

BOWDITCH'S JSJ TREE AND THE QUASI-ISOMETRY CLASSIFICATION OF CERTAIN COXETER GROUPS

PALLAVI DANI AND ANNE THOMAS

ABSTRACT. Bowditch's JSJ tree for splittings over 2-ended subgroups is a quasi-isometry invariant for 1-ended hyperbolic groups which are not cocompact Fuchsian [9]. Our main result gives an explicit "visual" construction of this tree for certain hyperbolic right-angled Coxeter groups. As an application of our construction we identify a large class of such groups for which the JSJ tree, and hence the visual boundary, is a complete quasi-isometry invariant. We also show that among the Coxeter groups we consider, the cocompact Fuchsian groups form a rigid quasi-isometry class.

1. INTRODUCTION

In this paper we use Bowditch's JSJ tree to investigate the quasi-isometry classification of certain right-angled Coxeter groups. After describing our results, we review previous work on this classification, about which surprisingly little is known.

Given a finite simplicial graph Γ , the associated *right-angled Coxeter group* W_Γ has generating set S equal to the vertices of Γ , and relations $s^2 = 1$ for all $s \in S$, and $st = ts$ whenever $s, t \in S$ are adjacent vertices. We call Γ the *defining graph* of W_Γ . We assume throughout that W_Γ is infinite, or equivalently that Γ is not a complete graph. We also assume that W_Γ is *2-dimensional*, that is, that Γ has no triangles. This is equivalent to the Davis complex for W_Γ being 2-dimensional (see Section 2.1 for background on the Davis complex).

Recall that a *special subgroup* of W_Γ is one generated by a subset of the generating set S . For each $T \subseteq S$, the special subgroup $\langle T \rangle$ is also a right-angled Coxeter group, whose defining graph has vertex set T and edge set consisting of all edges of Γ which have both endpoints in T . By Proposition 8.8.2 of [15], if W_Γ has infinitely many ends then W_Γ has a nontrivial decomposition as a finite tree of groups, in which each vertex group is a finite or 1-ended special subgroup and each edge group is a finite special subgroup. It then follows from Theorems 0.3 and 0.4 of Papasoglu and Whyte [26] that for W_Γ and W_Λ with infinitely many ends, W_Γ is quasi-isometric to W_Λ if and only if their respective tree-of-groups decompositions have the same set of quasi-isometry types of 1-ended vertex groups (without multiplicities). In light of these results, we focus on the case that W_Γ is 1-ended in this paper.

In [9], Bowditch gives a construction of a canonical JSJ splitting over 2-ended subgroups of any 1-ended hyperbolic group G which is not cocompact Fuchsian. (A finitely generated group is *Fuchsian* if it is non-elementary and acts properly discontinuously on the hyperbolic

Date: January 14, 2016.

The first author was supported in part by NSF Grant No. DMS-1207868. This research of the second author was supported in part by ARC Grant No. DP110100440. The second author was supported in part by an Australian Postdoctoral Fellowship.

plane.) The JSJ tree \mathcal{T}_G for this splitting is a quasi-isometry invariant, as it is defined in terms of the local cut point structure of the boundary. The group G acts on its JSJ tree \mathcal{T}_G with finite quotient. We recall the construction of Bowditch’s JSJ tree in Section 2.3.

Our main result gives an explicit construction of Bowditch’s JSJ tree \mathcal{T}_{W_Γ} for certain W_Γ ; see Theorem 1.2 below. We restrict our attention to the class of 2-dimensional, 1-ended, hyperbolic right-angled Coxeter groups W_Γ , which are not cocompact Fuchsian and which admit a non-trivial splitting over a 2-ended subgroup (if there is no such splitting then \mathcal{T}_{W_Γ} consists of a single vertex). Each of these assumptions corresponds to certain properties of the (triangle-free) graph Γ , as follows. First, by [15, Theorem 8.7.2], the group W_Γ is 1-ended if and only if Γ is connected and has no separating vertices or edges. Also, by [15, Corollary 12.6.3], the group W_Γ is hyperbolic if and only if Γ has no embedded cycles of length four.

A Coxeter group W (not necessarily right-angled) is cocompact Fuchsian if and only if W is either a hyperbolic polygon reflection group, or the direct product of a hyperbolic polygon reflection group with a finite Coxeter group. This result will not surprise experts, but as we could not find it stated explicitly in the literature, we record a proof in the Appendix, as Theorem A.2. For Γ triangle-free, this implies that W_Γ is cocompact Fuchsian if and only if Γ is a cycle of length ≥ 5 .

As for splittings, Papasoglu showed in [25] that among 1-ended, finitely presented groups that are not commensurable to surface groups, having a splitting over a 2-ended subgroup is a quasi-isometry invariant. In [24] Mihalik and Tschantz characterised the Coxeter groups which admit such splittings in terms of their defining graphs. A pair of vertices $\{a, b\}$ in Γ is called a *cut pair* if $\{a, b\}$ separates Γ , meaning that $\Gamma \setminus \{a, b\}$ has at least two components, each of which contains at least one vertex. In our setting, the main result of [24] states that W_Γ splits over a 2-ended subgroup if and only if there is a cut pair $\{a, b\}$ in Γ (see Section 2.2 for more details).

Correspondingly, we restrict to graphs Γ satisfying the following conditions.

Standing Assumptions 1.1. The graph Γ :

- (1) has no triangles (W_Γ is 2-dimensional);
- (2) is connected and has no separating vertices or edges (W_Γ is 1-ended);
- (3) has no squares (W_Γ is hyperbolic);
- (4) is not a cycle of length ≥ 5 (W_Γ is not cocompact Fuchsian); and
- (5) has a cut pair of vertices $\{a, b\}$ (W_Γ splits over a 2-ended subgroup).

The main result of this paper is a “visual” description of Bowditch’s JSJ tree for such W_Γ :

Theorem 1.2. *For Γ satisfying Standing Assumptions 1.1, the graph Γ visually determines Bowditch’s JSJ tree $\mathcal{T}_\Gamma = \mathcal{T}_{W_\Gamma}$, in the sense that the W_Γ -orbits of vertices and edges of \mathcal{T}_Γ are in bijection with subsets of vertices of the defining graph Γ which satisfy certain explicit graph-theoretic conditions. Moreover, the stabilisers of the vertices and edges of \mathcal{T}_Γ are conjugates of certain special subgroups of W_Γ .*

We prove Theorem 1.2 in Section 3, and an explicit statement of this theorem appears as Theorem 3.36. A consequence of our description is that there exists an algorithm to compute the JSJ tree of such a right-angled Coxeter group.

In order to prove Theorem 1.2, we relate the separation properties of subsets of vertices of Γ , of geodesics in the Davis complex, and of endpoints of geodesics in the visual boundary ∂W_Γ . The assumption that W_Γ is 2-dimensional is essential here, since geodesics can only separate the Davis complex if it is 2-dimensional. We also rely on Γ being triangle-free for many preliminary graph-theoretic results.

We are able to use a few arguments from Lafont’s papers [20, 21] in our proofs, but substantial additional work is required. The vertices of Bowditch’s JSJ tree \mathcal{T}_Γ are of three types. The first two types are defined in terms of subsets of ∂W_Γ , and include the so-called \approx -pairs and \sim -classes. The vertices of the third type are called stars, and are defined as subsets of vertices of the first two types. As we explain further in Remark 3.33, the \approx -pairs and certain \sim -classes in our setting correspond to key structures in the spaces investigated by Lafont in [20, 21]. Similar spaces have also been considered by Crisp–Paoluzzi [13], Malone [22] and Stark [27]. However the situation for \sim -classes is more varied in our setting. Moreover, there is no analogue of stars in [20, 21] and these other works. The subtleties of \sim -classes and the identification of stars are the most difficult parts of the proof of Theorem 1.2.

We give some applications of our main result in Section 4. First, we show that the JSJ tree is a complete invariant for a large subclass of groups. In the following statement, K_4 is the complete graph on four vertices.

Theorem 1.3. *Let \mathcal{G} be the class of graphs satisfying Standing Assumptions 1.1 which have no induced subgraphs which are subdivided copies of K_4 . Then Bowditch’s JSJ tree is a complete invariant for W_Γ with $\Gamma \in \mathcal{G}$. More precisely, given $\Lambda, \Gamma \in \mathcal{G}$, the groups W_Γ and W_Λ are quasi-isometric if and only if there is a type-preserving isomorphism from \mathcal{T}_Γ to \mathcal{T}_Λ .*

A result of Cashen [11] says that the JSJ tree is a complete invariant for any class of groups in which the JSJ tree has no stars. In Section 4.1, we prove Theorem 1.3 by characterising the groups W_Γ for which \mathcal{T}_Γ has no stars as exactly those for which the corresponding graph Γ has no subdivided K_4 subgraphs. We do not expect the JSJ tree to be a complete invariant in general (i.e. for all graphs satisfying Standing Assumptions 1.1).

As a consequence of Bowditch’s construction of the JSJ tree and Theorem 1.3, we also obtain that the visual boundary is a complete invariant for the same subclass of groups.

Corollary 1.4. *Let $\Gamma, \Lambda \in \mathcal{G}$ where \mathcal{G} is as defined in the statement of Theorem 1.3. Then W_Γ and W_Λ are quasi-isometric if and only if ∂W_Γ and ∂W_Λ are homeomorphic.*

The visual boundaries of hyperbolic groups have been investigated by Świątkowski and coauthors in a series of papers [23, 29, 30, 31, 32]. In particular, in [23] Martin and Świątkowski consider visual boundaries of fundamental groups of graphs of groups with finite edge groups. Suppose now that W_Γ and W_Λ are 2-dimensional hyperbolic right-angled Coxeter groups with infinitely many ends. When applied to the above-mentioned tree-of-groups decompositions for W_Γ and W_Λ , the results of [23] imply that ∂W_Γ and ∂W_Λ are homeomorphic if and only if the corresponding sets of homeomorphism types of visual boundaries of 1-ended vertex groups are the same. Our results thus have applications to determining the homeomorphism type of ∂W_Γ when W_Γ has infinitely many ends.

Our last application shows that among 2-dimensional right-angled Coxeter groups, the cocompact Fuchsian groups form a rigid quasi-isometry class.

Theorem 1.5. *Any 2-dimensional right-angled Coxeter group which is quasi-isometric to a cocompact Fuchsian group is cocompact Fuchsian.*

We give a precise statement as Theorem 4.2 and prove this result in Section 4.2. Although this result may already be known to experts, we could not find it in the literature.

We now discuss several quasi-isometry invariants that have previously been considered for right-angled Coxeter groups. In [14], we investigated *divergence*. Up to an equivalence relation which identifies polynomials of the same degree, the rate of divergence is a quasi-isometry invariant [18]. We characterised those W_Γ with linear and quadratic divergence by properties of their defining graphs and showed that for every positive integer d , there is a W_{Γ_d} with divergence polynomial of degree d .

Thickness, a quasi-isometry invariant related to divergence, was introduced by Behrstock, Druţu and Mosher in [4]. Behrstock–Druţu [3] show that thickness provides a polynomial upper bound on divergence. In [5], Behrstock, Hagen, Sisto and Caprace give an effective characterisation of the right-angled Coxeter groups W_Γ (not just the 2-dimensional ones) which are thick, and show that if W_Γ is not thick then it is hyperbolic relative to a collection of thick special subgroups, in which case its divergence is exponential. As discussed in [5], this relative hyperbolicity result combined with results from [4] and [17] can be used for quasi-isometry classification.

By Moussong’s Theorem [15, Theorem 12.3.3], right-angled Coxeter groups are CAT(0) groups. Although the visual boundary is not a quasi-isometry invariant for CAT(0) groups, Charney–Sultan [12] have introduced a boundary for CAT(0) spaces, called the *contracting boundary*, which is a quasi-isometry invariant. In Section 5 of [12], the contracting boundary is used to distinguish the quasi-isometry classes of two 1-ended non-hyperbolic right-angled Coxeter groups.

Finally, we remark that although every right-angled Artin group is a finite-index subgroup of some right-angled Coxeter group [16], there exist right-angled Coxeter groups which are not quasi-isometric to any right-angled Artin group. For example, there are no 1-ended hyperbolic right-angled Artin groups, since a right-angled Artin group is hyperbolic if and only if it is a free group. In fact, by considering divergence we see that there are infinitely many quasi-isometry classes of right-angled Coxeter groups which are not the quasi-isometry class of any right-angled Artin group: in [14] we constructed right-angled Coxeter groups with divergence polynomial of any degree, but the divergence of right-angled Artin groups can be only linear or quadratic [1, 2]. Moreover, the explicit construction given in [16] does not yield all of the right-angled Coxeter groups we identified as having quadratic divergence in [14]. Thus while results on the quasi-isometry classification of right-angled Artin groups (for instance [6, 7, 8, 19]) yield some information about right-angled Coxeter groups with certain specific defining graphs, they don’t readily lead to a quasi-isometry classification of right-angled Coxeter groups in general.

Acknowledgements. We thank Jason Behrstock, Chris Hruska, Misha Kapovich, Mike Mihalik and Kim Ruane for helpful conversations, and an anonymous referee for comments on an earlier version of this paper. We are grateful to Anthony Henderson for helping to support a visit by the first author to the University of Sydney, and to the University of

Glasgow for travel support for the second author. Some of the research for this paper was carried out at the MSRI Summer School in Geometric Group Theory in June 2015, and we thank MSRI for travel support for both authors.

2. BACKGROUND

In this section, we recall the Davis complex for W_Γ (Section 2.1), a splitting result due to Mihalik and Tschantz [24] (Section 2.2), and the construction and key properties of Bowditch’s JSJ tree (Section 2.3). We continue notation from the introduction and assume that the defining graph Γ satisfies the Standing Assumptions 1.1.

2.1. The Davis complex. A reference for the material in this section is [15].

Let Γ' be the barycentric subdivision of the defining graph Γ . Then the *chamber* K may be defined as the simplicial complex obtained by coning on Γ' . Since Γ is triangle-free, K is 2-dimensional. Denote by σ_0 the cone point of K . Since all vertices of Γ have valence at least two, the boundary of K may be identified with Γ (or with Γ'). We can then naturally define the *radial projection* of a path in $K \setminus \{\sigma_0\}$ to ∂K , and so obtain a path in Γ . Similarly, any point in K can be coned to σ_0 .

For each $s \in S$, the *mirror (of type s)* of the chamber K is the subset K_s of $\partial K = \Gamma'$ given by the star of the vertex s in the graph Γ' . The mirror K_s is thus the star graph of valence n , where $n = \text{card}\{t \in S \mid st = ts, t \neq s\} \geq 2$. Two mirrors K_s and K_t intersect (in a point) if and only if $st = ts$, and the boundary ∂K may be viewed as the union of mirrors $\cup_{s \in S} K_s$. For each $x \in K$ let $S(x) = \{s \in S \mid x \in K_s\} \subset S$. Denote by $W_{S(x)}$ the subgroup of W generated by the set $S(x)$, with $W_{S(x)}$ trivial if $S(x) = \emptyset$.

The *Davis complex* $\Sigma = \Sigma_\Gamma$ is defined to be the following 2-dimensional simplicial complex:

$$\Sigma = W \times K / \sim$$

where $(w, x) \sim (w', x')$ if and only if $x = x'$ and $w^{-1}w' \in W_{S(x)}$. We define a *chamber* of Σ to be any copy of K in Σ , and a *panel (of type s)* in Σ to be any copy of the mirror K_s . We write $w\sigma_0$ or just w for the image of the cone point σ_0 of the chamber $K = (e, K)$ under $w \in W$. Two chambers in Σ are then adjacent along a panel of type s if and only if their cone points w and w' satisfy $w^{-1}w' = s$.

Since Γ is triangle-free, we can now re-cellulate Σ so that its vertex set is W , its 1-skeleton is the Cayley graph \mathcal{C}_Γ of W with respect to the generating set S , and all 2-cells in Σ are squares with boundary word $stst = 1$, where s and t are commuting generators. We call this cellulation of Σ the *cellulation by big squares*, with the *big squares* being the 2-cells. We also define the *cellulation by small squares* of Σ to be the first square subdivision of the cellulation by big squares, with the *small squares* being the squares obtained on subdividing each big square into four. In the cellulation by small squares, the chamber K is then the “cubical cone” on the barycentric subdivision Γ' of Γ .

The Davis complex Σ , with the cellulation by either big or small squares, may now be metrisised so that each big square is a unit Euclidean square. By [15, Theorem 12.2.1], this piecewise Euclidean structure is CAT(0). In this metrisation, a geodesic γ in the Cayley graph $\mathcal{C} = \mathcal{C}_\Gamma$ is a geodesic in the Davis complex Σ if and only if no two successive labels a and b of γ are adjacent vertices in Γ . Also, each induced n -cycle in Γ , where $n \geq 5$ since Γ

is triangle- and square-free, corresponds to a family of convex subcomplexes of Σ which are planes tiled by big squares so that n big squares meet at each vertex. Since $n \geq 5$, each such subcomplex is quasi-isometric to the hyperbolic plane.

2.2. Splittings of right-angled Coxeter groups over 2-ended subgroups. Bowditch's JSJ tree has at least one edge if and only if G admits a splitting over a 2-ended subgroup. As mentioned in the introduction, Mihalik and Tschantz characterised the Coxeter groups which admit such splittings in [24]. We now state their result in our setting.

Let Γ be a finite simplicial graph which is triangle-free and has no separating vertices or edges, i.e. Γ satisfies (1) and (2) from Standing Assumptions 1.1. It follows that if $\{a, b\}$ is a cut pair of Γ , then a and b are non-adjacent, and therefore generate a 2-ended special subgroup. For such Γ , the relevant result from [24] can be stated as follows:

Theorem 2.1. [24] *The right-angled Coxeter group W_Γ defined by Γ as above splits over a 2-ended subgroup H if and only if there is a cut pair $\{a, b\}$ in Γ ; moreover, there is a unique cut pair $\{a, b\}$ so that some conjugate of H contains the special subgroup $\langle a, b \rangle$, necessarily with finite index, and W_Γ also splits over $\langle a, b \rangle$.*

Thus up to conjugacy and finite index, all splittings of W_Γ over 2-ended subgroups are splittings over special subgroups generated by cut pairs.

2.3. Bowditch's JSJ tree. Let G be a 1-ended hyperbolic group which is not cocompact Fuchsian (as defined in the introduction). For such G , Bowditch [9] uses the structure of local cut points of its boundary $M = \partial G$ to define a canonical JSJ tree $\mathcal{T} = \mathcal{T}_G$ associated to G . We now recall the construction and key properties of this tree \mathcal{T} from [9].

We begin with some terminology from [9]. Given $x \in M$, define the *valency* $\text{val}(x)$ of x to be the number of ends of the locally compact space $M \setminus \{x\}$. A priori $\text{val}(x) \in \mathbb{N} \cup \{\infty\}$, but Bowditch shows [9, Proposition 5.5] that if M is the boundary of a 1-ended hyperbolic group, then $\text{val}(x)$ is finite for all $x \in M$. The point x is a *local cut point* if $\text{val}(x) \geq 2$.

Now given $n \in \mathbb{N}$, let $M(n) = \{x \in M \mid \text{val}(x) = n\}$ and $M(n+) = \{x \in M \mid \text{val}(x) \geq n\}$. Bowditch defines relations \sim and \approx on $M(2)$ and $M(3+)$ respectively, as in Definitions 2.2 and 2.3 below. For $x, y \in M$, let $N(x, y)$ be the number of components of $M \setminus \{x, y\}$.

Definition 2.2. [The relation \sim] Given $x, y \in M(2)$, let $x \sim y$ if and only if either $x = y$ or $N(x, y) = 2$.

The following are some properties of \sim .

- (1) The relation \sim is an equivalence relation on $M(2)$ [9, Lemma 3.1].
- (2) Every point of $M(2)$ is in some \sim -class by definition.
- (3) The \sim -equivalence classes are of two types: either they are pairs or they are infinite. Moreover, the infinite ones are Cantor sets [9, Corollary 5.15 and Proposition 5.18].

Definition 2.3. [The relation \approx] Given $x, y \in M(3+)$, let $x \approx y$ if $x \neq y$ and $N(x, y) = \text{val}(x) = \text{val}(y) \geq 3$.

The following are some properties of \approx .

- (1) If $x \approx y$ and $x \approx z$ then $y = z$ [9, Lemma 3.8]. Thus a priori, a subset of $M(3+)$ is partitioned into pairs of the form $\{x, y\}$ where $x \approx y$.

(2) Lemma 5.12 and Proposition 5.13 of [9] show that in fact, all of $M(3+)$ is partitioned into \approx -pairs.

Remark 2.4. By the above properties, it is evident that $M(2+)$ is the disjoint union of the \sim -classes and the \approx -pairs in M .

Outline of the construction. The JSJ tree \mathcal{T} is produced as follows. Denote by T the set of \sim -classes and \approx -pairs.

Definition 2.5. [Betweenness] Given three classes $\eta, \theta, \zeta \in T$, the class θ is *between* η and ζ if there exist points $x \in \eta$ and $y \in \zeta$, and distinct points $a, b \in \theta$, so that x and y are separated in M by $\{a, b\}$.

Remark 2.6. By Lemma 3.18 of [9], if η and θ are distinct classes in T , then for all distinct pairs of points $a, b \in \theta$, the set η is contained in a single component of $M \setminus \{a, b\}$. Hence θ is between η and ζ if and only if there exist distinct points $a, b \in \theta \subset M$ and distinct components U and V of $M \setminus \{a, b\}$ such that $\eta \subseteq U$ and $\zeta \subseteq V$.

Bowditch shows in [9] that the set T forms a *pretree* (i.e. a set with the betweenness condition satisfying certain axioms) which is *discrete* (i.e. intervals between points are finite). In [10], Bowditch proves that every discrete pretree can be embedded in a *discrete median pretree* by adding in all possible *stars* of size at least 3.

Definition 2.7. [Stars] A *star* is a subset X of T with the property that no element of T is between any pair of elements in X , and which is maximal with respect to this property. Thus if $\eta \in T \setminus X$, then there exist $\zeta \in X$ and $\theta \in T$ such that θ is between η and ζ . The *size* of a star is its cardinality.

Any discrete median pretree can be realised as the vertex set of a simplicial tree [10]. The simplicial tree obtained in this way from T is essentially the JSJ tree \mathcal{T} , although certain additional vertices are added at the midpoints of some of the edges in order to get a cleaner statement about stabilisers.

The JSJ tree. The properties of this tree \mathcal{T} are summarised in Theorems 0.1 and 5.28 of [9] and are explained in more detail in Sections 3 and 5 of that paper. The group G acts minimally, simplicially and without edge inversions on \mathcal{T} with finite quotient graph.

2.3.1. *Vertices.* The tree \mathcal{T} has vertices of three types: $V_1(\mathcal{T})$, $V_2(\mathcal{T})$ and $V_3(\mathcal{T})$.

(1) *The vertices $V_1(\mathcal{T})$.* The set of Type 1 vertices consists of:

- the \sim -classes in $M(2)$ consisting of exactly two elements;
- all the \approx -pairs; and
- the extra vertices mentioned at the end of the previous section, added at the midpoints of any edges between stars of size at least 3 and infinite \sim -classes.

For all $v \in V_1(\mathcal{T})$, the stabiliser $G(v)$ of v is a maximal 2-ended subgroup of G . If v comes from a \sim - or a \approx -pair $\{x, y\}$, its degree in \mathcal{T} is equal to $N(x, y) < \infty$. The degree of the added vertices is 2 since they subdivide edges.

- (2) *The vertices $V_2(\mathcal{T})$.* The set of Type 2 vertices is in bijective correspondence with the collection of infinite \sim -classes in $M(2)$.

The stabiliser $G(v)$ of a vertex $v \in V_2(\mathcal{T})$ is a *maximal hanging Fuchsian (MHF)* subgroup of G . This means that there is a properly discontinuous action of $G(v)$ on the hyperbolic plane \mathbb{H}^2 , without parabolics, such that the quotient $\mathbb{H}^2/G(v)$ is non-compact and further, there is an equivariant homeomorphism from the \sim -class corresponding to v onto the limit set of the G -action in $\partial\mathbb{H}^2$. The degree of a vertex in $V_2(\mathcal{T})$ is infinite.

- (3) *The vertices $V_3(\mathcal{T})$.* The set of Type 3 vertices consists of the added stars of size at least 3 mentioned in the outline of the construction. The stabiliser of a vertex of this type is an infinite non-elementary group which is not a maximal hanging Fuchsian subgroup [9, Lemma 5.27]. The degree of a vertex of this type is infinite. From the description of the edges below, one sees that the elements in the star are in bijection with the edges incident to the corresponding vertex. In particular, stars of size at least 3 are infinite.

2.3.2. *Edges.* Given $v_1 \in V_1(\mathcal{T})$ coming from an \approx -pair ζ , and $v_2 \in V_2(\mathcal{T})$ coming from an infinite \sim -class θ , there is an edge between v_1 and v_2 if $\zeta \subset \bar{\theta}$. Every edge of the pretree T is of this form. The tree \mathcal{T} has the following additional edges. Given $v_3 \in V_3(\mathcal{T})$, there is an edge between v_3 and each vertex of $V_1(\mathcal{T}) \cup V_2(\mathcal{T})$ in the star corresponding to v_3 . Finally, if there are edges between vertices of $V_3(\mathcal{T})$ and $V_2(\mathcal{T})$, they are subdivided by adding a valence 2 vertex in the middle. It is evident from this construction that the two endpoints of an edge of \mathcal{T} are never of the same type, and that \sim -pair vertices (which are in $V_1(\mathcal{T})$) are necessarily connected to vertices in $V_3(\mathcal{T})$. The stabiliser of every edge is a 2-ended subgroup, and is the intersection of the stabilisers of the vertex groups incident to this edge.

Remark 2.8. (Quasi-isometry invariance) Let G and H be 1-ended hyperbolic groups which are not cocompact Fuchsian. By the construction above if G and H are quasi-isometric, then since ∂G and ∂H are homeomorphic, the JSJ trees \mathcal{T}_G and \mathcal{T}_H are equal as coloured trees. That is, the JSJ tree is a quasi-isometry invariant.

3. JSJ TREE FOR CERTAIN RIGHT-ANGLED COXETER GROUPS

In this section, we prove our main result, Theorem 1.2, by giving an explicit construction of the Bowditch JSJ tree $\mathcal{T} = \mathcal{T}_\Gamma$ associated to a right-angled Coxeter group $W = W_\Gamma$ such that Γ satisfies our Standing Assumptions 1.1. That is, we show that for each class of vertices in \mathcal{T} , the W -orbits of this class are in bijection with subsets of vertices of Γ which satisfy certain explicit graph-theoretic properties.

Throughout this section, we continue the notation of Section 2, and we assume that Γ satisfies the Standing Assumptions 1.1. We begin with some conventions and definitions in Section 3.1 and preliminary results in Section 3.2. In Section 3.3, we identify the \approx -pairs and their stabilisers. We construct certain \sim -classes in Section 3.4. In Section 3.5 we show that we have identified all \sim -classes, and so in fact constructed all the vertices of the pre-tree T . We also determine the stabilisers of \sim -classes. In Section 3.6 we construct certain stars,

and in Section 3.7 prove that we have identified all the stars and determine their stabilisers. Finally, we summarise the description of \mathcal{T} in terms of Γ in Theorem 3.36 of Section 3.8.

3.1. Conventions and definitions. Throughout Section 3, we denote the Davis complex of W by Σ . Recall that the 1-skeleton of the cellulation of Σ by big squares can be identified with the Cayley graph of W . We use \mathcal{C} to denote this Cayley graph, together with a choice of base point in Σ , to be labelled by the identity $e \in W$. Now we identify $\partial W = \partial \Sigma$ with the set of geodesic rays in \mathcal{C} emanating from e . However we also freely think of endpoints of bi-infinite geodesics in \mathcal{C} not passing through e as points of ∂W . Given any bi-infinite geodesic γ in \mathcal{C} , we denote its endpoints in ∂W by γ^+ and γ^- , and put $\partial\gamma = \{\gamma^+, \gamma^-\}$.

Recall the definition of a *cut pair* from the introduction. A vertex of Γ is *essential* if it has valence at least three, and a cut pair is *essential* if its two vertices are essential. A *reduced path* is a path in Γ or in \mathcal{C} which does not contain any loops or any backtracking. (Given any path in Γ or in \mathcal{C} , if such loops or backtracking exist, they can be cut out, eventually leading to a reduced path.) Note that a reduced path does not have to be a shortest path. Given a path τ in Γ or \mathcal{C} , the subpath between two vertices p and q on τ is denoted by $\tau_{[p,q]}$.

A geodesic γ in \mathcal{C} is *bicoloured (by a and b)* if it is labelled alternately by a pair of (nonadjacent) vertices a, b of Γ . If a geodesic γ is bicoloured by a and b we may say that it is *(a, b) -bicoloured*.

3.2. Preliminary results. In this section we show that endpoints of bicoloured geodesics are distinguished elements of ∂W . We also establish two results which relate separation properties of subsets of vertices in Γ , of subcomplexes of Σ , and of subsets of ∂W . The proofs in this section use some standard facts about reduced words in (right-angled) Coxeter groups, from [15].

For the proofs of Lemma 3.1 and Corollary 3.2 below, we do not need all of the Standing Assumptions 1.1. Indeed, in Lemma 3.1 the conclusion that η is eventually (a, b) -bicoloured holds for all right-angled Coxeter groups, and the entire conclusion of Lemma 3.1 and of Corollary 3.2 holds for all hyperbolic right-angled Coxeter groups.

Lemma 3.1. *Let γ be an (a, b) -bicoloured geodesic in \mathcal{C} , and let η be any other geodesic in \mathcal{C} . If $\gamma^+ = \eta^+$, then in the direction of η^+ , the geodesic η is eventually (a, b) -bicoloured, and the geodesic η eventually coincides with either γ or γc , with the latter case occurring only if there is a (unique) vertex c of $\Gamma \setminus \{a, b\}$ which is adjacent to both a and b .*

Proof. Let μ be a shortest path in \mathcal{C} from γ to η , and let w be the (possibly empty) word labelling μ . Then μ meets γ only at its starting point. We may assume, without loss of generality, that this point is e (so in particular, γ passes through e). Write $|w|$ for the length of w .

Since $\gamma^+ = \eta^+$, there is a $D > 0$ so that for all $n \geq 1$, there is a shortest path μ_n in \mathcal{C} from the vertex $(ab)^n$ of γ (without loss of generality) to some vertex of η , such that μ_n has length at most D . Let w_n be the word labelling the path μ_n . Now only the endpoints of the paths μ and μ_n lie on η , so if w'_n is the word labelling the subpath of η from the endpoint of μ to the endpoint of μ_n , the concatenation ww'_n is reduced.

By construction, the words ww'_n and $(ab)^nw_n$ are the labels of paths in \mathcal{C} from e to the same vertex (the endpoint of μ_n). Hence $ww'_n = (ab)^nw_n$ in W . The right-hand concatenation

$(ab)^n w_n$ might not be reduced, but we have $(ab)^n w_n = z_n \overline{w_n}$, where z_n is an initial subword of $(ab)^n$, $\overline{w_n}$ is a subword of w_n , and $z_n \overline{w_n}$ is reduced. Thus ww'_n and $z_n \overline{w_n}$ are reduced words representing the same element of W .

We first claim that in the direction of η^+ , infinitely many edge-labels of η are drawn from the set $\{a, b\}$. If not, then each word w'_n contains at most N instances of a or b , say. Hence for all n , the number of instances of a or b in the reduced word ww'_n is bounded above by $|w| + N$. However for n large enough, there are more than $|w| + N$ instances of a and b in z_n . Any two reduced words for the same element of a Coxeter group have the same letters, so this is a contradiction. This proves the claim.

Now suppose that η in the direction of η^+ is not eventually (a, b) -bicoloured. Then in the direction of η^+ , starting from the endpoint of μ , the labels on η are $g_1 h_1 g_2 h_2 \dots$, where each g_i is a reduced word in a and b , with g_i nonempty for $i \geq 2$, and each h_i is a nonempty reduced word in $\Gamma \setminus \{a, b\}$. For n large enough, the word w'_n contains $g_1 h_1 \dots g_{D+1} h_{D+1}$, and so the reduced word ww'_n contains more than D letters in $\Gamma \setminus \{a, b\}$. However the reduced word $z_n \overline{w_n}$ contains at most D letters in $\Gamma \setminus \{a, b\}$, and again we obtain a contradiction.

We have shown that in the direction of η^+ , the geodesic η is eventually (a, b) -bicoloured. Assume that in this direction, the geodesics η and γ do not eventually coincide. Then for all large enough n , the path μ_n is not contained in the geodesic γ , and this path ends in the (a, b) -bicoloured subray of η in the direction of η^+ . Thus the word $\overline{w_n}$ is nonempty and its first and last letters are not a or b . Fix such a large enough n , and choose $m > n$ so that w'_m has w'_n as an initial subword. Then the word $w'_{n,m}$ obtained by cancelling w'_n from the start of w'_m is a word in a and b , and for large enough m , we may assume that $w'_{n,m}$ contains both a and b . Similarly, if $z_{n,m}$ denotes the word obtained by cancelling z_n from the start of z_m , then $z_{n,m}$ is a word in a and b which we may assume contains both a and b . Denote by g_n the vertex of γ at which the subpath of μ_n labelled by $\overline{w_n}$ begins. By considering paths in \mathcal{C} from g_n to the endpoint of μ_m , we obtain that $\overline{w_n} w'_{n,m} = z_{n,m} \overline{w_m}$ in W , and both of these products are reduced.

Now we use the fact that any two reduced words representing the same element in a right-angled Coxeter group are related by a sequence of moves replacing a subword st by a subword ts , where s and t are commuting generators. Since both of the words $w'_{n,m}$ and $z_{n,m}$ contain both a and b , it follows that every letter in $\overline{w_n}$ commutes with both a and b , or is equal to a or b . The graph Γ is square-free so there is at most one vertex c of $\Gamma \setminus \{a, b\}$ which commutes with both a and b . As $\overline{w_n}$ is nonempty and reduced, and cannot start or end with a or b , we deduce that $\overline{w_n} = c$. Note that this holds for all large enough n .

To complete the proof, observe that for all large enough n , the geodesic γc is the unique (a, b) -bicoloured geodesic passing through the vertex $g_n c$. Since the vertices g_n approach $\eta^+ = \gamma^+$, and $\overline{w_n} = c$ for all large enough n , it follows that the geodesic η in the direction of η^+ eventually coincides with γc , as required. \square

Corollary 3.2. *Let γ be an (a, b) -bicoloured geodesic in \mathcal{C} , and let γ' be an (a', b') -bicoloured geodesic in \mathcal{C} . If $\partial\gamma = \partial\gamma'$ then $\{a, b\} = \{a', b'\}$, and either $\gamma = \gamma'$ or there is a (unique) vertex c of $\Gamma \setminus \{a, b\}$ which is adjacent to both a and b , and $\gamma' = \gamma c$.*

The next lemma produces separating pairs of points in ∂W from separating pairs of vertices in Γ . Note that since Γ is square-free, given a cut pair $\{a, b\}$ in Γ there is at most one

component of $\Gamma \setminus \{a, b\}$ consisting of a single vertex c which commutes with both a and b . For the remainder of this section, if Υ is a subcomplex of Σ , we denote by $N(\Upsilon)$ the union of the chambers which have nontrivial intersection with Υ .

Lemma 3.3. *Let $\{a, b\}$ be a cut pair in Γ and let γ be the (unique) (a, b) -bicoloured geodesic passing through the identity. Suppose that $\Gamma \setminus \{a, b\}$ has k components (with $k \geq 2$).*

(1) *If no component of $\Gamma \setminus \{a, b\}$ consists of a single vertex, then $\Sigma \setminus \gamma$, and consequently $\partial W \setminus \partial\gamma$, has exactly k components.*

More precisely, if $\Gamma \setminus \{a, b\}$ has k components $\Lambda_1, \dots, \Lambda_k$, none of which is a single vertex, then $\Sigma \setminus \gamma$ has k components Y_1, \dots, Y_k . For $1 \leq i \leq k$ the component Y_i contains exactly the Cayley graph vertices which can be written in the form $w_{ab}\lambda_i v$, where w_{ab} is a reduced word in a and b , λ_i is a generator corresponding to a vertex of Λ_i , and v is any element of W so that the word $w_{ab}\lambda_i v$ is reduced.

(2) *If $\Gamma \setminus \{a, b\}$ has a component consisting of a single vertex c , let γ' be the geodesic $\gamma' = \gamma c = c\gamma$ and let Υ be the subcomplex of Σ bounded by γ and γ' . Then $\Sigma \setminus \Upsilon$, and consequently $\partial W \setminus \partial\gamma$, has exactly $2(k - 1)$ components.*

More precisely, if $\Gamma \setminus \{a, b\}$ has k components $\Lambda_1, \dots, \Lambda_k$, with $\Lambda_1 = \{c\}$, then $\Sigma \setminus \Upsilon$ has $2(k - 1)$ components $Y_2, \dots, Y_k, Y'_2, \dots, Y'_k$. For $2 \leq i \leq k$ the component Y_i (respectively, Y'_i) contains exactly the Cayley graph vertices which can be written in the form $w_{ab}\lambda_i v$ (respectively, $w_{ab}c\lambda_i v$) where w_{ab} is a reduced word in a and b , λ_i is a generator corresponding to a vertex of Λ_i , and v is any element of W so that the word $w_{ab}\lambda_i v$ (respectively, $w_{ab}c\lambda_i v$) is reduced.

Proof. Observe first that in both cases, since a and b are nonadjacent vertices in Γ , the Cayley graph geodesic γ is also a geodesic in Σ .

To prove (1), assume that no component of $\Gamma \setminus \{a, b\}$ consists of a single vertex. Recall from Section 2.1 that we may identify the boundary of each chamber of Σ with the barycentric subdivision Γ' of Γ , and that Γ' may be viewed as the union of all mirrors. Now the complement of the mirrors K_a and K_b in Γ' has exactly k components, with the i th component being the union of the mirrors of type a vertex of Λ_i . The (infinite) union of chambers $N(\gamma)$ is obtained by gluing together chambers along panels of types a and b , and this gluing is type-preserving. Hence the complement of γ in $N(\gamma)$ has exactly k components, N_1, \dots, N_k , so that the boundary of N_i is the union of γ together with all panels in the boundary of $N(\gamma)$ with type a vertex of Λ_i .

We now define subsets Y_1, \dots, Y_k of $\Sigma \setminus \gamma$ as follows. Given a point $x \in \Sigma \setminus \gamma$ let $\pi(x)$ denote its closest-point projection to the geodesic γ . There is then a unique geodesic from x to $\pi(x)$, and this geodesic passes through exactly one of N_1, \dots, N_k . Define $x \in Y_i$ if this geodesic passes through N_i . The uniqueness of the projection geodesic implies that the Y_i are well-defined and disjoint. It is clear that each Y_i is connected and that $\cup_{1 \leq i \leq k} Y_i$ is $\Sigma \setminus \gamma$. Thus $\Sigma \setminus \gamma$ has k components. By similar arguments to those in the proof of Lemmas 2.1 and 2.3 of Lafont [20], it follows that $\partial\gamma$ separates ∂W into exactly k components as well.

Now, given a Cayley graph vertex w in Y_i , we construct a Cayley graph path from the identity to w as follows. Consider the geodesic η in Σ from w to the projection $\pi(w) \in \gamma$. We may choose a sequence of chambers K_1, \dots, K_n intersecting η so that consecutive chambers

in this sequence meet along a panel and the union $\cup_{i=1}^n K_i$ contains η . Since chambers are convex in Σ , we may assume that the chambers in this sequence are pairwise distinct. Also, if two chambers K_j and $K_{j'}$ with $j < j'$ meet along a panel, then $K_j \cup K_{j'}$ is convex, so we may assume that this occurs only if $j' = j + 1$. Let the centres of these chambers be $w_1 = w, \dots, w_n$ respectively. Then w_i and w_{i+1} are adjacent vertices of \mathcal{C} , so we may approximate the geodesic in Σ from w to $\pi(w)$ by the path in the Cayley graph with vertex set w_1, \dots, w_n . The w_1, \dots, w_n are pairwise distinct, since the corresponding chambers are pairwise distinct, and the label on this path does not contain any subword sts where s and t are commuting generators, since two chambers in this sequence meet along a panel only if they are consecutive. Hence this path in the Cayley graph is a geodesic in the Cayley graph.

Now concatenate the reverse of this geodesic with a Cayley graph geodesic from the identity to w_n . The geodesic from e to w_n in \mathcal{C} is labelled by a reduced word of the form w_{ab} as in the statement of the lemma. The geodesic from w_n to w leaves $N(\gamma)$ through a panel of type a vertex $\lambda_i \in \Lambda_i$. Thus the concatenated path is labelled by a word of the form $w_{ab}\lambda_i v$ as in the statement. It is clear that if a Cayley graph vertex w can be written as a reduced word $w_{ab}\lambda_i v$, then w is in Y_i . Thus Y_i consists of exactly the vertices described.

In case (2), since γ' is, like γ , bicoloured by a and b , both γ and γ' are geodesics in Σ . Moreover the subcomplex Υ is convex in Σ . (The subcomplex Υ consists of a band of big squares bounded by γ and γ' , and is isometric to $\mathbb{R} \times [0, 1]$.) By a similar argument to that in case (1), we see that $N(\Upsilon) \setminus \Upsilon$ has exactly $2(k-1)$ components, $N_2, \dots, N_k, N'_2, \dots, N'_k$, so that the boundary of N_i (respectively N'_i) is the union of γ (respectively γ') together with all panels in the boundary of $N(\Upsilon)$ with type a vertex of Λ_i .

We define subsets $Y_2, \dots, Y_k, Y'_2, \dots, Y'_k$ of $\Sigma \setminus \Upsilon$ as follows. Given a point $x \in \Sigma \setminus \Upsilon$ let $\pi(x)$ denote its closest-point projection to Υ . There is a unique geodesic from x to $\pi(x)$, which passes through exactly one of $N_2, \dots, N_k, N'_2, \dots, N'_k$. Define $x \in Y_i$ if this geodesic passes through N_i (and so $\pi(x) \in \gamma$) and $x \in Y'_i$ if it passes through N'_i (and so $\pi(x) \in \gamma'$). An argument similar to the proof of (1) implies that $\partial\gamma$ separates ∂W into exactly $2(k-1)$ components, and the description of these components is also similar to case (1). \square

The next result considers the separation properties of geodesics in the Cayley graph \mathcal{C} . Since an arbitrary geodesic γ in \mathcal{C} need not be a geodesic in Σ , we cannot use projections to γ in the following proof as we did in Lemma 3.3 above. Similarly, $N(\gamma)$ need not be convex, so we cannot use projections to $N(\gamma)$. Instead, we use paths in \mathcal{C} to approximate projections.

Lemma 3.4. *Let γ be a geodesic in \mathcal{C} . If γ separates $N(\gamma)$ then γ separates Σ .*

Proof. Suppose that $\Sigma \setminus \gamma$ is connected, and let p and q be distinct points in $N(\gamma) \setminus \gamma$. We will show that $N(\gamma) \setminus \gamma$ is connected by constructing a path from p to q in $N(\gamma) \setminus \gamma$.

Since $\Sigma \setminus \gamma$ is connected, there is a path η in $\Sigma \setminus \gamma$ from p to q . Denote by p' the first point of intersection of η with the boundary of $N(\gamma)$ and by q' the last point of intersection of η with the boundary of $N(\gamma)$. We may assume without loss of generality that the interiors of $\eta_{[p,p']}$ and $\eta_{[q',q]}$ lie in $N(\gamma) \setminus \gamma$ and that the interior of $\eta_{[p',q']}$ lies in $\Sigma \setminus N(\gamma)$, and that these subpaths contain no loops or backtracking.

Let K_1, \dots, K_n be a sequence of pairwise adjacent chambers in the closure of $\Sigma \setminus N(\gamma)$ through which the interior of $\eta_{[p',q']}$ passes, so that p' is in a panel shared by K_1 and some

chamber in $N(\gamma)$ and q' is in a panel shared by K_n and some chamber in $N(\gamma)$. Then similarly to the proof of Lemma 3.3, we may approximate $\eta_{[p',q']}$ by a reduced path τ in the Cayley graph \mathcal{C} with vertex set w_1, \dots, w_n , where w_1, \dots, w_n are the centres of the chambers K_1, \dots, K_n respectively. We construct another reduced path σ in \mathcal{C} from w_1 to w_n as follows. Let u be the vertex of γ which is adjacent in \mathcal{C} to w_1 , so that the midpoint of u and w_1 is the centre of a panel of K_1 which contains p' , and let v be the vertex of γ which is adjacent in \mathcal{C} to w_n , so that the midpoint of v and w_n is the centre of a panel of K_n containing q' . Then define σ to be the concatenation of the edge in \mathcal{C} between w_1 and u , the segment of γ between u and v , and the edge in \mathcal{C} between v and w_n . Note that since γ is a Cayley graph geodesic, the path σ is reduced.

Now we have that τ and σ are two distinct reduced paths in \mathcal{C} from w_1 to w_n , which intersect only at w_1 and w_n . So there is a filling by big squares of the loop in \mathcal{C} obtained by concatenating τ with σ . Denote by S_1, \dots, S_N the pairwise adjacent big squares in this filling which have at least one edge or vertex on $\gamma_{[u,v]}$, so that w_1 and u are adjacent vertices of S_1 and w_n and v are adjacent vertices of S_N . Then the boundary of $\cup_{1 \leq i \leq N} S_i$ is the union of σ with a path σ' in \mathcal{C} from w_1 to w_n so that σ' is entirely contained in chambers in the closure of $\Sigma \setminus N(\gamma)$ which share a panel or a vertex with $\partial N(\gamma)$. Now by concatenating the intersections of the squares S_1, \dots, S_N with these panels we obtain a path in the intersection of Σ with the boundary of $N(\gamma)$ which goes from the midpoint of u and w_1 to the midpoint of v and w_n . This path may be extended and/or restricted within panels to obtain a path from p' to q' which lies in $N(\gamma) \setminus \gamma$. Then by concatenating $\eta_{[p,p']}$, this path and $\eta_{[q',q]}$, we obtain a path from p to q in $N(\gamma) \setminus \gamma$. Hence $N(\gamma) \setminus \gamma$ is connected. \square

3.3. Identification of the \approx -pairs and their stabilisers. We now begin the construction of the JSJ tree \mathcal{T} , and the determination of stabilisers of vertices in \mathcal{T} . Recall that one class of finite-valence vertices in \mathcal{T} is given by \approx -pairs, which are pairs of points in ∂W that separate ∂W into at least three components. The following immediate corollary of Lemma 3.3 yields a collection of \approx -pairs. See Figure 3.1 for an example.

Corollary 3.5 (\approx -pairs). *Let $\{a, b\}$ be a pair of essential vertices in Γ such that $\Gamma \setminus \{a, b\}$ has $k \geq 3$ components. Then $\{a, b\}$ corresponds to a W -orbit of \approx -pairs in \mathcal{T} . Each \approx -pair in this orbit consists of the endpoints of a geodesic bicoloured by a and b , and yields a vertex of Type 1 in \mathcal{T} of valence:*

- (1) $k \geq 3$ if no component of $\Gamma \setminus \{a, b\}$ consists of a single vertex; and
- (2) $2(k - 1) \geq 4$ if $\Gamma \setminus \{a, b\}$ has a component consisting of a single vertex.



FIGURE 3.1. In both graphs, $\{a, b\}$ gives an orbit of \approx -pairs. On the left the corresponding Type 1 vertices of \mathcal{T} have valence $6 = 2(4 - 1)$ and on the right, they have valence 4.

We now show that every \approx -pair vertex of \mathcal{T} arises as in Corollary 3.5.

Lemma 3.6. *If ξ and ξ' are points of ∂W such that $\xi \approx \xi'$, then there is an essential cut pair $\{a, b\}$ of Γ such that $\Gamma \setminus \{a, b\}$ has at least three components, and a geodesic γ bicoloured by a and b such that $\partial\gamma = \{\xi, \xi'\}$.*

Proof. By Bowditch's construction, the pair $\{\xi, \xi'\}$ corresponds to a vertex v of Type 1 of the JSJ tree for W , the stabiliser H of v is 2-ended, and W splits over H . Now by Theorem 2.1, there exists a unique cut pair $\{a, b\}$ such that some conjugate of $\langle a, b \rangle$ is a finite-index subgroup of H . It follows that there is a geodesic γ bicoloured by a and b such that $\partial\gamma = \{\xi, \xi'\}$. Suppose $\Gamma \setminus \{a, b\}$ had exactly two components. Then by applying the appropriate case of Lemma 3.3, we see that $\partial W \setminus \{\xi, \xi'\}$ also has two components, which contradicts the assumption that $\xi \approx \xi'$. Thus $\Gamma \setminus \{a, b\}$ has at least three components.

We claim that a and b are both essential. By Standing Assumptions 1.1(2), Γ has no valence one vertices. Suppose one of a and b , say a , has valence two, and let Λ be a component of $\Gamma \setminus \{a, b\}$ which does not contain either of the two edges attached to a . Then a and Λ are necessarily in different components of $\Gamma \setminus \{b\}$, which contradicts Standing Assumptions 1.1(2). This proves the claim. Thus $\{a, b\}$ satisfies the hypothesis of Corollary 3.5. Therefore the \approx -pair $\{\xi, \xi'\}$ comes from the construction in Corollary 3.5. \square

We next determine the stabilisers of \approx -pairs.

Lemma 3.7. *Let ξ and ξ' be points of ∂W such that $\xi \approx \xi'$, and let γ be an (a, b) -bicoloured geodesic so that $\partial\gamma = \{\xi, \xi'\}$ (as guaranteed by Lemma 3.6). Then the stabiliser of the \approx -pair $\{\xi, \xi'\}$ is a conjugate of either $\langle a, b \rangle$, if there is no vertex c of Γ adjacent to both a and b , or of $\langle a, b, c \rangle$, if there is such a vertex c .*

Proof. The action of W on Σ preserves the labels of edges in \mathcal{C} , so that if $g \in W$ then $g\gamma$ is also an (a, b) -bicoloured geodesic. Conjugates of $\langle a, b \rangle$ obviously stabilise endpoints of (a, b) -bicoloured geodesics, and by Corollary 3.2 the only other elements of W which can stabilise these endpoints are those in conjugates of the subgroup $\langle a, b, c \rangle$, in the cases where there exists a (unique) vertex c which commutes with both a and b . \square

The following corollary of the results in this section will be used in the next section.

Corollary 3.8. *Let A be a set of vertices of Γ and let \mathcal{C}_A be a copy of the Cayley graph of $\langle A \rangle$ in \mathcal{C} . Then $\partial\mathcal{C}_A$ is an \approx -pair if and only if the following holds: there is an essential cut pair $\{a, b\}$ in Γ so that $\Gamma \setminus \{a, b\}$ has at least three components, and $A = \{a, b\}$ or $A = \{a, b, c\}$, with the latter occurring only if there is a vertex c of Γ adjacent to both a and b .*

3.4. Construction of certain \sim -classes. Recall from Section 2.3 that the \sim -classes are equivalence classes of points in $\partial W(2)$ such that any pair of points in the set separates ∂W into exactly two components, and the set is maximal with respect to this property. In Proposition 3.9 below, we describe a construction which yields \sim -classes. We show in Section 3.5 that every \sim -class arises in this way.

Let $|\Gamma|$ denote the geometric realisation of Γ . Note the distinction between $\Gamma \setminus \{a, b\}$ and $|\Gamma| \setminus \{a, b\}$, for an arbitrary pair of vertices a and b of Γ : a component of the former must have a vertex, while a component of the latter is allowed to be an arc between adjacent vertices a and b . In particular, if a and b are adjacent vertices of Γ , they do not separate Γ

(by Standing Assumptions 1.1(2)) but they do separate $|\Gamma|$. A *branch* of Γ is a subgraph of Γ consisting of a (closed) reduced path between a pair of essential vertices, which does not contain any essential vertices in its interior. For example, both of the graphs in Figure 3.1 above contain 4 branches between a and b .

Proposition 3.9 (\sim -classes). *Let A be a set of (not necessarily essential) vertices of Γ such that $\langle A \rangle$ is infinite and:*

- (A1) *elements of A pairwise separate $|\Gamma|$;*
- (A2) *given any subgraph Λ of Γ which is a subdivided copy of K_4 , if Λ contains at least three vertices of A then all the vertices of A lie on a single branch of Λ ; and*
- (A3) *the set A is maximal among all sets satisfying both (A1) and (A2).*

Let \mathcal{C}_A be any copy of the Cayley graph of $\langle A \rangle$ in \mathcal{C} , and let \mathcal{A}_\approx be the set of geodesics $\gamma \subset \mathcal{C}_A$ which are bicoloured by pairs $\{a, b\}$ as in Corollary 3.5 (so that $\partial\gamma$ is an \approx -pair). Assume that $\partial\mathcal{C}_A$ is not an \approx -pair. Then $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ contains at least two points, and:

- (1) *If $\langle A \rangle$ is 2-ended, $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx = \partial\mathcal{C}_A$ is a \sim -pair.*
- (2) *Otherwise, $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ is an infinite \sim -class.*

Remark 3.10. Let A be a subset of vertices of Γ . Then $\langle A \rangle$ is infinite if and only if A contains a pair of non-adjacent vertices; in particular, if $\langle A \rangle$ is infinite then $\text{card}(A) \geq 2$. Also, we can recognise whether $\partial\mathcal{C}_A$ is an \approx -pair using the graph-theoretic criteria on A given by Corollary 3.8. Finally, $\langle A \rangle$ is 2-ended if and only if either $A = \{a, b\}$, where a and b are non-adjacent vertices of Γ , or $A = \{a, b, c\}$, where a and b are non-adjacent and c is adjacent to both a and b . Thus all hypotheses of Proposition 3.9 can be verified in the graph Γ .

Example 3.11. We now discuss some examples which illustrate the statement of Proposition 3.9. In the graphs in Figure 3.1 there are no subgraphs which are subdivided copies of K_4 , so a set of vertices A satisfying (A1)–(A3) is just required to be maximal with respect to (A1). In the left-hand graph in this figure, if A is the union of $\{a, b\}$ with the vertex adjacent to both a and b , then A is maximal with respect to (A1). However by Corollary 3.8, $\partial\mathcal{C}_A$ is an \approx -pair, so the set A does not provide a \sim -class in this case. Now suppose that in either graph, A is a branch from a to b of length at least three. Then A is maximal with respect to property (A1), hence properties (A1)–(A3) hold, and since $\text{card}(A) \geq 4$ we have by Corollary 3.8 that $\partial\mathcal{C}_A$ is not an \approx -pair. Thus A provides an infinite \sim -class.

Now consider the examples in Figure 3.2. In this figure, the sets $\{a, b, c, d\}$ and $\{a, d, f, g, e\}$ on the left, the set $\{a_1, a_2, a_3, a_4, a_5\}$ in the centre and the set $\{p, q\}$ on the right all satisfy properties (A1)–(A3) in Proposition 3.9 and do not give \approx -pairs. The set $\{p, q\}$ thus yields a \sim -pair, while the other three correspond to infinite \sim -classes. On the left and in the centre of Figure 3.2 there is no subdivided K_4 subgraph, so (A3) just requires maximality with respect to (A1). In the right-hand graph, the situation is more subtle. The set $\{p, q\}$ satisfies (A1) and (A2), but is not maximal with respect to (A1). Adding any non-essential vertex to $\{p, q\}$ means that (A1) fails, and adding any essential vertex to $\{p, q\}$ means that (A2) fails, so (A3) does hold for the set $\{p, q\}$. Now consider the set $\{u, q\}$. This satisfies (A1) and (A2) but is maximal with respect to neither, since adding any vertex on the branch between u and q results in a larger set also satisfying both (A1) and (A2). So $\{u, q\}$ does not

satisfy (A3) (but the branch between u and q does). Finally, the set $\{p, q, r, s\}$ is maximal with respect to (A1) but fails (A2), and so does not give a \sim -class.

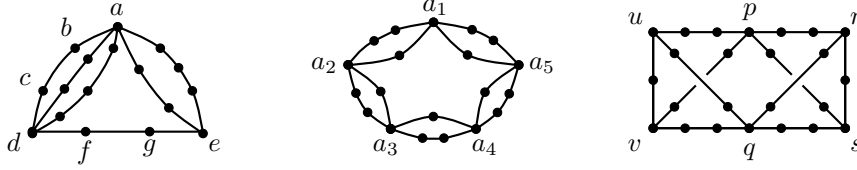


FIGURE 3.2. Examples to illustrate properties (A1)–(A3) in Proposition 3.9.

For the proof of Proposition 3.9, we need several preliminary results concerning the structure of a set A satisfying properties (A1), (A2) and/or (A3). Lemma 3.12 shows that subsets A which satisfy property (A1) have a “cyclic” configuration in Γ . In Lemma 3.14 we prove that if in addition A satisfies (A2), then the cyclic ordering on A is well-defined, and paths between pairs of points on induced cycles which contain A are strongly restricted. We use this in Corollary 3.15 to show that if A also satisfies (A3), then for any vertex $b \notin A$, there is an induced cycle in Γ which contains A but does not contain b .

Lemma 3.12. *Let A be a set of vertices of Γ with $\langle A \rangle$ infinite and satisfying property (A1). Then the elements of A lie on an induced cycle α in Γ .*

Proof. The proof is by induction on $\text{card}(A) = n \geq 2$. Since Γ is connected and has no separating vertices or edges, the statement is satisfied when A has 2 elements.

Now let A be a set of $n+1$ elements satisfying property (A1). By the induction hypothesis, after removing one element, say b , the remaining n elements lie on an induced cycle α . If b also lies on α we are done, otherwise label the elements of $A \setminus \{b\}$ going around α cyclically as a_1, \dots, a_n and let β be the subpath of α from a_1 to a_n which contains all of a_2, \dots, a_{n-1} .

Since the pair $\{a_n, b\}$ separates $|\Gamma|$ there is a path η in Γ from a_n to b which meets α only at a_n . Similarly, there is a path ζ in Γ from b to a_1 which meets $\alpha \cup \beta$ only at a_1 and b . Then $\alpha \cup \beta \cup \zeta$ is an embedded cycle containing all elements of A , as desired. It is clear from the construction that this cycle can be chosen to be an induced cycle. \square

Remark 3.13. Let A be as in Lemma 3.12 and let α be an induced cycle in Γ which contains all elements of A . Consider any copy of the subcomplex Σ_α of Σ corresponding to the special subgroup generated by all vertices in α . Since α contains at least 5 vertices and a copy of the Cayley graph \mathcal{C}_A is contained in Σ_α , we see that any copy of \mathcal{C}_A embeds in a subcomplex of Σ which is quasi-isometric to the hyperbolic plane.

Lemma 3.14. *Let A be a set of vertices of Γ with $\langle A \rangle$ infinite and satisfying properties (A1) and (A2). Let α be an induced cycle in Γ containing all elements of A (as guaranteed by Lemma 3.12) and label the elements of A going around α cyclically as a_1, \dots, a_n . Let σ be any other induced cycle containing all elements of A . Then:*

- (1) *The cycle σ can be oriented so as to induce this same order on A . Thus up to orientation, there is a well-defined cyclic ordering on the elements of A .*

- (2) Assuming $\text{card}(A) \geq 3$, if x and y are points on σ (not necessarily vertices of Γ) such that $\sigma \setminus \{x, y\}$ partitions A into two nonempty sets $\{a_i, a_{i+1}, \dots, a_{j-1}\}$ and $\{a_{i-1}, \dots, a_1, a_n, \dots, a_{j+1}, a_j\}$, where $i < j$ and x (respectively y) lies between a_i and a_{i-1} (respectively a_j and a_{j-1}), then any reduced path in Γ from x to y passes through either all of $a_i, a_{i+1}, \dots, a_{j-1}$, or all of $a_{i-1}, \dots, a_1, a_n, \dots, a_{j+1}, a_j$.

Proof. Part (1) is immediate when A has 2 or 3 elements. If $\text{card}(A) = n \geq 4$ suppose there is an induced cycle σ containing all elements of A which cannot be oriented to induce the same ordering. Then without loss of generality there is a vertex a_i with $i \neq 2, n$ so that a_1 is adjacent to a_i on σ . Since σ is an induced cycle and $n \geq 4$ there is also some a_k with $1 < k < i$ and some a_l with $i < l \leq n$ so that a_k and a_l are adjacent on σ . Let σ_{1i} be the subpath of σ connecting a_1 to a_i and containing no other vertices of A , and similarly define σ_{kl} . Notice that σ_{1i} connects the two components of $\alpha \setminus \{a_k, a_l\}$ and σ_{kl} connects the two components of $\alpha \setminus \{a_1, a_i\}$.

Now there is a subpath σ'_{1i} of σ_{1i} which connects the two components of $\alpha \setminus \{a_k, a_l\}$ and intersects α only at its endpoints, and there is a similar subpath σ'_{kl} of σ_{kl} . Consider the graph $\Lambda = \alpha \cup \sigma'_{1i} \cup \sigma'_{kl}$. This graph is a subdivided K_4 , and since it contains α it contains all vertices of A . But by construction, at most two of the four vertices a_1, a_i, a_k, a_l can lie on any branch of Λ . This contradicts (A2), and so completes the proof of part (1).

For (2), assume $\text{card}(A) = n \geq 3$ and let x and y be points as in the statement. Assume by way of contradiction that there is a reduced path ξ in Γ from x to y which fails the condition. Suppose that ξ misses some vertex, say a_k , in the set $\{a_i, a_{i+1}, \dots, a_{j-1}\}$, and some vertex, say a_l , in the set $\{a_{i-1}, \dots, a_1, a_n, \dots, a_{j+1}, a_j\}$. Then since a_k and a_l separate $|\Gamma|$, there is a reduced path β connecting a_k and a_l which meets $\sigma \cup \xi$ only at its endpoints. Now ξ connects the two components of $\sigma \setminus \{a_k, a_l\}$, so ξ contains a subpath ξ' which intersects σ only at its endpoints and connects the two components of $\sigma \setminus \{a_k, a_l\}$, as shown in Figure 3.3. Then $\sigma \cup \xi' \cup \beta$ forms a subdivided K_4 in Γ , which contains all $n \geq 3$ vertices of A . But the branch β of this K_4 contains only the two vertices a_k and a_l , and this contradicts (A2). So ξ must contain all of the vertices in at least one of the sets in the induced partition of A . \square

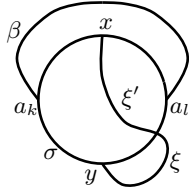


FIGURE 3.3

Corollary 3.15. Let A be a set of vertices of Γ satisfying properties (A1), (A2) and (A3), such that $\langle A \rangle$ is infinite and $\partial \mathcal{C}_A$ is not an \approx -pair. Let b be a vertex of Γ which is not in A .

- (1) If $\langle A \rangle$ is 2-ended, let $\{a_1, a_2\}$ be the unique cut pair in A . Then there is an induced cycle α in Γ which contains a_1 and a_2 and contains a vertex from both components of $\Gamma \setminus \{a_1, a_2\}$, but does not contain b .

(2) If $\langle A \rangle$ is not 2-ended, there is an induced cycle α in Γ which contains all elements of A but does not contain b .

Proof. Assume first that $\langle A \rangle$ is 2-ended. Note that $\Gamma \setminus \{a_1, a_2\}$ must have exactly two components, say Λ_1 and Λ_2 , otherwise $\partial\mathcal{C}_A$ would be an \approx -pair. Suppose $b \in \Lambda_1$ and let μ be a reduced path from a_2 to a_1 in Λ_2 . To obtain the desired cycle α , it suffices to concatenate μ with a reduced path from a_1 to a_2 in Λ_1 which does not contain b . We assume by contradiction that every reduced path from a_1 to a_2 in Λ_1 passes through b . It follows that property (A1) holds for the set $\{a_1, a_2, b\}$. Now suppose that there is a subdivided K_4 and contains all three vertices a_1, a_2 and b . Then as every reduced path from a_1 to a_2 in Λ_1 contains b , the vertices a_1, a_2 and b lie on the same branch of this K_4 . Thus property (A2) also holds for the set $\{a_1, a_2, b\}$. Hence $\{a_1, a_2, b\}$ is contained in a maximal set which satisfies both (A1) and (A2). If $A = \{a_1, a_2\}$ this contradicts (A3), since $b \notin A$.

Since $\langle A \rangle$ is 2-ended the only other possibility is that $A = \{a_1, a_2, c\}$ with $c \neq b$ and c adjacent to both a_1 and a_2 . Notice that as every reduced path from a_1 to a_2 in Λ_1 passes through b , the vertex c lies in Λ_2 . In particular, b and c are non-adjacent. Now the set $\{a_1, a_2, b, c\}$ must fail at least one of (A1) and (A2). Suppose first that (A1) fails for this set. Since (A1) does hold for $\{a_1, a_2, b\}$ and $\{a_1, a_2, c\}$, the pair $\{b, c\}$ is not a cut pair. Thus there is a reduced path η from a_1 to a_2 in $\Gamma \setminus \{b, c\}$. If η is in Λ_1 this contradicts our assumption that every reduced path from a_1 to a_2 in Λ_1 contains b , so η is in Λ_2 . Now $c \in \Lambda_2$ is adjacent to both a_1 and a_2 , so there must be a reduced path in Λ_2 , say β , from c to some point in η . Notice that β does not contain either a_1 or a_2 , and η does not contain c . Let Λ be the union of η, β , the edges between a_i and c for $i = 1, 2$, and a path in Λ_1 from a_1 to a_2 (via b). Then Λ is a subdivided K_4 subgraph which has a_1, a_2 and c as three of its essential vertices. Hence these three vertices are not contained in a single branch of Λ . This contradicts property (A2) for the set $\{a_1, a_2, c\}$. Therefore property (A1) must hold for the set $\{a_1, a_2, b, c\}$.

As (A1) holds for the set $\{a_1, a_2, b, c\}$, property (A2) must fail. Thus there is a subgraph Λ of Γ which is a subdivided K_4 and contains at least three vertices of $\{a_1, a_2, b, c\}$, so that these vertices are not all on the same branch of Λ . Since (A2) holds for the sets $\{a_1, a_2, b\}$ and $\{a_1, a_2, c\}$, we may assume that the vertices a_1, b and c all lie on Λ , but they are not all on the same branch. Note that then a_2 cannot lie on Λ . However $\Lambda \setminus \{a_1\}$ contains both b and c , and is contained in a single component of $\Gamma \setminus \{a_1, a_2\}$. This contradicts b and c being in different components of $\Gamma \setminus \{a_1, a_2\}$. We conclude that there is a reduced path from a_1 to a_2 in Λ_1 which does not pass through b , and so the required cycle α can be obtained.

Now suppose $\langle A \rangle$ is not 2-ended. Then $\text{card}(A) \geq 3$. Consider an induced cycle α containing all elements of A , and inducing the cyclic order a_1, \dots, a_n on the elements of A . Suppose b is on the cycle α , say between a_i and a_{i+1} . Since b is not in A , by property (A3) the set $A \cup \{b\}$ fails either (A1) or (A2).

If (A1) fails, there is some $a_k \in A$ (possibly equal to a_i or a_{i+1}) so that $\{b, a_k\}$ does not separate $|\Gamma|$. Then there is a path μ connecting the two components of $\alpha \setminus \{b, a_k\}$ and meeting α only at its endpoints x and y . We claim that both x and y lie on the (closed) subpath of α between a_i and a_{i+1} containing b . If not, then one possibility is that both x and

y both lie on the (open) component of $\alpha \setminus \{a_i, a_{i+1}\}$ which does not contain b (in this case, a_k must be distinct from both a_i and a_{i+1}). Then a_k and a_i (say) lie in different components of $\alpha \setminus \{x, y\}$. By slightly extending μ if either x or y is in A , we obtain a reduced path μ' connecting two points x' and y' in $\alpha \setminus A$, so that a_k and a_i lie in different components of $\alpha \setminus \{x', y'\}$, but neither a_i nor a_k lies on μ' . This contradicts Lemma 3.14(2). By a similar argument we can rule out exactly one of x and y lying on the (closed) subpath of α between a_i and a_{i+1} containing b . This proves the claim. Since μ does not contain b , we can then use μ to replace the subpath of α from x to y via b , and so obtain the required cycle.

Now suppose (A1) holds for the set $A \cup \{b\}$, but (A2) fails. Then there is a subgraph Λ of Γ which is a subdivided copy of K_4 , and distinct vertices $a_k, a_l \in A$ such that a_k, a_l and b are not contained in the same branch of Λ . If a_k and a_l are on different branches of Λ , then by (A1) there is a path η from a_k to a_l which meets Λ only at its endpoints. Then $\Lambda \cup \eta$ contains a subgraph, say Λ' , which is a subdivided K_4 , and such that a_k and a_l are on the same branch of Λ' . Moreover, Λ' can be chosen so that $b \in \Lambda'$, on a branch other than the one containing a_k and a_l . Thus we may assume that a_k and a_l are on the same branch of Λ .

We now have that a_k and a_l lie on the same branch of Λ , say β , and that b is not on β . Then there is an induced cycle σ in Λ so that σ contains β but not b . If all elements of A lie on β then σ is an induced cycle containing all elements of A but not b , so suppose that there is some $a_m \in A$ which does not lie on β . By property (A2) for A , the vertex a_m then cannot lie on Λ . Recall that by assumption, the cycle α which contains all of A also contains b . Then it is not hard to see that $\alpha \cup \Lambda$ contains a subdivided K_4 subgraph in which a_k, a_l and a_m do not lie on the same branch. This contradicts property (A2) for A . Thus all elements of A lie on the branch β , which completes the proof. \square

Finally, we consider separation properties of geodesics in \mathcal{C} labelled by elements of a set A satisfying (A1) and (A2).

Lemma 3.16. *Let A be a set of vertices of Γ with $\langle A \rangle$ infinite satisfying properties (A1) and (A2). Let γ be a bi-infinite geodesic in \mathcal{C}_A labelled by elements of A . Then γ separates Σ and $\partial\gamma$ separates ∂W .*

Proof. If $\text{card}(A) = 2$, with $A = \{a, b\}$, then a and b are not adjacent in Γ (since $\langle A \rangle$ is infinite), and so $\{a, b\}$ is a cut pair by (A1). Hence by Lemma 3.3, γ separates Σ and $\partial\gamma$ separates ∂W .

Now assume $\text{card}(A) \geq 3$. We first show that γ separates Σ . By Lemma 3.4, it is enough to show that γ separates $N(\gamma)$. If not, then given any pair of points x and y in $N(\gamma) \setminus \gamma$, there is a path η in $N(\gamma) \setminus \gamma$ connecting them. Since η misses γ , and in particular, the cone points of the chambers that make up $N(\gamma)$, we may assume that η (as well as x and y) lie in $\partial N(\gamma)$. We show below that this leads to a contradiction.

Write $N(\gamma)$ as the union of chambers K_i , with $-\infty \leq i \leq \infty$, such that for each i , the chambers K_{i-1} and K_i intersect in the panel P_i , of type g_i . Recall that each panel P is a star, and denote by ∂P the set of endpoints of the spokes of this star. The boundary of any chamber K can be identified with the graph Γ , and under this identification, the cone point of a panel P of type g corresponds to the vertex g in Γ , while the points of ∂P correspond to midpoints of the edges emanating from g in Γ .

Identify the boundary of K_0 with Γ as above. Fix an embedded cycle α containing A so that α induces a cyclic ordering a_1, \dots, a_n on the elements of A with $a_1 = g_0$. This order on A induces a partition of ∂P_0 into two sets P_0^+ and P_0^- , as follows. Let $x \in \partial P_0$, so that x may be identified with the midpoint of an edge emanating from a_1 in Γ . Let α_{1n} be the subpath of α connecting a_1 to a_n and containing no other vertices of α , and similarly define α_{12} . If x lies on α , then we put x in P_0^- if x lies on α_{1n} , and x in P_0^+ if x lies on α_{12} .

If x does not lie on α , let b be the other vertex of the edge of Γ which has x as its midpoint. As Γ has no separating vertices or edges, there is a reduced path η connecting b to some vertex b' of α other than a_1 , so that η intersects α only at b' . An argument similar to the proof of Lemma 3.14(2), applied to the union of η with the edge containing x , implies that b' must lie on either α_{1n} or α_{12} . Put $x \in P_0^-$ if b' lies on α_{1n} , and $x \in P_0^+$ if b' lies on α_{12} .

To see that this is well-defined, suppose that there are reduced paths ξ and ξ' connecting a_1 via b to vertices $b' \neq a_1$ on α_{1n} and $b'' \neq a_1$ on α_{12} respectively, so that ξ and ξ' meet α only at their endpoints. Then the graph $\Lambda = \alpha \cup \xi \cup \xi'$ is a subdivided K_4 containing all of A , and its branch ξ contains at most two vertices of A . This contradicts (A2). Thus P_0^- and P_0^+ are well-defined, and we have completed the partition of ∂P_0 .

Now we inductively define a partition of ∂P_i into P_i^+ and P_i^- for each $i \neq 0$ such that:

- (1) If ν is a path from ∂P_i to ∂P_{i+1} in ∂K_i whose interior contains no points of $\partial P_i \cup \partial P_{i+1}$, then ν either connects P_i^+ to P_{i+1}^+ or P_i^- to P_{i+1}^- .
- (2) If ν is a path from P_i^+ to P_i^- in $\partial K_i \setminus P_i$, then ν passes through the cone point of P_{i+1} .

We first prove property (2) for $i = 0$. Suppose ν is a path from P_0^+ to P_0^- in $\partial K_0 \setminus P_0$. Consider ν under the identification of ∂K_0 with Γ and call its endpoints x and y . Assume ν is reduced. By our assumption, ν does not pass through the vertex of Γ labelled $g_0 = a_1$. Then by Lemma 3.14(2), ν must pass through all the other vertices of A . Translating back to K_0 , this means, in particular, that ν passes through the cone point of P_1 .

Now suppose that $P_1^\pm, \dots, P_{i-1}^\pm$ have been defined satisfying (2), with $P_1^\pm, \dots, P_{i-2}^\pm$ satisfying (1). Identify ∂K_i with Γ , and let x correspond to a point of ∂P_i . Then there is a path ν in Γ between the vertex labelled g_i and the vertex labelled g_{i-1} passing through x . Put x in P_i^+ if ν passes through P_{i-1}^+ and in P_i^- otherwise. If μ and ν are two such paths, entering P_{i-1} through P_{i-1}^- and P_{i-1}^+ respectively, then $\mu \cup \nu$ reduces to a path connecting P_{i-1}^+ and P_{i-1}^- , which does not pass through the cone point of P_i . However, this contradicts property (2) for P_{i-1} . This shows that the partition on ∂P_i is well defined, and the partition on ∂P_{i-1} satisfies property (1). The proof that the partition on ∂P_i satisfies property (2) is similar to the proof in the base case. The definition of P_i^\pm for $i < 0$ is similar.

Observe that an argument similar to the base case of (2) shows for all i

- (3) If ν is a path from P_i^+ to P_i^- in $\partial K_{i-1} \setminus P_i$, then ν passes through the cone point of P_{i-1} .

Now assume that the path η in $\partial N(\gamma)$ constructed above connects some point $x \in P_0^+$ to some point $y \in P_0^-$. Then η is a union of segments, each of which connects some point of ∂P_i to some point of ∂P_i or ∂P_{i+1} for some i , and does not pass through the cone points of the panels P_i and $P_{i\pm 1}$. Now since η connects x to y , it must (without loss of generality) have some segment which connects a point of P_j^+ to either P_j^- or P_{j+1}^- for some j . However this contradicts one of properties (1), (2), and (3) above.

It follows that γ does separate $N(\gamma)$, and therefore γ separates Σ as required. By similar arguments to those in Lemma 2.3 of Lafont [20], we obtain that $\partial\gamma$ separates ∂W as well. \square

We are now ready to prove Proposition 3.9.

Proof of Proposition 3.9. Let \mathcal{C}_A be some copy of the Cayley graph of $\langle A \rangle$ in \mathcal{C} . Assume that $\langle A \rangle$ is infinite and that $\partial\mathcal{C}_A$ is not an \approx -pair.

We first show that $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ contains at least two points. Since $\langle A \rangle$ is infinite, the set $\partial\mathcal{C}_A$ has at least two points. If $\langle A \rangle$ is 2-ended, then $\partial\mathcal{C}_A$ has exactly two points, and by the assumption that $\partial\mathcal{C}_A$ is not an \approx -pair, it follows that no bicoloured geodesic in \mathcal{C}_A has endpoints an \approx -pair. So the set \mathcal{A}_\approx is empty and $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx = \partial\mathcal{C}_A$ has exactly two points. If $\langle A \rangle$ is not 2-ended, we may construct a geodesic γ in \mathcal{C}_A so that neither direction of γ is eventually bicoloured by a pair of vertices in A . Thus by Lemma 3.1, neither γ^+ nor γ^- is the endpoint of a bicoloured geodesic. So γ^+ and γ^- are two distinct points in $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$.

We next show that any two points in $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ are in the same \sim -class. For this, let $\xi, \eta \in \partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$, with $\xi \neq \eta$. To show that $\xi \sim \eta$, we must show that $\partial W \setminus \{\xi, \eta\}$ has exactly two components. Since $\langle A \rangle$ is convex in W and $\xi, \eta \notin \partial\mathcal{A}_\approx$, the boundary points ξ and η are the endpoints of some geodesic γ in $\mathcal{C}_A \setminus \mathcal{A}_\approx$. By Lemma 3.16, $\partial\gamma = \{\xi, \eta\}$ separates ∂W . Thus ξ and η are in $\partial W(2+)$. On the other hand, since ξ and η are not in $\partial\mathcal{A}_\approx$, by our characterisation of \approx -pairs in Lemma 3.5 we know that neither of them is part of an \approx -pair. It follows that neither ξ nor η is in $\partial W(3+)$, by property (2) following Definition 2.3. So ξ and η are both in $\partial W(2)$, and they separate ∂W into exactly two components. Thus $\xi \sim \eta$.

To complete the proof of Proposition 3.9, we show that $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ is a full \sim -class, i.e., that no point of $\partial W(2)$ outside the set $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ is equivalent to a point inside it. Assume by contradiction that there are points $\xi \in \partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ and $\eta \notin \partial\mathcal{C}_A$ so that $\xi \sim \eta$. Then since $\xi \sim \eta$, we have that $\{\xi, \eta\}$ separates ∂W into exactly two components, say U_1 and U_2 . Now $\xi \in \partial W(2)$, so there is a neighbourhood V of ξ in ∂W so that $V \setminus \{\xi\}$ has exactly two components, say V_1 and V_2 . Suppose V intersects only one of U_1 and U_2 , say U_1 . Then $\partial W \setminus \eta$ is the union of the open sets $V \cup U_1$ and U_2 , and hence η is a cut point of ∂W . This contradicts the result from [28] that the boundary of a 1-ended hyperbolic group has no global cut points, as discussed in the introduction to [9]. Thus V intersects both U_1 and U_2 nontrivially. On the other hand, if V_1 has nonempty intersection with both U_1 and U_2 , this contradicts the connectedness of V_1 . So without loss of generality $V_1 = V \cap U_1$, and similarly $V_2 = V \cap U_2$. We claim that there is a path in $\partial W \setminus \{\xi, \eta\}$ which connects a point in V_1 to a point in V_2 . This yields a contradiction, since this path connects U_1 to U_2 in $\partial W \setminus \{\xi, \eta\}$.

Let ξ' be a point of $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ with $\xi' \neq \xi$. We showed above that $\xi \sim \xi'$. Let γ be a geodesic in \mathcal{C} connecting ξ to η . Since $\xi \in \partial\mathcal{C}_A$, in the direction of ξ the geodesic γ is eventually labelled by elements of A . Now since $\eta \notin \partial\mathcal{C}_A$, not all of γ is contained in \mathcal{C}_A . Starting at ξ and travelling back along γ , let b be the first label of γ which is not in A .

We first consider the case that $\langle A \rangle$ is 2-ended. Then $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx = \partial\mathcal{C}_A$ consists of the two endpoints ξ and ξ' of an (a_1, a_2) -bicoloured geodesic, where $a_1, a_2 \in A$ and $\{a_1, a_2\}$ is a cut pair in Γ . Let Λ_1 and Λ_2 be the two components of $\Gamma \setminus \{a_1, a_2\}$. Then by Corollary 3.15(1), there is an induced cycle α in Γ which contains both a_1 and a_2 , does not contain b , and has nonempty intersection with both Λ_1 and Λ_2 . Let Σ_α be the corresponding subcomplex of

Σ containing our chosen copy of \mathcal{C}_A , as in Remark 3.13. Then Σ_α is quasi-isometric to the hyperbolic plane, and $\xi, \xi' \in \partial\Sigma_\alpha$ but $\eta \notin \partial\Sigma_\alpha$ since b is not on α . Let C_1 and C_2 be the two components of $\partial W \setminus \{\xi, \xi'\}$. Since the cycle α contains vertices in both components of $\Gamma \setminus \{a_1, a_2\}$, Lemma 3.3 implies that the circle $\partial\Sigma_\alpha$ has nonempty intersection with both C_1 and C_2 . Now if both C_1 and C_2 have nonempty intersection with V_1 this contradicts the connectedness of V_1 , and so without loss of generality $V \cap C_i = V_i$ for $i = 1, 2$. Hence $\partial\Sigma_\alpha$ has nonempty intersection with both V_1 and V_2 . Let $p_i \in \partial\Sigma_\alpha \cap V_i$ for $i = 1, 2$. Then the arc of $\partial\Sigma_\alpha$ from p_1 to p_2 which misses ξ gives a path in $\partial W \setminus \{\xi, \eta\}$ from V_1 to V_2 , proving the claim in this case.

Now suppose that $\langle A \rangle$ is not 2-ended. By Corollary 3.15(2), there is an induced cycle α in Γ which contains all elements of A but does not contain b . Let Σ_α be the corresponding subcomplex of Σ which contains \mathcal{C}_A and is quasi-isometric to the hyperbolic plane. Then $\xi \in \partial\Sigma_\alpha$ but $\eta \notin \partial\Sigma_\alpha$. Similarly to the previous case, to prove the claim it suffices to show that the sets $\partial\Sigma_\alpha \cap V_1$ and $\partial\Sigma_\alpha \cap V_2$ are both nonempty.

Since $\langle A \rangle$ is not 2-ended, there is a point $\xi'' \in \partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ with ξ'' distinct from both ξ and ξ' . Let C_1 be the component of $\partial W \setminus \{\xi, \xi'\}$ which does not contain ξ'' and let C_2 be the component of $\partial W \setminus \{\xi, \xi''\}$ which does not contain ξ' . Then similarly to the previous case, we may assume that $V \cap C_i = V_i$ for $i = 1, 2$. Now the component C_1 contains the interior of the arc of $\partial\Sigma_\alpha$ from ξ to ξ' which misses ξ'' , hence there is a point $p_1 \in \partial\Sigma_\alpha \cap V_1$. Similarly, we obtain a point $p_2 \in \partial\Sigma_\alpha \cap V_2$, which completes the proof of the claim in this case.

We have shown that $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ is a \sim -class in ∂W . Now this \sim -class is a \sim -pair if and only if $\partial\mathcal{C}_A$ contains exactly two points, which occurs if and only if $\langle A \rangle$ is 2-ended. This completes the proof of Proposition 3.9. \square

Remark 3.17. We can now construct some edges in the pre-tree T . Let A be a set of vertices of Γ so that $\langle A \rangle$ is infinite and not 2-ended and properties (A1)–(A3) from Proposition 3.9 hold. Let \mathcal{C}_A be a copy of the Cayley graph of $\langle A \rangle$. Then since $\langle A \rangle$ is not 2-ended, $\partial\mathcal{C}_A$ is not a \approx -pair. Let $v = \partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ be the corresponding infinite \sim -class. Then for each $\gamma \in \mathcal{A}_\approx$, the endpoints $\{\gamma^+, \gamma^-\}$ form an \approx -pair which is in the closure of the set v . Hence there is an edge in the pre-tree T connecting v to this \approx -pair vertex.

3.5. Identification of the \sim -classes and their stabilisers. We now show that we have already constructed all the \sim -classes in ∂W . Together with our identification of the \approx -pairs in Lemma 3.6 and the relation of betweenness, this means we have constructed the pretree T . This section also identifies the stabilisers of \sim -classes.

Bowditch shows that $\partial W(2+)$ is equal to the disjoint union of the \approx -pairs and \sim -classes in ∂W . (See Remark 2.4.) Observe that the constructions in Corollary 3.5 and Proposition 3.9 use exactly those points in $\partial W(2+)$ which correspond to geodesic rays that are eventually labelled by elements of some set A satisfying conditions (A1) and (A2) in Proposition 3.9. (Given a \approx -pair associated to an (a, b) -bicoloured bi-infinite geodesic, the set $\{a, b\}$ satisfies (A1) and (A2).) The following proposition shows that such points in $\partial W(2+)$ exhaust all of $\partial W(2+)$. It then follows that there are no additional \sim -classes.

Proposition 3.18. *Let $\gamma^+ \in \partial W$ be represented by a geodesic ray γ in \mathcal{C} based at e . Then $\gamma^+ \in \partial W(2+)$ if and only if γ is eventually labelled by the elements of a set A satisfying properties (A1) and (A2) in Proposition 3.9.*

Proof. The “if” direction follows from the discussion in the paragraph above the statement.

For the converse, let C be the set of generators of W that occur infinitely often as labels of γ (so that γ is eventually labelled by C), and suppose that C fails either (A1) or (A2). To establish that γ^+ is not in $\partial W(2+)$, it is enough to show that $U_n \setminus \{\gamma^+\}$ is connected for all n , where U_n is the subset of ∂W represented by geodesic rays in \mathcal{C} based at e which agree with γ for the first n edges. Given $\eta^+, \mu^+ \in U_n \setminus \{\gamma^+\}$, represented by geodesic rays η and μ in \mathcal{C} which agree with γ for the first n edges, we will construct a path between η^+ and μ^+ in $U_n \setminus \{\gamma^+\}$.

We begin with two special cases. First, suppose that η and μ are both labelled, starting at $\gamma(n)$, exclusively by vertices of an induced cycle α of Γ . If Σ_α denotes the subcomplex corresponding to α based at $\gamma(n)$, then the parts of η and μ beyond $\gamma(n)$ lie in Σ_α . Thus η^+ and μ^+ are in $\partial\Sigma_\alpha$, which is homeomorphic to S^1 (see Remark 3.13). There are two arcs connecting η^+ and μ^+ in $\partial\Sigma_\alpha$, and at least one of these is contained in U_n ; call this arc ν .

If ν does not pass through γ^+ , then it is the desired path between η^+ and μ^+ . Now suppose γ^+ lies on ν . This means that if $V \subseteq U_n \cap \partial\Sigma_\alpha$ is such that $V \setminus \{\gamma^+\}$ has two components V_1 and V_2 , then ν intersects both V_1 and V_2 . We will obtain a detour to ν by constructing a path between V_1 and V_2 in $U_n \setminus \{\gamma^+\}$. Observe that since γ^+ lies on $\nu \subset \partial\Sigma_\alpha$, the set C defined above, of eventual labels of γ , is a subset of the vertices of α .

If C fails (A1), then there exists a pair $\{c_1, c_2\} \subseteq C$ which does not separate $|\Gamma|$. Then c_1 and c_2 are nonadjacent, so the pair $\{c_1, c_2\}$ does not separate Γ . It follows that there is a path β in Γ which connects the two components of $\alpha \setminus \{c_1, c_2\}$ and meets α only at its endpoints. Let τ be the cycle obtained by taking the union of β with the arc of α between the endpoints of β which contains c_1 . The path β can be chosen so that τ is induced. Now for $m \geq n$, let $\Sigma_{\tau, m}$ be the subcomplex corresponding to τ based at $\gamma(m)$. Since c_2 is not on τ , it follows that γ^+ does not lie on $\partial\Sigma_{\tau, m}$. Let x and y be two vertices of α which lie on $\tau \setminus \{c_1\}$ and are such that c_1 and c_2 lie in distinct components of $\alpha \setminus \{x, y\}$. Then $\Sigma_{\tau, m} \cap \Sigma_\alpha$ contains (at least) an (x, y) -bicoloured geodesic ζ_m passing through $\gamma(m)$. We may choose m large enough so that $\partial\zeta_m \subset V$, and at least one arc of $\partial\Sigma_{\tau, m}$ connecting ζ_m^+ and ζ_m^- is in U_n . This arc will yield the required path between V_1 and V_2 , provided we can show that m can be chosen so that $\partial\zeta_m$ has non-trivial intersection with V_1 as well as V_2 .

Assume m is large enough that the labels of γ beyond $\gamma(m-1)$ are in C . Now suppose the two edges of γ incident to $\gamma(m)$ are labelled c and c' (so in particular, c and c' lie on α), and ξ_1, ξ_2 are geodesic rays based at $\gamma(m)$ such that $\xi_1^+, \xi_2^+ \in V$. Observe that by similar arguments to Lemma 3.3, carried out in the subcomplex Σ_α , if the first letters labelling ξ_1 and ξ_2 are in different components of $\alpha \setminus \{c, c'\}$ then one of ξ_1^+ and ξ_2^+ lies in V_1 , and the other lies in V_2 . We use this fact to choose m as follows.

Let C_1 denote the elements of C which lie on the arc of $\alpha \setminus \{x, y\}$ containing c_1 , and let C_2 denote the remainder of C (which may include one or both of x and y). It follows from the definition of C that γ is (eventually) labelled alternately by nonempty words in C_1 and C_2 . Now if m is chosen as in the previous paragraphs with the additional property that at $\gamma(m)$,

the label of γ transitions from a word in C_1 to a word in C_2 , then the above observation can be used to show that $\partial\zeta_m$ intersects both V_1 and V_2 , and therefore, as described above, $\partial\Sigma_{\tau,m}$ contains the desired detour. Together with ν , this yields a path between η^+ and μ^+ in $U_n \setminus \{\gamma^+\}$ when C fails (A1).

If C fails (A2), then there is a subdivided K_4 subgraph Λ of Γ (with $\alpha \subset \Lambda$) and elements $c_1, c_2 \in C$ which lie on distinct branches of Λ . Let β_1 and β_2 be the branches of Λ containing c_1 and c_2 respectively. Without loss of generality, we may assume c_1 is an interior vertex of the branch β_1 . We now apply a similar argument as in the case that C fails (A1), with τ taken to be a cycle in Λ which contains β_1 but not β_2 , and x and y the endpoints of the branch β_1 . This completes the proof of the first special case.

The second special case we consider is when η (respectively μ) is labelled, starting at $\gamma(n)$, exclusively by the vertices of an induced cycle α (respectively β), with $\alpha \neq \beta$. By Standing Assumptions 1.1 and an elementary graph-theoretic argument, there exists a sequence of induced cycles $\alpha = \sigma_1, \sigma_2, \dots, \sigma_k = \beta$ such that every consecutive pair of cycles intersects in at least a pair of non-adjacent vertices. We claim that the σ_i can be chosen so that for all $1 \leq i < k$, the intersection of σ_i and σ_{i+1} generates a group with infinitely many ends (that is, strictly speaking, this intersection is the defining graph of a special subgroup with infinitely many ends).

To see this claim, observe that if the intersection of consecutive cycles generates a 2-ended group, then the intersection is exactly a pair of non-adjacent vertices or a pair of adjacent edges. For each pair σ_i, σ_{i+1} which intersect in a pair of non-adjacent vertices, add an extra cycle between σ_i and σ_{i+1} consisting of one arc from σ_i connecting these vertices and one from σ_{i+1} . Then the intersection of this new cycle with σ_i or σ_{i+1} either generates a group with infinitely many ends, or consists of a pair of adjacent edges. Thus after reindexing we have a sequence in which we only need to deal with the latter case. Now if σ_i and σ_{i+1} intersect in a subpath ϵ consisting of a pair of edges, then adding the cycle obtained by deleting the interior of ϵ from $\sigma_i \cup \sigma_{i+1}$ results in the desired intersections between successive cycles in the sequence.

Assume that $\sigma_1, \dots, \sigma_k$ have been chosen to satisfy the above claim, and let Σ_i be the subcomplexes corresponding to σ_i based at $\gamma(n)$. Now for all $1 < i < k$ choose a geodesic ray η_i in \mathcal{C} based at e , which agrees with γ for the first n edges, lies in $\Sigma_i \cap \Sigma_{i+1}$ beyond $\gamma(n)$, and is not equal to γ (this last criterion is possible because by construction $\Sigma_i \cap \Sigma_{i+1}$ contains a tree with infinitely many ends). Finally, define $\eta_1 = \eta$ and $\eta_k = \mu$. Observe that $\eta_i^+ \in U_n \setminus \{\gamma^+\}$ for all i , and moreover, η_i^+ and η_{i+1}^+ satisfy the hypotheses of the first special case above. Now we can construct the desired path between η^+ and μ^+ by concatenating the paths obtained above between η_i^+ and η_{i+1}^+ for all i .

In the general case, we consider arbitrary $\mu_1^+, \mu_2^+ \in U_n \setminus \{\gamma^+\}$. Write μ^+ for either μ_1^+ or μ_2^+ . It is enough to construct a path in $U_n \setminus \{\gamma^+\}$ from μ^+ to some $\eta^+ \in U_n \setminus \{\gamma^+\}$, where η is labelled, starting at $\gamma(n)$, exclusively by vertices of an induced cycle α . Suppose μ is labelled, starting at $\gamma(n)$, by $w_1 w_2 w_3 \dots$ such that w_i is a word in the letters contained in some cycle β_i for all i . We define geodesic rays η_i interpolating between η and μ as follows: $\eta_0 = \eta$ and for $i > 0$, η_i agrees with γ until $\gamma(n)$, and is labelled $w_1 w_2 \dots w_i u_i$ beyond $\gamma(n)$, where u_i is an infinite word in β_i chosen so that the result is a geodesic ray not equal to γ .

Then $\eta_i^+ \in U_n \setminus \{\gamma^+\}$ for all i . Beyond $\gamma(n)$, the geodesics η_i and η_{i+1} agree on a segment labelled $w_1 w_2 \dots w_i$, after which η_i is labelled by β_i and η_{i+1} is labelled by β_{i+1} . An argument similar to the second special case (with an appropriate change of base point) shows that we can construct a path in $U_n \setminus \gamma^+$ between η_i^+ and η_{i+1}^+ for all i . This completes the proof of the proposition, as concatenating the paths obtained between η_i^+ and η_{i+1}^+ for each i results in a continuous path in $U_n \setminus \{\gamma^+\}$ between $\eta^+ = \eta_0^+$ and μ^+ . \square

Example 3.19. We can now describe the \sim -classes for Figures 3.1 and 3.2, and the \approx -pairs to which they are adjacent in the pre-tree T (see also Example 3.11 and Remark 3.17).

First consider Figure 3.1. In the left-hand graph, there are three W -orbits of infinite \sim -classes, corresponding to the three branches between a and b which are of length at least three. In the right-hand graph, there are four W -orbits of infinite \sim -classes, corresponding to the four branches between a and b . In both cases, there are no \sim -pairs, all infinite \sim -classes are adjacent in T to \approx -pairs, and all \approx -pairs correspond to (a, b) -bicoloured geodesics.

On the left of Figure 3.2, the W -orbits of infinite \sim -classes correspond to the three branches between a and d , the two branches between a and e , and the set $\{a, d, f, g, e\}$, and there are no \sim -pairs. The infinite \sim -classes for branches between a and d (respectively, a and e) are adjacent in T to \approx -pairs of valence 4 corresponding to (a, d) -bicoloured geodesics (respectively, valence 3 corresponding to (a, e) -bicoloured geodesics), while the infinite \sim -classes for the set $\{a, d, f, g, e\}$ are adjacent to \approx -pairs of both kinds.

In the centre of Figure 3.2, there is a W -orbit of infinite \sim -classes for each of the five branches of length three, and a W -orbit of infinite \sim -classes corresponding to the set $\{a_1, a_2, a_3, a_4, a_5\}$. There are no \sim -pairs. All \approx -pairs correspond to (a_i, a_{i+1}) -bicoloured geodesics. For each i , the infinite \sim -classes corresponding to the branch between a_i and a_{i+1} are adjacent in T to \approx -pairs of valence 4 corresponding to (a_i, a_{i+1}) -bicoloured geodesics. The infinite \sim -classes for $\{a_1, a_2, a_3, a_4, a_5\}$ are adjacent in T to all five kinds of \approx -pairs.

On the right of Figure 3.2, there is a W -orbit of infinite \sim -classes corresponding to each branch of length three, as well as three W -orbits of \sim -pairs, corresponding to the set $\{p, q\}$, the branch between u and v , and the branch between r and s . There are no \approx -pairs.

The following results and observations concerning sets A which correspond to \sim -classes will be used to identify the stabilisers of \sim -classes, and in later sections. In Lemma 3.20 we prove that if A satisfies (A1)–(A3) and A contains an interior vertex of a branch of Γ , then A contains all vertices of this branch. This is used to establish Corollary 3.21, which says that A contains an essential cut pair $\{a, b\}$ such that a and b are consecutive in the cyclic order on A . Remark 3.22 then records some implications for Σ .

Lemma 3.20. *Let A be a set of vertices of Γ satisfying properties (A1), (A2) and (A3), and such that $\langle A \rangle$ is infinite. Suppose an element $a \in A$ is an interior vertex of a branch of Γ . Then A contains all vertices of this branch, including its endpoints.*

Proof. Let α be an induced cycle in Γ containing all elements of A . Let β be the branch of Γ containing a and let b and b' be the essential vertices of Γ which are its endpoints.

First consider a non-essential vertex $c \neq a$ which also lies between b and b' on α . Then $\{a, c\}$ separates $|\Gamma|$, and for all $a' \in A \setminus \{a\}$, the pair $\{a, a'\}$ separates $|\Gamma|$ if and only if the pair $\{c, a'\}$ separates $|\Gamma|$. Thus (A1) holds for $A \cup \{c\}$.

Now suppose there is a subgraph Λ of Γ which is a subdivided K_4 and contains at least three vertices of $A \cup \{c\}$. Then Λ contains at least three vertices of A , or Λ contains c (these cases are not mutually exclusive). In the former case, by (A1) all vertices of A lie on the same branch of Λ . Since a and c are both interior vertices of β , the vertex c must lie on this branch of Λ as well, and so in this case (A2) holds for $A \cup \{c\}$.

In the case that Λ contains c , we claim that all elements of A lie on the same branch of Λ as c . Since c is non-essential, Λ must contain the entire branch β , and so Λ contains a as well. Now suppose there exists $a' \in A$ which does not lie on the branch of Λ containing a and c . Then $\Lambda \setminus \{a, a'\}$ is connected but $\{a, a'\}$ separates $|\Gamma|$, so there is a path between a and a' which intersects Λ only at a , and possibly a' , which contradicts a being non-essential. This proves the claim. Thus (A2) holds for $A \cup \{c\}$ in this case as well. We have shown that both (A1) and (A2) hold for the set $A \cup \{c\}$, and so by (A3) we conclude that $c \in A$.

To show that b and b' are both in A , notice first that $\{b, b'\}$ separates $|\Gamma|$, and for all vertices c in the interior of the branch β , the pairs $\{b, c\}$ and $\{b', c\}$ separate $|\Gamma|$. Consider the case that every vertex of A lies on the branch β . Then (A1) holds for $A \cup \{b, b'\}$ by the previous observations. For (A2), if any subgraph Λ which is a subdivided copy of K_4 contains a triple of vertices of A then Λ must contain a non-essential vertex of β . It follows that a branch of Λ contains all of β , and so all vertices of $A \cup \{b, b'\}$ lie on the same branch of Λ . Thus (A2) also holds for $A \cup \{b, b'\}$ in this case, and so by (A3) we have $b, b' \in A$.

Now suppose there is some $a' \in A$ which does not lie on the branch β . If $\{a', b\}$ does not separate $|\Gamma|$ then a' and b are non-adjacent and there must be a path η in Γ which connects the two components of $\alpha \setminus \{a', b\}$. Since $\{a, a'\}$ separates $|\Gamma|$, there is a path between a and a' which intersects $\alpha \cup \eta$ only at its endpoints. This contradicts the fact that a is not essential. So $\{a', b\}$, and similarly $\{a', b'\}$, separates $|\Gamma|$. Thus $A \cup \{b, b'\}$ satisfies (A1).

For (A2), to avoid trivial cases we may assume by contradiction that Γ has a subdivided K_4 subgraph Λ which contains three vertices of $A \cup \{b, b'\}$, so that for some vertex a' of A in this triple, a' and b lie in different branches of Λ . If b' is also on Λ , then $\Lambda \cup \beta \setminus \{a, a'\}$ is connected, and as before, this contradicts the fact that a is not essential. Finally, if b' is not on Λ , then there is another vertex a'' of A which is on Λ . If a'' is between a' and b on $\alpha \setminus \beta$, then $\Lambda \cup \alpha \setminus \{a, a''\}$ is connected, and by the same argument as before, this contradicts the fact that a is not essential. If a' is between a'' and b on $\alpha \setminus \beta$, then using a' instead of a'' in the previous sentence, we again have a contradiction. This completes every case, and shows that $A \cup \{b, b'\}$ satisfies (A2). By (A3) it follows that $b, b' \in A$ in this case as well. \square

Corollary 3.21. *Let A be a set of vertices of Γ satisfying properties (A1), (A2) and (A3), and such that $\langle A \rangle$ is infinite. Then A contains an essential cut pair $\{a, b\}$ such that a and b are consecutive in the cyclic order on A .*

Proof. We first show that A contains an essential cut pair. If A contains an interior vertex of a branch of Γ , then by Lemma 3.20, A contains the endpoints of this branch. These are an essential cut pair. Otherwise, A consists only of essential vertices. Then as $\langle A \rangle$ is infinite A must contain a pair of non-adjacent essential vertices, which by (A1) are a cut pair.

If $\text{card}(A) = 2$ or 3 , let $\{a, b\}$ be an essential cut pair in A . It is then immediate that a and b are consecutive in any cyclic order on A . Now assume that $\text{card}(A) = n \geq 4$. Let α be an induced cycle containing all of A , and inducing the cyclic order a_1, \dots, a_n on A .

We first show that α must contain a vertex which is not in A . If every vertex of α is in A , then as Γ is not a cycle and has no separating vertices or edges, there is a path η which connects a non-adjacent pair a_i and a_j in A , so that η intersects the cycle α only at a_i and a_j . By slightly extending η to the midpoints of edges of α incident to a_i and a_j , we obtain a path which contradicts Lemma 3.14(2). Thus α must contain a vertex which is not in A .

We now have that α contains a vertex $c \notin A$. Then without loss of generality c lies on the subpath of α between a_1 and a_2 which has no other elements of A . If a_1 is non-essential then by Lemma 3.20, A and thus α contains all vertices of the branch on which a_1 lies. Using Lemma 3.20 again, this contradicts either $c \notin A$ or there being no vertex of A between c and a_1 on α . Thus a_1 is essential, and similarly a_2 is essential. By (A1), it now suffices to show that a_1 and a_2 are not adjacent in Γ . If there is an edge ϵ of Γ with endpoints a_1 and a_2 , then since ϵ doesn't separate Γ , there is a path η connecting the two components of $\alpha \setminus \{a_1, a_2\}$ which meets α only at its endpoints. Now $\alpha \cup \eta \cup \epsilon$ is a subdivided K_4 which does not contain all elements of A on a single branch. This contradicts (A2). Therefore a_1 and a_2 are not adjacent in Γ , and so $\{a_1, a_2\}$ is the desired essential cut pair. \square

Remark 3.22. Let A be a set of vertices of Γ with $\langle A \rangle$ infinite and satisfying properties (A1), (A2) and (A3). Write K_A for the Davis complex chamber for the special subgroup $\langle A \rangle$. Developing this chamber and then imposing the cellulation by big squares gives a subcomplex Σ_A of Σ with 1-skeleton the copy of the Cayley graph \mathcal{C}_A of A which contains the identity. For example, in Figure 3.2, if $A = \{a, b, c, d\}$ then Σ_A contains big squares with edges labelled by commuting generators, and so Σ_A properly contains \mathcal{C}_A , while if $A = \{a_1, a_2, a_3, a_4, a_5\}$ then K_A is a star graph of valence 5, and $\mathcal{C}_A = \Sigma_A$ is a tree. Each coset of $\langle A \rangle$ in W also corresponds to some copy of Σ_A and to some copy of \mathcal{C}_A , with \mathcal{C}_A the 1-skeleton of Σ_A .

Now fix a copy of Σ_A , and let α be an induced cycle in Γ containing all vertices of A . We define \mathcal{A}^+ to be the set of geodesics in $\mathcal{C}_A \subseteq \Sigma_A$ which are bicoloured by essential cut pairs $\{a, b\}$ in A so that a and b are consecutive in the cyclic order on A . By Corollary 3.21, the set \mathcal{A}^+ is nonempty. Let γ be a geodesic in \mathcal{A}^+ , with γ bicoloured by a and b . Then by definition of \mathcal{A}^+ , if $A \setminus \{a, b\}$ is nonempty then all vertices of $A \setminus \{a, b\}$ lie in the same component of $\alpha \setminus \{a, b\}$, and thus in the same component of $\Gamma \setminus \{a, b\}$. It follows from Lemma 3.3 that $\Sigma_A \setminus \gamma$, if nonempty, is contained in a single component of $\Sigma \setminus \gamma$. We thus refer to the elements of \mathcal{A}^+ as the *frontier geodesics* of Σ_A . Now observe that for any component U of $\Sigma \setminus \Sigma_A$, there is a unique frontier geodesic $\gamma \in \mathcal{A}^+$ such that U is a component of $\Sigma \setminus \gamma$. We then say that γ is the frontier geodesic *corresponding* to the component U of $\Sigma \setminus \Sigma_A$.

We finish this section by determining the stabilisers of \sim -classes.

Proposition 3.23. *Let v be a \sim -class in T , and let A be the corresponding set of vertices of Γ as in Proposition 3.9, so that $v = \partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx$ for some copy \mathcal{C}_A of the Cayley graph of $\langle A \rangle$.*

- (1) *If v is a \sim -pair, let γ be an (a, b) -bicoloured geodesic such that $v = \partial\gamma = \partial\mathcal{C}_A$, hence $a, b \in A$. Then the stabiliser of v is a conjugate of $\langle a, b \rangle$, if there is no vertex c of Γ adjacent to both a and b , or of $\langle a, b, c \rangle$, if there is such a vertex c .*
- (2) *If v is an infinite \sim -class, then the stabiliser of v is a conjugate of $\langle A \rangle$.*

Proof. The proof of (1) is similar to that for stabilisers of \approx -pairs in Lemma 3.7. Now suppose v is an infinite \sim -class, so $\langle A \rangle$ is not 2-ended, and let \mathcal{C}_A be the copy of the Cayley graph of $\langle A \rangle$ which contains the identity. It suffices to show that the stabiliser of v is $\langle A \rangle$.

Suppose that some $g \in W \setminus \langle A \rangle$ stabilises v . Then g stabilises $\partial\mathcal{C}_A = \partial\Sigma_A$ and the set of endpoints of the frontier geodesics \mathcal{A}^+ , since these bound components of $\Sigma \setminus \Sigma_A$. Let γ be an (a, b) -bicoloured geodesic in \mathcal{A}^+ which passes through the identity. Then $g\gamma$ is not in \mathcal{A}^+ since $g \notin \langle A \rangle$. Hence there is a geodesic $\gamma' \in \mathcal{A}^+$ so that γ' and $g\gamma$ have the same endpoints. By Corollary 3.2, γ' is also (a, b) -bicoloured, and $g\gamma = \gamma'c$ where c is a (unique) vertex of Γ adjacent to both a and b . If there is no such vertex c we are done, otherwise let $w \in \langle A \rangle$ be the label on a shortest path from e to γ' , so that $\gamma' = w\gamma$. Then $g\gamma = w\gamma c$. Since γ passes through the identity $\gamma c = c\gamma$, so $g\gamma = wc\gamma$. As the stabiliser of γ is $\langle a, b \rangle$, it follows that $g = wcz$ where $w \in \langle A \rangle$ and $z \in \langle a, b \rangle \leq \langle A \rangle$. Since we assumed $g \notin \langle A \rangle$, we have $c \notin A$. However g , w and z all stabilise v , so c must as well.

Now let $a' \in A \setminus \{a, b\}$ be such that a' does not commute with both a and b (since $\langle A \rangle$ is not 2-ended, such an a' exists). Using properties (A1) and (A2), we see that a' cannot be adjacent to c . Consider the geodesic $ca'\gamma$. Since $a'\gamma \in \mathcal{A}^+$, $c \notin A$ and c stabilises v , by similar reasoning to the previous paragraph there is a $\gamma'' \in \mathcal{A}^+$ with the same endpoints as $ca'\gamma$, and a vertex x of Γ which is adjacent to both a and b so that $ca'\gamma = \gamma''x$. Now Γ has no squares, so $x = c$ and thus $ca'\gamma = \gamma''c$. As $\gamma'' \in \mathcal{A}^+$ and is (a, b) -bicoloured, we can write $\gamma'' = w'\gamma$ where $w' \in \langle A \rangle$ is the label on a shortest path from γ to γ'' . Then $ca'\gamma = w'\gamma c = w'c\gamma$, hence $c(w')^{-1}ca' \in \langle a, b \rangle$. It follows that w' contains an instance of a' , and every letter in w' commutes with c . But a' does not commute with c , so we have obtained a contradiction. We conclude that the stabiliser of v is $\langle A \rangle$. \square

3.6. Construction of certain stars. Recall that a star is a maximal set X of vertices in the pretree T such that given any two vertices in X , no vertex in T is between them, in the sense of Definition 2.5. Bowditch shows that stars of size at least 3 are in fact infinite. In Proposition 3.24 below, we construct certain infinite stars. We will show in Section 3.7 that we have identified all of the stars of size at least 3, and therefore all of the Type 3 vertices.

Proposition 3.24 (Stars of size at least 3). *Let B be a set of essential vertices of Γ satisfying the following properties:*

- (B1) *if $C = \{c_1, c_2\}$ is any pair of essential vertices of Γ (which possibly intersects B), then $B \setminus C$ is contained in a single component of $\Gamma \setminus C$;*
- (B2) *the set B is maximal with respect to (B1); and*
- (B3) *$\text{card}(B) \geq 4$.*

Then B corresponds to a W -orbit of stars of size at least 3 as follows: let \mathcal{G}_B denote the set of (b_i, b_j) -bicoloured geodesics in \mathcal{C} passing through e , where $\{b_i, b_j\} \subseteq B$ is a cut pair. Let \mathcal{B} be the $\langle B \rangle$ -orbit of \mathcal{G}_B . Then for each $\eta \in \mathcal{B}$, the endpoints $\partial\eta$ are contained in a unique \approx -pair or \sim -class, and this containment induces a map $\Phi : \mathcal{B} \rightarrow T$ so that $\Phi(\mathcal{B})$ corresponds to a star of size at least 3 in T .

Example 3.25. In Figure 3.2, the set $\{a, d, e\}$ in the left-most graph satisfies (B1) and (B2) but not (B3), and since there are only three essential vertices condition (B3) will never be satisfied. Any subset of $\{a_1, a_2, a_3, a_4, a_5\}$ in the middle graph which contains at least 4

vertices fails (B1). The sets $\{u, v, p, q\}$ and $\{p, q, r, s\}$ in the right-most graph satisfy (B1), (B2) and (B3), and each pair of vertices within these sets is a cut pair in Γ .

We begin with some preliminaries for the proof of Proposition 3.24.

Observation 3.26. Let B be a set of essential vertices of Γ which satisfies property (B1) of Proposition 3.24, and y be an essential vertex of Γ . If there exist paths η, μ and ν in Γ from y to B which intersect only at y , then $B \cup \{y\}$ satisfies property (B1). This is because given any essential vertices $c_1, c_2 \neq y$, at least one of the paths η, μ and ν misses $\{c_1, c_2\}$, so that y is in the same component as $B \setminus \{c_1, c_2\}$ in $\Gamma \setminus \{c_1, c_2\}$.

Lemma 3.27. Let B be a set of essential vertices satisfying (B1), (B2) and (B3). Then there exists a cut pair in Γ consisting of vertices $\{b, b'\} \subset B$.

Proof. By items (4) and (5) in Standing Assumptions 1.1, we have that Γ is not a cycle, and has at least one cut pair. Since we may replace each vertex in the pair with a nearest essential vertex if necessary, it follows that Γ has at least one essential cut pair. Thus if B contains all the essential vertices of Γ , the lemma follows.

We may now assume that $\Gamma \setminus B$ contains an essential vertex. Let x be an essential vertex in $\Gamma \setminus B$ “nearest to B ” in the sense that there exists $b \in B$ and a path τ in Γ between x and b which does not contain any essential vertices in its interior. Since $x \notin B$, there exists, by property (B2), a pair of essential vertices which separate x from B . One of the vertices in the pair must be b , for otherwise τ would connect x to B . Let c be the other vertex, so that $\Gamma \setminus \{b, c\}$ contains disjoint components Λ_x and Λ_B , with $x \in \Lambda_x$ and $B \setminus \{b, c\} \subset \Lambda_B$. If $c \in B$, then put $b' = c$ to obtain the desired cut pair $\{b, b'\} \subset B$.

If $c \notin B$, choose a reduced path η from c to B through Λ_B , which meets B only at its endpoint $b' \in B$. Since $b \notin \Lambda_B$, we know that $b \neq b'$. We claim that $\{b, b'\}$ is the desired pair. If every path from c to B enters through either b or b' , then clearly $\{b, b'\}$ separates Γ (putting c and $B \setminus \{b, b'\}$ in different components). Assume, therefore, that there is a reduced path ν from c to B , which intersects B only in its endpoint $b'' \neq b, b'$. It follows that ν does not intersect Λ_x .

There exists a path μ from c to b in Λ_x which is disjoint from η and ν . Since $c \notin B$, Observation 3.26 implies that the paths η and ν must intersect. Let y be the last point of $\eta \cap \nu$ encountered while traveling along η from c to b' (with y possibly equal to c), so that the segment $\eta_{[yb']}$ meets ν only at y , as in Figure 3.4.

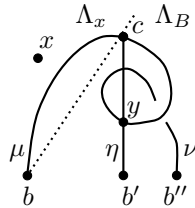


FIGURE 3.4.

We now have three paths from y to B which meet only at y : $\eta_{[yb']}$, $\nu_{[yb']}$, and the concatenation of $\nu_{[yc]}$ with $\mu = \mu_{[cb]}$. Then by Observation 3.26, we have $y \in B$, which contradicts our choice of η . It follows that $\{b, b'\}$ separates Γ . \square

Lemma 3.28. *Let B be a set of essential vertices satisfying (B1), (B2) and (B3). Suppose that B contains three distinct vertices b_1, b_2 and b_3 which lie on an induced cycle σ in Γ . Then b_1, b_2 and b_3 lie on a subdivided K_4 subgraph of Γ . Moreover, the set $\{b_1, b_2, b_3\}$ intersects at least two branches of this subgraph.*

Proof. By property (B3), $\text{card}(B) \geq 4$, so there is $b_4 \in B$ distinct from b_1, b_2 and b_3 .

If b_4 lies on σ , assume without loss of generality that b_1 and b_3 separate σ into two components, one containing b_2 and the other containing b_4 . By property (B1), b_2 and b_4 lie in the same component of $\Gamma \setminus \{b_1, b_3\}$, so there is a path η which connects the two components of $\sigma \setminus \{b_1, b_3\}$. We may assume η is reduced and meets σ only in a pair of essential vertices x and y . Now b_1 and b_3 lie in different components of $(\sigma \cup \eta) \setminus \{x, y\}$, but by property (B1), they lie in the same component of $\Gamma \setminus \{x, y\}$. Thus there is a reduced path τ connecting the component of $\sigma \setminus \{x, y\}$ containing b_1 with that containing b_3 . Then if τ and η are disjoint, $\sigma \cup \tau \cup \eta$ is a subdivided K_4 subgraph containing b_1, b_2 and b_3 , and otherwise $\sigma \cup \tau \cup \eta$ contains such a subgraph. It is clear from the construction that b_1, b_2 , and b_3 lie on at least two distinct branches of this subgraph.

On the other hand, if b_4 is not on σ then since Γ is connected, there is at least one (reduced) path τ connecting b_4 to σ . If τ first meets σ at $z \neq b_1, b_2, b_3$, then by Observation 3.26 and property (B2) we conclude that $z \in B$, and apply the previous paragraph with $b_4 = z$. So we may assume that every reduced path from b_4 to τ first meets σ at b_1, b_2 or b_3 . Now using the same argument as in the preceding paragraph, we conclude that for $1 \leq i \leq 3$, there is a reduced path τ_i connecting b_4 to b_i in the complement of b_{i+1} and b_{i+2} (indices taken mod 3). It is easy to see that $\sigma \cup \tau_1 \cup \tau_2 \cup \tau_3$ either is or contains a subdivided K_4 subgraph. Again, it is also clear that b_1, b_2 , and b_3 lie on at least two distinct branches of this subgraph. \square

Lemma 3.29. *Let A and B be sets of vertices of Γ as in the statements of Propositions 3.9 and 3.24 respectively. Then $\text{card}(A \cap B) \leq 2$.*

Proof. Suppose there are three distinct elements b_1, b_2 and b_3 in $A \cap B$. By Lemma 3.12 there is an induced cycle α containing b_1, b_2 and b_3 . Then by Lemma 3.28, b_1, b_2 and b_3 lie on at least two distinct branches of a subdivided K_4 subgraph of Γ . This is a contradiction, since A satisfies property (A2) of Proposition 3.9. \square

Lemma 3.30. *If B is a set of essential vertices satisfying (B1) and (B2) with at least two elements, and if z is an arbitrary vertex in $\Gamma \setminus B$, then (at least) one of the following is true:*

- (1) *there exists a pair of essential vertices $\{c_1, c_2\}$ and an element $b \in B$ such that z and b lie in different components of $\Gamma \setminus \{c_1, c_2\}$; or*
- (2) *the vertex z is not essential and lies on a branch between a pair of vertices b_1 and b_2 in B .*

Proof. If z is essential, then (1) follows easily from (B1) and (B2). Otherwise z lies on a branch between two essential vertices c_1 and c_2 . If $\{c_1, c_2\} \subseteq B$ then (2) holds. If not, then since $\text{card}(B) \geq 2$, there exists $b \in B$ such that $b \notin \{c_1, c_2\}$, and (1) holds, because b is necessarily essential, and so cannot lie on the branch between c_1 and c_2 containing z . \square

We are now ready to prove Proposition 3.24.

Proof of Proposition 3.24. Given $\eta \in \mathcal{B}$, by definition of the sets \mathcal{G}_B and \mathcal{B} and Lemma 3.3, the endpoints $\partial\eta$ separate ∂W , so that $\partial\eta \subset \partial W(2+)$. We can thus define a map $\Phi : \mathcal{B} \rightarrow T$ by letting $\Phi(\eta)$ be the unique \approx -pair or \sim -class containing $\partial\eta$. This map is well-defined since Bowditch shows that the \approx -pairs and \sim -classes partition $\partial W(2+)$.

We claim that:

- (a) The map $\Phi : \mathcal{B} \rightarrow T$ satisfies $\text{card } \Phi^{-1}(v) \leq 2$ for all $v \in T$.
- (b) No vertex in T is between any pair of vertices in $\Phi(\mathcal{B})$.
- (c) The set $\Phi(\mathcal{B})$ is maximal with respect to condition (b).

Note that \mathcal{G}_B is nonempty by Lemma 3.27, and since $\text{card}(B) \geq 4$ it follows that $\text{card } \mathcal{B} = \infty$. This, together with conditions (a), (b) and (c), implies that $\Phi(\mathcal{B})$ corresponds to a star of size at least 3 in T .

Proof of (a): Suppose $\Phi(\eta) = v \in T$ where $\eta \in \mathcal{B}$ is labelled by b_1 and b_2 . If v is a \approx -pair or a \sim -pair, and $\eta' \in \mathcal{B}$ with $\Phi(\eta') = v$, then $\partial\eta = \partial\eta'$. As both η and η' are bicoloured, it follows from Corollary 3.2 that there is at most one possibility for η' distinct from η .

Now suppose v is an infinite \sim -class, and $\Phi(\eta') = v$. By left-multiplying by an element of $\langle B \rangle$ if necessary, we may assume that η passes through the identity, and that $\eta' = w\gamma$, where $w \in \langle B \rangle$ and $\gamma \in \mathcal{G}_B$ passes through the identity. Since $\partial\eta$ and $\partial\eta' = \partial(w\gamma)$ are part of the same \sim -class, the supports of w, η_1 and γ are all part of the same set A as in Proposition 3.9 giving rise to this \sim -class. Then by Lemma 3.29 the union of these supports contains at most two elements. It follows that $\eta = \eta'$.

Thus $\text{card } \Phi^{-1}(v) \leq 2$ for all $v \in T$.

Proof of (b): Let $\eta, \eta' \in \mathcal{B}$ be such that $\Phi(\eta) = v$, $\Phi(\eta') = v'$ and $v \neq v'$. Suppose by way of contradiction that there is a vertex v'' of T between the vertices v and v' .

First assume v'' is either a \sim - or \approx -pair. Then v'' corresponds to $\partial\mu$, where μ is a (c_1, c_2) -bicoloured bi-infinite geodesic in \mathcal{C} , so that $\Sigma \setminus \mu$ has at least two components. Since v'' is between $v = \Phi(\eta)$ and $v' = \Phi(\eta')$, by Remark 2.6 the sets $\partial\eta$ and $\partial\eta'$ are in different components of $\partial W \setminus \partial\mu$. Thus η and η' are in different components of $\Sigma \setminus \mu$.

Assume without loss of generality that η passes through e , and that η' is of the form $w\gamma$, where w is a (possibly empty) reduced word in B labelling a shortest path ν from e to η' , and γ is a bi-infinite geodesic passing through e . Now ν necessarily crosses μ , and could potentially travel along it for some distance. Let p and p' be the points where it first meets μ and leaves μ , respectively, and let u (respectively u') be the reduced word labelling the path starting at p (respectively p') and ending at η (respectively η') along ν . Then every vertex on η (respectively η') can be reached from μ by a path labelled by u (respectively u'), followed by an alternating word in B . (This allows the possibility that u or u' is empty.) By the description of the components of $\Sigma \setminus \mu$ in Lemma 3.3, we see that all such paths lie in the same component of $\Sigma \setminus \mu$, since the initial letter labelling any such path is in B , and by property (B1), $B \setminus \{c_1, c_2\}$ lies in a single component of $\Gamma \setminus \{c_1, c_2\}$. This is a contradiction. Thus the vertex v'' cannot be either a \sim -pair or a \approx -pair.

If v'' is an infinite \sim -class, then it comes from a set A as in Proposition 3.9 so that $\langle A \rangle$ is not 2-ended. More precisely, \mathcal{C} contains a copy \mathcal{C}_A of the Cayley graph of $\langle A \rangle$ so that v'' is the set $\partial\mathcal{C}_A \setminus \partial\mathcal{A}_\approx = \partial\mathcal{C}_A \setminus \partial W(3+)$. As discussed in Remark 3.22, this copy of \mathcal{C}_A

is contained in a copy of Σ_A , the subcomplex of Σ corresponding to the special subgroup generated by A . Recall from Remark 3.22 that \mathcal{A}^+ denotes the set of frontier geodesics of Σ_A , and that there is a unique frontier geodesic corresponding to any component of $\Sigma \setminus \Sigma_A$.

We claim that v'' being between v and v' implies that either η and η' are both in \mathcal{A}^+ , or there exists $\mu \in \mathcal{A}^+$ such that $\partial\mu$ separates $\partial\eta$ from $\partial\eta'$. For this, assume that $\eta \notin \mathcal{A}^+$. By [9, Lemma 3.18] (see Remark 2.6), η then cannot intersect any geodesic in \mathcal{A}^+ . Thus η lies entirely in a single component of $\Sigma \setminus \Sigma_A$. If η' is also in this component, or η' is the frontier geodesic for this component, then no pair of points in $\partial\mathcal{C}_A \setminus \partial W(3+)$ would separate $\partial\eta$ from $\partial\eta'$, contradicting the fact that v'' is between v and v' . Thus η and η' are either in distinct components of $\Sigma \setminus \Sigma_A$, or if $\eta' \in \mathcal{A}^+$ it corresponds to a different component of $\Sigma \setminus \Sigma_A$ to that containing η . Then if μ is the geodesic in \mathcal{A}^+ corresponding to the component containing η , it is clear that $\partial\mu$ separates $\partial\eta$ from $\partial\eta'$. This proves the claim.

In the case that there exists μ in \mathcal{A}^+ such that $\partial\mu$ separates $\partial\eta$ from $\partial\eta'$, the proof proceeds as in the case that v'' is in $V_1(\mathcal{T})$. Otherwise, η and η' are both in $\mathcal{A}^+ \subset \mathcal{C}_A$. By construction η and η' are in $\mathcal{B} \subset \mathcal{C}_B$. Since \mathcal{C}_A (respectively \mathcal{C}_B) is convex, the shortest path between η and η' is labelled by a word in A (respectively B). Thus the support of this path, together with the supports of η and η' , are in $A \cap B$. But the union of these supports must have at least three elements, which contradicts Lemma 3.29.

Thus we have shown that no vertex of T is between v and v' .

Proof of (c): Consider an arbitrary vertex $v \in T \setminus \Phi(\mathcal{B})$. We wish to show that $\{v\} \cup \Phi(\mathcal{B})$ fails (b), i.e. that there is some vertex $v_\eta = \Phi(\eta)$ in $\Phi(\mathcal{B})$ such that there exists a vertex $v' \in V_1(\mathcal{T}) \cup V_2(\mathcal{T})$ between v and v_η .

Since v is in T , there is a bi-infinite geodesic μ which is bicoloured by a cut pair $\{x, y\}$ in Γ , such that μ separates Σ and $\partial\mu \subseteq v \subset \partial W$. Let w be a (possibly empty) reduced word labelling a shortest path from e to μ , so that every vertex along μ in \mathcal{C} is labelled by a group element consisting of w followed by a word in x and y . Since v is not in $\Phi(\mathcal{B})$, either a letter in w or x or y is not in B . Starting from e along the chosen path to μ , let z denote the first letter encountered that is not in B . Thus either z is some letter used in w or it is one of x and y . Let u denote the (possibly empty) subword of w before the occurrence of z .

By Lemma 3.30, there exists an essential cut pair $\{c_1, c_2\}$ (possibly intersecting B) which separates z from $B \setminus \{c_1, c_2\}$. Let ν be the (c_1, c_2) -bicoloured geodesic passing through the vertex of \mathcal{C} labelled u (or e if u is empty). By Lemma 3.3, ν separates Σ , and so $\partial\nu$ is part of the set defining some vertex v' in T .

Now by Lemma 3.27, there exists a cut pair $\{b_1, b_2\} \subset B$. We will construct a (b_1, b_2) -bicoloured bi-infinite geodesic η based at the vertex of \mathcal{C} labelled u' , where u' is chosen so that the length of a shortest path between ν and η is at least 2, as follows. Since $\text{card}(B) \geq 4$, there exist $b_3, b_4 \in B$ distinct from b_1, b_2 . If u is empty, then $u' = b_3b_4$, and if u has length 1, then u' is one of b_3 or b_4 , chosen to be distinct from u (so that uu' is reduced). Otherwise, u' is empty (and η is based at e). It is easy to verify that the shortest path from ν to η passes through e and has uu' as its label. Now η defines a vertex $v_\eta = \Phi(\eta)$ corresponding to $\partial\eta$.

We claim that v' and v_η are distinct. Suppose not. Then if $v' = v_\eta$ is a \sim - or \approx -pair, the geodesics ν and η are parallel, that is, $\{c_1, c_2\} = \{b_1, b_2\}$ and uu' consists of a single letter, which is a contradiction. Thus $v' = v_\eta$ is an infinite \sim -class, in which case the labels of

ν, η, u , and u' are contained in a set A as in Proposition 3.9. Now by Lemma 3.29, we have that $\text{card}(A \cap B) \leq 2$, which is a contradiction, since this intersection contains b_1, b_2, b_3 , and b_4 . This proves the claim.

Finally, we observe that v' is between v_η (which is in $\Phi(\mathcal{B})$) and v . This is a consequence of Definition 2.5 and Lemma 3.3, because at least one half of μ can be reached by a path with first letter z , and at least one half of η can be reached by a path with first letter in $B \setminus \{c_1, c_2\}$. It follows that $\Phi(\mathcal{B})$ is maximal with respect to property (b), as desired.

The above shows that the set $\Phi(\mathcal{B}) \subset V_1(\mathcal{T}) \cup V_2(\mathcal{T})$ defined above yields one star. By acting on this set on the left by elements of W , one obtains a W -orbit of stars. \square

3.7. Identification of all stars and their stabilisers. In this section we prove that we have identified all of the stars in the JSJ tree \mathcal{T} , and determine their stabilisers.

Proposition 3.31. *Every star of size at least 3 in T comes from a set B satisfying conditions (B1), (B2) and (B3) in the statement of Proposition 3.24.*

Proof. Suppose that v is a vertex of \mathcal{T} corresponding to a star of size at least 3 in T . Then v is an infinite collection of vertices in T . We now describe how to construct an associated set of geodesics \mathcal{B}_v by choosing, for each vertex in this collection, a bi-infinite geodesic which is bicoloured by an essential cut pair.

For each \approx -pair or \sim -pair in the collection v , we add to the set \mathcal{B}_v a (separating) bi-infinite geodesic γ whose boundary is this \approx -pair or \sim -pair. We claim that γ may be chosen so that the labels of this geodesic are essential vertices. This is immediate from Lemma 3.6 for \approx -pairs. For \sim -pairs, we have by Corollary 3.21 that the corresponding set A contains an essential cut pair. Now since $\langle A \rangle$ is 2-ended, there is only one essential cut pair contained in A . If $\text{card}(A) = 2$ the geodesic γ must be bicoloured by this cut pair, and if $\text{card}(A) = 3$ then $A = \{a, b, c\}$ with $\{a, b\}$ an essential cut pair and c commuting with both a and b ; in this case, we may take γ to be one of the two geodesics bicoloured by a and b with $\partial\gamma = v$.

Given an infinite \sim -class v' in v which comes from a set A as in Proposition 3.9, by Remark 3.22, for each component of $\Sigma \setminus \Sigma_A$ there is a unique corresponding geodesic in the set \mathcal{A}^+ . Now for any $v'' \neq v'$ in the collection v , the subset of ∂W corresponding to v'' lies in a single component of $\Sigma \setminus \Sigma_A$, as shown in the proof of Proposition 3.24(b). Moreover, there is a unique component U of $\Sigma \setminus \Sigma_A$ which contains the subsets of ∂W corresponding to all of the vertices in the collection v other than v' itself, for otherwise, v' would be between some pair of vertices in v , contradicting the fact that v is a star. Let $\mu_{v'} \in \mathcal{A}^+$ be the frontier geodesic corresponding to the component U . By the definition of \mathcal{A}^+ , μ is bicoloured by essential vertices. For each such v' in v , add $\mu_{v'}$ to the collection \mathcal{B}_v .

Thus for each vertex in the star v , the set \mathcal{B}_v contains a unique separating bi-infinite geodesic bicoloured by a pair of essential vertices. We claim that \mathcal{B}_v arises from the construction given in Proposition 3.24 associated to some set B satisfying (B1), (B2) and (B3).

Let B be the support of the geodesics in the set \mathcal{B}_v , that is the union of all the pairs labelling the geodesics in the set. We will show that B satisfies (B1), (B2) and (B3).

If (B1) fails, then there is a pair of vertices $\{c_1, c_2\}$ and $b_1, b_2 \in B$ which are in different components of $\Gamma \setminus \{c_1, c_2\}$. For $i = 1, 2$, let $b'_i \in B$ be such that there is a (b_i, b'_i) -bicoloured geodesic γ_i in \mathcal{B}_v . Without loss of generality γ_1 passes through e . Let w be the label on

a shortest path from e to γ_2 . Write w as w_1w_2 , where w_1 is the maximal initial subword consisting of generators from the same component of $\Gamma \setminus \{c_1, c_2\}$ as b_1 . (The subwords w_1 and/or w_2 could be empty.) Let γ be the (c_1, c_2) -bicoloured geodesic through the vertex w_1 . Then γ separates Σ by Lemma 3.3, and since the first letter of w_2 (if it is nonempty) or the generator b_2 (if w_2 is empty) lie in a separate component of $\Gamma \setminus \{c_1, c_2\}$ from b_1 , we get that at least one endpoint of the geodesic γ_2 is separated by $\partial\gamma$ from at least one endpoint of the geodesic γ_1 . Since the endpoints of γ_1 and γ_2 are contained in vertices say v_1 and v_2 of T , and γ is also contained in a vertex v_γ of T , this implies that v_γ lies between v_1 and v_2 . (Note that v_γ is a distinct vertex from v_1 and v_2 since our procedure for choosing geodesics in \mathcal{B}_v involved choosing the geodesic closest to the star in each infinite \sim -class.) This contradicts v_1 and v_2 being in the same star. We have shown that no $\{c_1, c_2\}$ as above can exist.

Thus B satisfies property (B1) of Proposition 3.24. It is therefore contained in a maximal set say B' satisfying (B1). Now if we run the procedure from Proposition 3.24 on the set B we will recover all of the pre-tree vertices in v . But v is maximal since it is a star, so $B = B'$ and hence B satisfies property (B2).

To show that (B3) holds, we need to rule out $\text{card}(B) = 2, 3$. In the first case, suppose b_1 and b_2 are the only labels of geodesics in \mathcal{B}_v . Since stars are infinite, there are infinitely many geodesics in \mathcal{B}_v with these labels. If γ and γ' are two such geodesics, they are necessarily disjoint. Let w be the (non-empty) word labelling a shortest path ν between γ and γ' and let z be the first letter in w . Assume without loss of generality that γ and ν intersect at e .

We now use Lemma 3.30 to define a geodesic μ passing through some point along ν . If Case (1) of the lemma holds, then define μ to be the (c_1, c_2) -bicoloured geodesic passing through e . (Observe that in this case b from the lemma is necessarily one of b_1 and b_2 .) If Case (2) of the lemma holds, then let u be the longest subword of w labelled by generators from the branch β of Γ containing z guaranteed by the lemma. If $u \neq w$, then define μ to be the (b_1, b_2) -bicoloured geodesic through the vertex of \mathcal{C} labelled u . In both these cases, we arrive at a contradiction by a proof similar to that of Proposition 3.24(c).

Finally, if $u = w$, then the vertices of β (which include b_1, b_2 and the labels of u) are contained in a set A as in Proposition 3.9, and it follows that γ and γ' are either both in \mathcal{A}_\sim (which is a contradiction because then the vertex corresponding to A in T is between the vertices corresponding to γ and γ') or they are both in \mathcal{A}^+ (which is a contradiction, since \mathcal{B}_v only contains one geodesic from each such set). This completes the proof that $\text{card}(B) \neq 2$.

Now suppose $\text{card}(B) = 3$, with $B = \{b_1, b_2, b_3\}$. Since B comes from bicoloured geodesics which separate Σ , there are at least two cut pairs in B , say $\{b_1, b_2\}$ and $\{b_2, b_3\}$.

Suppose that b_1 and b_3 don't separate Γ . Starting with a path μ between them in Γ , and using the separation properties of $\{b_1, b_2\}$ and $\{b_2, b_3\}$, we can complete μ to a cycle σ containing all elements of B . Since $\{b_1, b_3\}$ doesn't separate, there is a path connecting μ to the other component of $\sigma \setminus \{b_1, b_3\}$ which meets μ at an essential vertex $c \neq b_1, b_3$, and the other component at a single point. Then there are three paths from c to B which meet only at c , so by Observation 3.26, we have $c \in B$, which is a contradiction.

Thus we may assume that $\{b_1, b_3\}$ also separates Γ , so that B satisfies (A1). If B also satisfies (A2), then it is contained in a set A which is maximal among all sets satisfying both (A1) and (A2). That is, B is contained in a set A which satisfies (A1), (A2) and (A3),

and therefore defines an infinite \sim -class. However, \mathcal{B}_v only contains one geodesic from each \sim -class in v , so this is a contradiction.

Thus B does not satisfy (A2). If there were no subdivided K_4 subgraph containing B , then (A2) would be vacuously true. It follows that there is a subdivided K_4 subgraph Λ of Γ containing B , and moreover, that b_1, b_2 and b_3 lie on at least two distinct branches of Λ .

Now if there is an essential vertex x of $\Lambda \setminus B$ which admits three paths to B meeting only at x , then by Observation 3.26, the maximality of B is violated, and this is a contradiction. Thus no essential vertex of $\Lambda \setminus B$ admits three such paths to B . This can only happen if B is contained in a single branch of Λ . This is a contradiction.

We have shown that the set B satisfies properties (B1), (B2) and (B3) in Proposition 3.24. It is then easy to verify that the geodesics in \mathcal{B}_v , together with any geodesics parallel to them, form the set \mathcal{B} defined in the statement of Proposition 3.24, and that $\Phi(\mathcal{B}) = \Phi(\mathcal{B}_v)$. Hence $v = \Phi(\mathcal{B})$ is one of the stars we have already constructed. This completes the proof. \square

Example 3.32. We now describe the stars (if any) for the graphs in Figures 3.1 and 3.2, and their adjacencies in \mathcal{T} . There are no sets of vertices B satisfying conditions (B1), (B2) and (B3) in either graph in Figure 3.1, or in the left-hand and central graphs in Figure 3.2. Hence by Proposition 3.31, the corresponding JSJ trees have no stars, and so \mathcal{T} coincides with the pre-tree T ; compare Example 3.19, which describes all \approx -pairs and \sim -classes, and their adjacencies in T . The left-hand (respectively, right-hand) graph in Figure 3.1 has JSJ tree which is particularly easy to describe: it is biregular with vertices alternating between \approx -pairs of valence 6 (respectively, 4) and infinite \sim -classes.

On the right of Figure 3.2, there are two W -orbits of stars, corresponding to the sets $\{u, v, p, q\}$ and $\{p, q, r, s\}$. There are no \approx -pairs. Each infinite \sim -class corresponds to a branch of this graph of length three, and so is adjacent in \mathcal{T} to a Type 1 vertex which is added in to subdivide the edge between this infinite \sim -class (which is a Type 2 vertex) and a representative of exactly one of these W -orbits of stars (which are Type 3 vertices). The \sim -pairs for the branch between u and v and the branch between r and s are adjacent in \mathcal{T} to a representative of exactly one W -orbit of stars, and the \sim -pairs which correspond to the set $\{p, q\}$ are adjacent in \mathcal{T} to representatives of both W -orbits of stars.

Remark 3.33. Having identified all vertices of the JSJ tree \mathcal{T} , we compare with the work of Lafont [20, 21]. It is easy to see that the \approx -pairs in our setting are analogous to the endpoints of branching geodesics in [21] (which are studied using some results from [20]). Now suppose A is a set of vertices of Γ corresponding to an infinite \sim -class such that A consists of all vertices in a single branch of Γ of length at least three. Then the \sim -classes corresponding to A are analogous to the visual boundaries of chambers in [21]. The spaces considered in [21] consist entirely of chambers glued together along branching geodesics, and so when \mathcal{T} consists entirely of “single branch” infinite \sim -classes alternating with \approx -pairs, as in the case of the graphs in Figure 3.1, the situation is very similar to that in [21]. However in our setting there may not be any \approx -pairs or any infinite \sim -classes, and there may also be \sim -pairs, infinite \sim -classes which do not come from a single branch, and stars.

We finish this section by determining the stabilisers of stars.

Proposition 3.34. *Let v be a vertex of \mathcal{T} corresponding to a star of size at least 3 in T , and let B be the corresponding subset of vertices of Γ satisfying (B1), (B2) and (B3) in Proposition 3.24. Then the stabiliser of v is a conjugate of $\langle B \rangle$.*

Proof. Let \mathcal{G}_B , \mathcal{B} and $\Phi : \mathcal{B} \rightarrow T$ be as in the statement of Proposition 3.24. It suffices to show that the stabiliser of the star $\Phi(\mathcal{B})$ is $\langle B \rangle$. Let $g \in W$ be in this stabiliser.

Let η be a geodesic in \mathcal{G}_B and let v' be the \approx -pair or \sim -class $\Phi(\eta)$. Since g stabilises v , we have that $g(v') = \Phi(\xi)$ for some $\xi \in \mathcal{B}$. Note that there are at most two choices for ξ , and these choices have the same endpoints, by (a) in the proof of Proposition 3.24. We will need the following result.

Lemma 3.35. *The geodesics ξ and $g\xi$ have the same endpoints.*

Proof. If v' is an \approx -pair or a \sim -pair then $v' = \partial\eta$ and $g(v') = \partial\xi$, hence $\partial\xi = g(\partial\eta) = \partial(g\eta)$ as required. Now suppose v' is an infinite \sim -class and that ξ and $g\xi$ do not have the same endpoints. By Lemma 3.1, since ξ and $g\xi$ are bicoloured, the sets $\partial\xi$ and $\partial(g\xi)$ are disjoint.

Let A be the set of vertices of Γ corresponding to v' as in Proposition 3.9, so that $v' = \partial\mathcal{C}_A \setminus \partial\mathcal{A}_{\approx}$ for some copy \mathcal{C}_A of the Cayley graph of $\langle A \rangle$. Then by similar arguments to those in the proof of Proposition 3.31, η is the frontier geodesic of $\Sigma_A \supseteq \mathcal{C}_A$ corresponding to the component of $\Sigma \setminus \Sigma_A$ which contains $\mathcal{B} \setminus \{\eta\}$. Similarly, writing Σ'_A for the copy of Σ_A for the infinite \sim -class $g(v')$, we have that ξ is the frontier geodesic of Σ'_A corresponding to the component of $\Sigma \setminus \Sigma'_A$ which contains $\mathcal{B} \setminus \{\xi\}$. By Proposition 3.23 the stabiliser of v' is $\langle A \rangle$ and the stabiliser of $g(v')$ is $g\langle A \rangle g^{-1}$. It follows that g takes the frontier geodesic η of Σ_A to a frontier geodesic $g\eta$ of Σ'_A . The endpoints of $g\eta$ are in $g(v')$, so $g\eta$ is the frontier geodesic of Σ'_A corresponding to the component of $\Sigma \setminus \Sigma'_A$ which contains $g\mathcal{B} \setminus \{g\eta\}$.

Now v is a star of size at least 3, so there is a $v'' \in v$ with v'' distinct from both v' and $g(v')$. Note that $g(v'') \neq g(v')$. Let $\mu \in \mathcal{B}$ be such that $\Phi(\mu) = v''$. Then μ and $g\eta$ lie in distinct components of $\Sigma \setminus \xi$. Now the endpoints of $g\mu$ are contained in $g(v'')$, and $g(v'')$ is in the star v . However $g\mu$ lies in $g\mathcal{B} \setminus \{g\eta\}$, so $g\mu$ and μ are in distinct components of $\Sigma \setminus \Sigma'_A$. This means that $g(v')$ separates $g(v'')$ from v'' in T , which contradicts v being a star. Thus ξ and $g\xi$ have the same endpoints as required. \square

Assume now that there is a $g \in W \setminus \langle B \rangle$ which stabilises v . By Lemma 3.27, there exists a cut pair $\{b_i, b_j\} \subset B$. Let η be the (unique) geodesic in \mathcal{G}_B which is (b_i, b_j) -bicoloured. By similar arguments to those in the proof of Proposition 3.23, since $g \notin \langle B \rangle$ there is a (unique) vertex x of $\Gamma \setminus B$ so that x is adjacent to both b_i and b_j , and x stabilises v .

Since $\text{card}(B) \geq 4$, there is an element $b_k \in B \setminus \{b_i, b_j\}$. Consider the geodesic $b_k\eta \in \mathcal{B}$. The group element x stabilises v but $x \notin \langle B \rangle$, so we have by the same arguments as for g and η above that the geodesic $xb_k\eta$, which is not in \mathcal{B} , has the same endpoints as some geodesic, say ξ' , which is in \mathcal{B} . Corollary 3.2 then implies that the geodesic ξ' is (b_i, b_j) -bicoloured and $xb_k\eta = \xi'x$. Now since ξ' is (b_i, b_j) -bicoloured and is in \mathcal{B} , we have $\xi' = h\eta$ for some $h \in \langle B \rangle$. Thus $xb_k\eta$ equals $h\eta x = hx\eta$ for some $h \in \langle B \rangle$. As the stabiliser of η is $\langle b_i, b_j \rangle$, it follows that $b_k^{-1}x^{-1}hx = b_kxhx \in \langle b_i, b_j \rangle$. Now x and b_k are not in $\langle b_i, b_j \rangle$, so this means x commutes with b_k . Thus x is adjacent to all three elements of $\{b_i, b_j, b_k\}$. Hence by Observation 3.26, $x \in B$ as well. This contradicts $x \in \Gamma \setminus B$, and so completes the proof. \square

3.8. **The tree \mathcal{T} .** We now list all vertices and edges in \mathcal{T} , and describe their stabilisers.

Theorem 3.36. *The vertices of \mathcal{T} and their stabilisers are as follows.*

- (1) *The W -orbits in $V_1(\mathcal{T})$ are in bijective correspondence with the following four families.*
 - (a) *(\approx -pairs, valence $k \geq 3$) Pairs $\{a, b\}$ such that $\Gamma \setminus \{a, b\}$ has $k \geq 3$ components and $\Gamma \setminus \{a, b\}$ has no component consisting of a single vertex. The stabilisers of these vertices are the conjugates of either $\langle a, b \rangle$, if there is no vertex c of Γ adjacent to both a and b , or of $\langle a, b, c \rangle$, if there is such a vertex c .*
 - (b) *(\approx -pairs, valence $2(k - 1) \geq 4$) Pairs $\{a, b\}$ such that $\Gamma \setminus \{a, b\}$ has $k \geq 3$ components and $\Gamma \setminus \{a, b\}$ has a component consisting of a single vertex c . The stabilisers of these vertices are the conjugates of $\langle a, b, c \rangle$.*
 - (c) *(\sim -classes of size 2, valence 2) Sets A satisfying properties (A1), (A2) and (A3) in Proposition 3.9 with $\langle A \rangle$ being 2-ended and which are not as in (1)(a) or (1)(b). The stabilisers of these vertices are the conjugates of either $\langle A \rangle$, if $\text{card}(A) = 3$ or $\text{card}(A) = 2$ and there is no vertex c of Γ adjacent to both elements of A , or of $\langle A \cup \{c\} \rangle$, if $\text{card}(A) = 2$ and there is such a vertex c .*
 - (d) *Valence 2 vertices added between Type 2 and Type 3 vertices (if any). The stabilisers of these vertices are the intersections of stabilisers of their endpoints.*
- (2) *The W -orbits in $V_2(\mathcal{T})$ are in bijective correspondence with sets A satisfying properties (A1), (A2) and (A3) in Proposition 3.9 with $\langle A \rangle$ infinite but not 2-ended. The stabilisers of these vertices are conjugates of $\langle A \rangle$.*
- (3) *The W -orbits in $V_3(\mathcal{T})$ are in bijective correspondence with sets B satisfying properties (B1), (B2) and (B3) in Proposition 3.24. The stabilisers of these vertices are conjugates of $\langle B \rangle$.*

The edges of \mathcal{T} are as described in Section 2.3.2. Furthermore, there is an edge between vertices v and v' if and only if the corresponding stabilisers intersect.

Proof. The identification of all vertices in \mathcal{T} follows from the construction of the JSJ tree together with Corollary 3.5 and Lemma 3.6 for the \approx -pairs, Propositions 3.9 and 3.18 for the \sim -classes, and Propositions 3.24 and 3.31 for the stars of size at least 3. We determined the stabilisers of \approx -pairs in Lemma 3.7, of \sim -classes in Proposition 3.23 and of stars in Proposition 3.34. Finally the stabilisers of any valence 2 vertices added between Type 2 vertices and stars are the intersections of the stabilisers of the endpoints, since these vertices are just subdividing edges of \mathcal{T} . \square

4. APPLICATIONS

In this section we give some applications of our main result. In Section 4.1 we prove Theorem 1.3 and in Section 4.2 we provide a precise statement for and proof of Theorem 1.5.

4.1. A collection for which \mathcal{T} is a complete invariant. In [22], Malone uses the quasi-isometries of “fattened trees” from Behrstock–Neumann [7] to construct a quasi-isometry between any pair of “geometric amalgams of free groups” which have the same Bowditch JSJ tree. Cashen [11] observes that this construction works for any pair of 1-ended hyperbolic groups, so long as their corresponding JSJ trees do not have any stars. In light of this, to establish Theorem 1.3 it is enough to prove the following.

Proposition 4.1. *Let Γ be a graph satisfying Standing Assumptions 1.1. Then the JSJ tree $\mathcal{T} = \mathcal{T}_\Gamma$ has no stars if and only if Γ has no subdivided K_4 subgraphs.*

Proof. First we show that if Γ has a subdivided K_4 subgraph Λ then there is a star in \mathcal{T} . Let b_1, b_2, b_3, b_4 be the essential vertices of the K_4 , considered as vertices in Λ . If c_1, c_2 is an arbitrary pair of essential vertices of Γ , then the subgraph $\Lambda \setminus \{c_1, c_2\}$ of $\Gamma \setminus \{c_1, c_2\}$ is a (subdivided) K_4 with at most two points missing, and is therefore connected. It follows that $\{b_1, b_2, b_3, b_4\} \setminus \{c_1, c_2\}$ is contained in a single component of $\Gamma \setminus \{c_1, c_2\}$. Thus $\{b_1, b_2, b_3, b_4\}$ satisfies properties (B1) and (B3) of Proposition 3.24. It is therefore contained in a set satisfying (B2), yielding an orbit of stars.

Now suppose \mathcal{T} has a star, which comes from a set of essential vertices B as in Proposition 3.24. By Lemma 3.28 it suffices to show that some three elements of B lie on an induced cycle in Γ . By Lemma 3.27, there is a cut pair $\{b_1, b_2\} \subseteq B$. Then we may choose $b_3 \neq b_1, b_2$ in B , and reduced paths τ and μ such that τ connects b_1 and b_2 in the component of $\Gamma \setminus \{b_1, b_2\}$ which does not contain b_3 , and μ connects b_2 and b_3 and misses b_1 . By construction, $\tau \cap \mu = \{b_2\}$. Now let η be a reduced path from b_3 to b_1 which misses b_2 . Clearly $\eta \cap \tau = \{b_1\}$. If $\eta \cap \mu = \{b_3\}$ then $\tau \cup \mu \cup \eta$ is an induced cycle containing b_1, b_2 , and b_3 . If $\eta \cap \mu$ contains points other than b_3 , then let x be the first point of $\eta \cap \mu$ that η meets, starting from b_1 , so that $\eta_{[b_1, x]} \cap \mu = \{x\}$. Then x is connected to b_1, b_2 and b_3 by the disjoint paths $\eta_{[b_1, x]}$, $\mu_{[b_2, x]}$ and $\mu_{[b_3, x]}$. Now by Observation 3.26, $x \in B$. Then $\sigma = \eta_{[b_1, x]} \cup \mu_{[b_2, x]} \cup \tau$ is an induced cycle containing $b_1, b_2, x \in B$, as desired. \square

4.2. Quasi-isometric rigidity. In this section we state and prove Theorem 4.2, which gives a precise version of Theorem 1.5 of the introduction. As we discuss in the Appendix, Theorem A.2 of the Appendix provides an alternative proof of some parts of Theorem 4.2.

Recall the definition of a Fuchsian group from the introduction. Note that if G is Fuchsian then so is $G \times F$ for any finite group F . Denote by Λ_n the n -cycle of length $n \geq 5$ and put $W_n = W_{\Lambda_n}$. Then W_n is the group generated by reflections in the sides of a right-angled hyperbolic n -gon.

A *generalised Θ graph* is defined as follows. Let $k \geq 3$ and $n_1, \dots, n_k \geq 0$ be integers. The graph $\Theta(0, 0, \dots, 0)$ is the graph with two essential vertices a and b each of valence k , and k edges e_1, e_2, \dots, e_k connecting a and b . Then the graph $\Theta(n_1, n_2, \dots, n_k)$ is obtained by, for each $1 \leq i \leq k$, subdividing the edge e_i into $n_i + 1$ edges by inserting n_i new vertices of valence 2. For example, the graphs in Figure 3.1 are $\Theta(1, 2, 2, 3)$ and $\Theta(2, 2, 2, 3)$.

Theorem 4.2. *Let W_Γ be 2-dimensional. The following are equivalent:*

- (1) W_Γ is cocompact Fuchsian.
- (2) W_Γ is quasi-isometric to W_n for some $n \geq 5$.
- (3) $\Gamma = \Lambda_n$ for some $n \geq 5$.

Proof. The implications (1) \implies (2), (3) \implies (1) and (3) \implies (2) are obvious. In order to establish (2) \implies (3) and so complete the proof, we will use the following lemma, whose proof is an elementary exercise in graph theory.

Lemma 4.3. *Suppose Γ is a triangle-free simplicial graph which has no separating vertices or edges. If Γ is not a cycle graph or a single edge then Γ contains an induced subgraph $\Theta(n_1, n_2, n_3)$ with $n_1, n_2, n_3 \geq 1$. \square*

Now suppose that W_Γ is 2-dimensional and quasi-isometric to W_n for some $n \geq 5$. Then W_Γ is 1-ended, so Γ has no separating vertices or edges. Also Γ is not a single edge, as then W_Γ would be finite. Assume that Γ is not a cycle graph. Then by Lemma 4.3, Γ contains an induced subgraph $\Theta = \Theta(n_1, n_2, n_3)$. If more than one n_i is 1, then Θ and hence Γ contains a square, thus W_Γ is not hyperbolic. This is a contradiction, so we may assume $n_2, n_3 \geq 2$.

Since W_Θ is a 1-ended special subgroup of W_Γ , its visual boundary ∂W_Θ is a closed, connected subset of ∂W_Γ . Now ∂W_Γ is homeomorphic to S^1 , so ∂W_Θ is homeomorphic to either S^1 or to a closed interval in S^1 . Let a and b be the two essential vertices of Θ and let γ be an (a, b) -bicoloured geodesic in the Cayley graph of W_Θ . Since $\Theta \setminus \{a, b\}$ has three components, by Corollary 3.5 the endpoints of γ are in $\partial W_\Theta(3+)$. Thus ∂W_Θ cannot be homeomorphic to either S^1 or an interval in S^1 . Hence Γ is a cycle graph Λ_n with $n \geq 5$. \square

APPENDIX A. CHARACTERISATION OF COCOMPACT FUCHSIAN COXETER GROUPS

In this appendix we prove Theorem A.2, which characterises cocompact Fuchsian groups among all Coxeter groups. The proof of Theorem A.2 uses results from [15] which were suggested to us by an anonymous referee.

In this appendix only, we consider general Coxeter groups. Recall that a *Coxeter group* is a group with presentation

$$(A.1) \quad W = \langle S \mid s_i^2 = 1 \forall s_i \in S, (s_i s_j)^{m_{ij}} = 1 \forall \text{ distinct } s_i, s_j \in S \rangle$$

where $S = \{s_i\}_{i=1}^n$ is a finite set and $m_{ij} = m_{ji} \in \{2, 3, 4, \dots\} \cup \{\infty\}$, with $m_{ij} = \infty$ meaning that there is no relation between s_i and s_j . The pair (W, S) is called a *Coxeter system*. The *nerve* of a Coxeter system (W, S) is the simplicial complex $N(W, S)$ with vertex set S and a k -simplex σ_T with vertex set $T \subset S$ for each subset T of S such that $\langle T \rangle$ is finite and $|T| = k + 1$. The Coxeter system (W, S) is *right-angled* if each $m_{ij} \in \{2, \infty\}$, and if $W = W_\Gamma$ is right-angled then its defining graph Γ is precisely the 1-skeleton of the nerve $N(W, S)$. Moreover, Γ is triangle-free if and only if Γ is equal to the nerve of W_Γ .

Example A.1. Let P be a compact convex hyperbolic polygon with edges labelled cyclically e_1, \dots, e_n , so that for $1 \leq i \leq n$ the dihedral angle between e_i and e_{i+1} is $\pi/m_{i,i+1}$ with $m_{i,i+1}$ an integer ≥ 2 . Put $m_{i+1,i} = m_{i,i+1}$ and if edges e_i and e_j are nonadjacent, put $m_{ij} = m_{ji} = \infty$. Let $S = \{s_i\}_{i=1}^n$ where s_i is the reflection of \mathbb{H}^2 in the geodesic containing e_i . Let $W = W(P)$ be the group generated by this set of reflections S . Then W is cocompact Fuchsian, and is a Coxeter group with presentation as in (A.1).

A Coxeter group W is a *hyperbolic polygon reflection group* if $W = W(P)$ for some hyperbolic polygon P as in Example A.1. Following Definition 10.6.2 of [15] we say that a Coxeter system (W, S) is of *type* HM^2 if its nerve is a generalised homology 1-sphere (see Definition 10.4.5 of [15] for the definition of this latter term.)

Theorem A.2. *A Coxeter group W is cocompact Fuchsian if and only if W is either a hyperbolic polygon reflection group, or the direct product of a hyperbolic polygon reflection group with a finite Coxeter group.*

Proof of Theorem A.2. Suppose that a Coxeter group W is cocompact Fuchsian. Then W has a finite index torsion-free subgroup which is the fundamental group of a closed hyperbolic surface. Hence by [15, Theorem 10.9.2], either (W, S) is of type HM^2 , or (W, S) has a direct product decomposition $(W, S) = (W_0, S_0) \times (W_1, S_1)$ where W_1 is a finite (nontrivial) Coxeter group and (W_0, S_0) is of type HM^2 . It thus suffices to show that a Coxeter group of type HM^2 must be a hyperbolic polygon reflection group. Now as remarked on p. 194 of [15], every generalised homology 1-sphere is homeomorphic to the circle S^1 . Therefore the nerve of (W, S) is an n -cycle for some $n \geq 3$. Since W is hyperbolic, it follows that W is a hyperbolic polygon reflection group as required. \square

Corollary A.3. *Let Γ be a finite, simplicial, triangle-free graph. Then the associated right-angled Coxeter group W_Γ is cocompact Fuchsian if and only if Γ is an n -cycle with $n \geq 5$.*

Theorem A.2 provides an alternative proof of $(1) \implies (3)$ in Theorem 4.2, however we cannot see how to use Theorem A.2 for either $(2) \implies (3)$ or $(2) \implies (1)$, one of which would be necessary in order to obtain the quasi-isometric rigidity result.

REFERENCES

- [1] A. ABRAMS, N. BRADY, P. DANI, M. DUCHIN, AND R. YOUNG, *Pushing fillings in right-angled Artin groups*, J. Lond. Math. Soc. (2), 87 (2013), pp. 663–688.
- [2] J. BEHRSTOCK AND R. CHARNEY, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann., 352 (2012), pp. 339–356.
- [3] J. BEHRSTOCK AND C. DRUŢU, *Divergence, thick groups, and short conjugators*, Illinois J. Math., 58 (2014), pp. 939–980.
- [4] J. BEHRSTOCK, C. DRUŢU, AND L. MOSHER, *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*, Math. Ann., 344 (2009), pp. 543–595.
- [5] J. BEHRSTOCK, M. F. HAGEN, A. SISTO, AND P.-E. CAPRACE, *Thickness, relative hyperbolicity, and randomness in Coxeter groups*, (2013). arxiv:1312.4789.
- [6] J. A. BEHRSTOCK, T. JANUSZKIEWICZ, AND W. D. NEUMANN, *Quasi-isometric classification of some high dimensional right-angled Artin groups*, Groups Geom. Dyn., 4 (2010), pp. 681–692.
- [7] J. A. BEHRSTOCK AND W. D. NEUMANN, *Quasi-isometric classification of graph manifold groups*, Duke Math. J., 141 (2008), pp. 217–240.
- [8] M. BESTVINA, B. KLEINER, AND M. SAGEEV, *The asymptotic geometry of right-angled Artin groups. I*, Geom. Topol., 12 (2008), pp. 1653–1699.
- [9] B. H. BOWDITCH, *Cut points and canonical splittings of hyperbolic groups*, Acta Math., 180 (1998), pp. 145–186.
- [10] B. H. BOWDITCH, *Treelike structures arising from continua and convergence groups*, Mem. Amer. Math. Soc., 139 (1999), pp. viii+86.
- [11] C. H. CASHEN, *Quasi-isometry invariants from decorated trees of cylinders of two-ended JSJ decompositions*, (2014). arxiv:1403.2865.
- [12] R. CHARNEY AND H. SULTAN, *Contracting boundaries of $\text{CAT}(0)$ spaces*, J. Topol., 8 (2015), pp. 93–117.
- [13] J. CRISP AND L. PAOLUZZI, *Commensurability classification of a family of right-angled Coxeter groups*, Proc. Amer. Math. Soc., 136 (2008), pp. 2343–2349.

- [14] P. DANI AND A. THOMAS, *Divergence in right-angled Coxeter groups*, Trans. Amer. Math. Soc., 367 (2015), pp. 3549–3577.
- [15] M. W. DAVIS, *The geometry and topology of Coxeter groups*, vol. 32 of London Mathematical Society Monographs Series, Princeton University Press, Princeton, NJ, 2008.
- [16] M. W. DAVIS AND T. JANUSZKIEWICZ, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra, 153 (2000), pp. 229–235.
- [17] C. DRUȚU, *Relatively hyperbolic groups: geometry and quasi-isometric invariance*, Comment. Math. Helv., 84 (2009), pp. 503–546.
- [18] S. M. GERSTEN, *Quadratic divergence of geodesics in CAT(0) spaces*, Geom. Funct. Anal., 4 (1994), pp. 37–51.
- [19] J. HUANG, *Quasi-isometry rigidity of right-angled artin groups i: the finite out case*, (2014). arxiv:1410.8512.
- [20] J.-F. LAFONT, *Rigidity result for certain three-dimensional singular spaces and their fundamental groups*, Geom. Dedicata, 109 (2004), pp. 197–219.
- [21] ———, *Diagram rigidity for geometric amalgamations of free groups*, J. Pure Appl. Algebra, 209 (2007), pp. 771–780.
- [22] W. MALONE, *Topics in geometric group theory*, ProQuest LLC, Ann Arbor, MI, 2010. Thesis (Ph.D.)—The University of Utah.
- [23] A. MARTIN AND J. ŚWIĄTKOWSKI, *Infinitely-ended hyperbolic groups with homeomorphic Gromov boundaries*, J. Group Theory, 18 (2015), pp. 273–289.
- [24] M. MIHALIK AND S. TSCHANTZ, *Visual decompositions of Coxeter groups*, Groups Geom. Dyn., 3 (2009), pp. 173–198.
- [25] P. PAPASOGLU, *Quasi-isometry invariance of group splittings*, Ann. of Math. (2), 161 (2005), pp. 759–830.
- [26] P. PAPASOGLU AND K. WHYTE, *Quasi-isometries between groups with infinitely many ends*, Comment. Math. Helv., 77 (2002), pp. 133–144.
- [27] E. STARK, *Abstract commensurability and quasi-isometry classification of hyperbolic surface group amalgams*, (2014). arxiv:1404.6239.
- [28] G. A. SWARUP, *On the cut point conjecture*, Electron. Res. Announc. Amer. Math. Soc., 2 (1996), pp. 98–100 (electronic).
- [29] J. ŚWIĄTKOWSKI, *Trees of manifolds as boundaries of spaces and groups*, (2013). arxiv:1304.5067.
- [30] ———, *Trees of metric compacta and trees of manifolds*, (2013). arxiv:1304.5064.
- [31] ———, *The dense amalgam of metric compacta and topological characterization of boundaries of free products of groups*, (2014). arxiv:1410.4989.
- [32] ———, *Hyperbolic Coxeter groups with Sierpiński carpet boundary*, (2015). arxiv:1502.02044.