LINEAR-DEGREE COVERS OF SALVETTI COMPLEXES WHERE SPECIFIED LOOPS DON'T LIFT

MARK HAGEN

In this note, we give a slightly streamlined proof of [HP16, Theorem A]. The idea is essentially the same, but the construction has been modified slightly to rely less on features specific to Salvetti complexes and avoid technicalities about fundamental domains, in the hope of being able to more easily sharpen the argument for other purposes.

Let X be the Salvetti complex of a right-angled Artin group presented by a finite graph. Let $x_0 \in X$ denote the 0–cube.

Theorem 0.1. *Let Z be a compact special cube complex, let $z \in Z^{(0)}$, and let $f : (Z, z) \to (X, x_0)$ be a based local isometry. Let $A = \pi_1(X, x_0)$, let $H = \pi_1(Z, z)$, and let $B = f_*(H)$ (note that f_* is injective). Suppose that $g \in A - B$. Then there exists a cover $(Y, y) \to (X, x_0)$ such that*

- *$|Y^{(0)}| \leq |Z^{(0)}| \cdot (\|g\| + 1)$, where $\|-\|$ is the word metric on A with respect to the standard generators (in particular, $Y \to X$ has degree at most $|Z^{(0)}| \cdot (\|g\| + 1)$.*
- *Let C be the image of $\pi_1(Y, y)$ in A. Then $B \leq C$ but $g \notin C$.*

Proof. Let $p : (\widetilde{Z}, \tilde{z}) \to (Z, z)$ be the universal cover, so that $f \circ p$ lifts to a map $\tilde{f} : (\widetilde{Z}, \tilde{z}) \to (\widetilde{X}, \tilde{x}_0)$. Since f is a local isometry, $\tilde{f}$ is an embedding whose image is a convex subcomplex of the CAT(0) cube complex $\widetilde{X}$. Let $\widetilde{Z}$ denote the image of $\tilde{f}$. Let $\mathfrak{g} : \widetilde{X} \to \widetilde{Z}$ be the gate map.

Claim 1. Suppose that $\mathfrak{g}(g\tilde{x}_0) = g\tilde{x}_0$; equivalently, $g\tilde{x}_0 \in \widetilde{Z}$. Then the finite cover $Y \to X$ as in the statement of the theorem exists.

Proof of Claim 1. Let $\tilde{\gamma}$ be a combinatorial geodesic from $\tilde{x}_0$ to $g\tilde{x}_0$. Then $\tilde{\gamma} \subset \widetilde{Z}$, by convexity of $\widetilde{Z}$. Let γ be the closed path in X based at x to which $\tilde{\gamma}$ descends, so g is the path-homotopy class of γ. But $g \notin C$, by assumption, so the lift of γ to Z based at z is not a closed path. Let $Y \to X$ be the degree–$|Z^{(0)}|$ cover obtained by applying canonical completion to $Z \to X$, and using [BRHP15, Lemma 2.8]. Then Y has all of the required properties. □

In view of Claim 1, we can now assume that $g\tilde{x}_0 \notin \widetilde{Z}$. By the construction of the gate, the hyperplanes separating $g\tilde{x}_0$ from $\mathfrak{g}(g\tilde{x}_0)$ are precisely those separating $g\tilde{x}_0$ from $\widetilde{Z}$, and by our assumption there is at least one such hyperplane.

Claim 2. There exists $n \geq 1$ and hyperplanes $h_1, \ldots, h_n$ of $\widetilde{X}$ with the following properties:

- h_i separates $g\tilde{x}_0$ from $\widetilde{Z}$, for all i.
- h_i separates h_{i-1} from h_{i+1}, for $2 \leq i \leq n - 1$.
- h_i osculates with h_{i+1} for all i, and no hyperplane separates h_1 from $\widetilde{Z}$ or h_n from $g\tilde{x}_0$.

Proof of Claim 1. Let $\mathcal{H}$ be the set of hyperplanes separating $g\tilde{x}_0$ from $\mathfrak{g}(g\tilde{x}_0)$. There is a standard way to partially order the elements of $\mathcal{H}$: we say that $v \in \mathcal{H}$ precedes $h \in \mathcal{H}$ if v separates $\mathfrak{g}(g\tilde{x}_0)$ from h. Let $h_1, \ldots, h_n$ be a maximal chain in this partial order, numbered so that h_i precedes h_{i+1}. Then $h_1, \ldots, h_n$ have the first property in the claim because of the construction of the gate. They have the second property just because of the way they were chosen. Now, if h_i, h_{i+1} do not osculate, then they are separated by a hyperplane h, which is necessarily in $\mathcal{H}$. But then $h_1, \ldots, h_i, h, h_{i+1}, \ldots, h_n$ is a chain, contradicting maximality. So

h_i, h_{i+1} osculate. Similarly, if h separates $\widetilde{Z}$ from h_1 or $g\tilde{x}_0$ from h_n, then we could lengthen the chain. This proves the third assertion. $\square$

For each i, let h_i^+ be the combinatorial hyperplane in the carrier $\mathcal{N}(h_i)$ parallel to h_i and in the same h_i–halfspace as $g\tilde{x}_0$; the other combinatorial hyperplane h_i^- bounding $\mathcal{N}(h_i)$ is on the $\widetilde{Z}$ side. The preceding claim implies that $g\tilde{x}_0 \in \mathcal{N}(h_n^+)$ and $\widetilde{Z} \cap h_1^- \neq \emptyset$.

Here is the main construction in the proof. Let $Q_1 = \widetilde{Z} \cap h_1^-$. As the intersection of two convex subcomplexes, Q_1 is a convex subcomplex.

Claim 3. Let $K_1 = B \cap \mathrm{Stab}_A(h_1)$. Then K_1 acts on Q_1 cocompactly.

Proof of Claim 3. Since Z is compact, B acts on $\widetilde{Z}$ cocompactly. Next, let $b \in B$ and suppose that $bQ_1 \cap Q_1 \neq \emptyset$. Then $bh_1^- \cap h_1^- = \emptyset$. But immersed combinatorial hyperplanes in the Salvetti complex X are embedded, so $b \in \mathrm{Stab}_A(h_1^-) = \mathrm{Stab}_A(h_1)$. Hence, by a standard argument, e.g. [HS20, Lemma 2.3], the action of $B \cap \mathrm{Stab}_A(h_1)$ on Q_1 is cocompact. $\square$

Let $\bar{Q}_1 = K_1 \backslash Q_1$, which is a nonpositively-curved cube complex since Q_1 is convex in $\widetilde{X}$ and therefore CAT(0), and K_1 acts freely. Moreover, $\bar{Q}_1$ is compact by the preceding claim. Finally, the inclusions $Q_1 \hookrightarrow \widetilde{Z}$ and $Q_1 \hookrightarrow h_1^-$ descend to local isometries $\bar{Q}_1 \to Z$ and $\bar{Q}_1 \to \mathcal{N}(\bar{h}_1^-)$, where $\bar{h}_1^-$ is the immersed combinatorial hyperplane in X to which h_1^- projects.

Applying canonical completion (see [BRHP15, Lemma 2.8]), we have a finite cover $\hat{h}_1^- \to \bar{h}_1^-$ to which the local isometry $\bar{Q}_1 \to \bar{h}_1^-$ lifts as an embedding. Moreover, $|(\hat{h}_1^-)^{(0)}| = |\bar{Q}_1^{(0)}|$.

Note that since Q_1 is disjoint from its translates, the local isometry $\bar{Q}_1 \to Z$ is an embedding.

Let Z_1' be constructed as follows. Start with Z and $\mathcal{N}(\hat{h}_1)$, and glue them together along the common subcomplex $\bar{Q}_1$, where $\bar{Q}_1$ is regarded as a subcomplex of $\hat{h}_1^- \subset \mathcal{N}(\hat{h}_1)$.

Note that Z_1' is locally CAT(0), because it is obtained by gluing nonpositively curved cube complexes along a common locally convex subcomplex.

There is a natural locally injective cubical map $(Z_1', z) \to X$ extending the map $Z \to X$ and the map $\mathcal{N}(\hat{h}_1) \to X$.

Indeed, each cube of Z_1' belonging to Z maps via $Z \to X$, and the cubes of $\mathcal{N}(\hat{h}_1)$ are sent via the carrier immersion $\mathcal{N}(\hat{h}_1) \to \mathcal{N}(\bar{h}_1) \to X$. These maps agree on $\bar{Q}_1$, so we get a well-defined cubical map.

The map is locally injective on Z and $\mathcal{N}(\hat{h}_1)$. Let e be an oriented edge of Z and let f be an oriented edge of $\mathcal{N}(\hat{h}_1)$ such that e, f have a common vertex u. We need to show that the images of e, f are distinct (as oriented edges), which will verify local injectivity. Note that $u \in \bar{Q}_1$, and let $\tilde{u}$ be a lift of u to Q_1, so that there are edges $\tilde{e} \subset \widetilde{Z}$ and $\tilde{f} \subset \mathcal{N}(h_1)$ incident to $\tilde{u}$. If $\tilde{e}$ and $\tilde{f}$ are distinct as oriented edges, then their images in X are distinct (as oriented edges), so e, f do not witness a failure of local injectivity. Otherwise, $\tilde{e} = \tilde{f}$, so e, f are edges of $\bar{Q}_1$ and were hence identified.

We now add cubes to Z_1' to obtain a local isometry $Z_1 \to X$. Recall that Z_1' is locally CAT(0). Moreover, the map $Z_1' \to X$ is injective on links, by construction.

Let $\Lambda_0 = Z_1'$. Let χ_0 be the sum, over vertices in Λ_0, of the number of cubes containing the given vertex.

If our map $\Lambda_0 \to X$ is not a local isometry, then there exist edges $e_1, \ldots, e_m$ of Λ_0 with a common vertex, whose images in X span an m–cube. Now, since $Z \to X$ is a local isometry, and $\bar{Q}_1 \to X$ is a local isometry, after relabelling we have $k \leq m$ such that $e_1, \ldots, e_k$ are edges of Z (which already span a k–cube of Z), and $e_{k+1}, \ldots, e_m$ are edges of of $\hat{h}_1^-$ (which already span a cube). Given $i \leq k, j > k$, if e_i, e_j do not already span a square in Λ_0, we add such a square s.

Now, e_j is an edge of $\hat{h}_1^-$ not belonging to $\bar{Q}_1$. So, since every vertex of $\hat{h}_1^-$ lies in $\bar{Q}_1$, both ends of e_j are in $\bar{Q}_1$, and hence in Z. There are two cases:

- e_j is a loop in $\hat{h}_1^-$. Let u, v be the vertices of e_i, with $v \in \hat{h}_1^-$. By the definition of Q_1, there is no edge in Z containing u and mapping to the same edge in X as e_j, so when adding s, we add such an edge, as a loop at u. (This loop is part of the attaching map of s, along with e_j; the attaching map traverses e_i once in each direction.)
- e_j joins vertices u, v that are also joined by a nontrivial geodesic in $\bar{Q}_1$ and, say, $u \in e_i$. Recall how $\hat{h}_1^-$ was constructed from $\bar{Q}_1$ (i.e. canonical completion of Haglund-Wise, see e.g. [BRHP15, Theorem 2.6]): for each generator ℓ, each embedded ℓ–labelled path in $\bar{Q}_1$ is closed by the addition of an ℓ–labelled edge, and at each vertex with no ℓ edge, an ℓ loop is added. In particular, letting ℓ be the label of e_j, we have that u, v are joined by an ℓ–labelled geodesic γ in $\bar{Q}_1$. Since ℓ commutes with the label of e_i, since the images of e_i, e_j span a square, and $Z \to X$ is a local isometry, Z contains an edge e_i' incident to v with the same label as e_i. We attach the square using e_j, e_i, e_i', and a new edge parallel to e_j. So we have not added 0–cubes.

Hence, adding s allows us to extend the map $\Lambda_0 \to X$ to a local injection, and does not add any new 0–cubes. We add such a square for all e_i, e_j as above for which it is needed. Note that any 1–cube added is parallel to a 1–cube in $\hat{h}_1^-$. After filling in higher-dimensional cubes (which does not add 0–cells), we obtain a locally injective cubical map $\Lambda_1 \to X$ such that

- Λ_1 is locally CAT(0);
- $\Lambda_1 \to X$ extends the map $\Lambda_0 \to X$;
- $|\Lambda_1^{(0)}| = |\Lambda_0^{(0)}|$;
- every 1–cube of Λ_1 either lies in Z, or is a loop parallel to a 1–cube of $\hat{h}_1^-$, or is parallel to an edge of $\hat{h}_1^-$ that joins the endpoints of a path in Z so that the resulting cycle consists of edges all with the same label (when labels are pulled back from X);
- $\chi_1 \geq \chi_0$, with equality only if $\Lambda_0 \to X$ was a local isometry.

The above properties allow us to repeat the construction, which is guaranteed to terminate in a local isometry $Z_1 \to X$ because $|\Lambda_0^{(0)}|$ is finite, χ_* increases at each step, and local injectivity guarantees that χ_* is bounded (in terms of the number of cubes in X and $|\Lambda_0^{(0)}|$).

Now we can do the main iterative step. So far, we have constructed a local isometry $Z_1 \to X$ such that:

- $Z \subset Z_1$;
- $|Z_1^{(0)}| = |Z^{(0)}| + |\bar{Q}_1^{(0)}|$.

Moreover, by construction, the universal cover of Z_1 lifts to an embedding $\widetilde{Z}_1 \to \widetilde{X}$ such that $\widetilde{Z} \cup \mathcal{N}(h_1) \subset \widetilde{Z}_1$.

Claim 4. The orbit $\pi_1(Z_1, z) \cdot \tilde{x}_0$ is separated from $g\tilde{x}_0$ by h_1. In particular, the image of $\pi_1(Z_1, z)$ under the homomorphism to A induced by $Z \to X$ does not contain g.

Proof of Claim 4. In building Z_1', the complex Z was attached to $\mathcal{N}(\hat{h}_1) \cong \hat{h}_1 \times [-\frac{1}{2}, \frac{1}{2}]$ along a subcomplex of, say, $\hat{h}_1 \times \{-\frac{1}{2}\}$. The additional cubes added to form Z_1 each had an edge in Z, and no edge dual to $\hat{h}_1$. So Z_1 strongly deformation retracts to the subspace $Z_1 - (\hat{h}_1 \times \{-\frac{1}{2}\})$. So, any loop in Z_1 based at z is path homotopic to a loop whose lift to $\widetilde{X}$ at $\tilde{x}_0$ does not cross h_1. This implies the first part of the claim, and the second follows since h_1 separates $\widetilde{Z}$, and hence $\tilde{x}_0$, from $g\tilde{x}_0$. □

We remark that we cannot conclude at this point simply by applying canonical completion to Z_1, any more than we could have by applying canonical completion to Z: completing to a finite cover may yield a cover in which the base lift of g is a loop. However, $Z_1 \to X$ is a local isometry such that the induced map on fundamental group has image not containing g, by the preceding claim. But, the complexity of $Z_1 \to X$ relative to g is lower than that of $Z \to X$, in the following sense. Every hyperplane separating $\widetilde{Z}_1$ from $g\tilde{x}_0$ must separate $\widetilde{Z}$ from $g\tilde{x}_0$, by the preceding claim. But h_1 does not separate $\widetilde{Z}_1$ from $g\tilde{x}_0$. So $\mathsf{d}_{\widetilde{X}}(g\tilde{x}_0, \widetilde{Z}_1) < \mathsf{d}_{\widetilde{X}}(g\tilde{x}_0, \widetilde{Z})$.

We now iterate the above procedure, with Z_1 playing the role of Z and h_2 playing the role of h_1, which is possible since h_1 osculates with h_2 and therefore $\widetilde{Z}_1 \cap h_2^- = h_1^+ \cap h_2^-$. Letting $Q_2 = h_1^+ \cap h_2^-$, and letting L be the finite-index subgroup of $\pi_1 h_1$ corresponding to $\hat{h}_1$, we have that $L \cap \mathrm{Stab}_A(h_2)$ acts on Q_2 cocompactly, and the quotient embeds in $\hat{h}_1^+$. So we can repeat the above argument, gluing Z_1 to $\mathcal{N}(\hat{h}_2)$ along $\bar{Q}_2$, where $\bar{Q}_2$ has at most $|\hat{h}_1^{(0)}| \leq |Z^{(0)}|$ vertices, and the same is true of $\hat{h}_2$ (which is obtained as above using canonical completion).

So, repeating n times, we obtain a local isometry $Z_n \to X$ such that:

- $|Z_n^{(0)}| \leq (n+1)|Z^{(0)}| \leq (\|g\|+1)|Z^{(0)}|$;
- g is not in the image of $\pi_1(Z_n, z)$ under the homomorphism induced by $Z_n \to X$;
- $g\tilde{x}_0 \in \widetilde{Z}_n$, since no hyperplane separates g, h_n and $\mathcal{N}(h_n) \subset \widetilde{Z}_n$.

Now we can apply Claim 1 to Z_n to obtain a cover $Y \to X$ such that Y contains Z_n and g does not lift to a closed path at the basepoint in Y. Moreover, the degree of $Y \to X$ is bounded by $|Z_n^{(0)}| \leq (\|g\|+1)|Z^{(0)}|$, as required. □

References

[BRHP15] Khalid Bou-Rabee, Mark F. Hagen, and Priyam Patel. Residual finiteness growths of virtually special groups. *Math. Z.*, 279(1-2):297–310, 2015.

[HP16] Mark F. Hagen and Priyam Patel. Quantifying separability in virtually special groups. *Pacific J. Math.*, 284(1):103–120, 2016.

[HS20] Mark F. Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. *Israel J. Math.*, 236(1):45–89, 2020.