The Bestvina-Brady Construction Revisited – Geometric Computation of Σ-Invariants for Right Angled Artin Groups

Kai-Uwe Bux   Carlos Gonzalez

Revised 16th October 1999

The starting point of our investigation is the remarkable paper [**2**] where M. Bestvina and N. Brady gave an example of an infinitely related group of type FP_2. The result about right angled Artin groups behind their example is best intepreted by means of the Bieri-Strebel-Neumann-Renz Σ-invariants.

For a group G the invariants $\Sigma^n(G)$ and $\Sigma^n(G,\mathbb{Z})$ are sets of non-trivial homomorphisms $\chi : G \to \mathbb{R}$. They contain full information about finiteness properties of subgroups of G with abelian factor group. The main result of [**2**] determines for the canonical homomorphism χ, taking each generator of the right angled Artin group G to 1, the maximal n with $\chi \in \Sigma^n(G)$ respectively $\chi \in \Sigma^n(G,\mathbb{Z})$.

In [**6**] J. Meier, H. Meinert and L. VanWyk completed the picture by computing the full Σ-invariants of right angled Artin groups using as well the result of Bestvina and Brady as algebraic techniques from Σ-theory. Here we offer a new account on their result which is totally geometric. In fact, we return to the

1991 Mathematics Subject Classification: 20F32, 20F36.
Key words and phrases: Artin Groups, Graph Groups, Geometric Group Theory, Finiteness Properties.

Bestvina-Brady construction and simplify their argument considerably by bringing a more general notion of links into play. At the end of the first section we reprove their main result. By recomputing the full Σ-invariants, we show in the second section that the simplification even adds some power to the method. The criterion we give provides new insight on the geometric nature of the "n-domination" condition employed in [**6**].

The invariants $\Sigma^n(G)$ respectively $\Sigma^n(G,\mathbb{Z})$ are defined if G is a group of *type* F_n (i.e. G admits a $\mathrm{K}(G,1)$-complex with finite n-skeleton) respectively of *type* FP_n (i.e. the trivial $\mathbb{Z}G$-module $\mathbb{Z}$ admits a projective resolution which is finitely generated up to dimension n).

We will not give the most general definition of Σ here. So assume G is of type F_n and take X to be the universal cover of a $\mathrm{K}(G,1)$-complex with finite n-skeleton. Let G act non-trivially on $\mathbb{R}$ by translations and let $\chi : G \to \mathbb{R}$ be a homomorphism describing the action. Then there always exists a G-equivariant continuous map $h : X \to \mathbb{R}$ which we call a *height function*. This induces a filtration of X by closed subspaces $h^{-1}([r,\infty))$. By definition $\chi \in \Sigma^n(G,\mathbb{Z})$ (respectively $\chi \in \Sigma^n(G)$) if and only if the directed system of reduced homology groups $\left(\widetilde{\mathrm{H}}_i(h^{-1}([r,\infty)))\right)_{r\in\mathbb{R}}$ is essentially trivial for all $i < n$ (respectively if the directed system of homotopy groups $(\pi_i(h^{-1}([r,\infty))))_{r\in\mathbb{R}}$ is essentially trivial for all $i < n$). Here we follow Brown (cf. [**5**]) and call a system of groups $(H_r)_{r\in\mathbb{R}}$ *essentially trivial* if for every $a \in \mathbb{R}$ there is an element $b \leq a$ such that $H_a \to H_b$ is the trivial morphism. It is a crucial but non-obvious fact that the definition above does not depend on the choices of X and h. A survey is given in [**3**].

Let

- L be a finite *flag complex* (i.e. a simplicial complex wherein a collection of vertices forms a simplex if and only if they are pairwise joined by an edge) and
- $G = \left\langle L^{(0)} \mid vv' = v'v \text{ if } \{v,v'\} \in L^{(1)} \right\rangle$ be the corresponding Artin group generated by all vertices in $L^{(0)}$ subjected to the

relations that generators commute if they are joined by an edge in $L^{(1)}$.

Further let us fix

- $\chi : G \to \mathbb{R}$, a non-trivial homomorphism describing an action of G on $\mathbb{R}$ by translations,
- $L^\dagger \leq L$, the full subcomplex spanned by all vertices $v \in L^{(0)}$ with $\chi(v) = 0$ (*dead* vertices), and
- $L^* \leq L$, the full subcomplex spanned by the remaining (*living*) vertices.

Call a complex n-*acyclic* ($n \in \mathbb{Z}$) if its reduced homology groups vanish in dimensions $\leq n$. Note that for $n \geq -1$ this implies that the complex is non-empty. We adopt the corresponding convention for n-*connectivity*.

We aim to prove:

Theorem. *The following are equivalent:*

1. $\chi \in \Sigma^{n+1}(G, \mathbb{Z})$ *respectively* $\chi \in \Sigma^{n+1}(G)$.
2. *For every possibly empty dead simplex $\sigma \in L^\dagger$ the living link $\mathrm{Lk}_{L^*}(\sigma) := L^* \cap \mathrm{Lk}_L(\sigma)$ is $(n - \dim(\sigma) - 1)$-acyclic respectively L^* is, additionally, n-connected.*

Recall that L is the link of its empty simplex.

1 Links

We aim at a criterion for $\chi \in \Sigma^n(G)$ formulated entirely in terms of the complex L. The first step, however, will be to reduce our task to determine the acyclicity respectively connectivity length of a certain infinite complex.

For every vertex $v \in L^{(0)}$, let

- $\mathbb{S}^1_v$ denote a circle with base point and $\mathbb{S}^0_v := \{v_+, v_-\}$ the Link of the base point in $\mathbb{S}^1_v$. For each simplex $\sigma \in L$ there is a corresponding torus

- $\mathbb{T}_\sigma := \times_{v\in\sigma} \mathbb{S}^1_v \subseteq \times_{v\in L^{(0)}} \mathbb{S}^1_v$. Set

- $\mathbb{T}_L := \bigcup_{\sigma\in L} \mathbb{T}_\sigma \subseteq \times_{v\in L^{(0)}} \mathbb{S}^1_v$ where we glue the tori along the obvious identifications. Set

- $\mathbb{S}_L := \bigcup_{\sigma\in L} \mathbb{S}_\sigma$ where the spheres $\mathbb{S}_\sigma := \ast_{v\in\sigma} \mathbb{S}^0_v$ are glued together analogously. $\mathbb{S}_L$ is clearly the link of the vertex in $\mathbb{T}_L$. Note that the map $v \mapsto v_+$ induces an embedding $L \hookrightarrow \mathbb{S}_L$ and that the map $v_+, v_- \mapsto v$ leads to a retraction of $\mathbb{S}_L$ onto L.

  We pass to

- X_L, the universal cover of $\mathbb{T}_L$. Regarding each torus $\mathbb{T}_\sigma$ as a cube with opposite faces identified we think of X_L as a cube complex.

The link $\mathbb{S}_L$ is a CAT(1) "all right spherical complex" which is shown for instance in [**2**, Lemmas 5.6 and 5.8]. This implies

Proposition 1. *$\mathbb{T}_L$ is a locally CAT(0) geodesic space and X_L is a globally CAT(0) geodesic cube complex. In particular, $\mathbb{T}_L$ is a $\mathrm{K}(G,1)$-complex.* **q.e.d.**

Note that each cube is indexed by a simplex in L, its *type*. Vertices correspond to the empty simplex. An embedding $L' \leq L$ of a subcomplex induces an embedding $\mathbb{T}_{L'} \leq \mathbb{T}_L$. Choosing corresponding basepoints we obtain an embedding $X_{L'} \leq X_L$ where $X_{L'}$ is a component of the preimage of $\mathbb{T}_{L'}$ in X_L. For a simplex $\sigma \in L$ the corresponding subcomplexes X_σ are called *sheets*. They are components of the preimage of the torus $\mathbb{T}_\sigma$.

Lemma 2. *If L' is a full subcomplex of L then $\mathbb{T}_{L'}$ is a locally convex subspace of $\mathbb{T}_L$ and $X_{L'}$ is a convex subspace of X_L (i.e. geodesic segments between two points of $X_{L'}$ within X_L do not leave $X_{L'}$).*

In particular any sheet is isometric to a Euclidean space.

Proof. The first claim is obvious and implies the second in virtue of the Theorem of Cartan-Hadamard. Take two points in $X_{L'}$ and a path in $X_{L'}$ connecting them. Consider its image in $\mathbb{T}_{L'}$. By

Cartan-Hadamard [**1**, Theorem 13] there is a unique locally geodesic curve in $\mathbb{T}_{L'}$ representing its relative homotopy class. Since $\mathbb{T}_{L'}$ is locally convex in $\mathbb{T}_L$ this curve is a locally geodesic curve in $\mathbb{T}_L$. Hence its lift is the geodesic path joining the two points and it does not leave $X_{L'}$ by construction. **q.e.d.**

The main idea is to assign to a convex subspace a natural link that should behave like a "normal unit bundle". Let Y be a closed convex subspace of X, a globally CAT(0) geodesic space with proper metric (i.e. closed balls are compact). Then for each point in X there is exactly one point in Y with minimal distance. This induces an *orthogonal projection* $o_Y^X : X \to Y$ and we define the *link* of Y to be the "bundle of directions within fibres". To be more precise, let

- $\mathrm{Lk}_X(Y) := \{[\gamma] \mid \gamma : [0,r] \to X$ geodesic, $r > 0$, $o_Y^X(\gamma) = \gamma(0) \in Y\}$ be the set of all geodesics "orthonormal to Y" where those geodesics are identified that have an initial segment in common defining this way the same germ. This is a bundle over Y; the *bundle projection*

- $\pi_Y^X : \mathrm{Lk}_X(Y) \to Y$ maps each geodesic $[\gamma]$ to its initial point $\gamma(0)$.

  For each point $x \in X \setminus Y$ there is a unique shortest geodesic segment γ_x connecting Y with x. This geodesic segment represents an element $[\gamma_x] \in \mathrm{Lk}_X(Y)$. Let us call the map

- $\widetilde{o}_Y^X : X \setminus Y \to \mathrm{Lk}_X(Y)$, $x \mapsto [\gamma_x]$ the *geodesic projection*. Obviously $o_Y^X |_{X\setminus Y} = \pi_Y^X \circ \widetilde{o}_Y^X$.

Let us return to the complex X_L. The homomorphism χ induces a G-equivariant height function

- $h : X_L \to \mathbb{R}$ which is given by $X_L \to \mathbb{R}^{L^{(0)}} := \bigoplus_{v\in L^{(0)}} \mathbb{R}v \to \mathbb{R}$ where the first arrow is a lift of $\mathbb{T}_L \hookrightarrow \times_{v\in L^{(0)}} \mathbb{S}^1_v$ and the second arrow is the linear map induced by $v \mapsto \chi(v)$. Note that $h |_{X_\sigma}$ is an affine map for each sheet X_σ.

The level set at height r is the preimage $h^{-1}(r)$. A level set $h^{-1}(r)$ cuts every cube containing an interior point of its height r into two pieces.

For a vertex $x \in X_L$ consider the level set at height $h(x)$. The union of all cubes in X_L contained in the level set forms a subcomplex of X_L whose components are precisely of the form $X_{L^\dagger}$. Call these components level complexes.

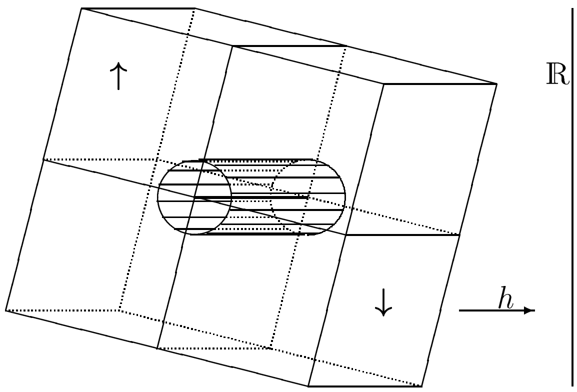

Figure 1

Since we know that level complexes are convex subcomplexes we can consider their links. So consider a geodesic γ representing an element $[\gamma] \in \mathrm{Lk}_{X_L}(X_{L^\dagger})$. There is a unique minimal cube c carrying the initial segment of γ and this cube only depends on $[\gamma]$. Three cases can occur. First, c might be cut into two pieces by the level set containing the level complex $X_{L^\dagger}$. Secondly c might meet the level set at a bottom face. In this case we say that $[\gamma]$ belongs to the ascending link

- $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$.

Finally c might meet the level set at a top face; this defines the descending link

- $\mathrm{Lk}^{\downarrow}_{X_L}(X_{L^\dagger})$. See Figure 1.

Observation 3. *The cubical structure of X_L induces a cellular structure on $\mathrm{Lk}_{X_L}(X_{L^\dagger})$, $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ and $\mathrm{Lk}^{\downarrow}_{X_L}(X_{L^\dagger})$. Note that two points in $X_{L^\dagger}$ that belong to the same open cube have isomorphic fibres in $\mathrm{Lk}_{X_L}(X_{L^\dagger})$ the isomorphism being given by parallelism of geodesic segments. Hence $\mathrm{Lk}_{X_L}(X_{L^\dagger})$, $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$, and $\mathrm{Lk}^{\downarrow}_{X_L}(X_{L^\dagger})$ have a natural product structure over open cubes. The fibres are given in the following way. Let $\sigma \in L^\dagger$ be the type of an open cube in $X_{L^\dagger}$. Then the fibre over any point in this cube is $\mathbb{S}_{\mathrm{Lk}_{L^*}(\sigma)}$. The fibre of* *the ascending link and the descending link are just the living links* $\mathrm{Lk}_{L^*}(\sigma)$. *Hence* $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ *is a retract of* $\mathrm{Lk}_{X_L}(X_{L^\dagger})$. **q.e.d.**

Observation 4. *Given a geodesic segment* γ *contained in a closed cube* c *of* X_L *we can naturally extend* γ *to an infinite geodesic because* c *lies in a sheet which is isometric to a Euclidean space.*

Hence we can represent elements of $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ *by infinite geodesic rays starting in* $X_{L^\dagger}$. *Along such a ray, the height* h *increases to infinity. Considered this way,* $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ *"intersects" each level set above* $h(X_{L^\dagger})$. **q.e.d.**

Corollary 5. *If the homotopy group* $\pi_n(\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger}))$ *is non-trivial then for all real numbers* $a \leq b \leq c$ *the inclusion* $h^{-1}([b,c]) \hookrightarrow h^{-1}([a,c])$ *induces a non-trivial map in homotopy* π_n. *The analogue for homology groups is also true.*

Proof. We will treat the homotopical statement only. But the same proof aplies in the homological case. Pick a level complex $X_{L^\dagger}$ below a and an n-sphere in its ascending link $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ representing a non-trivial element of $\pi_n(\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger}))$. Regarding the points in $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ as infinite geodesic rays as in Observation 4, this sphere maps to the level set at height b becoming a "high sphere" in this way.

Now observe that the geodesic projection $X_L \setminus X_{L^\dagger} \to \mathrm{Lk}_{X_L}(X_{L^\dagger})$ maps the high sphere to the sphere we started with. Hence the high sphere cannot be contracted in $X_L \setminus X_{L^\dagger}$ whence it cannot be contracted in $h^{-1}([a,c])$ provided that the sphere in $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ cannot be contracted in $\mathrm{Lk}_{X_L}(X_{L^\dagger})$. However that follows from the fact that $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ is a retract of $\mathrm{Lk}_{X_L}(X_{L^\dagger})$. **q.e.d.**

We now generalize the Morse argument of [**2**] to level complexes. Since we do not confine ourselfes to discrete homomorphisms the key observation is slightly more involved but still obvious.

Observation 6. *Given a convex polyhedron* C *and a codimension-*1*-face* F, *let* $\tilde{F}$ *be the union of all faces that do not meet* F. *Then every* *deformation retraction of $\partial C \setminus F$ to $\tilde{F}$ extends to a deformation retraction of C to $\tilde{F}$.* **q.e.d.**

Let $a < b$ be real numbers such that $b - a < |\chi(v)|$ for all living vertices. Note that no two distinct level complexes in $h^{-1}([a,b])$ are joined by an edge in $X_L^{(1)}$. Hence the ascending links of these level complexes may be identified with disjoint subspaces of the level set at height b. Here, again, we use Observation 4.

Lemma 7 (Morse-Lemma). *Let $a < b < c$ be real numbers such that $b - a < |\chi(v)|$ for all living vertices. For each level complex $X_{L^\dagger}$ with $h(X_{L^\dagger}) \in [a,b)$, let M be the mapping cylinder of the bundle projection $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger}) \to X_{L^\dagger}$ where we regard $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ as a subspace of the level set at height b. Let $\widetilde{M}$ be the union of all mapping cylinders M obtained that way. Then $\widetilde{M} \cup h^{-1}([b,c])$ is a deformation retract of $h^{-1}([a,c])$.*

In particular, since all level complexes are contractable, the preimage $h^{-1}([a,c])$ is homotopy equivalent to $h^{-1}([b,c])$ with the ascending links coned off.

Proof. We can regard $\widetilde{M} \cup h^{-1}([b,c])$ as a subspace of X_L. It is nothing but the subcomplex of all cubes above height a cut at height c. So we have to get rid of all cubes intersecting the level set at height a. A deformation from $h^{-1}([a,c])$ to $\widetilde{M} \cup h^{-1}([b,c])$ can thus be constructed inductively on skeleta using Observation 6 — just take F to be the intersection of a given cube with the level set at height a. **q.e.d.**

Of course there is an analogue concerning descending links.

Corollary 8. *If all ascending and descending links of level complexes in X_L are n-connected then for every closed interval J the preimage $h^{-1}(J)$ is n-connected. The analogous statement for n-acyclicity is also true.*

Proof. Using the Morse-Lemma we observe first that for two closed intervals $I \subseteq J$ the inclusion $h^{-1}(I) \hookrightarrow h^{-1}(J)$ induces isomorphisms in the homotopy respectively reduced homology groups in dimensions $\leq n$. Now the claim follows from the fact that X_L is contractible and is the direct limit of the subspaces $h^{-1}(J)$. **q.e.d.**

Observing that ascending respectively descending links are all isomorphic we may rephrase the results obtained so far in the following

Theorem 9. *For any non-trivial homomorphism χ of G the following are equivalent:*

1. $\chi \in \Sigma^{n+1}(G, \mathbb{Z})$ *respectively* $\chi \in \Sigma^{n+1}(G)$.
2. $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ *is n-acyclic respectively n-connected.* **q.e.d.**

Remark 10. It is well known (cf. [**4**, Theorem B]) that the kernel of a discrete homomorphism χ is of type FP_n if and only if $\{\chi, -\chi\} \subseteq \Sigma^{n+1}(G, \mathbb{Z})$. The homotopical counterpart is also true. However, in the setting of Artin groups one does not need the general result. Observe that for discrete χ the kernel $\ker\chi$ acts cocompactly on $h^{-1}(J)$ for any closed compact interval J. Furthermore, Corollary 8 implies that the system of homotopy respectively reduced homology groups is essentially trivial up to to the required dimension. Corollary 5 implies that it is not essentially trivial in higher dimensions. Now Brown's criteria (cf. [**5**, Theorems 2.2 and 3.2]) apply. So we conclude that $\chi \in \Sigma^{n+1}(G, \mathbb{Z})$ respectively $\chi \in \Sigma^{n+1}(G)$ if and only if $\ker\chi$ is of type FP_n respectively F_n.

Remark 11. If $L^\dagger = \emptyset$ as in [**2**] a level complex $X_{L^\dagger}$ is simply a vertex. Then we have $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger}) \cong L$ whence Theorem 9 immediately implies the main result of [**2**].

2 Relating links to L and $L^\dagger$

Theorem 9 provides a criterion for $\chi \in \Sigma^n(G)$ in terms of the ascending link $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$. So we have to relate $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ to the pair $(L, L^\dagger)$.

Lemma 12. $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ *is* n*-acyclic/connected if and only if for each closed cube in* $X_{L^\dagger}$ *its preimage in* $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ *under the bundle projection is* n*-acyclic/connected.*

Proof. Assume first that all cube preimages are n-acyclic. Since they cover $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ we can compute its homology using the spectral sequence for the homology of a union. Technically, it is slightly more convenient to use results of D. Quillen. Note that $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ and $X_{L^\dagger}$ are complexes whose cells are convex polyhedra and whose cells intersect in cells. Hence both are determined by the corresponding posets of closed cells. In fact, their barycentric subdivisions are the geometric realisations of these posets. The bundle projection $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger}) \to X_{L^\dagger}$ induces an order preserving map between the associated posets. Hence the properties of $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ follow from [**7**, Proposition 7.6 and proof], where the π_1-case is discussed as well.

For the converse, note that each closed cube $c \in X_{L^\dagger}$ is a convex subspace of $X_{L^\dagger}$ whence we have a retraction from $X_{L^\dagger}$ to c by orthogonal projection. We only have to show that this retraction lifts to a retraction $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger}) \to \tilde{c}$ where $\tilde{c}$ is the preimage of c. So take a geodesic γ connecting a point $\gamma(0) \in X_{L^\dagger}$ with $\gamma(r) = o_c^{X_{L^\dagger}}(\gamma(0)) \in c$. The only problem we may encounter is that the fibre in $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ over $\gamma(0)$ may not be a subcomplex of the fibre over its endpoint $\gamma(r)$. But this problem does not occur as follows from the following argument of Bestvina and Brady.

Consider the minimal open face of c containing the endpoint $\gamma(r)$ and let σ be its type. Then the geodesic ends perpendicular to that face coming from an open cube indexed by a simplex $\tau \supseteq \sigma$. Induction shows that all cubes met by γ are labeled by supersimplices of σ. Recall that by Observation 3 the fibre over an open cube of type τ is $\mathrm{Lk}_{L^*}(\tau)$. Hence the fibre over $\gamma(0)$ naturally embeds in the fibre $\mathrm{Lk}_{L^*}(\sigma)$ over $\gamma(r)$.

So we can define a retraction from $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ to $\tilde{c}$ whence the homotopy and reduced homology groups of $\tilde{c}$ embed into the homotopy and reduced homology groups of $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$. **q.e.d.**

The final step is to reformulate the criterion in terms of the pair $\left(L, L^{\dagger}\right)$.

Lemma 13. *The following are equivalent:*

1. *For every closed cube in $X_{L^\dagger}$, its preimage in $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ under the bundle projection is n-acyclic respectively n-connected.*
2. *For every possibly empty dead simplex $\sigma \in L^{\dagger}$, the living link $\mathrm{Lk}_{L^*}(\sigma)$ is $(n-\dim(\sigma)-1)$-acyclic respectively L^* is, additionally, n-connected.*

Proof. Recall the structure of $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ exhibited in Observation 3.

We will consider a family of spaces $\tilde{c}_{\sigma,\tau}$ where σ and τ are disjoint possibly empty dead simplices such that $\sigma \cup \tau$ is a dead simplex. Take a cube c of type $\sigma \cup \tau$. Let c_σ and c_τ be faces of the corresponding type. Take $\tilde{c}_{\sigma,\tau}$ to be the preimage in the ascending link of a c_σ-translate $c_{\sigma,\tau}$ in c meeting c_τ in its centre. See Figure 2. Note that $\tilde{c}_{\sigma,\tau}$ does not depend on the choices up to isomorphy. Further check that if v is a dead vertex outside $\sigma \cup \tau$ such that $\sigma \cup \tau \cup \{v\}$ is a dead simplex we have $\tilde{c}_{\sigma\cup\{v\},\tau} = \tilde{c}_{\sigma,\tau\cup\{v\}} \times [0,1] \cup \tilde{c}_{\sigma,\tau} \times \{0,1\}$. Here one should think of $[0,1]$ as an edge of type v.

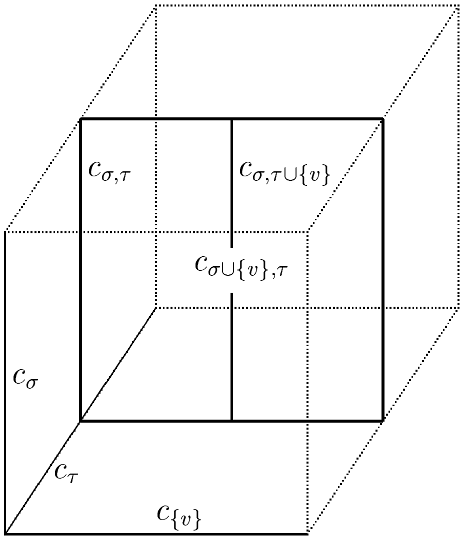

Figure 2

Hence we have the following homotopy equivalence

$$\tilde{c}_{\sigma\cup\{v\},\tau} \equiv \tilde{c}_{\sigma,\tau} \cup_{\tilde{c}_{\sigma,\tau\cup\{v\}}} \tilde{c}_{\sigma,\tau}$$

where $X \cup_Y X$ denotes the union of two copies of X identified along the comon subspace $Y \subseteq X$. Note that the Mayer-Vietoris sequence implies

> If X is m-acyclic then $X \cup_Y X$ is m-acyclic if and only if Y is $(m-1)$-acyclic.

Using this we can treat the homological part. Since we have $\tilde{c}_{\emptyset,\tau} = \mathrm{Lk}_{L^*}(\tau)$ for all dead simplices τ we can use induction on $\dim(\sigma)$ to infer from condition 2 that $\tilde{c}_{\sigma,\tau}$ is $(n - \dim(\tau) - 1)$-acyclic. Hence condition 1 holds because $c_{\sigma,\emptyset}$ is nothing but a cube of type σ.

On the other hand assume there is a dead simplex $\sigma \in L^\dagger$ such that $\mathrm{Lk}_{L^*}(\sigma)$ is not $(n - \dim(\sigma) - 1)$-acyclic. Then we may assume its dimension to be minimal. Now we can use induction to show that the preimage of a cube of type σ is not n-acyclic.

Finally we have to deal with the homotopical part. Assuming that preimages in $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ of closed cubes in $X_{L^\dagger}$ are n-connected, we obtain that L^* is n-connected because it is the fibre over 0-dimensional cubes. On the other hand we can combine the induction above with the Seifert-van Kampen-Theorem to obtain the converse. **q.e.d.**

Let us summarize.

Theorem 14. *For any non-trivial homomorphism $\chi : G \to \mathbb{R}$ the following are equivalent:*

1. $\chi \in \Sigma^{n+1}(G, \mathbb{Z})$ *respectively* $\chi \in \Sigma^{n+1}(G)$.
2. $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ *is n-acyclic respectively n-connected.*
3. *The preimage of every closed cube in $\mathrm{Lk}^{\uparrow}_{X_L}(X_{L^\dagger})$ under the bundle projection is n-acyclic respectively n-connected.*
4. *For every possibly empty dead simplex $\sigma \in L^\dagger$ the living link $\mathrm{Lk}_{L^*}(\sigma)$ is $(n - \dim(\sigma) - 1)$-acyclic respectively L^* is, additionally, n-connected.*

If the homomorphism χ is discrete then these conditions hold if and only if $\ker\chi$ is of type FP_n respectively F_n.

Proof. Put together Theorem 9, Lemma 12, Lemma 13, and Remark 10. **q.e.d.**

In **[6]** the Σ-invariants are described in terms of a rather involved domination condition. The notion n-domination is defined recursively. We call a subcomplex $L' \leq L$ (-1)-dominating if it is non-empty. We call it n-dominating for $n > -1$ if for every vertex $v \in L \setminus L'$ the *restricted link* $\mathrm{Lk}_{L'}(v) := L' \cap \mathrm{Lk}_L(v)$ is $(n-1)$-acyclic and $(n-1)$-dominating as a subcomplex of the entire link $\mathrm{Lk}_L(v)$. We end our tour with the following lemma showing our conditions to be equivalent to the one employed in [**6**].

Proposition 15. *The following are equivalent:*

1. *$L^* \leq L$ is n-dominating and L^* is n-acyclic.*
2. *For every possibly empty dead simplex σ in $L^\dagger$ the living link $\mathrm{Lk}_{L^*}(\sigma)$ is $(n - \dim(\sigma) - 1)$-acyclic.*

Proof. Consider the following hypotheses:

(dom$_i$) For every possibly empty dead simplex σ in $L^\dagger$ of dimension i the living link $\mathrm{Lk}_{L^*}(\sigma)$ is a $(n-i-1)$-dominating subcomplex of the full link $\mathrm{Lk}_L(\sigma)$.

(ac$_i$) For every possibly empty dead simplex σ in $L^\dagger$ of dimension i the living link $\mathrm{Lk}_{L^*}(\sigma)$ is $(n-i-1)$-acyclic.

Note the following equivalences

- (ac$_n$) $\Longleftrightarrow$ (dom$_n$)
- 1 $\Longleftrightarrow$ (dom$_{-1}$) and (ac$_{-1}$)
- 2 $\Longleftrightarrow$ (ac$_i$) holds for all $i \leq n$

For two dead simplices $\sigma \neq \tau = \sigma \cup \{v\}$, the identities $\mathrm{Lk}_{L^*}(\tau) \cong \mathrm{Lk}_{\mathrm{Lk}_{L^*}(\sigma)}(v)$ and $\mathrm{Lk}_L(\tau) \cong \mathrm{Lk}_{\mathrm{Lk}_L(\sigma)}(v)$ hold since L^* is a full subcomplex of L. Using these, one can prove the following additional equivalence:

- (dom$_i$) $\Longleftrightarrow$ (dom$_{i+1}$) and (ac$_{i+1}$)

From this the claim follows easily. **q.e.d.**

References

1. W. Ballmann, 'Singular spaces of non-positive curvature', *Sur les Grupes Hyperboliques d'après Mikhael Gromov* (ed. E. Ghys and P. de la Harpe, Birkhäuser, Boston, 1990) p. 189 – 201.

2. M. Bestvina, N. Brady, 'Morse Theory and Finiteness Properties of Groups', *Inventiones mathematicae* 129 (1997) p. 445 – 470.

3. R. Bieri, 'Finiteness length and connectivity length for groups', *Geometric Group Theory down under; Proc. of the special year on geometric group theory, Australia Nat. Univ. Canberra, 1996* (ed. J. Cossey, C.F. Miller, W. Neumann and M. Shapiro, de Gruyter, 1999) p. 9 – 22.

4. R. Bieri, B. Renz, 'Valuations on free resolutions and higher geometric invariants of groups', *Commentarii Mathematicae Helvetici* 63 (1988) p. 464 – 497.

5. K.S. Brown, 'Finiteness Properties of Groups', *Journal of Pure and Applied Algebra* 44 (1987) p. 45 – 75.

6. J. Meier, H. Meinert, L. VanWyk, 'Higher Generation Subgroup Sets and the Σ-Invariants of Graph Groups', *Commentarii Mathematici Helvetici* 73 (1998) p. 22 – 44.

7. D. Quillen, 'Homotopy Properties of the Poset of Nontrivial p-Subgroups of a Group', *Advances in Mathematics* 28 (1978) p. 101 – 128.

Kai-Uwe Bux
Department of Mathematics
University of Utah
155 S. 1400 E., RM. 233
Salt Lake City, UT 84112-0090, USA
kubux@gmx.net

Carlos Gonzalez
Fachbereich Mathematik der
Johann Wolfgang Goethe-Universität
D-60054 Frankfurt am Main, Germany
gonzalez@secude.de