

COCOMPACTLY CUBULATED GRAPH MANIFOLDS

MARK F. HAGEN* AND PIOTR PRZYTYCKI†

ABSTRACT. Let M be a graph manifold. We show that $\pi_1 M$ is the fundamental group of a compact nonpositively curved cube complex if and only if M is chargeless. We also prove that in that case $\pi_1 M$ is virtually 3-dimensional compact special.

1. INTRODUCTION

A *graph manifold* is a compact oriented aspherical 3-manifold M that has only Seifert-fibred blocks in its JSJ decomposition. We say that a torsion-free group is (*cocompactly cubulated*) if it is the fundamental group of a (compact) nonpositively curved cube complex. A (cocompactly) cubulated group is (*compact special*) if the complex is (*compact special*), i.e. admits a local isometry into the Salvetti complex of a right-angled Artin group. Moreover, if the compact special cube complex is 3-dimensional, then its fundamental group is 3-dimensional compact special.

Liu proved in [Liu13] that if a graph manifold M admits a nonpositively curved Riemannian metric, then $\pi_1 M$ is virtually cubulated (and in fact special). Under the stronger hypothesis that M has nonempty boundary, the same conclusion was obtained in [PW13]. However, the resulting cube complex was in general not compact.

The main goal of this paper is to answer Question 9.4 of Aschenbrenner, Friedl, and Wilton [AFW12] by characterizing graph manifolds M with $\pi_1 M$ virtually cocompactly cubulated, i.e. having a finite-index subgroup that acts freely and cocompactly on a CAT(0) cube complex. We show, moreover, that whenever this is the case, $\pi_1 M$ is virtually 3-dimensional compact special.

We note that if M has no JSJ tori, i.e. M is Seifert-fibred, then by [BH99, Thm 6.12] the group $\pi_1 M$ is virtually cocompactly cubulated only if the Euler number of the Seifert fibration vanishes. In this situation M is virtually a product of a circle with a surface that is a special square complex. Hence $\pi_1 M$ is virtually 3-dimensional compact special. If M is a Sol manifold, then $\pi_1 M$ is not cocompactly cubulated.

Throughout the article, we therefore assume that M is not a Sol manifold and has at least one JSJ torus, so that its underlying graph $\bar{\Gamma} = (\bar{V}, \bar{E})$ has at least one edge. We also assume that M does not contain π_1 -injective Klein bottles, so that the base orbifolds of all Seifert-fibred blocks are oriented and hyperbolic. For each $\bar{v} \in \bar{V}$, we denote by $B_{\bar{v}} \subset M$ the corresponding Seifert-fibred block, and for each edge $\bar{e} \in \bar{E}$, we denote by $T_{\bar{e}}$ the corresponding JSJ torus. For an edge \bar{e} incident to \bar{v} , let $Z_{\bar{v}}^{\bar{e}} \subset T_{\bar{e}}$ be an embedded circle that is a fiber in $B_{\bar{v}}$.

* This material is based upon work partially supported by the National Science Foundation under Grant Number NSF 1045119.

† Partially supported by the Foundation for Polish Science and National Science Centre DEC-2012/06/A/ST1/00259.

Definition 1.1. A graph manifold M is *chargeless* if for every block $B_{\bar{v}}$ disjoint from ∂M we can assign integers $n_{\bar{e}}$ to all edges $\bar{e} = (\bar{v}, \bar{v}')$, so that in integral homology $H_1(B_{\bar{v}})$ we have $\sum_{\bar{e}} n_{\bar{e}} [Z_{\bar{v}'}^{\bar{e}}] = 0$.

In other words, a graph manifold is chargeless if in each block disjoint from ∂M there is a horizontal surface whose boundary circles are vertical in adjacent blocks. For a block intersecting ∂M there is such a surface as well, since we can choose its boundary slope on ∂M , which is the required degree of freedom in $H_1(B_{\bar{v}})$, see e.g. [WY97, Lem 1.1]. Note that if M' is a finite cover of a graph manifold M , then M' is chargeless if and only if M is chargeless. Our first result is the following.

Theorem A. *Let M be a chargeless graph manifold. Then $\pi_1 M$ is virtually 3-dimensional compact special.*

Our main theorem is the following converse.

Theorem B. *Let M be a graph manifold. If $\pi_1 M$ is virtually cocompactly cubulated, then M is chargeless.*

Theorem B is one of few results giving an obstruction to being cocompactly cubulated for a specific class of groups. Another notable result of this type is Wise's characterization of tubular groups that are cocompactly cubulated [Wis12, Thm 5.8].

Organization. In Section 2 we introduce notation used to study graph manifolds. We construct an efficient family of surfaces in a graph manifold. The CAT(0) cube complex dual to the resulting system of walls in the universal cover is cocompact, as required in Theorem A. In Section 3 we study how images of isometric embeddings of \mathbb{E}^2 in a CAT(0) cube complex may intersect hyperplanes. We describe intersection patterns that allow one to recognize a combinatorial flat or a half-flat of dimension > 2 in the complex. In Section 4 we show that such objects are not allowed in cocompact cubulations of graph manifolds. This enables us to find cocompact cubical convex cores for $\pi_1 B_{\bar{v}}$. We then use Caprace–Sageev's rank-rigidity result to limit the possible cubulations of the blocks. The cubulations of adjacent blocks interact along what we call *straight subgroups*, and analyzing them is the last step in the proof of Theorem B.

Prerequisites. We assume basic knowledge of CAT(0) cube complexes and hyperplanes. We consider two metrics: the combinatorial metric on the 1-skeleton with combinatorial geodesics and isometric embeddings on 1-skeleta; and the CAT(0) metric on the entire complex, with CAT(0) geodesics and isometrically embedded copies \mathcal{E} of \mathbb{E}^2 . For an introduction to CAT(0) cube complexes, we recommend Sageev's survey article [Sag12].

Acknowledgements. We thank Henry Wilton and Daniel T. Wise for encouragement, Jordan Sahattchiev for discussions, and the referee for corrections. M.F.H. thanks the Institute of Mathematics of the Polish Academy of Sciences for its hospitality during the period in which most of this work was completed.

2. CHARGELESS GRAPH MANIFOLDS

In this section we prove Theorem A.

2.1. Preliminaries. Let M be a compact oriented irreducible 3-manifold that is not a Sol manifold and does not contain π_1 -injective Klein bottles. The manifold M contains a minimal collection of incompressible tori, called JSJ *tori*, all of whose complementary components, called *blocks* are Seifert-fibred or atoroidal [Bon02, Thm 3.4]. This decomposition is unique up to isotopy. We say that M is a *graph manifold* if all of the blocks are Seifert-fibred, and there is at least one JSJ torus.

The *underlying graph* $\bar{\Gamma} = (\bar{V}, \bar{E})$ of M is the graph dual to the JSJ decomposition. For each $\bar{v} \in \bar{V}$, we denote by $B_{\bar{v}}$ the corresponding block, and for each edge $\bar{e} \in \bar{E}$, we denote by $T_{\bar{e}}$ the corresponding JSJ torus. Let $F_{\bar{v}}$ be the base orbifold of $B_{\bar{v}}$, which is oriented and hyperbolic.

Let $G = \pi_1 M$, which we identify with the group of covering transformations of the universal cover \widetilde{M} of M . Let $\Gamma = (V, E)$ be the underlying Bass–Serre tree of \widetilde{M} arising from the JSJ decomposition of M . The vertices $v \in V$ of Γ correspond to the components $\widetilde{B}_v \subset \widetilde{M}$ of the preimages of blocks in M , which we also call *blocks*. The edges $e \in E$ of Γ correspond to the components $\widetilde{T}_e \subset \widetilde{M}$ of the preimages of JSJ tori in M , which we call JSJ *planes*. We denote by $G_v < G$ the stabilizer of the block $\widetilde{B}_v \subset \widetilde{M}$ and by $G_e < G$ the stabilizer of the JSJ plane \widetilde{T}_e .

An immersed surface in a block is *horizontal* if it is transverse to the fibers and *vertical* if it is a union of fibers.

2.2. Efficient collection. We will cubulate a chargeless graph manifold using a family of surfaces of the following type.

Definition 2.1 (Turbine collection). Let M be a chargeless graph manifold. A *turbine* collection \mathcal{S}^{tur} is a collection of surfaces that are immersed in M and have the following form. For each block $B_{\bar{v}}$ consider an embedded surface $S_{\bar{v}} \subset B_{\bar{v}}$ such that $\partial S_{\bar{v}} \cap T_{\bar{e}}$ is a union of $n_{\bar{e}}$ parallel copies of $Z_{\bar{v}}^{\bar{e}}$ from Definition 1.1. Let $A_{\bar{e}}$ be a non- ∂ -parallel vertical annulus in $B_{\bar{v}}$, both of whose boundary components are homotopic to $Z_{\bar{v}}^{\bar{e}}$ in $T_{\bar{e}}$. Gluing $2n_{\bar{e}}$ parallel copies of the annuli $A_{\bar{e}}$ with 2 parallel copies of $S_{\bar{v}}$ yields a surface $S_{\bar{v}}^{\text{tur}}$ properly immersed in M . We let \mathcal{S}^{tur} denote the collection of $S_{\bar{v}}^{\text{tur}}$ for all \bar{v} , together with some non- ∂ -parallel vertical annuli A with $\partial A \subset T$ for all boundary tori T of M .

Remark 2.2. Let $\Gamma' \rightarrow \bar{\Gamma}$ be a cover of the underlying graph $\bar{\Gamma}$ of M . Let $M' \rightarrow M$ be the induced cover of M , i.e. the graph of spaces whose underlying graph is Γ' and whose vertex spaces are isomorphic to the various blocks of M . If Γ' is simple, i.e. has no loops or double edges, then every surface in \mathcal{S}^{tur} lifts to an embedding in M' .

Definition 2.3 (Torus collection, efficient collection). Each block $B_{\bar{v}}$ of a graph manifold M has a finite cover $B'_{\bar{v}}$ that is a product of a circle with a surface $F'_{\bar{v}}$. A *torus* collection $\mathcal{S}_{\bar{v}}^{\text{tor}}$ is a family of vertical tori in $B'_{\bar{v}}$ whose base curves *fill* $F'_{\bar{v}}$. With respect to some hyperbolic metric on $F'_{\bar{v}}$ this means that the complementary components of the union of the geodesic representatives of the base curves are discs or annular neighborhoods of the boundary. Each torus $T \in \mathcal{S}_{\bar{v}}^{\text{tor}}$ is equipped with a map to M factoring through $B'_{\bar{v}}$. Let \mathcal{S}^{tor} be the union of the tori in all $\mathcal{S}_{\bar{v}}^{\text{tor}}$ together with all the JSJ tori of M .

The union $\mathcal{S}^{\text{eff}} = \mathcal{S}^{\text{tur}} \cup \mathcal{S}^{\text{tor}}$ of a turbine collection and a torus collection for a chargeless graph manifold is called an *efficient* collection. We also assume that \mathcal{S}^{eff} is in general position. Let $\widetilde{\mathcal{S}}^{\text{eff}}$ be the collection of all connected surfaces in the universal cover \widetilde{M} of M covering the surfaces of \mathcal{S}^{eff} .

Each surface in $\widetilde{\mathcal{S}}^{\text{eff}}$ cuts \widetilde{M} into two halfspaces and the collection of such pairs endows \widetilde{M} with a Haglund–Paulin wallspace structure (see e.g. [HW12, Sec 2.1]).

Construction 2.4. Sageev’s construction yields the CAT(0) cube complex \mathcal{X} dual to $\widetilde{\mathcal{S}}^{\text{eff}}$. A vertex x of \mathcal{X} is a choice of a closed halfspace $x(S)$ for each $S \in \widetilde{\mathcal{S}}^{\text{eff}}$ with $x(S) \cap x(S') \neq \emptyset$ for all $S, S' \in \widetilde{\mathcal{S}}^{\text{eff}}$. Cubes of dimension k are spanned by the sets of 2^k vertices differing on k surfaces S . If that complex is disconnected, we restrict to the connected component containing the vertices x for which there is a point $m \in \widetilde{M}$ with $m \in x(S)$ for all S .

The group $\pi_1 M$ acts naturally on \mathcal{X} . Hyperplanes of \mathcal{X} correspond to surfaces in $\widetilde{\mathcal{S}}^{\text{eff}}$ and their stabilizers coincide. Maximal cubes of \mathcal{X} correspond to maximal collections of pairwise intersecting surfaces in $\widetilde{\mathcal{S}}^{\text{eff}}$.

Proposition 2.5. *Let M be a chargeless graph manifold with an efficient collection \mathcal{S}^{eff} . Let \mathcal{X} be the CAT(0) cube complex dual to $\widetilde{\mathcal{S}}^{\text{eff}}$. Then the action of $\pi_1 M$ on \mathcal{X} is free and cocompact. Moreover, $\pi_1 M \backslash \mathcal{X}$ is virtually special.*

Before the proof, note that Proposition 2.5 implies Theorem A, except for 3–dimensionality. This last property will be obtained in Remark 2.7.

Lemma 2.6. *Let S, S' be distinct surfaces in $\widetilde{\mathcal{S}}^{\text{eff}}$ intersecting a block $\widetilde{B}_v \subset \widetilde{M}$. If $S \cap S' \neq \emptyset$, then $S \cap S' \cap \widetilde{B}_v \neq \emptyset$. Moreover, at most one of $S \cap \widetilde{B}_v, S' \cap \widetilde{B}_v$ is horizontal.*

Proof. If one of $S \cap \widetilde{B}_v, S' \cap \widetilde{B}_v$ is horizontal and the other vertical, then they intersect.

If they are both horizontal, then they cover the same unique surface in \mathcal{S}^{eff} that has horizontal intersection with the block of M covered by \widetilde{B}_v . By Remark 2.2 applied to the universal cover $\Gamma' \rightarrow \overline{\Gamma}$, through which the map $\Gamma \rightarrow \overline{\Gamma}$ factors, the surfaces S, S' are disjoint.

If $S \cap \widetilde{B}_v, S' \cap \widetilde{B}_v$ are both vertical and disjoint, but S and S' intersect, then there is an adjacent block $\widetilde{B}_{v'}$ with both $S \cap \widetilde{B}_{v'}, S' \cap \widetilde{B}_{v'}$ nonempty and horizontal. As before, we obtain that S and S' are disjoint, which is a contradiction. \square

Proof of Proposition 2.5. The action is free by [PW13, Sec 3] since \mathcal{S}^{eff} is *sufficient* in the sense of [PW13, Def 1.4].

For cocompactness it suffices to show that, for any collection of pairwise intersecting surfaces $\mathcal{S} \subset \widetilde{\mathcal{S}}^{\text{eff}}$, there is a point of \widetilde{M} uniformly close to each $S \in \mathcal{S}$. First note that by the Helly property for trees, there is a block $\widetilde{B}_v \subset \widetilde{M}$ intersecting all the surfaces in \mathcal{S} . By Lemma 2.6 the surfaces $S \cap \widetilde{B}_v$ for $S \in \mathcal{S}$ pairwise intersect. Moreover, the collection \mathcal{S} contains at most one surface S^{hor} with $S^{\text{hor}} \cap \widetilde{B}_v$ horizontal and we can assume without loss of generality that there is such a surface in \mathcal{S} . Then the family of lines $S \cap S^{\text{hor}}$ with $S \in \mathcal{S} - \{S^{\text{hor}}\}$ is a family of uniform quasi-geodesics in the quasi-tree $S^{\text{hor}} \cap \widetilde{B}_v$, pairwise at finite distance. Then, for example by [Sag97, Lem 3.4], there is a point in $S^{\text{hor}} \cap \widetilde{B}_v$ uniformly close to all these lines. Hence $G = \pi_1 M$ acts cocompactly on \mathcal{X} .

We verify specialness using [HW08, Thm 9.19], saying that for a free and cocompact action of a group G on a CAT(0) cube complex \mathcal{X} the following are equivalent:

- (1) G has a finite-index subgroup G' such that $G' \backslash \mathcal{X}$ is special.
- (2) For each hyperplane \mathfrak{h} of \mathcal{X} , the subgroup $\text{Stab } \mathfrak{h} \leq G$ is separable, and for intersecting hyperplanes $\mathfrak{h}, \mathfrak{h}'$, the double coset $\text{Stab } \mathfrak{h} \text{Stab } \mathfrak{h}' \subset G$ is separable.

Each hyperplane in \mathcal{X} has stabilizer $\text{Stab } S$ for some $S \in \widetilde{\mathcal{S}}^{\text{eff}}$. Since \mathcal{S}^{eff} is in general position, each $\text{Stab } S$ coincides with a conjugate of the fundamental group of a surface in \mathcal{S}^{eff} .

Since the surfaces in \mathcal{S}^{eff} are virtually embedded, the groups $\text{Stab } S$ are separable [PW13, Thm 1.1], and double cosets $\text{Stab } S \text{ Stab } S'$ are separable for intersecting surfaces $S, S' \in \widetilde{\mathcal{S}}^{\text{eff}}$, by [PW13, Thm 1.2]. Hence, by the above criterion, the complex $\pi_1 M \setminus \mathcal{X}$ is virtually special. \square

We now complete the proof of Theorem A.

Remark 2.7. While $\pi_1 M \setminus \mathcal{X}$ in Proposition 2.5 might not be 3-dimensional, we can arrange it to be by enhancing the construction in the following way. First, by [LW97, Prop 4.4] by passing to a finite cover of M we can assume that every block is a product of a circle with a surface of genus ≥ 1 . Such a surface admits a filling family \mathcal{L} of essential curves and arcs whose dual cube complex is 2-dimensional, such that the boundary of each arc lies in a boundary circle and for each boundary circle there is exactly one such arc.

One way to construct \mathcal{L} is to start with a family of disjoint arcs as above, which is possible by the assumption on the genus, and complete it with a maximal family of disjoint pairwise non-homotopic curves to \mathcal{L}_1 . Then embed the graph dual to \mathcal{L}_1 in the surface and define \mathcal{L}_2 to be the curves in the boundary of a regular neighborhood of that graph. The union $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2$ is as required.

We use the curves and arcs in \mathcal{L} as bases for the vertical tori in the torus collection and the vertical annuli used in the construction of the turbine collection. Then any collection of pairwise intersecting surfaces in $\widetilde{\mathcal{S}}^{\text{eff}}$ all intersecting a block $\widetilde{B}_v \subset \widetilde{M}$ consists of at most 3 surfaces, one of which might have horizontal intersection with \widetilde{B}_v and at most two might have vertical intersection.

3. CONSTRUCTING FLATS FROM PARALLEL FAMILIES

We now begin to establish statements needed in the proof of Theorem B. In this section we give methods of recognizing flats in a CAT(0) cube complex using parallel families, defined below.

Definition 3.1 (Parallel family of lines). Let \mathcal{S} be a locally finite *multiset* of geodesic lines in the Euclidean plane \mathbb{E}^2 , i.e. a family of lines that are allowed to occur multiple times. We say that \mathcal{S} is a *parallel family* if the lines are pairwise parallel and for each $S \in \mathcal{S}$ each of the two components of $\mathbb{E}^2 - S$ contains another line from \mathcal{S} . We say that \mathcal{S} *consists of n parallel families* if \mathcal{S} can be partitioned into n parallel families that are maximal, or equivalently, pairwise transverse.

Definition 3.2 (Combinatorial flat). A *combinatorial flat* \mathbf{E}_n is a CAT(0) cube complex that is the standard tiling of \mathbb{E}^n by unit n -cubes.

Remark 3.3. Let \mathcal{X} be the CAT(0) cube complex from Construction 2.4, dual to a collection \mathcal{S} consisting of n parallel families. If none of the lines in \mathcal{S} coincide, then \mathcal{X} is a combinatorial n -flat. Otherwise \mathcal{X} is the product of n complexes, each of which is built from a sequence $(Q_k)_{k \in \mathbb{Z}}$ of cubes by gluing them along vertices so that the vertex $Q_{k-1} \cap Q_k$ is opposite to $Q_k \cap Q_{k+1}$ in Q_k . Each cube Q_k corresponds to a collection of coinciding lines, and $\dim Q_k$ is equal to the number of such lines.

We would now like to recognize a combinatorial flat inside an ambient CAT(0) cube complex.

Remark 3.4. Any hyperplane \mathfrak{h} in a CAT(0) cube complex \mathcal{X} has a metric neighborhood of the form $\mathfrak{h} \times I$. Hence for any CAT(0) geodesic ε in \mathcal{X} that is neither disjoint from nor

contained in \mathfrak{h} , the intersection $\varepsilon \cap \mathfrak{h}$ is a point separating the intersections of ε with the two open halfspaces of \mathfrak{h} . Consequently, for any isometrically embedded copy $\mathcal{E} \subset \mathcal{X}$ of \mathbb{E}^2 that is neither disjoint from nor contained in \mathfrak{h} , the intersection $\mathcal{E} \cap \mathfrak{h}$ is a line separating the convex intersections of \mathcal{E} with these halfspaces.

Definition 3.5 (Parallel family of hyperplanes). Let \mathcal{X} be a CAT(0) cube complex containing an isometrically embedded copy \mathcal{E} of \mathbb{E}^2 . We do not require that \mathcal{E} is a combinatorial flat, only that it is isometrically embedded with respect to the CAT(0) metric of \mathcal{X} . Let $\text{Ess}(\mathcal{E})$ denote the collection of hyperplanes that are neither disjoint from nor contain \mathcal{E} . We say that \mathcal{E} has n parallel families if the collection of lines $\mathfrak{h} \cap \mathcal{E}$, with $\mathfrak{h} \in \text{Ess}(\mathcal{E})$, consists of n parallel families. The family of hyperplanes intersecting \mathcal{E} along one of these families is called *parallel* as well.

Note that parallelism of $\mathfrak{h} \cap \mathcal{E}$, $\mathfrak{h}' \cap \mathcal{E}$ does not imply that the hyperplanes \mathfrak{h} , \mathfrak{h}' are disjoint; these hyperplanes may intersect outside of \mathcal{E} .

Let \mathcal{X} be finite-dimensional. If a group J of automorphisms of \mathcal{X} acts cocompactly on \mathcal{E} , then for each hyperplane $\mathfrak{h} \in \text{Ess}(\mathcal{E})$, the family of lines $\mathfrak{h}' \cap \mathcal{E}$, with $\mathfrak{h}' \in \text{Ess}(\mathcal{E})$, parallel to $\mathfrak{h} \cap \mathcal{E}$ is a parallel family. Consequently, there exists n such that \mathcal{E} has n parallel families.

Lemma 3.6. *Let \mathcal{X} be a finite-dimensional CAT(0) cube complex with a free action of a group $J = \mathbb{Z}^2$. Let $\mathcal{E} \subset \mathcal{X}$ be a J -cocompact isometrically embedded copy of \mathbb{E}^2 . Then one of the following holds:*

- (1) *For some $n \geq 3$, and some finite-index subgroup $J' \leq J$, the complex \mathcal{X} contains a J' -invariant combinatorial n -flat \mathcal{Y} such that $\mathcal{Y}^{(1)} \subset \mathcal{X}^{(1)}$ is isometrically embedded.*
- (2) *The plane \mathcal{E} has two parallel families.*

Proof. Denote by n the number of parallel families of \mathcal{E} . We first prove $n \geq 2$. Each complementary component in \mathcal{E} of the union of all the hyperplanes in $\text{Ess}(\mathcal{E})$ is contained in the cubical star of a vertex. Since each cube is bounded, each such component must be bounded. Hence there are hyperplanes $\mathfrak{a}, \mathfrak{b} \in \text{Ess}(\mathcal{E})$ with lines $\mathfrak{a} \cap \mathcal{E}$ and $\mathfrak{b} \cap \mathcal{E}$ intersecting transversely.

We shall now construct a combinatorial n -flat $\mathcal{Y} \subset \mathcal{X}$. For a vertex $x \in \mathcal{X}$ and a hyperplane \mathfrak{h} , let $x(\mathfrak{h})$ be the closed halfspace of \mathfrak{h} containing x . If $\mathfrak{h} \notin \text{Ess}(\mathcal{E})$, so that \mathcal{E} is contained in one or two of the closed halfspaces of \mathfrak{h} , we define $\mathcal{E}(\mathfrak{h})$ to be one such fixed halfspace. Consider the set $\mathcal{Y}^{(0)}$ of vertices $x \in \mathcal{X}$ satisfying the following conditions:

- (a) For $\mathfrak{h} \notin \text{Ess}(\mathcal{E})$, we have $x(\mathfrak{h}) = \mathcal{E}(\mathfrak{h})$.
- (b) For $\mathfrak{h}, \mathfrak{h}' \in \text{Ess}(\mathcal{E})$, we have $x(\mathfrak{h}) \cap x(\mathfrak{h}') \cap \mathcal{E} \neq \emptyset$.

Let \mathcal{Y} be the subcomplex spanned by $\mathcal{Y}^{(0)}$, i.e. the union of the closed cubes all of whose vertices are contained in $\mathcal{Y}^{(0)}$. By condition (a), each vertex $x \in \mathcal{Y}^{(0)}$ is uniquely determined by the choice of halfspaces $x(\mathfrak{h})$ with $\mathfrak{h} \in \text{Ess}(\mathcal{E})$ satisfying condition (b). By condition (b), for all \mathfrak{h} from the same parallel family we have $x(\mathfrak{h}) \ni e$ for some $e \in \mathcal{E}$. Thus \mathcal{Y} is isomorphic to the CAT(0) cube complex dual to the intersection lines in \mathcal{E} of $\text{Ess}(\mathcal{E})$. Hence \mathcal{Y} is a combinatorial n -flat or possibly a complex described in Remark 3.3. A pair of points $e, e' \in \mathcal{E}$ is separated by $\mathfrak{h} \cap \mathcal{E}$ in \mathcal{E} if and only if it is separated by \mathfrak{h} in \mathcal{X} . Hence $\mathcal{Y}^{(1)} \subset \mathcal{X}^{(1)}$ is isometrically embedded.

Let J' be the finite-index subgroup of J that acts trivially on the finite set of hyperplanes containing \mathcal{E} . The complex \mathcal{Y} is J' -invariant since J preserves parallel families and J' preserves the collection of $\mathcal{E}(\mathfrak{h})$. In the case where \mathcal{Y} is not a combinatorial n -flat, we replace J' with a further finite-index subgroup whose elements do not map a hyperplane in $\text{Ess}(\mathcal{E})$ to

a distinct hyperplane with the same intersection line with \mathcal{E} . We can then replace \mathcal{Y} with a J' -invariant subcomplex that is a combinatorial n -flat isometrically embedded on 1-skeleton. If $n \geq 3$, case (1) of the lemma is verified. Otherwise, we have $n = 2$ and hence case (2). \square

The following complements Lemma 3.6.

Definition 3.7 (Combinatorial half-plane). Let \mathbf{E}_2 be a combinatorial 2-flat with a bi-infinite combinatorial geodesic $\gamma \subset \mathbf{E}_2^{(1)}$. The CAT(0) cube complex that is the closure of a component of $\mathbf{E}_2 - \gamma$ is a *combinatorial half-plane bounded by γ* .

Lemma 3.8. *Let \mathcal{X} be a finite-dimensional CAT(0) cube complex with a free action of a group $J = \mathbb{Z}^2$. Let $\mathcal{E} \subset \mathcal{X}$ be a J -cocompact isometrically embedded copy of \mathbb{E}^2 that has two parallel families. Then one of the following holds:*

- (1) *For some finite-index subgroup $J' \leq J$, the complex \mathcal{X} contains a J' -invariant subcomplex $\mathcal{F} \times \mathbf{E}_1$ where \mathcal{F} is a combinatorial half-plane, such that $(\mathcal{F} \times \mathbf{E}_1)^{(1)} \subset \mathcal{X}^{(1)}$ is isometrically embedded.*
- (2) *There exists a J -cocompact convex subcomplex $\widehat{\mathcal{Y}} \subset \mathcal{X}$.*

We first reduce the proof of Lemma 3.8 to the following.

Lemma 3.9. *Let \mathcal{X} be a finite-dimensional CAT(0) cube complex and let $g \in \text{Aut}(\mathcal{X})$ act by a translation on a bi-infinite combinatorial geodesic $\gamma \subset \mathcal{X}^{(1)}$. Then one of the following holds:*

- (1) *There exists $m \geq 1$ and a g^m -invariant combinatorial half-plane $\mathcal{F} \subset \mathcal{X}$, bounded by γ and contained in its cubical convex hull, such that $\mathcal{F}^{(1)} \subset \mathcal{X}^{(1)}$ is isometrically embedded.*
- (2) *There exists R such that if intersecting hyperplanes $\mathfrak{h}, \mathfrak{h}'$ intersect γ , then they intersect it in points at distance $\leq R$.*

Here the *cubical convex hull* $\widehat{\mathcal{Y}}$ of a subcomplex \mathcal{Y} of a CAT(0) cube complex \mathcal{X} is the subcomplex of \mathcal{X} spanned by the vertices that are not separated from $\mathcal{Y}^{(0)}$ by any hyperplane.

Proof of Lemma 3.8. Let \mathcal{Y} be the combinatorial 2-flat $\alpha \times \beta$ from the proof of Lemma 3.6, where the factors come from the two parallel families. We choose embeddings of the combinatorial geodesics α, β into \mathcal{Y} with arbitrary fixed coordinates. As before, let J' be a finite-index subgroup of J preserving \mathcal{Y} . Consider the J' -invariant cubical convex hull $\widehat{\mathcal{Y}}$ of \mathcal{Y} , which is the product of the cubical convex hulls $\widehat{\alpha}$ and $\widehat{\beta}$ of α, β . Let $\text{Stab}_{J'} \alpha = \langle a \rangle$, $\text{Stab}_{J'} \beta = \langle b \rangle$ and note that $\langle a, b \rangle \leq J'$ has finite index.

By Lemma 3.9, there are two possibilities. The first is that one of the subcomplexes $\widehat{\alpha}$ or $\widehat{\beta}$ contains a combinatorial half-plane \mathcal{F} , bounded by α or β , whose 1-skeleton is isometrically embedded in that of $\widehat{\alpha}$ or $\widehat{\beta}$. Moreover, \mathcal{F} is a^m - or b^m -invariant for some $m > 0$. Then after possibly replacing J' with a further finite-index subgroup, we have that $\widehat{\mathcal{Y}}$ contains a J' -invariant subcomplex $\mathcal{F} \times \mathbf{E}_1$ isometrically embedded on 1-skeleton. This verifies case (1) of the lemma. The second possibility is that there exists R such that any pair of intersecting hyperplanes from $\text{Ess}(\mathcal{E})$ intersects \mathcal{Y} in lines at distance $\leq R$. Then $\widehat{\mathcal{Y}}$ contains finitely many J' -orbits of vertices, and hence J' acts cocompactly on $\widehat{\mathcal{Y}}$.

In order to find a subcomplex that is J -invariant, we relax condition (a) from the proof of Lemma 3.6, so that $x(\mathfrak{h})$ can be arbitrary for \mathfrak{h} containing \mathcal{E} . The cubical convex hull of such a set of vertices has the form $\widehat{\mathcal{Y}} \times I^k$, where k is the number of such \mathfrak{h} , hence satisfies case (2) of the lemma. \square

The following completes the proof of Lemma 3.8.

Proof of Lemma 3.9. All the hyperplanes discussed in this proof are assumed to intersect γ . Finite-dimensionality implies that there exists $m \geq 1$ such that any hyperplane \mathfrak{h} is disjoint from and separates $g^{\pm m}\mathfrak{h}$. We order the hyperplanes so that $\mathfrak{h} < \mathfrak{h}'$ if the direction of the subpath of γ from $\mathfrak{h} \cap \gamma$ to $\mathfrak{h}' \cap \gamma$ agrees with the direction of the translation g . Note that if $\mathfrak{h} < \mathfrak{h}''$ intersect, then each $\mathfrak{h} < \mathfrak{h}' < \mathfrak{h}''$ intersects \mathfrak{h} or \mathfrak{h}'' . In particular, if $\mathfrak{h}'' = g^{mj}\mathfrak{h}'$, then \mathfrak{h}' intersects \mathfrak{h} .

If case (2) of the lemma does not hold, then there exist intersecting hyperplanes $\mathfrak{h} < \mathfrak{h}'$ with $\mathfrak{h} \cap \gamma$, $\mathfrak{h}' \cap \gamma$ at arbitrary distance. Since there are finitely many $\langle g^m \rangle$ -orbits of hyperplanes, there exist hyperplanes $\mathfrak{a} < \mathfrak{b}$ such that $g^{-mi}\mathfrak{a}$ intersects $g^{mj}\mathfrak{b}$ for infinitely many $i, j \geq 0$. By the previous paragraph, $g^{-mi}\mathfrak{a}$ intersects $g^{mj}\mathfrak{b}$ for all $i, j \geq 0$.

We shall now partition the set of hyperplanes into two families $\mathfrak{A} \supset \{g^{mi}\mathfrak{a}\}$, $\mathfrak{B} \supset \{g^{mj}\mathfrak{b}\}$, with $i, j \in \mathbb{Z}$, such that any $\mathfrak{h} \in \mathfrak{A}$ and $\mathfrak{h}' \in \mathfrak{B}$ satisfying $\mathfrak{h} < \mathfrak{h}'$ intersect. For every hyperplane \mathfrak{h} , for sufficiently large $i, j \geq 0$ we have $g^{-mi}\mathfrak{a} < \mathfrak{h} < g^{mj}\mathfrak{b}$. Observe that if \mathfrak{h} intersects $g^{mj}\mathfrak{b}$ and $g^{mj''}\mathfrak{b}$ for some $j < j''$, then \mathfrak{h} intersects all $g^{mj'}\mathfrak{b}$ with $j < j' < j''$. Hence \mathfrak{h} intersects all $g^{-mi}\mathfrak{a}$ for i sufficiently large or all $g^{mj}\mathfrak{b}$ for j sufficiently large. If the former property holds, or if both properties hold simultaneously, then let $\mathfrak{h} \in \mathfrak{B}$; otherwise let $\mathfrak{h} \in \mathfrak{A}$ and note that then \mathfrak{h} is disjoint from all $g^{-mi}\mathfrak{a}$ for i sufficiently large. Now for any $\mathfrak{h} \in \mathfrak{A}$ and $\mathfrak{h}' \in \mathfrak{B}$ satisfying $\mathfrak{h} < \mathfrak{h}'$, there is $i \geq 0$ with $g^{-mi}\mathfrak{a} < \mathfrak{h}$. After possibly increasing i , the hyperplane \mathfrak{h}' intersects $g^{-mi}\mathfrak{a}$, but \mathfrak{h} does not intersect $g^{-mi}\mathfrak{a}$, whence \mathfrak{h} intersects \mathfrak{h}' , as required.

Let $p = (x, x')$ be a pair of vertices of γ such that either $x = x'$ or the direction of the subpath $xx' \subset \gamma$ agrees with the direction of the translation g . For every hyperplane \mathfrak{h} we choose the halfspace $p(\mathfrak{h})$ to be $x(\mathfrak{h})$ if $\mathfrak{h} \in \mathfrak{A}$ or $x'(\mathfrak{h})$ if $\mathfrak{h} \in \mathfrak{B}$. By the previous paragraph, any such halfspaces have non-empty intersection, and hence define a vertex in the cubical convex hull of γ . The set $\mathcal{F}^{(0)}$ of such vertices is g^m -invariant. Let \mathcal{F} be the subcomplex spanned by $\mathcal{F}^{(0)}$. The 1-skeleton of \mathcal{F} is obtained by adding edges between vertices determined by p that differ by moving x through a hyperplane in \mathfrak{A} or x' through a hyperplane in \mathfrak{B} . The squares of \mathcal{F} are obtained by performing two such moves simultaneously. It is easy to see that \mathcal{F} is a combinatorial half-plane bounded by γ and that $\mathcal{F}^{(1)} \subset \mathcal{X}^{(1)}$ is isometrically embedded. \square

4. COCOMPACTLY CUBULATED GRAPH MANIFOLDS

In this section, we prove Theorem B. As in Section 2, we denote by M a graph manifold with universal cover \widetilde{M} and Bass-Serre tree Γ . We pull back a fixed Riemannian metric on M to \widetilde{M} . Finally, we assume that $G = \pi_1 M$ acts freely and cocompactly on a CAT(0) cube complex \mathcal{X} .

4.1. Straight subgroups.

Lemma 4.1. *The inclusion $\widetilde{T}_e \subset \widetilde{M}$ does not extend to a G'_e -equivariant quasi-isometric embedding $\phi: \widetilde{T}_e \times [0, \infty) \rightarrow \widetilde{M}$ for any finite-index subgroup $G'_e \leq G_e$.*

Proof. Let $e = (v, v')$. If such ϕ exists, then let L be its additive constant, so that the L -neighborhood of the image under ϕ of any path-connected set is path-connected. We first show that for any such ϕ and for any $s \in [0, \infty)$ there is $t \in \widetilde{T}_e$ with $\phi(t, s) \in \mathcal{N}_L(\widetilde{B}_v \cup \widetilde{B}_{v'})$. Otherwise, fix any $t \in \widetilde{T}_e$. Without loss of generality $\phi(t, s)$ lies in a component of $\widetilde{M} - \mathcal{N}_L(\widetilde{B}_v)$ that does not contain $\widetilde{B}_{v'}$. Let $g \in G'_e$ be non-central in G_v . Then g does not stabilize any of the edges of Γ incident to v except for e . Thus $\phi(t, s)$ and $\phi(gt, s)$ lie in different components

of $\widetilde{M} - \mathcal{N}_L(\widetilde{B}_v)$. Consider any path τ in \widetilde{T}_e joining t with gt . Then $\mathcal{N}_L(\phi(\tau, s))$ joins two different components of $\widetilde{M} - \widetilde{B}_v$ but is disjoint from \widetilde{B}_v . Thus $\mathcal{N}_L(\phi(\tau, s))$ is disconnected, which is a contradiction. Since G'_e acts cocompactly on \widetilde{T}_e , it follows that the entire image of ϕ is contained in a uniform neighborhood of $\widetilde{B}_v \cup \widetilde{B}_{v'}$.

The quasi-isometric embedding ϕ induces a bi-Lipschitz embedding of asymptotic cones $\mathbf{Cone}(\widetilde{T}_e \times [0, \infty)) \rightarrow \mathbf{Cone}(\widetilde{B}_v \cup \widetilde{B}_{v'})$. Here we choose an arbitrary ultrafilter, scaling constants, and observation points in $\widetilde{T}_e \times \{0\}$. Then $\mathbf{Cone}(\widetilde{T}_e \times [0, \infty)) = \mathbb{E}^2 \times [0, \infty)$ which has topological dimension 3. On the other hand $\mathbf{Cone}(\widetilde{B}_v \cup \widetilde{B}_{v'})$ is obtained by gluing two copies of the product of an \mathbb{R} -tree with \mathbb{E}^1 along \mathbb{E}^2 ; the result has dimension 2, a contradiction. \square

Definition 4.2 (Essential hyperplane, essential core). Let Y be a subspace of a metric space X . We say that $Z \subset X$ is Y -essential if Z is not contained in any k -neighborhood $\mathcal{N}_k(Y)$ of Y .

Let \mathcal{V} be a subspace of a CAT(0) cube complex \mathcal{X} . Let $\text{Ess}_{\mathcal{X}}(\mathcal{V})$ be the set of hyperplanes \mathfrak{h} of \mathcal{X} such that the intersections of \mathcal{V} with both halfspaces of \mathfrak{h} are \mathfrak{h} -essential. We usually omit the subscript \mathcal{X} , and this agrees with the notation $\text{Ess}(\mathcal{E})$, which we have used so far.

Similarly, if H acts on \mathcal{X} , a hyperplane \mathfrak{h} is called H -essential if the intersections of an (hence any) H -orbit with both halfspaces of \mathfrak{h} are \mathfrak{h} -essential. If $\mathcal{V} \subset \mathcal{X}$ is H -cocompact, then $\text{Ess}(\mathcal{V})$ is the set of H -essential hyperplanes.

If \mathcal{V} is a convex subcomplex, then each $\mathfrak{h} \in \text{Ess}_{\mathcal{X}}(\mathcal{V})$ intersects \mathcal{V} , these intersections form $\text{Ess}_{\mathcal{V}}(\mathcal{V})$, and thus we can identify $\text{Ess}_{\mathcal{X}}(\mathcal{V})$ with $\text{Ess}_{\mathcal{V}}(\mathcal{V})$ and omit the subscript. The essential core \mathcal{V}^{ess} is the CAT(0) cube complex dual to $\text{Ess}(\mathcal{V})$, which is a quotient of \mathcal{V} . For example, if $\mathcal{V} = \mathbf{E}_1 \times I$, then $\text{Ess}(\mathcal{V})$ consists of all the hyperplanes except $\mathbf{E}_1 \times \{\frac{1}{2}\}$, hence $\mathcal{V}^{ess} = \mathbf{E}_1$.

If a group H acts on \mathcal{V} freely and cocompactly, then it acts on \mathcal{V}^{ess} freely and cocompactly as well. If $J \leq H$ and a hyperplane \mathfrak{h} of \mathcal{X} is J -essential, then we have $\mathfrak{h} \in \text{Ess}(\mathcal{V})$ and the corresponding hyperplane in \mathcal{V}^{ess} is J -essential as well.

Definition 4.3 (Straight subgroup). Let e be an edge of Γ . A maximal cyclic subgroup of G_e is *straight* if it contains a nontrivial element stabilizing a G_e -essential hyperplane of \mathcal{X} .

Proposition 4.4. *For each edge e of Γ , the group G_e has exactly two straight subgroups.*

Proof. By [Hag07, Thm 1.4], G acts semisimply on \mathcal{X} . By the Flat Torus Theorem [BH99, Thm II.7.1], \mathcal{X} contains an isometrically embedded copy \mathcal{E} of \mathbb{E}^2 on which G_e acts cocompactly. We apply Lemma 3.6 with $J = G_e$. Note that the set of G_e -essential hyperplanes coincides with the union of parallel families $\text{Ess}(\mathcal{E})$. In case (1) of Lemma 3.6, a G -equivariant quasi-isometry $\mathcal{X} \rightarrow \widetilde{M}$ maps \mathcal{V} to $\widetilde{T}_e \times (-\infty, \infty)^{n-2}$ contradicting Lemma 4.1. Hence we have case (2), which says that \mathcal{E} has two parallel families. Consider the two maximal cyclic subgroups $\mathbb{Z}_a, \mathbb{Z}_b < G_e$ stabilizing lines in these families. Then $\mathbb{Z}_a, \mathbb{Z}_b$ are straight, since they have finite-index subgroups stabilizing all hyperplanes in one of the corresponding families of $\text{Ess}(\mathcal{E})$. On the other hand, each G_e -essential hyperplane intersects \mathcal{E} along one of these lines, so these are the only straight subgroups. \square

4.2. Cocompact cores.

Lemma 4.5. *For each edge e of Γ , there exists a convex G_e -cocompact subcomplex $\widehat{\mathcal{Y}}_e \subset \mathcal{X}$.*

Proof. As in the proof of Proposition 4.4, the Flat Torus Theorem yields a G_e -cocompact copy \mathcal{E} of \mathbb{E}^2 isometrically embedded in \mathcal{X} . We apply Lemma 3.6 as before and conclude that we have case (2), hence we can apply Lemma 3.8. Case (1) of that lemma leads to a contradiction with Lemma 4.1. Hence we have case (2) saying that there exists a G_e -cocompact convex subcomplex $\widehat{\mathcal{Y}}_e \subset \mathcal{X}$. \square

Proposition 4.6. *For each vertex v of Γ , there exists a convex G_v -cocompact subcomplex $\mathcal{V}_v \subset \mathcal{X}$.*

Given a subcomplex $\mathcal{Y} \subset \mathcal{X}$, let $\mathcal{Y}^{+0} = \mathcal{Y}$ and let the *cubical k -neighborhood* \mathcal{Y}^{+k} be the union of all closed cubes intersecting $\mathcal{Y}^{+(k-1)}$. If $\mathcal{Y} \subset \mathcal{X}$ is convex, then \mathcal{Y}^{+k} is convex as well.

Proof. Fix a base vertex $x \in \mathcal{X}$. By Lemma 4.5, there is a G -equivariant family of G_e -cocompact convex subcomplexes $\widehat{\mathcal{Y}}_e \subset \mathcal{X}$, where e varies over the set of edges incident to v . Let $\phi: \widetilde{M} \rightarrow \mathcal{X}$ be a G -equivariant quasi-isometry. Each $\widetilde{M} - \mathcal{N}_k(\widetilde{T}_e)$ has two \widetilde{T}_e -essential components. Let \overleftarrow{M}_e^k be the \widetilde{T}_e -essential component disjoint from \widetilde{B}_v and let $\overrightarrow{M}_e^k = \widetilde{M} - \overleftarrow{M}_e^k$. In other words, the space \widetilde{M} has two *poles* with respect to each \widetilde{T}_e (see [CP11, Appendix A]). Poles are quasi-isometry invariants [CP11, Lem A.2], so by [CP11, Lem A.4], there is k such that $\mathcal{X} - \widehat{\mathcal{Y}}_e^{+k}$ has two $\widehat{\mathcal{Y}}_e$ -essential components, one of which, denoted by $\overleftarrow{\mathcal{X}}_e^k$, is disjoint from the orbit $G_v x \subset \mathcal{X}$. Moreover, there exists K such that $\phi(\overleftarrow{M}_e^K) \subset \overleftarrow{\mathcal{X}}_e^k$; see [CP11, Sublem A.3]. Let $\overrightarrow{\mathcal{X}}_e^k = \mathcal{X} - \overleftarrow{\mathcal{X}}_e^k$. Let $\mathcal{V}_v \subset \mathcal{X}$ be the convex G_v -invariant subcomplex $\bigcap_e \overrightarrow{\mathcal{X}}_e^k$, which contains x . Let $K' \geq K$ be such that $\mathcal{N}_{K'}(\widetilde{B}_v)$ contains all complementary components of $\mathcal{N}_K(\widetilde{T}_e)$ that are not \widetilde{T}_e -essential. Then we have $\phi^{-1}(\mathcal{V}_v) \subset \bigcap_e \overrightarrow{M}_e^{K'} \subset \mathcal{N}_{K'}(\widetilde{B}_v)$. Since $\mathcal{N}_{K'}(\widetilde{B}_v)$ is G_v -cocompact, the complex \mathcal{V}_v is G_v -cocompact as well. \square

The final preparatory lemma requires using two results from [CS11].

Lemma 4.7. *Consider the product of the free cyclic and a finitely generated free non-abelian group $H = \mathbb{Z} \times \mathbb{F}$. Suppose that H acts freely and cocompactly on a CAT(0) cube complex \mathcal{V} . Then the following hold:*

- (1) *The essential core \mathcal{V}^{ess} of \mathcal{V} is a product $\mathcal{V}_a \times \mathcal{V}_b$, where $\mathcal{V}_a, \mathcal{V}_b$ are unbounded.*
- (2) *The group H has a finite-index subgroup $H' = H_a \times H_b$ that preserves the above decomposition, where H_a acts trivially on \mathcal{V}_b and H_b acts trivially on \mathcal{V}_a .*
- (3) *We have $H_a = \mathbb{Z} \cap H'$ and the group H_b embeds as a finite-index subgroup of the free group $H/\mathbb{Z} \cong \mathbb{F}$ under the natural quotient.*

Proof. Since H is a direct product with infinite factors, none of its elements is rank 1 in the action on \mathcal{V}^{ess} . Hence [CS11, Cor 6.4(iii)] implies assertion (1). By [CS11, Prop 2.6] the group $\text{Aut}(\mathcal{V}^{ess})$ contains $\text{Aut}(\mathcal{V}_a) \times \text{Aut}(\mathcal{V}_b)$ as a finite-index subgroup, from which assertion (2) follows. Since each of the subcomplexes $\mathcal{V}_a, \mathcal{V}_b$ is unbounded, H_a, H_b are non-trivial.

Any finite-index subgroup $H' \leq H$ has the form $\mathbb{Z}' \times \mathbb{F}'$, where $\mathbb{Z}' = \mathbb{Z} \cap H'$. We will identify one of H_a, H_b as \mathbb{Z}' . First note that either H_a or H_b has trivial center. Otherwise we would have \mathbb{Z}^2 in the center of $\mathbb{Z}' \times \mathbb{F}'$. Secondly, consider $z \in \mathbb{Z}' = \mathbb{Z} \cap H'$. Then $z \in H_a$ or $z \in H_b$, since otherwise the projections of z to H_a, H_b would both be central. Without loss of generality we assume $\mathbb{Z}' \leq H_a$. Conversely, an element $h \in H_a$ commutes with both \mathbb{Z}' and H_b . Since the only elements with non-cyclic centralizer in $\mathbb{Z}' \times \mathbb{F}'$ are in \mathbb{Z}' , we have $h \in \mathbb{Z}'$, establishing $H_a = \mathbb{Z}'$. Hence H_b embeds in the quotient H/\mathbb{Z} . Finally, H_b has finite index in H/\mathbb{Z} since H' has finite index in H . This establishes assertion (3). \square

4.3. Vanishing of charges.

Proof of Theorem B. Suppose $G = \pi_1 M$ acts freely and cocompactly on a CAT(0) cube complex \mathcal{X} . By [LW97, Prop 4.4], after passing to a finite cover we can assume that the blocks $B_{\bar{v}}$ of M have no singular fibers. In other words their base orbifolds $F_{\bar{v}}$ are surfaces and consequently their fundamental groups are free. Hence $G_v = \mathbb{Z}_v \times \mathbb{F}_v$, where \mathbb{Z}_v is the cyclic fiber group and \mathbb{F}_v is a free group.

By Proposition 4.6, there is a convex G_v -cocompact subcomplex $\mathcal{V}_v \subset \mathcal{X}$. We consider the essential core \mathcal{V}_v^{ess} of \mathcal{V}_v , which decomposes as $\mathcal{V}_a \times \mathcal{V}_b$ by Lemma 4.7(1) applied to $H = G_v$. Let $H' = H_a \times H_b$ be the subgroups provided by Lemma 4.7(2).

Let \bar{v} be the image of v in $\bar{\Gamma}$ and let $F'_{\bar{v}} \rightarrow F_{\bar{v}}$ be the finite cover of the base surface $F_{\bar{v}}$ corresponding to the embedding $H_b \leq G_v/\mathbb{Z}_v \cong \pi_1 F_{\bar{v}}$ coming from Lemma 4.7.(3). Let \mathcal{C} be the collection of boundary curves of $F'_{\bar{v}}$. For each $c \in \mathcal{C}$, choose an edge e of Γ incident to v so that $G_e \cap H_b < H_b \cong \pi_1 F'_{\bar{v}}$ is a conjugate of $\pi_1 c$. The inclusion $G_e \cap H_b < G_e$ can be considered as inclusion on homology $\mathbb{H}_1(G_e \cap H_b) < \mathbb{H}_1(G_e) = \mathbb{H}_1(T_{\bar{e}})$. Choose an embedded curve $\check{c} \subset T_{\bar{e}}$, non-vertical in $B_{\bar{v}}$, and an integer n_c such that $n_c[\check{c}]$ is the image of $[c] \in \mathbb{H}_1(G_e \cap H_b)$ under this inclusion. Let $\phi: \mathbb{H}_1(H_b) \rightarrow \mathbb{H}_1(G_v) = \mathbb{H}_1(B_{\bar{v}})$ be induced by $H_b < G_v$. Hence $\phi(\sum_{c \in \mathcal{C}} [c]) = \sum_{c \in \mathcal{C}} n_c[\check{c}]$. Since \mathcal{C} is the collection of boundary curves of a surface, we have $\sum_c [c] = 0$ in $\mathbb{H}_1(H_b)$, whence $\sum_c n_c[\check{c}] = 0$ in $\mathbb{H}_1(B_{\bar{v}})$. To show that M is chargeless, that is to verify the analogous formula from Definition 1.1, it remains to show that \check{c} is homotopic to $Z_{\bar{v}'}^{\bar{e}}$.

Claim. *Let $\mathbb{Z}_a, \mathbb{Z}_b < G_e$ be the straight subgroups coming from Proposition 4.4. Then the groups $H_a, G_e \cap H_b$ are finite-index subgroups of $\mathbb{Z}_a, \mathbb{Z}_b$ or vice-versa.*

Before we justify the claim, we explain how to use it to complete the proof. The claim yields $H_a \leq \mathbb{Z}_a$ and since $H_a \leq \mathbb{Z}_v$ by Lemma 4.7(3), we obtain $\mathbb{Z}_a = \mathbb{Z}_v$. Moreover, the claim gives that the image of $[c] \in \mathbb{H}_1(G_e \cap H_b)$ in $\mathbb{H}_1(G_e) = G_e$ lies in \mathbb{Z}_b , whence $\langle [\check{c}] \rangle = \mathbb{Z}_b$. If $e = (v, v')$ and we apply the claim to e and v' , we similarly obtain that either \mathbb{Z}_a or \mathbb{Z}_b coincides with $\mathbb{Z}_{v'}$. Since $\mathbb{Z}_{v'} \cap \mathbb{Z}_v = \{1\}$, we have $\mathbb{Z}_b = \mathbb{Z}_{v'}$. For Definition 1.1 the circle $Z_{\bar{v}'}^{\bar{e}}$ was defined to satisfy $\langle [Z_{\bar{v}'}^{\bar{e}}] \rangle = \mathbb{Z}_{v'}$ in $\mathbb{H}_1(T_{\bar{e}}) = G_e$. Thus we have $\langle [\check{c}] \rangle = \mathbb{Z}_b = \mathbb{Z}_{v'} = \langle [Z_{\bar{v}'}^{\bar{e}}] \rangle$, which ends the proof of the theorem.

It remains to justify the claim. By Lemma 4.7(3) we have $H_a = \mathbb{Z}_v \cap H'$ and hence $G_e \cap H' = H_a \times (G_e \cap H_b)$, which is a product of cyclic groups. Thus it suffices to find nontrivial elements in $\mathbb{Z}_a \cap H_a, \mathbb{Z}_b \cap (G_e \cap H_b)$. Since \mathbb{Z}_a is straight, it contains a nontrivial element h stabilizing a G_e -essential hyperplane of \mathcal{X} . Since $H' \leq H$ has finite index, after passing to a power we can assume $h \in H_a \times (G_e \cap H_b)$. Since $G_e < G_v$, the corresponding hyperplane in the G_v -essential core \mathcal{V}_v^{ess} of \mathcal{V}_v is G_e -essential as well. Without loss of generality assume that h stabilizes a hyperplane $\mathcal{V}_a \times \mathfrak{h}$ where \mathfrak{h} is a hyperplane in \mathcal{V}_b . Then $h = a \times b$, where b stabilizes \mathfrak{h} . Since $1 \times b$ and $H_a \times 1$ stabilize the G_e -essential hyperplane $\mathcal{V}_a \times \mathfrak{h}$, they cannot generate a finite-index subgroup of G_e , whence $b = 1$. Thus $h \in H_a$, as desired, and using the same argument we find a nontrivial element in $\mathbb{Z}_b \cap (G_e \cap H_b)$. \square

REFERENCES

- [AFW12] Matthias Aschenbrenner, Stefan Friedl, and Henry Wilton. 3-manifold groups. 2012. ArXiv e-print 1205.0202.
- [BH99] Martin R. Bridson and Andr e Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.

- [Bon02] Francis Bonahon. Geometric structures on 3-manifolds. In *Handbook of geometric topology*, pages 93–164. North-Holland, Amsterdam, 2002.
- [CP11] Pierre-Emmanuel Caprace and Piotr Przytycki. Bipolar Coxeter groups. *J. Algebra*, 338(1):35–55, 2011.
- [CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21:851–891, 2011.
- [Hag07] Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. 2007. ArXiv e-print 0705.3386.
- [HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1 551–1620, 2008.
- [HW12] G. Christopher Hruska and Daniel T. Wise. Finiteness properties of cubulated groups. 2012. ArXiv e-print 1209.1074.
- [Liu13] Yi Liu. Virtual cubulation of nonpositively curved graph manifolds. *Journal of Topology*, 6(4):793–822, 2013.
- [LW97] John Luecke and Ying-Qing Wu. Relative Euler number and finite covers of graph manifolds. In *Geometric topology (Athens, GA, 1993)*, volume 2 of *AMS/IP Stud. Adv. Math.*, pages 80–103. Amer. Math. Soc., Providence, RI, 1997.
- [PW13] Piotr Przytycki and Daniel T. Wise. Graph manifolds with boundary are virtually special. *Journal of Topology*, 2013. To appear.
- [Sag97] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.
- [Sag12] Michah Sageev. CAT(0) complexes and groups. 2012. Notes to PCMI lecture course, to appear.
- [Wis12] Daniel T. Wise. Cubular tubular groups. 2012. Preprint available at <http://www.math.mcgill.ca/wise/papers.html>.
- [WY97] Shicheng Wang and Fengchun Yu. Graph manifolds with non-empty boundary are covered by surface bundles. *Math. Proc. Cambridge Philos. Soc.*, 122(3):447–455, 1997.

DEPT. OF MATH., UNIVERSITY OF MICHIGAN, ANN ARBOR, MI, USA
E-mail address: markfhagen@gmail.com

INSTITUTE OF MATHEMATICS, POLISH ACADEMY OF SCIENCES, WARSAW, POLAND
E-mail address: pprzytyc@mimuw.edu.pl