \mathbb{Z}A \longrightarrow M \longrightarrow 0,$$

where A and B are free Q -sets with finitely many orbits. It is convenient to put an ordering on A and B that is compatible with the Q -action. Such an ordering can be obtained from an ordering of Q and an ordering on some fixed finite sets of orbit representatives.

Let Z be the subset of the product of spheres $\prod_{a \in A} S_a^1$ consisting of all tori $S_{a_1}^1 \times \dots \times S_{a_k}^1$ with $a_1 < \dots < a_k$. Now Q acts in the obvious way on the product of spheres and makes Z into a cubical Q -complex. Note that Z is a $K(\mathbb{Z}A, 1)$ -complex with a single vertex $*$, and that the 1-cells of Z are in one-to-one correspondence with the elements of A .

For every $b \in B$ we attach to Z a 2-cell e_b^2 as follows. If $\partial(b) = \sum_{i=1}^k n_i a_i$, where $a_1 < \dots < a_k$, use the attaching path $w = a_1^{n_1} \dots a_k^{n_k}$ (we use the ordering on A to obtain a uniquely defined edge path from a sum on the a_i). Q acts on the set of new 2-cells by $qe_b^2 = e_{qb}^2$.

For every pair (b, a) we furthermore attach a 3-cell $e_{(b,a)}^3$ (shown in Figure 3.1). This 3-cell is the product of the 2-cell e_b^2 and the 1-cell a . We refer to such 3-cells as prisms.

Again, Q acts on this set of 3-cells by acting on the labels: $qe_{(b,a)}^3 = e_{(qb,qa)}^3$. We call the resulting Q -complex Y . The action induces Q -module structures on all homotopy groups and turns the fundamental group into a Q -module isomorphic to M .




FIGURE 3.1. Prism 3-cell.




FIGURE 3.2. Connection of pair in 1-skeleton of S .

LEMMA 3.1. *The second homotopy module $\pi_2(Y)$ is a finitely generated Q -module.*

Proof. Consider the composition of Q -module maps

$$\pi_2(Y) \xrightarrow{h} H_2(Y) \xrightarrow{j} H_2(Y, Z) \approx \mathbb{Z}B,$$

where h is the Hurewicz map and j is induced by inclusion. It suffices to show that this composition $j \circ h$ is injective because the image is a submodule of the finitely generated Q -module $\mathbb{Z}B$ and hence is finitely generated (the group ring $\mathbb{Z}Q$ is noetherian).

Note first that a bouquet of spherical diagrams S representing an element $\beta \in \pi_2(Y)$ in the kernel has to contain e_b^2 cells in oppositely oriented pairs. Connect such a pair by a path uvw in the 1-skeleton of S as shown in Figure 3.2. Here u and w are paths in the 1-skeleton of spherical diagrams, whereas v is contained in the stems of the bouquet. We can add a stack of prime 3-cells along the path uvw as shown in Figure 3.3. Note first that the boundary of this configuration is a diagram L that represents an element of $\pi_2(Y)$ that agrees with the element β represented by S , as we altered S only by boundaries of 3-cells. Note next that L contains two fewer e_b^2 -faces than S . To obtain a bouquet of spheres S' from L that represents β , simply open up L along u and w . If we continue this process, we arrive at a bouquet of spherical diagrams that represents β but does not contain e_b^2 -faces. This bouquet of spheres over Z represents the trivial element of $\pi_2(Y)$, as the subcomplex Z of Y is aspherical. Thus β is trivial. \square




FIGURE 3.3. Adding a stack of prisms.

Lemma 3.1 shows that we can attach finitely many Q -orbits of 3-cells to Y to obtain a Q -complex with trivial second homotopy module. Denote by X the 3-skeleton of this complex. Note that X satisfies condition (a) in Section 3.1. The cells of X are the cubes of Z^3 , the 2-cells e_b^2 , and the prisms $e_{(b,a)}^3$, together with finitely many Q -orbits of additional 3-cells. However, X/Q is not finite.

4. Support geometry on X

By fixing Q -orbit representatives $A_0 \subseteq A$, we can define the support of an element $a = qa_0 \in QA_0 = A$ to be the set $\text{supp}(a) = \{q\}$. As the 1-cells in X are in one-to-one correspondence with the elements from A , this also gives us a way to speak of the support of a 1-cell in X .

If $f : K \rightarrow X$ is a map of a finite complex K into X , then we define

$$\text{supp}(f) = \bigcup \text{supp}(a),$$

where the union is taken over all 1-cells contained in the image $f(K)$.

If σ is a cell in X , then we define

$$\text{supp}(\sigma) = \text{supp}(f_\sigma),$$

where f_σ is an attaching map for σ . If $c = \sum_{i=1}^k n_i \sigma_i \in C_n(X)$ is a cellular n -chain, then we define the support of c to be the union of the supports of the cells σ_i .

If D is a diagram over X , then it represents a map f_D from a finite complex into X , and we can define

$$\text{supp}(D) = \text{supp}(f_D).$$

Of course $\text{supp}(D)$ is just the union of the support of the edge labels in D .

We define the diameter $\text{diam}(C)$ of a finite subset C of \mathbb{R}^m to be the diameter of the smallest ball containing it. As $Q \subseteq \mathbb{R}^m$, we can speak of the diameter of a map $f : K \rightarrow X$ as above by taking

$$\text{diam}(f) = \text{diam}(\text{supp}(f)).$$

Using this we can define the diameter of cells, chains and diagrams as above. Note that the diameter of a 1-cell is 0. The support of a square consists of two points, and the diameter of the square is the distance between these points. The support of a cube consists of three points. The support of a prism $e_{(b,a)}^3$ is $\text{supp}(e_b^2) \cup \text{supp}(a)$.

If ρ is a positive real number, then we write $X_\rho \subseteq X$ for the subcomplex consisting




FIGURE 4.1. Prism of diameter ρ and its support.

of all cells with diameter less than or equal to ρ . Similarly define Z_ρ . It is clear that X_ρ/Q is finite, and that X_ρ contains the 1-skeleton of X . We aim to show that there is some ρ with $\pi_1(X_\rho) = M$ and $\pi_2(X_\rho) = 0$ so that X_ρ satisfies both condition (a) and condition (b) in Section 3.1. The first step towards this is the next lemma. (See Figure 4.1.)

LEMMA 4.1. *Let S be a spherical diagram over X of diameter ρ that contains e_b^2 -faces only in oppositely oriented pairs. Then S represents the trivial element of $\pi_2(X_\rho)$.*

Proof. Proceed as in the proof of Lemma 3.1. Consider a pair of oppositely oriented e_b^2 -faces. Connect their base points by a path u . We can use a stack of prisms to replace the pair by squares. The support of the stack of prisms is $\text{supp}(e_b^2) \cup \text{supp}(u) \subseteq \text{supp}(S)$, which has diameter less than or equal to ρ , so each prism used is contained in X_ρ . If we continue this process, then we end up with a spherical diagram S' over Z that represents the same element of $\pi_2(X_\rho)$ as S and $\text{supp}(S') = \text{supp}(S)$. If a_1, \dots, a_n are edge labels occurring in S' , then S' is a spherical diagram over $S_{a_1}^1 \times \dots \times S_{a_n}^1$. Then 3-skeleton of this n -torus is a subcomplex of X_ρ with trivial second homotopy group; thus S' represents the trivial element of $\pi_2(X_\rho)$. \square

5. Entrance of n -tameness

Now we have to impose the assumption that M is n -tame. This means that for each set of n -points $u_1, \dots, u_n \in S^{m-1}$ that is not contained in an open hemisphere, there is a centralizer $\lambda \in C(M)$ such that, for at least one u_j , we have

$$\chi_{u_j}(\lambda) = \min\langle u_j, \text{supp}(\lambda) \rangle > 0.$$

Using compactness of $S^{m-1} \times \dots \times S^{m-1}$, one easily finds the following.

LEMMA 5.1. *If M is n -tame, then there is a finite subset $\Lambda \subseteq C(M)$ and a number $\epsilon > 0$ with the property that, for each n points subset $\{u_1, \dots, u_n\}$ of S^{m-1} , not contained in an open hemisphere, there is a centralizer $\lambda \in \Lambda$ such that, for at least one u_j , we have $\chi_{u_j}(\lambda) > \epsilon$.*

6. Pushing cells

6.1.

For each $(\lambda, a) \in \Lambda \times A$, the difference $\lambda a - a$ is an element of the kernel of the augmentation map $\mathbb{Z}A \rightarrow M$. Using the fact that $\Lambda \times A$ is Q -finite, we may thus assume that there is some $b \in B$ with $\partial(b) = \lambda a - a$. We denote the corresponding 2-cell e_b^2 of X by $e_{(\lambda,a)}^2$, and call this a pushing-cell. Note that the boundary of $e_{(\lambda,a)}^2$ is of the form $ap(\lambda a)$, where $p(\lambda a)$ is a path with support $q + \text{supp}(\lambda)$, where $\{q\}$ is the support of a .

We use the cell $e_{(\lambda,a)}^2$ to push the edge a across $e_{(\lambda,a)}^2$ to the path $p(\lambda a)$. (See Figure 6.1.) This has the effect that the support of a is pushed in direction $u \in S^{m-1}$ if $\chi_u(\lambda) > 0$.

We will have to push not only edges but also 2-cells e_b^2 across (not yet specified) 3-cells e^3 in direction u . In the algebraic setting of the free resolution

$$\mathcal{F} = (\mathbb{Z}C \rightarrow \mathbb{Z}B \rightarrow \mathbb{Z}A \rightarrow M),$$

this is easy. Multiplication by λ defines a chain endomorphism of \mathcal{F} , and so the assignment $b \mapsto \lambda b$ can be viewed as ‘pushing the basis element $b \in B$ to λb ’ along a given chain homotopy $\lambda(-) \simeq \text{id}_M$. We will mimic the same procedure in the homotopical setting of X . However, first we need more notation. For a given direction $u \in S^{m-1}$ and a cell σ of X , define

$$\chi_u(\sigma) = \min\langle u, \text{supp}(\sigma) \rangle.$$

Let X_u^+ be the subcomplex of X consisting of cells σ such that $\chi_u(\sigma) \geq 0$.

LEMMA 6.1. *Suppose that $u \in \Sigma(M)$. Then the inclusion induced map $\pi_1(X_u^+) \rightarrow \pi_1(X)$ is injective.*

Proof. As both groups $\pi_1(X_u^+)$ and $\pi_1(X)$ are abelian, we can argue homologically. Suppose that $\alpha \in C_1(X_u^+)$ is a 1-cycle that is the boundary of a 2-chain $\beta \in C_2(X)$. As $u \in \Sigma(M)$, there exists a centralizer $\lambda \in C(M)$ such that $\chi_u(\lambda) > 0$. For s a positive integer, we have

$$\chi_u(\lambda^s \beta) \geq s\chi_u(\lambda) + \chi_u(\beta),$$

and so $\lambda^s \beta$ is a chain in $C_2(X_u^+)$ for s large enough. The difference $\lambda^s \alpha - \alpha$ is the boundary of a 2-chain γ made up of pushing 2-cells from X_u^+ (if $s = 1$ and $\alpha = \sum_{i=1}^k n_i a_i$, $n_i \in \mathbb{Z}$, then $\gamma = \sum_{i=1}^k n_i e_{(\lambda,a_i)}^2$). Thus

$$\alpha = \lambda^s \alpha - \partial(\gamma) = \partial(\lambda^s \beta - \gamma),$$

which shows that α is the boundary of a 2-chain in $C_2(X_u^+)$. □




FIGURE 6.1. Pushing 2-cell.




FIGURE 6.2. Collar diagram.

For the construction of the pushing 3-cells, we need Corollary 6.2 and Figure 6.2.

COROLLARY 6.2. *Let Λ and ϵ be as in Lemma 5.1. Consider the boundary path $w = a_1 \dots a_k$ of a 2-cell e^2_b of X . For each edge a_j , let $\lambda_j \in \Lambda$. We consider the collar diagram over X (shown in Figure 6.2).*

If $u \in S^{m-1}$ has the property that $\chi_u(\lambda_j) \geq \epsilon$ for all j , then the crosshatched region of the diagram can be filled in with a diagram D_u such that $\chi_u(D_u) \geq \chi_u(w) + \epsilon$.

Proof. Let w' be the boundary path of the crosshatched region of the diagram. Note that $\chi_u(w') \geq \chi_u(w) + \epsilon$. As $u \in \Sigma(M)$, Lemma 6.1 tells us that there exists a diagram D with $\chi_u(D) \geq \chi_u(w')$. □

6.2.

It is important to note that, because the set

$$\Delta = \{u \in S^{m-1} \mid \chi_u(\lambda_j) \geq \epsilon, j = 1, \dots, k\}$$

is compact, we can get away with a finite set of filling diagrams $\{D_u \mid u \in \Delta\}$ such that $\chi_u(D_u) \geq \chi_u(w) + \epsilon/2$.

Now a 2-cell e^2_b together with a filled in collar diagram, as in Corollary 6.2, make up a spherical diagram that we can use to attach a 3-cell to X . These cells will be our pushing 3-cells. As $B = QB_0$ is Q -finite and we can work with a finite set of filling diagrams for a given $e^2_b, b \in B_0$, the set of pushing 3-cells can be assumed to be Q -finite.

7. Proof of the Main Theorem

7.1.

Let $\{\rho_0 < \rho_1 < \rho_2 < \dots\}$ be the set of diameters of cells of X larger than or equal to some big real number ρ_0 (occurring as diameter) that we will specify as we go along. (As there are only Q -finitely many cells of diameter less than or equal to some fixed number ρ , the countable set of diameters forms a discrete subset of the reals.) We aim to construct a cellular map $r : X \rightarrow X$ that is the identity on the




FIGURE 7.1. Image of square.




FIGURE 7.2. The tinted region shows $q' + \text{supp}(\lambda)$ in the case $\langle u, \text{supp}(\lambda) \rangle > \epsilon$.

subspace X_{ρ_0} and the property that $r(X_{\rho_i}) \subseteq X_{\rho_{i-1}}$, $i = 1, 2, 3, \dots$. The Main Theorem follows quickly once the map has been constructed. Indeed, an iteration of r gives a retraction $r^\infty : X \rightarrow X_{\rho_0}$. Thus, as $\pi_2(X)$ is trivial, so is $\pi_2(X_{\rho_0})$. Furthermore, it follows that the inclusion induced map $\pi_1(X_{\rho_0}) \rightarrow \pi_1(X)$ is injective. As it is also surjective (both spaces have the same 1-skeleton), we can see that X_{ρ_0} fulfills conditions (a) and (b) of Section 3.1, and we are done.

Edges have diameter 0 and all cells other than cubes and prisms have bounded diameter, so we can choose ρ_0 big enough so that X_{ρ_0} contains all edges and all cells except large cubes and large prisms.

To define the map r on the 2-skeleton, it suffices to define it on squares e^2 of diameter ρ larger than ρ_0 . Let $\text{supp}(e^2) = \{q, q'\}$, where $\{q\}$ and $\{q'\}$ are the supports of the edges a and a' of the square e^2 . Let u be the unit vector

$$u = \frac{q - q'}{|q - q'|}.$$

As M is 2-tame, there exists a centralizer $\lambda \in \Lambda$ such that

$$\langle -u, \text{supp}(\lambda) \rangle > \epsilon \quad \text{or} \quad \langle u, \text{supp}(\lambda) \rangle > \epsilon.$$

The left-hand and right-hand diagrams in Figure 7.1 describe the map $r : e^2 \rightarrow X$ in the first and second case, respectively. (The left-hand one shifts q in direction $-u$, and the right-hand one shifts q' in direction u .) Note that the support of every square in $r(e^2)$ (in the second case, say) is contained in $\{q\} \cup q' + \text{supp}(\lambda)$, and hence has diameter strictly less than $\text{diam}(e^2) = \rho$, for ρ_0 big enough (see Figure 7.2).

7.2.

REMARK 7.1. At this stage we should note that we have shown that G is finitely presented as we constructed a map $r : X^2 \rightarrow X^2$ with the properties stated at the beginning of Section 7.1.




FIGURE 7.3. *Drum.*

7.3.

It remains to extend the map $r : X^2 \rightarrow X^2$ to the 3-cells of X . This is the core of the proof. Define r to be the identity on 3-cells of diameter less than or equal to ρ_0 . Assume inductively that r has been defined on $X_{\rho'}$, $\rho' = \rho_{n-1} > \rho_0$, such that $r(X_{\rho_i}) \subseteq X_{\rho_{i-1}}$, $i = 1, \dots, n-1$. By construction, it is clear that if $\rho = \rho_n = \text{diam}(e^3)$, then $r(\partial(e^3))$ is contained in $X_{\rho'}$. To extend r to e^3 , we have to show that $r(\partial(e^3))$ represents the trivial element of $\pi_2(X_{\rho'})$.

As only prisms and cubes can have diameter larger than ρ_0 , it suffices to consider those. We start with a prism $e_{(b,a)}^3$. If $\partial(e_b^2) = a_1 \dots a_n$ and $\text{supp}(a) = \{q\}$, $\text{supp}(a_i) = \{q_i\}$, then $\text{supp}(e_{(b,a)}^3) = \{q, q_1, \dots, q_n\}$. Let $\text{diam}(e_{(b,a)}^3) = \rho = \rho_n > \rho_0$. If ρ_0 is large, the distance between the q_i s is small compared with the distance from q (Figure 4.1 shows the prism and its support). Let $u \in S^{m-1}$ be the unit vector in direction $q - \sum_{i=1}^n q_i$. We identify three cases.

Case (a) (drums): Suppose that each edge a_i of the boundary of e_b^2 has a centralizer $\lambda_i \in \Lambda$ that shifts in direction $u_i = (q - q_i)/|q - q_i|$. In this case the diagram $r(\partial(e_{(b,a)}^3))$ is of the form shown in Figure 7.3, which we refer to as a *drum*. It is clear for elementary geometric reasons that, if only ρ_0 is sufficiently large, then each of the λ_i will also shift in direction u . Moreover, the sets $\text{supp}(q\lambda_i)$ are well inside the ball of radius ρ containing $\text{supp}(e_{(b,a)}^3)$. Using Corollary 6.2, we replace the top and bottom collar diagram by the remaining part of the boundary of its pushing 3-cell in direction u . We end up with a new spherical diagram with diameter smaller than or equal to $\rho' < \rho$ that contains non-square cells in oppositely oriented pairs. By Lemma 4.1, this diagram represents the trivial element of $\pi_2(X_{\rho'})$.

Case (b) (special prisms): We call the prism $e_{(b,a)}^3$ *special* if e_b^2 is of the form $e_{(\lambda,a')}^2$ for some $\lambda \in \Lambda$ and $a' \in A$, and $\chi_u(\lambda) > 0$ for $u = (q - q')/|q - q'|$, where $\{q'\}$ is the support of the edge a' . Figure 7.4 shows a special prism and its support.

It is important to note that only q and q' lie on the boundary of the ball of diameter ρ containing $\text{supp}(e_{(b,a)}^3)$; all other support points of this 3-cell are well inside this ball (for ρ_0 large enough). Furthermore, note that there are exactly two




FIGURE 7.4. Special prism (left) and its support (right).




FIGURE 7.5. Image of boundary of special prism under r .

edges in the boundary of the special prism labeled with a' , and these are the only edge labels with support q' .

Again we consider the spherical diagram $S = r(\partial(e^3_{(b,a)}))$. Because we have considered drums before, we can assume that it is not a drum. Note that the support of S satisfies the same properties as the support of the special prism: it lies well inside the ball of diameter ρ except for the external points q and q' and except that a' is the only edge label with support q' (see Figure 7.5).

We have $\max\{\text{diam}(\text{supp}(S) - \{q\}), \text{diam}(\text{supp}(S) - \{q'\})\} \leq \rho' < \rho$. Note that S is a spherical diagram over $X_{\rho'}$ of diameter ρ . We will show that S represents the trivial element of $\pi_2(X_{\rho'})$.

S not being a drum implies that the image $r([a_i, a])$ of one of the squares in $\partial(e^3_{(b,a)})$ contains a pushing 2-cell $e^2_{(\lambda_i, a)}$ so that $\chi_{-u}(\lambda_i) \geq \epsilon$.

Suppose first that $[a', a]$ is such a square. Then S contains a subdiagram that is part of the boundary of a stack of prisms. The situation is shown in Figure 7.6. The boundary of the subdiagram in Figure 7.6 is heavier. The stack of prisms that we are talking about is the product of the bottom 2-cell $e^2_{(\lambda, a')}$ with the edge path u running from the bottom 2-cell to the top 2-cell in S . The support of this stack is contained in the support of S , but it does not contain q (because the label a does not occur in the stack). Thus every prism in the stack has diameter smaller than or




FIGURE 7.6. Subdiagram and its boundary.




FIGURE 7.7. Subdiagram with boundary.

equal to ρ' . We can replace the subdiagram by the remaining boundary in the stack of prisms (consisting of squares) to obtain a spherical diagram S' that no longer contains a' as an edge label. The support of S' is contained in $\text{supp}(S) - \{q'\}$ and hence $\text{diam}(S') \leq \rho'$. Thus S' is a spherical diagram over $X_{\rho'}$ of diameter smaller than or equal to ρ' that contains e_b^2 -faces only in oppositely oriented pairs and hence represents the trivial element of $\pi_2(X_{\rho'})$ by Lemma 4.1. As S and S' represent the same element of $\pi_2(X_{\rho'})$ (we altered S only by boundaries of prisms in $X_{\rho'}$), we can see that S too represents the trivial element of $\pi_2(X_{\rho'})$.

The other case where $r([a, a'])$ is not of the form just considered can be treated similarly. In this case, S contains a subdiagram that is contained in a sum of two stacks of prisms. (See Figure 7.7, where the boundary of the subdiagram is heavier.) One stack is the product of $e_{(\lambda, a')}^2$ and the edge path u , and the other one is the product of some other pushing 2-cell $e_{(\lambda', a')}^2$ and u . The support of both stacks is contained in $\text{supp}(S) - \{q\}$ (neither stack has a as an edge label). We can replace the subdiagram by the remaining boundary in the sum of stacks (consisting of squares) to obtain a spherical diagram S' that no longer contains a' as an edge label and has diameter of at most ρ' . As before, S' represents the trivial element of $\pi_2(X_{\rho'})$ by Lemma 4.1, and, as S and S' are equivalent modulo boundaries of prisms in $X_{\rho'}$, so does S .




FIGURE 7.8. *Non-drum case.*




FIGURE 7.9. *Surrounding a prism with special prisms.*

We can now extend r to special prisms $e^3_{(b,a)}$ of diameter ρ such that $r(e^2_{(b,a)}) \subseteq X_{\rho'}$.

Case (c) (non-drums): Assume that $r(\partial(e^3_{(b,a)}))$ is not a drum (see Figure 7.8). The image of one of the squares $[a_j, a]$ under r then contains a pushing 2-cell $e^2_{(\lambda, a)}$ with $\chi_{-u_j}(\lambda) \geq \epsilon$ for $u_j = (q - q_j)/|q - q_j| \in S^{m-1}$. For all $i = 1, \dots, n$, we then consider the special prisms $e^3_{(b', a_i)}$, where $e^2_{b'}$ is the pushing 2-cell $e^2_{(\lambda, a)}$. Note that each of these special prisms has diameter of at most ρ .

We can surround $e^3_{(b,a)}$ with the special prisms $e^3_{(b', a_i)}$, as shown in Figure 7.9. The boundary of this union of 3-cells is the boundary of a stack of prisms $e^2_{(b, a'_k)}$, each of which has support well inside the ball of diameter ρ containing $\text{supp}(e^3_{(b,a)})$. Hence each of the prisms in the stack is of diameter at most ρ' .

Thus we have shown the boundary of the prism $e^3_{(b,a)}$ as the boundary of a union of 3-cells consisting of special prisms of diameter ρ and prisms of diameter ρ' . This shows that $\partial(e^3_{(b,a)})$ defines the trivial element of the second homotopy group of the subcomplex U of X consisting of the 2-skeleton of X together with $X_{\rho'}$ and special prisms of diameter at most ρ . Therefore we can define a map $e^3_{(b,a)} \rightarrow U$ that is the identity on the boundary. On U , the map $r : U \rightarrow X_{\rho'}$ has been defined already (by (a) and (b) above and induction hypothesis), and the composition

$$e^3_{(b,a)} \rightarrow U \xrightarrow{r} X_{\rho'}$$

in the desired extension of r to the prism $e^3_{(b,a)}$.




FIGURE 7.10. The tinted region is $q_1 + \text{supp}(\lambda)$.




FIGURE 7.11. Surrounding cube by prisms.

7.4.

So far we have extended r to prisms of diameter ρ . We now turn to cubes. Let $e^3_{(a_1, a_2, a_3)}$ be a cube with edges a_1, a_2, a_3 , and let $\{q_i\}$ be the support of a_i , $i = 1, 2, 3$. We assume that $\text{diam}(e^3_{(a_1, a_2, a_3)}) = \text{diam}(\{q_1, q_2, q_3\}) = \rho > \rho_0$. Let z be the center of the ball B of diameter ρ that contains $\text{supp}(e^3_{(a_1, a_2, a_3)})$, and define $u_i = (z - q_i)/|z - q_i| \in S^{m-1}$.

It is clear that the 3-point set $\{u_1, u_2, u_3\}$ is not contained in an open hemisphere (otherwise the ball B could be shrunk), and so there is a centralizer $\lambda \in \Lambda$ such that $\langle u_j, \text{supp}(\lambda) \rangle > \epsilon$ for some u_j , say u_1 . (This is the first time we have used the fact that M is 3-tame; the discussion on prisms only required 2-tameness.) In the case in which the set $\{u_1, u_2, u_3\}$ contains antipodal points, we can assume that u_1 and u_2 are antipodal. Thus u_2 and u_3 are not antipodal, and hence the distance between q_2 and q_3 is at most $\rho' < \rho$. (See Figure 7.10.)

We argue very much as in case (c) (non-drum prisms). First surround $e^3_{(a_1, a_2, a_3)}$ with prisms $e^3_{(b, a_i)}$, $i \in \{2, 3\}$, where $e^2_b = e^2_{(\lambda, a_1)}$ (see Figure 7.11). Note that the support of each of these prisms is contained in $\{q_1, q_2, q_3\} \cup q_1 + \text{supp}(\lambda)$ and hence

is of diameter at most ρ . The boundary of this union of 3-cells is the boundary of a stack of cubes $e_{(a_2, a_3, a'_i)}^3$. The diameter of each such cube is contained in $\{q_2, q_3\} \cup q_1 + \text{supp}(\lambda)$, and hence is of diameter at most $\rho' < \rho$.

Thus we have shown the boundary of the cube $e_{(a_1, a_2, a_3)}^3$ as the boundary of a union of 3-cells consisting of prisms of diameter at most ρ and cubes of diameter at most ρ' . As r is already defined on these prisms by our discussion of cases (a), (b) and (c), and on these cubes by induction hypothesis, we can extend r to our given cube $e_{(a_1, a_2, a_3)}^3$ by the same arguments as were used at the end of the discussion of case (c).

By what was said at the beginning of Section 7.1, this concludes the proof of the Main Theorem.

References

1. H. ÅBERG, 'Bieri–Strebel valuations (of finite rank)', *Proc. London Math. Soc.* (3) 52 (1986) 269–304.
2. R. BIERI and J. R. J. GROVES, 'Metabelian groups of type FP_∞ are virtually of type FP', *Proc. London Math. Soc.* (3) 45 (1982) 365–384.
3. R. BIERI and J. R. J. GROVES, 'The geometry of the set of characters induced by valuations', *J. Reine Angew. Math.* 347 (1984) 168–195.
4. R. BIERI and J. HARLANDER, 'A remark on the polyhedrality theorem for the Σ -invariants of modules over abelian groups', *Math. Proc. Cambridge Philos. Soc.* 131, 39–43.
5. R. BIERI and R. STREBEL, 'Valuations and finitely presented metabelian groups', *Proc. London Math. Soc.* (3) 41 (1980) 439–464.
6. W. A. BOGLEY and S. J. PRIDE, *Two-dimensional homotopy and combinatorial group theory*, London Mathematical Society Lecture Note Series 197 (ed. C. Hog-Angeloni, A. Sieradski and W. Metzler, Cambridge University Press, 1993) (Chapter 5).
7. K. S. BROWN, 'Finiteness properties of groups', *J. Pure Appl. Algebra* 44 (1987) 45–75.
8. K. U. BUX, 'Finiteness properties of certain metabelian arithmetic groups in the function field case', *Proc. London Math. Soc.* 75 (1997) 308–322.
9. D. J. COLLINS and J. HUEBSCHMANN, 'Spherical diagrams and identities among relations', *Math. Ann.* 261 (1982) 155–183.
10. R. C. LYNDON and P. E. SCHUPP, *Combinatorial group theory* (Springer).
11. D. H. KOCHLOUKOVA, 'The FP_m -conjecture for a class of metabelian groups', *J. Algebra* 184 (1996) 1175–1204.
12. D. H. KOCHLOUKOVA, 'The FP_m -conjecture for a class of metabelian groups and related topics', PhD Dissertation, University of Cambridge, 1997.
13. G. A. NOSKOV, 'Bieri–Strebel invariant and homological finiteness properties of metabelian groups', SFB-Preprint 93-028, Universität Bielefeld, 1993.

*Fachbereich Mathematik (Fach 187),
Johann Wolfgang Goethe-Universität
F-60054 Frankfurt-am-Main
Germany*

bieri@math.uni-frankfurt.de