

On the recognition problem for virtually special cube complexes

Martin R. Bridson and Henry Wilton

Abstract.

We address the question of whether the property of being virtually special (in the sense of Haglund and Wise) is algorithmically decidable for finite, non-positively curved cube complexes. Our main theorem shows that it cannot be decided by examining one hyperplane at a time. Specifically, we prove that there does not exist an algorithm that, given a compact non-positively curved squared 2-complex X and a hyperplane H in X , will decide whether or not there is a finite-sheeted cover of X in which no lift of H self-osculates.

§1. Introduction

Haglund and Wise’s concept of a *special cube complex* [6] has had profound ramifications in geometric group theory and low-dimensional topology, most notably in Agol’s solution to the Virtually Haken Conjecture [2]. A compact, non-positively curved cube complex is *special* if it admits a local isometry to the canonical cube complex associated to a right-angled Artin group. When phrased thus, this is not a condition that one expects to be able to check in a practical manner, given a compact cube complex X . But Haglund and Wise proved that being special is equivalent to a purely combinatorial condition which requires that each of the finitely many hyperplanes in X should not behave in one of a small number of forbidden ways: a hyperplane should not cross itself, it should be 2-sided, it should not self-osculte, and it should not inter-osculte with another hyperplane (see [6] for details). We say that a hyperplane is *clean* if it does not exhibit any of these forbidden behaviours.

Received October 1, 2014.

Revised May 4, 2016.

2010 *Mathematics Subject Classification*. 20F10, 20F67, 57M07.

Both authors are supported by the EPSRC. Bridson is also supported by a Wolfson Research Merit Award from the Royal Society.

Given the combinatorial nature of these conditions, it is easy to check if a given hyperplane is clean and hence whether a given compact cube complex is special or not. However, in applications, it is not specialness that matters so much as *virtual specialness*, i.e. whether the complex has a finite-sheeted covering space that is special. Agol [1, p. 33] has promoted the study of the following question.

Question 1. *Is there an algorithm that, given a finite, non-positively curved cube complex, will determine whether or not that complex is virtually special?*

This is equivalent to the question:

Question 1'. *Is there an algorithm that, given a finite, non-positively curved cube complex X (whose hyperplanes are H_1, \dots, H_n , say), will determine the existence of a finite-sheeted covering $\widehat{X} \rightarrow X$ in which each lift of each H_i is clean.*

One has to think clearly for a moment to see that this question is not equivalent to:

Question 2. *Is there an algorithm that, given a finite, non-positively curved cube complex X and a collection of hyperplanes H_1, \dots, H_n , will determine the existence of a finite-sheeted covering $\widehat{X} \rightarrow X$ in which each lift of each H_i is clean.*

One's first thought is that a positive solution to Question 1' should proceed via a positive solution to Question 2. But on reflection one realises that the latter would be significantly more powerful because certain of the complexes X for which one did establish the existence of \widehat{X} would not be virtually special. Focusing on this difference, we shall answer Question 2 in the negative. The following theorem is an immediate consequence of our main technical result, Theorem 2.4.

Theorem 1.1. *There does not exist an algorithm that, given a finite, non-positively curved cube complex X and a hyperplane H in X , can determine whether or not there exists a finite-sheeted covering $\widehat{X} \rightarrow X$ in which each lift of H is clean.*

Note that this theorem is equivalent to the corresponding statement with 'each' replaced by 'some': this is clear if $\widehat{X} \rightarrow X$ is normal, and since the lift of a clean hyperplane is clean, we are free to pass to a finite-sheeted normal covering.

The main ingredient in our proof of Theorem 2.4 (and hence Theorem 1.1) is our work in [3], in which we showed that there is no algorithm that determines whether or not a finite complex has a proper

finite-sheeted cover (see Theorem 3.2). In [3] we also proved that this undecidability persists in the setting of finite, non-positively curved cube complexes (see Theorem 3.1). The point of this paper is to note that, with care, the same procedure can be made to translate the non-existence of a proper finite-sheeted cover into the failure of a certain hyperplane to be virtually clean.

The obstruction to cleanness that we focus on to prove Theorem 2.4 is self-osculation. Our construction shows that the question of whether one can remove self-osculations in finite-sheeted covers is already undecidable in dimension 2, and this is the context in which we shall work for the remainder of this paper. Although this theorem does not settle Question 1, it strongly constrains the nature of any possible positive solution: a putative recognition algorithm would have to take account of all the hyperplanes of X simultaneously. We therefore believe that Question 1 has a negative solution:

Conjecture 1.2. *Virtual specialness is an undecidable property of non-positively curved cube complexes.*

§2. \mathcal{VH} -complexes

Rather than working with non-positively curved cube complexes in full generality, we restrict attention to the class of \mathcal{VH} -complexes.

Definition 2.1. A square complex is \mathcal{VH} if there is a partition of the edges of the 1-skeleton into *vertical* and *horizontal* subgraphs, such that the boundary cycle of each 2-cell crosses vertical and horizontal edges alternately.

We refer the reader to [4, 7] for the basic theory of \mathcal{VH} -complexes, but we recall some salient facts here. In common with all cube complexes, much of the geometry of a \mathcal{VH} -complex X is encoded in its set of *hyperplanes* $Y \subseteq X$. (A hyperplane is a maximal connected union of midcubes.) We are concerned with certain regularity hypotheses on the hyperplanes.

Definition 2.2. Let Y be a hyperplane embedded in a \mathcal{VH} -complex X . The union of all the open cells intersecting Y , together with the closest-point projection map, defines a bundle $\mathring{N} \rightarrow Y$ with fibre $(0, 1)$, which has a natural completion to an I -bundle denoted by $\mathring{N} \rightarrow Y$ equipped with a map $\iota : N \rightarrow X$ extending the inclusion of \mathring{N} into X . If N is a trivial bundle then Y is called *non-singular* (or *2-sided*). In this case, $N \cong Y \times [0, 1]$, and the two maps

$$\partial_t : y \mapsto \iota(y, t)$$

for $t = 0, 1$, are called the *pushing maps* associated to Y .

The hyperplane Y is called *clean* if the pushing maps ∂_0, ∂_1 are both embeddings.

If every hyperplane is non-singular then the complex X is *non-singular*, and if every hyperplane is clean then the whole complex X is called *clean*. The following theorem of Haglund and Wise reconciles this terminology with that of virtually special complexes.

Theorem 2.3. [6, Theorem 5.7] *A \mathcal{VH} -complex is virtually special if and only if it is virtually clean.*

Thus, we may specialize Question 1 to the case of \mathcal{VH} -complexes as follows.

Question 3. *Is there an algorithm that decides whether or not a given finite \mathcal{VH} -complex is virtually clean?*

As in the introduction, our focus here is a *local* version of this question. We call a hyperplane Y in a \mathcal{VH} -complex X *virtually clean* if there exists a finite-sheeted covering map $p : \tilde{X} \rightarrow X$ and a connected component $\hat{Y} \subseteq p^{-1}Y$ such that the hyperplane \hat{Y} of \tilde{X} is clean. Our main theorem can now be stated as follows.

Theorem 2.4. *There is a recursive sequence of pairs of compact \mathcal{VH} -complexes X_n and hyperplanes $Y_n \subseteq X_n$ such that the set of natural numbers n for which Y_n is virtually clean is recursively enumerable but not recursive.*

The complexes X_n that we shall construct to prove Theorem 2.4 are never (globally) virtually clean—indeed, $\pi_1 X_n$ is not residually finite—so they are not themselves candidates to answer Question 3.

§3. Undecidability and finite covers

The proof of Theorem 2.4 relies on our recent work on the triviality problem for profinite completions, in which we proved the following theorem [3, Theorem B].

Theorem 3.1. *There is a recursive sequence of compact, non-positively curved square complexes X_n such that the set*

$$\{n \in \mathbb{N} \mid \widehat{\pi_1 X_n} \cong 1\}$$

is recursively enumerable but not recursive.

Here, $\widehat{\Gamma}$ denotes the profinite completion of a group Γ . By definition, $\widehat{\Gamma}$ is trivial if and only if every finite quotient of Γ is trivial.

The proof of this theorem begins with a non-positively curved square complex J that has no proper finite-sheeted covers. With a careful choice of J , we can ensure that, in the output sequence (X_n) , each X_n is a \mathcal{VH} -complex. The main result of this section is Theorem 3.7, a slight refinement of Theorem 3.1. We will start by explaining the construction that proves Theorem 3.1, and then observe some additional properties of the examples that we use to deduce Theorem 3.7.

The first ingredient in the proof of Theorem 3.1 is the corresponding result for finitely presented groups (without any non-positive curvature condition).

Theorem 3.2. *There is a recursive sequence of finite presentations \mathcal{P}_n of groups Γ_n such that the set*

$$\{n \in \mathbb{N} \mid \widehat{\Gamma}_n \not\cong 1\}$$

is recursively enumerable but not recursive.

For J we will use a complex that was constructed by Wise.

Theorem 3.3 (Wise [7]). *There exists a compact \mathcal{VH} -complex J such that $\pi_1 J$ is infinite but has no proper finite quotients.*

Burger and Mozes exhibited compact \mathcal{VH} -complexes with the even more striking property that their fundamental groups are infinite and simple [5].

We next perform a construction reminiscent of the well known Kan-Thurston construction.

Definition 3.4. Let J be as above and fix a simple closed loop c in a vertical component of the 1-skeleton of J . Let K be the presentation 2-complex associated to a finite presentation $\mathcal{P} \equiv \langle a_1, \dots, a_n \mid r_1, \dots, r_m \rangle$ of a group Γ , and let $V := K^{(1)}$. We define $J_{\mathcal{P}}$ to be the space obtained by attaching m copies of J to V , with the j -th copy attached by a cylinder joining c to the edge-loop in V labelled r_j . More formally, writing $\rho_j : \mathbb{S}^1 \rightarrow V$ for this last loop, we define \sim to be the equivalence relation on

$$V \coprod (\mathbb{S}^1 \times [0, 1]) \times \{1, \dots, m\} \coprod J \times \{1, \dots, m\}$$

defined by

$$\forall t \in \mathbb{S}^1 \forall j \in \{1, \dots, m\} : \rho_j(t) \sim (t, 0, j) \text{ and } (t, 1, j) \sim (c(t), j),$$

and define $J_{\mathcal{P}}$ to be the quotient space.

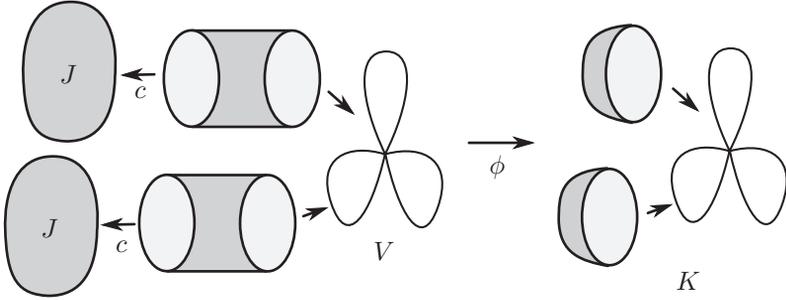


Fig. 1. The complex $J_{\mathcal{P}}$ and the map ϕ .

The construction of $J_{\mathcal{P}}$ is illustrated in Figure 1. The next proposition collects together some easy properties of this construction.

Proposition 3.5. *Let $J_{\mathcal{P}}$ be as above. Then we have the following properties:*

- (1) $J_{\mathcal{P}}$ has the structure of a \mathcal{VH} -complex;
- (2) there is a natural continuous map $\phi : J_{\mathcal{P}} \rightarrow K$, with the property that every homomorphism from $\pi_1 J_{\mathcal{P}}$ to a finite group factors through ϕ_* ;
- (3) the composition

$$\pi_1 V \rightarrow \pi_1 J_{\mathcal{P}} \xrightarrow{\phi_*} \Gamma$$

is surjective.

Proof. To prove (1), we subdivide the copies of J and V so that the attaching maps of the cylinders have the same length at either end. The gluing cylinders can then be subdivided into squares, giving the whole space the structure of a square complex, which is obviously \mathcal{VH} .

The map ϕ is defined to crush each copy of J to a point; since a cylinder with one end identified with a point is homeomorphic to a disc, the resulting space is homeomorphic to Y . The kernel of ϕ_* is normally generated by the copies of $\pi_1 J$. Since $\pi_1 J$ has no finite quotients, if f is a homomorphism from $\pi_1 J_{\mathcal{P}}$ to a finite group, then f kills the conjugacy classes defined by the loops in each copy of J , and therefore factors through ϕ_* . This proves (2).

Since V is the entire 1-skeleton of K , item (3) is immediate. Q.E.D.

Definition 3.6. A *pointed \mathcal{VH} -pair* is a triple $(L, V, *)$ where L is a compact \mathcal{VH} -complex, V is a vertical component of the 1-skeleton and $*$ is a vertex of V .

Theorem 3.7. *There is a recursive sequence of pointed \mathcal{VH} -pairs $(L_n, V_n, *_{n_1})$ such that the following set is recursively enumerable but not recursive:*

$$\{(n, \gamma) \mid n \in \mathbb{N}, \hat{\gamma} \neq 1 \text{ in } \widehat{\pi_1 L_n}\}$$

where γ runs through simple loops in V_n based at $*_{n_1}$ and $\hat{\gamma}$ denotes the element of $\widehat{\pi_1 L_n}$ defined by γ .

Proof. Let \mathcal{P}_n and Γ_n be as in Theorem 3.2 and let K_n be the presentation complex for \mathcal{P}_n . Set $L_n = J_{\mathcal{P}_n}$, let V_n be the vertical subgraph of $J_{\mathcal{P}_n}$ from Proposition 3.5 and let $*_{n_1}$ be the unique vertex of V_n . Note that the construction of $(L_n, V_n, *_{n_1})$ from \mathcal{P}_n is algorithmic, and therefore the sequence is recursive.

Consider now a finite set of simple based loops $\gamma_1, \dots, \gamma_k$ that generate $\pi_1(V_n, *_{n_1})$. If there were an algorithm to determine whether or not such a loop γ survives in some quotient of $\pi_1(L_n, *_{n_1})$ then, applying the algorithm to each γ_i in turn, we would obtain an algorithm to determine whether or not $\pi_1(V_n, *_{n_1})$ has non-trivial image in some finite quotient of $\pi_1(L_n, *_{n_1})$. But the latter occurs if and only if Γ_n has a non-trivial finite quotient (by items (2) and (3) of Proposition 3.5), which is undecidable.

This proves that the set of pairs (n, γ) for which γ survives in some finite quotient of $\pi_1(L_n, *_{n_1})$ is not recursive. However, if γ *does* survive in some finite quotient then a systematic search will eventually discover this. Therefore, the set of such (n, γ) is recursively enumerable. Q.E.D.

It will be convenient in what follows to have a fixed enumeration of the pairs (n, γ) with γ a simple loop in V_n based at $*_{n_1}$: denote this by $m \mapsto (n_m, \gamma_m)$. Correspondingly, we write $(L_m, V_m, *_{m_1})$ in place of $(L_{n_m}, V_{n_m}, *_{n_{m_1}})$, so γ_m is a simple loop in V_m based at $*_{m_1}$.

§4. Proof of Theorem 2.4

In this final section, we modify the output of Theorem 3.7 in order to deduce Theorem 2.4. First we construct L'_n from L_n by attaching a vertical loop α (consisting of a single edge) at $*_{n_1}$. We also let $V'_n = V_n \cup \alpha$. For each simple based loop γ in V_n we define γ' to be the concatenation $\gamma' \cdot \alpha$ in L'_n .

Each loop γ'_n is a local isometry

$$\gamma'_n : \mathbb{S}^1 \rightarrow L'_n$$

which induces a cellular structure on \mathbb{S}^1 , sending vertices to vertices and edges to edges. To construct the complex X_n , take two copies of L'_n , and attach an annulus $A_n \cong \mathbb{S}^1 \times I$, where the two ends of the cylinder

A_n are attached by the two copies of the map γ'_n . The construction is illustrated in Figure 2.

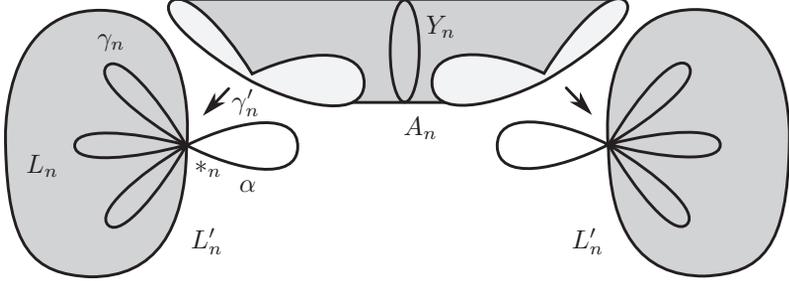


Fig. 2. The construction of the complex X_n .

Since the loops γ' are vertical and locally geodesic, we can extend the \mathcal{VH} -structures of the two copies of L'_n to a non-singular \mathcal{VH} -structure on X_n by dividing the annulus A_n into squares.

The hyperplane Y_n is defined to be the unique vertical hyperplane in the annulus A_n . It is homeomorphic to a circle, and two-sided (i.e. non-singular) by definition.

Our final lemma relates the geometry of Y_n in finite-sheeted covers of X_n to the group-theoretical properties of $\phi_*(\gamma_n) \in \Gamma_n$ (in the notation of Proposition 3.5 with $L_n = \mathcal{J}_{\mathcal{P}_n}$).

Lemma 4.1. *The following statements are equivalent.*

- (1) *The hyperplane $Y_n \subseteq X_n$ is virtually clean.*
- (2) *There is a based, finite-sheeted covering space $(R, \dagger) \rightarrow (L_n, *_n)$ so that the lift at \dagger of γ_n is not a loop.*
- (3) *There is a subgroup N of finite index in $\pi_1 L_n$ such that $\gamma_n \notin N$.*
- (4) *The group element $[\gamma_n]$ survives in some finite quotient of $\pi_1 L_n$.*
- (5) *The element $\phi_*([\gamma_n])$ survives in some finite quotient of Γ_n .*

Proof. We first show that (1) implies (2). Suppose that $p : Z \rightarrow X_n$ is a finite-sheeted covering map and that some component W of $p^{-1}Y_n$ is clean. Consider the pushing map $\partial_0 : W \rightarrow Z$. Choose some $\dagger \in p^{-1}(*_n)$ that lies in $\partial_0(W)$. Let R be the component of $p^{-1}L_n$ that contains \dagger . Because ∂_0 is an embedding, it follows that the lift of γ_n at \dagger is not a loop.

We next show that (2) implies (1); to this end, we define a certain retraction $\rho : X_n \rightarrow L_n$. Consider the natural retraction $X_n \rightarrow L'_n$ which identifies the two copies of L'_n and maps the gluing annulus via

γ'_n . Next, consider the retraction $L'_n \rightarrow L_n$ which maps $\alpha(t)$ to $\bar{\gamma}_n(t)$, where $\bar{\gamma}_n$ is the path γ_n in the opposite direction. The concatenation

$$X_n \rightarrow L'_n \rightarrow L_n$$

defines the retraction ρ . To show that (2) implies (1), we pull the covering map $R \rightarrow L_n$ back along ρ to obtain a covering space $Z \rightarrow X_n$. Let W be the unique component of the preimage of Y_n such that \dagger is contained in $\partial_t(W)$ for some $t = 0$ or 1 . By construction, W is a circle and ∂_t is given by the lift of $\gamma'_n = \gamma_n \cdot \alpha$. Since the lift of γ_n is an embedded arc with distinct endpoints, the lift of γ'_n is an embedded loop, as required.

Items (2) and (3) are equivalent by standard covering-space theory: in one direction, take $N = \pi_1(R, \dagger)$, and in the other take (R, \dagger) to be the based covering space corresponding to the subgroup N .

The equivalence of items (3) and (4) is also standard. To show that (4) implies (3), just take N to be the kernel of the finite quotient map. For the converse, the action of $\pi_1 L_n$ by left translation on $\pi L_n/N$ defines a homomorphism $\pi_1 L_n \rightarrow \text{Sym}(\pi L_n/N)$ in which γ_n survives.

Finally, that (4) and (5) are equivalent follows from item (2) of Proposition 3.5. Q.E.D.

Lemma 4.1 tells us that deciding if the hyperplane $Y_n \subseteq X_n$ is virtually clean is equivalent to deciding if (n, γ_n) lies in the set considered in Theorem 3.7; and this set is recursively enumerable but not recursive. This completes the proof of Theorem 2.4.

References

- [1] Ian Agol. Virtual properties of 3-manifolds. <https://dl.dropboxusercontent.com/u/8592391/virtuallspecialICM.pdf>.
- [2] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.
- [3] Martin R. Bridson and Henry Wilton. The triviality problem for profinite completions. *Invent. Math.*, 202(2):839–874, 2015.
- [4] Martin R. Bridson and Daniel T. Wise. $\mathcal{V}\mathcal{H}$ complexes, towers and subgroups of $F \times F$. *Mathematical Proceedings of the Cambridge Philosophical Society*, 126(3):481–497, 1999.
- [5] Marc Burger and Shahar Mozes. Finitely presented simple groups and products of trees. *C. R. Acad. Sci. Paris Sér. I Math.*, 324(7):747–752, 1997.
- [6] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geometric and Functional Analysis*, 17(5):1551–1620, 2008.

- [7] Daniel T. Wise. Complete square complexes. *Comment. Math. Helv.*, 82(4):683–724, 2007.

Martin R. Bridson, Mathematical Institute, University of Oxford, Andrew Wiles Building, Oxford, OX2 6GG, UK

E-mail address: `bridson@maths.ox.ac.uk`

Henry Wilton, DPMMS, Centre for Mathematical Sciences, Wilberforce Road, Cambridge, CB3 0WB, UK

E-mail address: `h.wilton@maths.cam.ac.uk`