

RIGHT-ANGLED COXETER GROUPS WITH NON-PLANAR BOUNDARY

PALLAVI DANI, MATTHEW HAULMARK, AND GENEVIEVE WALSH

ABSTRACT. We investigate the planarity of the boundaries of right-angled Coxeter groups. We show that non-planarity of the defining graph does not necessarily imply non-planarity of every boundary of the associated right-angled Coxeter group, although it does in many cases. Our techniques yield a characterization of the triangle-free defining graphs such that the associated right-angled Coxeter group has boundary a Menger curve.

1. INTRODUCTION

A group is CAT(0) if it acts geometrically (properly discontinuously, co-compactly and by isometries) on a CAT(0) space. Every CAT(0) metric space X has a well-defined visual boundary ∂X . We will denote a proper CAT(0) space on which G acts geometrically by X_G , or by X_Γ , when G is the right-angled Coxeter group defined by a graph Γ . See Section 2 for more detailed definitions.

When a group G acts geometrically on X_G , the topology of ∂X_G can provide information about the algebra of G , even though the boundary of G may not be well-defined. The dimension of ∂X_G is related to the cohomological dimension of G [16, 4], and in the case that G is torsion-free, G is a PD(3) group exactly when $\partial X_G \cong S^2$ [4]. Also, splittings of G are expressed as topological features in ∂X_G (see [6, 26, 18] among others).

An important question about the topology of boundaries is their planarity. We say that a topological space is *planar* if it can be embedded in S^2 . When G can be virtually realized as a geometrically finite Kleinian group, every CAT(0) boundary ∂X_G is planar. This can be seen as follows. The limit set of the Kleinian group is a subset of S^2 . Moreover, the CAT(0) boundary is well-defined, since G is either hyperbolic or CAT(0) with isolated flats [20]. This boundary is either the limit set itself, or the limit set with parabolic fixed points replaced by circles. In either case, it is planar. By a special case of a theorem of Bestvina–Kapovich–Kleiner [2], if G is the fundamental group of a 3-manifold, then no boundary of G contains a K_5 or a $K_{3,3}$ (though this doesn't immediately imply that every boundary is planar, as we discuss below).

Conjecture 1.1 below, which was asked as questions in [27, Questions 1.3 and 1.4], presents a sort of converse to the above statements.

This work of the first author was supported by a grant from the Simons Foundation (#426932, Pallavi Dani) and by NSF Grant No. DMS-1812061. The third author was supported by NSF Grant No. DMS 1709964. Some of this work was conducted at the Centre International de Rencontres Mathématiques in Luminy and we thank them for their hospitality. We thank Kevin Schreve for his comments on a draft of this paper and the referee for their helpful comments.

Conjecture 1.1. *Let G be a CAT(0) group with a planar visual boundary. Then every visual boundary of G is planar, and furthermore, G is virtually the fundamental group of a compact 3-manifold.*

This paper contains evidence for Conjecture 1.1 in the setting of right-angled Coxeter groups. See Corollary 1.5.

By work of Haïssinsky [17, Theorem 1.10], this conjecture is known to hold for hyperbolic groups which are CAT(0) cubed, and hence for hyperbolic right-angled Coxeter groups. We note that Conjecture 1.1 implies the Cannon Conjecture [8] for hyperbolic groups which are CAT(0). As we write there are no known examples of hyperbolic groups which are not CAT(0).

Conjecture 1.1 is more speculative than the analogous conjecture for hyperbolic groups in [17, Conjecture 1.6]. This is because CAT(0) boundaries, unlike the boundaries of hyperbolic groups, are not always well-defined [12] and not always locally connected [23]. This last property means that a result of Claytor [11, Theorem C] does not necessarily apply. In particular, when ∂X is not locally connected non-planarity of the boundary does not imply that there is a $K_{3,3}$ or K_5 in the boundary. Schreve and Stark [27] have an example of two homeomorphic CAT(0) complexes with two different boundaries, one of which contains an embedded $K_{3,3}$ and one of which does not. Both boundaries are non-planar.

In this paper we study the planarity of boundaries of right-angled Coxeter groups. Given a finite simplicial graph Γ we denote the associated right-angled Coxeter group by W_Γ . Every W_Γ is CAT(0), and in particular, W_Γ acts geometrically on a CAT(0) cube complex Σ_Γ called its *Davis–Moussong complex*.

It is tempting to conjecture that if Γ is non-planar then every CAT(0) boundary of W_Γ is non-planar, up to a finite subgroup. Indeed, Świątkowski speculates that planarity of the defining graph may be a necessary condition for a planar boundary. See [30, Remark 3]. However, we prove that non-planarity of the defining graph of a right-angled Coxeter group does not guarantee non-planarity of the boundary; there is a W_Γ with a non-planar defining graph Γ , which has a planar CAT(0) boundary, as we explain in Example 1.1 below. Indeed there are many.

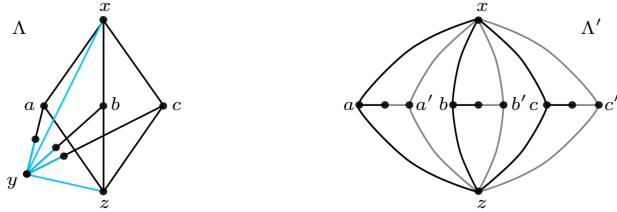


FIGURE 1.1. In the graph Λ on the left, each blue segment is an edge, while each black segment is a path which may or may not be subdivided. The graph Λ' on the right is the double of Λ over the vertex y , as defined in Section 2.2. The two copies of Λ minus the open star of y are shown in Λ' in black and grey respectively.

Example 1.1. Let Λ denote the graph on the left in Figure 1.1, which is non-planar. As we observe in Lemma 2.5, the group W_Λ contains an index two subgroup isomorphic to $W_{\Lambda'}$, where Λ' is the planar graph on the right of Figure 1.1. Any right-angled Coxeter group with planar defining graph is virtually the fundamental

group of a 3-manifold. Indeed, the defining graph can be embedded as an induced subgraph in a triangulation T of a 2-sphere. Now W_T , the right-angled Coxeter group defined by the one-skeleton of T , is virtually a closed 3-manifold group, since the Davis–Moussong complex Σ_T of W_T is a manifold. The Davis–Moussong complex $\Sigma_{\Lambda'}$ of $W_{\Lambda'}$ is a convex subcomplex of Σ_T . Consequently $\partial\Sigma_{\Lambda'}$ embeds in $\partial\Sigma_T \cong S^2$. Therefore, $\partial\Sigma_{\Lambda'}$ is planar. (Also see Theorem 11.4.1 of [15].) Now $W_{\Lambda'}$ acts on $W_{\Lambda'}$ by conjugation, and this induces a geometric action of $W_{\Lambda'}$ on $\Sigma_{\Lambda'}$. It follows that $W_{\Lambda'}$ has a planar boundary as well.

The graph Λ in Figure 1.1 is a subdivision of the graph in Figure 1.2. More generally, let Π denote any graph obtained by subdividing the black segments of the graph in Figure 1.2 enough to get a triangle-free graph. Then W_{Π} has a planar boundary by a similar argument. We show that the graph Π is the only obstruction:

Theorem 1.2. *Let Γ be a triangle-free non-planar graph, and let X_{Γ} be a proper CAT(0) space on which W_{Γ} acts geometrically. Then either ∂X_{Γ} is non-planar or W_{Γ} contains a finite-index special subgroup whose defining graph contains an induced copy of the graph Π in Figure 1.2.*

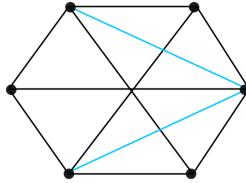


FIGURE 1.2. The graph Π is defined to be any graph as shown in this figure, such that the blue segments are edges, the black segments may or may not be edges, and enough of the black segments are subdivided to ensure that the result is triangle-free. A specific instance of such a subdivision appears on the left in Figure 1.1.

A graph is *inseparable* if it is connected, has no separating complete subgraph, no cut pair, and no separating complete subgraph suspension. When Γ is inseparable and ∂X_{Γ} is locally connected and planar, it follows from our Corollary 5.2, that ∂X_{Γ} cannot contain an induced copy of the graph in Figure 1.2. Thus we have the following:

Theorem 1.3. *Let Γ be a triangle-free inseparable graph and let X_{Γ} be a CAT(0) space on which W_{Γ} acts geometrically. If Γ is non-planar and ∂X_{Γ} is locally connected and contains no local cut points, then ∂X_{Γ} is non-planar.*

Since planar defining graph implies planar boundary, the following corollary is immediate.

Corollary 1.4. *Let Γ be a triangle-free inseparable graph and let X_{Γ} be a CAT(0) space on which W_{Γ} acts geometrically. Assume that ∂X_{Γ} is locally connected and has no local cut points. Then ∂X_{Γ} is planar if and only if Γ is planar.*

Using Theorem 1.3 we can now conclude that Conjecture 1.1 holds for a class of right-angled Coxeter groups:

Corollary 1.5. (to Theorem 1.3) *Let Γ be a graph with no triangles, and X_Γ a proper $\text{CAT}(0)$ space on which W_Γ acts geometrically. Then if ∂X_Γ is a Sierpinski carpet, W_Γ is virtually a 3-manifold group.*

The corollary holds as follows. Since ∂X is a Sierpinski carpet, it is planar, locally connected, and has no local cut points or cut pairs. Therefore in this setting, we can conclude from Theorem 1.3 that the defining graph Γ is planar as well. This implies that W_Γ is virtually a 3-manifold group by the argument in Example 1.1.

Theorem 1.3 also allows us to characterize the right-angled Coxeter groups with triangle-free defining graphs which have Menger curve boundaries.

Corollary 1.6. (to Theorem 1.3) *Let Γ be a triangle-free inseparable graph, such that W_Γ is either hyperbolic or $\text{CAT}(0)$ with isolated flats. Then the following are equivalent:*

- (1) Γ is non-planar.
- (2) Every $\text{CAT}(0)$ boundary ∂X_Γ is Menger curve.
- (3) Some $\text{CAT}(0)$ boundary ∂X_Γ is a Menger curve.

Proof. We claim that under the hypotheses of the corollary, ∂X_Γ is 1-dimensional for every $\text{CAT}(0)$ space X_Γ on which W_Γ acts geometrically. Since Γ is triangle-free, the Davis–Moussong complex Σ_Γ is 2-dimensional, and the virtual cohomological dimension (vcld) of W_Γ is equal to 1 or 2. By Stallings [28] a group G with $\text{vcld}(G) = 1$ is virtually free. Since Γ is inseparable, the group W_Γ is not virtually free, so $\text{vcld}(W_\Gamma) = 2$. Since Coxeter groups are virtually torsion-free, a theorem of Bestvina–Mess [4, Corollary 1.4], see also [1, Theorem 1.7], implies that the covering dimension of ∂X_Γ is equal to $\text{vcld}(W_\Gamma) - 1$. This proves the claim.

Now suppose Γ is non-planar. When W_Γ is hyperbolic, ∂X_Γ is locally connected (see [4, 29, 5]). Since $\text{vcld}(W_\Gamma) = 2$, the highest rank of a virtually abelian subgroup is 2, and we apply [21] to conclude that ∂X_Γ is locally connected in the $\text{CAT}(0)$ with isolated flats case. Since Γ is inseparable, W_Γ does not split over a 2-ended subgroup. Thus by [6, Theorem 6.2] (in the hyperbolic case) or by [18, Theorem 1.3] (in the isolated flats case) we conclude that ∂X_Γ has no local cut points. Now we may apply Theorem 1.3 to conclude that ∂X_Γ is non-planar. Then by [22, Theorem 4] (in the hyperbolic case) or by [18, Theorem 1.2] (in the isolated flats case), we conclude that ∂X_Γ is a Menger curve.

On the other hand if Γ is planar, then $\partial \Sigma_\Gamma$ is planar (as we showed in Example 1.1), and cannot be a Menger curve. Conditions (2) and (3) are equivalent because in this situation, the $\text{CAT}(0)$ boundary is well defined [20]. \square

We remark that [30, Theorem 1] implies (in the setting of triangle-free graphs) that if Γ is inseparable and W_Γ is hyperbolic, then the Gromov boundary ∂W_Γ is a Sierpinski carpet. Świątkowski suggests in Remark 3 of [30] that planarity of the nerve may be a necessary condition for a Coxeter group to have Sierpinski carpet boundary (up to a product with a finite Coxeter group). Corollary 1.6 shows that this is true in the case of hyperbolic right-angled Coxeter groups defined by triangle-free graphs.

In the setting of hyperbolic groups, Menger curve boundary is known to be generic [13]. Recently Haulmark–Hruska–Sathaye [19] provide examples of large

type (i.e. not right-angled) Coxeter groups which are not hyperbolic and have visual boundary homeomorphic to the Menger curve. Corollary 1.6 provides a technique for constructing large classes of examples of right-angled Coxeter groups with Menger curve boundary. The following example gives a concrete such class.

Example 1.2. One class of finite simplicial graphs which yield right-angled Coxeter groups with Menger curve boundary are the Möbius Ladders (see https://en.wikipedia.org/wiki/Möbius_ladder) These graphs are inseparable and non-planar. By a result of Caprace [9, 10] the right-angled Coxeter groups defined by these graphs have isolated flats; therefore, they satisfy the hypotheses of Corollary 1.6 and have Menger curve boundary.

1.1. Overview of the paper. In Section 2, we give some preliminaries on right-angled Coxeter groups and their boundaries. In Section 3 we use graph-theoretic techniques to show that if there is not an induced $K_{3,3}$ subdivision in our non-planar graph Γ , then by taking the double cover some vertex finitely many times, we arrive at a graph Γ' such that $W_{\Gamma'}$ is a finite-index subgroup of W_{Γ} , and Γ' either contains an induced $K_{3,3}$ subdivision or an induced subdivided copy of one of two specific graphs. (See Figure 3.3.) One of the two specific graphs is Π from Figure 1.2. In Section 4 we show if the defining graph of a right-angled Coxeter group contains either an induced $K_{3,3}$ subdivision, or an induced copy of one of the two specific graphs (the one not equal to Π), then any boundary ∂X_{Γ} is non-planar. Finally in Section 5 we deal with the case of the remaining graph Π . We show that if any visual boundary of a right-angled Coxeter group is connected, locally connected, has no local cut points and is planar, then the defining graph of that right-angled Coxeter group cannot contain a copy of Π . Theorem 1.2 is proven in Section 4, while Theorem 1.3 is proven in Section 5.

2. PRELIMINARIES

2.1. Boundaries of CAT(0) Spaces. Let X be a proper CAT(0) space. The *visual* or CAT(0) *boundary* of X , denoted ∂X , is the set of equivalence classes of geodesic rays, where two rays $c_1, c_2: [0, \infty) \rightarrow X$ are equivalent if there exists a constant $D \geq 0$ such that $d(c_1(t), c_2(t)) \leq D$ for all $t \in [0, \infty)$.

The boundary ∂X comes equipped with a natural topology called the *cone topology*. To define this topology, identify ∂X with the set of geodesic rays based at some fixed point p in X . Then if c is a geodesic ray based at p , a basic open set around c consists of geodesic rays based at p whose projection onto a ball of radius t around p is close to $c(t)$.

If G acts geometrically on X one would like to define ∂G to be ∂X . If G is a hyperbolic group, then X is a Gromov hyperbolic metric space and ∂X is the Gromov boundary of G . In particular, in this case the CAT(0) boundary ∂G is well-defined. For example, if G is virtually free, then the boundary of any CAT(0) space that G acts on geometrically is a Cantor set. Hruska–Kleiner have shown that ∂G is also well-defined in the setting of CAT(0) groups with isolated flats [20]. In general the homeomorphism type of the boundary is not well-defined for CAT(0) groups (see [12, 27]). However, for special subgroups of Coxeter groups, one can find a boundary for that special subgroup in any CAT(0) boundary for the Coxeter group, by the following lemma of Mihalik and Tschantz [24]. Suppose that (W, S) is a finitely generated Coxeter system, \mathcal{C} is the Cayley graph of W with respect to

S , and W acts geometrically on a CAT(0) space X . Fix a point $x \in X$, and define a graph $\mathcal{C}_x \subset X$ to have as vertices the orbit $W \cdot x$ and as edges the collection of CAT(0) geodesic paths connecting wx and wsx , for $w \in W$ and $s \in S$. Note that the collection of edges will be W -equivariant by uniqueness of CAT(0) geodesics.

Lemma 2.1 ([24] Corollary 6.6). *Suppose that (W, S) is a finitely generated Coxeter group with Cayley graph \mathcal{C} , acting geometrically on the CAT(0) space X , and take an $x \in X$, and $P_x : \mathcal{C} \rightarrow \mathcal{C}_x$ mapping \mathcal{C} quasi-isometrically and W -equivariantly into X . Then for each subset $A \subset S$, (the image of) the subgroup $\langle A \rangle$ is quasi-convex in X .*

We will use the lemma below often to find circles and Cantor sets in the boundary ∂X of some CAT(0) space for W_Γ , although it applies more generally. We will also use it to refer to points of the boundary. For example, if x and y are disjoint vertices on Γ , the special subgroup defined by x and y is virtually cyclic and we will refer to the points of its boundary as $(xy)^\infty$ and $(yx)^\infty$.

Definition. Let G act geometrically on a proper metric space X . We say that $H < G$ is *quasiconvex* for the action of G on X if, given a point $x_0 \in X$, any geodesic between points of the orbit Hx_0 lies in the K -neighborhood of Hx_0 , for some $K > 0$.

Definition. Let X be a proper CAT(0) space and suppose G acts properly by isometries on X . The *limit set*, $\Lambda(G)$, is the set of all $\xi \in \partial X$ such that there is a sequence (g_n) of elements of G with $(g_n x) \rightarrow \xi$ as $n \rightarrow \infty$ for some (hence any) $x \in X$.

Lemma 2.2. *Suppose that G acts geometrically on a CAT(0) space X and $H < G$ is a Gromov hyperbolic group which is quasiconvex for the action of G on X . Then ∂H is homeomorphic to $\Lambda(H)$, the limit set of H in ∂X . In particular, ∂H embeds in ∂X .*

Proof. We fix a point $x_0 \in X$, and use it to identify the CAT(0) visual boundary ∂X of X with the set of geodesic rays emanating from x_0 , as well as to define the orbit Hx_0 . We denote the CAT(0) metric on X by d .

Let K be the quasiconvexity constant of H , and define X_H to be the union of the closed balls of radius K around the points of the orbit Hx_0 . The action of H on X leaves X_H invariant, and we would like to identify ∂H with the boundary of X_H . However, X_H endowed with the metric d is not a geodesic space, and hence does not satisfy the definition of δ -hyperbolicity. Instead, we consider the length metric d' induced by d on X_H . To avoid confusion, for the rest of this proof we preface terms which depend on the metric being used (or on its associated topology) with the name of the metric (eg. d -geodesic, d -closed set, etc.).

In general, $d'(x_1, x_2) \geq d(x_1, x_2)$ for $x_1, x_2 \in X_H$. Any d -geodesic in X_H is a d' -geodesic (but not necessarily vice versa). Thus, the two metrics agree when restricted to the orbit Hx_0 (since the d -geodesic between a pair of orbit points is contained in X_H by quasiconvexity). Let $B_d(x, r)$ and $B_{d'}(x, r)$ denote the closed balls of radius r about x in the specified metrics. If $x \in X_H$, then for sufficiently small ϵ , the ball $B_d(x, \epsilon)$ intersects X_H in a single component, and $B_{d'}(x, \epsilon) = B_d(x, \epsilon) \cap X_H$. Note that $d(x, y) = d'(x, y)$ for each y in such a $B_{d'}(x, \epsilon)$. If x is in the d -interior of X_H , then ϵ can be chosen so that $B_{d'}(x, \epsilon) = B_d(x, \epsilon)$ and the two metrics agree on this entire set. Therefore, the topology on X_H induced by d' is the

same as its topology as a closed subset of (X, d) . The set X_H is a union of closed balls around each point of the orbit of H and is closed as H is finitely generated.

We claim that (X_H, d') is a proper geodesic space. It is enough to show that X_H is d' -complete and d' -locally compact, by [7, Proposition I.3.7] (Hopf–Rinow Theorem) and [7, Corollary I.3.8].

To prove that X_H is d' -complete, note that every d' -Cauchy sequence (x_n) in X_H is also d -Cauchy. Since X is d -complete and $X_H \subset X$ is d -closed, (x_n) d -converges to some $x \in X_H$. Now given any d' -open neighborhood U of x , there exists ϵ sufficiently small such that $B_d(x, \epsilon) \cap X_H = B_{d'}(x, \epsilon) \subset U$. Since (x_n) is contained in X_H and d -converges to x , it eventually lies in $B_d(x, \epsilon) \cap X_H \subset U$. This proves that it d' -converges to x as well.

Since the topology on X_H induced by d' is the same as the topology induced as a closed subset of (X, d) , (X_H, d') is homeomorphic to a closed subset of a locally compact Hausdorff space and hence locally compact [32, Theorem 18.4].

This proves the claim that (X_H, d') is a proper geodesic space.

The action of H on X_H preserves d' with quotient one closed ball. Thus, H acts geometrically on (X_H, d') , and is therefore quasi-isometric to it by the Švarc–Milnor Lemma. It follows that (X_H, d') is δ -hyperbolic for some $\delta > 0$ (fixed for the rest of the proof) and ∂H is homeomorphic to ∂X_H . To finish the proof, we will show that ∂X_H is homeomorphic to the limit set $\Lambda(H) \subset X$ (with the subspace topology).

Next, we define a map $\phi : \partial X_H \rightarrow \partial X$. As (X_H, d') is a geodesic hyperbolic space, we can identify ∂X_H with the set of equivalence classes of d' -geodesic rays in X_H emanating from x_0 . We assume all geodesics are parametrized by arclength. Let $[\alpha] \in \partial X_H$, where α is a d' -geodesic ray from x_0 . Choose a sequence $(h_n)_{n \geq 1}$ in H such that the map $\eta : [0, \infty) \rightarrow X_H$ which sends $[n-1, n)$ to the orbit point $h_n x_0$ is a d' -quasigeodesic that is d' -asymptotic to α (i.e. α and β have finite Hausdorff distance). Such quasigeodesics exist; the sequence obtained by choosing h_n such that $d(\alpha(n), h_n x_0) \leq K$ is an example. For $n > 0$, let β_n denote the d -geodesic segment from x_0 to $h_n x_0$. Since (X, d) is proper, the Arzela–Ascoli Theorem implies that there exists a subsequence (β_{n_i}) which converges to a d -geodesic ray $\beta : [0, \infty) \rightarrow X$. We define $\phi([\alpha]) = [[\beta]]$, where $[\cdot]$ and $[[\cdot]]$ denote the appropriate equivalence classes.

We must show that ϕ is well-defined. First we establish some useful facts about β . Since each β_i is in X_H , and X_H is d -closed, it follows that β is in X_H and (thinking of it as a d -geodesic) represents a point of ∂X_H . Moreover, it can be shown that (β_{n_i}) d' -converges to β . Now it is evident that the sequence $(h_{n_i} x_0)$ of endpoints of the β_{n_i} d' -converges to both $[\eta] = [\alpha]$ and $[\beta]$, so these are equal. It follows that α and β are d' -asymptotic, and therefore also d -asymptotic.

Now suppose β' is a d -geodesic ray obtained from the above construction by either choosing a different representative for $[\alpha]$, or a different quasigeodesic η , or a different convergent subsequence of (β_n) . Then the argument in the previous paragraph shows that β' is d -asymptotic to α , and therefore to β . It follows that $[[\beta]] = [[\beta']]$, i.e., that ϕ is well-defined.

To see that ϕ is injective, suppose β is a d -geodesic ray from x_0 representing $\phi(\alpha) = \phi(\alpha')$. Then the above argument shows that β is a d' -geodesic asymptotic to both α and α' , and it follows that $[\alpha] = [\alpha']$.

It is evident from the construction of ϕ that its image is in $\Lambda(H)$. Now let $[[\beta]] \in \Lambda(H)$. By definition, there is a sequence $(h_n x_0)$ whose d -limit is $[[\beta]]$. As

observed in the last paragraph of Section II.8.5 of [7], the geodesics β_n from x_0 to $h_n x_0$ converge to β . Arguing as before, β is in fact in X_H , and represents an element of ∂X_H . Applying the above construction to β yields a d -geodesic ray from x_0 which is d -asymptotic to β . Since β is the only such ray, we have that $\phi([\beta]) = [[\beta]]$, i.e., ϕ surjects onto $\Lambda(H)$.

It remains to prove the continuity of ϕ and ϕ^{-1} . A neighborhood basis for $[[\beta]] \in \partial X$ is given by sets $U(\beta, r, \epsilon)$, for $r, \epsilon > 0$, which is the set of d -geodesic rays β' from x_0 such that $d(\beta(r), \beta'(r)) < \epsilon$. A neighborhood basis for $[\alpha] \in \partial X_H$ is given by $V(\alpha, r, \rho)$, for $r > 0, \rho > 2\delta$ (where δ is the hyperbolicity constant of X_H), which is the set of equivalence classes of d' -geodesic rays α' from x_0 such that $d'(\alpha(r), \alpha'(r)) < \rho$.

Now suppose $\phi([\alpha]) = [[\beta]]$, for some d' -geodesic ray α from x_0 , and consider some $U = U(\beta, r, \epsilon)$. To choose a basis element about $[\alpha]$ mapping into U , fix $\rho = 3\delta$, and choose R large enough so that for every $p \in B_{d'}(\beta(R), 7\delta)$, the projection p' of p to the d -sphere of radius r centered at x_0 has the property that $d(\beta(r), p') < \epsilon$ (so that any d -geodesic passing through this ball is in U). Now if α' is in $V = V(\alpha, R, 3\delta)$, then by definition, $d'(\alpha(R), \alpha'(R)) < 3\delta$. Let β' be the d -geodesic ray from x_0 representing $\phi([\alpha'])$. Then as shown above, α' and β' are d' -geodesic rays that are d' -asymptotic, so, since X_H is proper, $d'(\alpha'(R), \beta'(R)) < 2\delta$ by [7, Lemma III.H.3.3]. Similarly $d'(\alpha(R), \beta(R)) < 2\delta$. It follows that $d(\beta(R), \beta'(R)) \leq d'(\beta(R), \beta'(R)) < 7\delta$. Thus $[[\beta']] \in U$, and so $\phi(V) \subset U$. This proves that ϕ is continuous at each point in ∂X_H .

To prove the continuity of ϕ^{-1} , consider $[\beta] = \phi^{-1}([[\beta]])$, and a neighborhood $V(\beta, r, \rho)$ of $[\beta]$ in ∂X_H . Choose ϵ sufficiently small so that $\epsilon < \rho$ and $B_{d'}(\beta(r), \epsilon) = B_d(\beta(r), \epsilon) \cap X_H$. Then for $\gamma \in U(\beta, r, \epsilon)$, we have $d'(\beta(r), \gamma'(r)) = d(\beta(r), \gamma(r)) < \epsilon < \rho$, and it follows that γ represents an element of $V(\beta, r, \rho)$. Moreover, $\phi^{-1}([[\gamma]]) = [\gamma]$. Thus $\phi^{-1}(U(\beta, r, \epsilon)) \subset V(\beta, r, \rho)$. Thus ϕ^{-1} is continuous at each point in its domain. \square

Corollary 2.3. *Let G act geometrically on a CAT(0) space X , and let H be a surface subgroup which is quasiconvex for this action. Then $\Lambda(H)$ is homeomorphic to ∂H , where ∂H is the CAT(0) boundary of H .*

Proof. Each surface group (fundamental group of a closed surface or free group of rank ≥ 1) is hyperbolic except for $Z + Z$. Each $Z + Z$ subgroup of G has limit set a circle by the Flat Torus Theorem. The hyperbolic cases follow from Lemma 2.2. \square

Corollary 2.4. *Let H and H' be surface subgroups of a CAT(0) group G which are quasiconvex for the action. Then $\Lambda(H \cap H') = \Lambda(H) \cap \Lambda(H')$.*

2.2. Right-angled Coxeter groups. Let Γ be a finite simplicial graph. The *right-angled Coxeter group* associated to Γ has generating set S equal to the vertices of Γ , relations $s^2 = 1$ for each s in S and relations $st = ts$ whenever s and t are adjacent vertices of Γ . Given a graph Γ we denote the associated right-angled Coxeter group by W_Γ . Right-angled Coxeter groups are canonical examples of groups with nice geometric structures. For example, the right-angled Coxeter group on a path of length at least 3 can be realized as a Fuchsian group which acts geometrically on a strict subset of \mathbb{H}^2 . Hence any CAT(0) boundary of such a right-angled Coxeter group is either 2 points or a Cantor set.

A technique for finding index two subgroups of a right-angled Coxeter group that features heavily in the current paper is doubling. Let $v \in S$ and define $D_v \Gamma$

to be the graph obtained from Γ by gluing two copies of Γ along the star of v then deleting v and its open star in the new graph. We call $D_v\Gamma$ the *double of Γ over v* . If a vertex s of Γ is not in the link $\text{Lk}(v)$ we use s' to denote its double in $D_v\Gamma$. The following lemma is a folk result analogous to Example 1.4 of Bestvina–Kleiner–Sageev [3] (which is in the setting of right-angled Artin groups). For the sake of brevity we will not include a proof here.

Lemma 2.5 (Doubling Lemma). *Assume Γ is a finite simplicial graph, let v be a vertex of Γ , and set $\Delta = D_v\Gamma$. Then W_Δ is an index two subgroup of W_Γ .*

Associated to Γ is proper piecewise Euclidean CAT(0) complex Σ_Γ called the *Davis–Moussong* complex on which Γ acts geometrically. The space Σ_Γ is canonically constructed based solely on the combinatorial data of Γ (See [14] for the details of this construction.)

3. GRAPH THEORETIC RESULTS

In this section we prove some graph theoretic results which are used in the next section to prove Theorem 1.2.

3.1. Graph terminology. Given a graph Γ , an edge subdivision operation consists of adding a valence two vertex in the interior of an edge of Γ . A Γ *subdivision* is a graph obtained from Γ by a (possibly trivial) sequence of edge subdivision operations. A subgraph Λ of Γ is said to be *induced* if every pair of vertices of Λ which are adjacent in Γ are also adjacent in Λ (i.e. if u and w are vertices of Λ , and Γ contains the edge $[u, v]$, then Λ does too). An essential vertex of Γ is any vertex of valence at least 3. Vertices of valence two are called *non-essential*. A *branch* of Γ is an embedded path between essential vertices of the graph. It contains its endpoints, but does not contain any other essential vertices. The branch between a pair of essential vertices x and y will be denoted by $[x, y]$. We will sometimes also use this interval notation for paths which are not necessarily branches, when there is no ambiguity.

A *cycle* in Γ is an embedded loop. We will denote cycles either by the essential vertices or by the paths that they visit. For example, if a cycle passes through the essential vertices v_1, \dots, v_n in order, (so that $[v_i, v_{i+1}]$ is a branch for each $i \pmod n$), then we will denote it by (v_1, \dots, v_n) . (We only use this notation for graphs in which every pair of essential vertices is connected by at most one branch.) On the other hand, if a cycle passes through embedded paths $\alpha_1, \dots, \alpha_n$ in Γ such that for each $i \pmod n$, the terminal vertex of α_i is equal to the initial point of α_{i+1} , then we denote it by $(\alpha_1, \dots, \alpha_n)$. (Here we are not assuming that the α_i are branches.)

A graph is *inseparable* if it is connected, has no separating complete subgraph, no cut pair, and no separating complete subgraph suspension. Obviously, a triangle-free graph is inseparable if and only if it is connected, has no separating vertex, no separating edge, no cut pair, and no separating vertex suspension. This is equivalent to the corresponding right-angled Coxeter group not splitting over a finite group or a virtually cyclic group [25, 28].

Let $\Lambda \subset \Gamma$ be a $K_{3,3}$ (resp. K_5) subdivision. We say that a vertex of Λ is Λ -*essential* if it has valence bigger than 2 in Λ , and Λ -*non-essential* if its valence in Λ is 2. Note that given a vertex of Λ , its valence in Γ could be higher than its valence in Λ (so in particular, a Λ -non-essential vertex could be an essential vertex of Γ). If $\Lambda \subset \Gamma$ is a $K_{3,3}$ subdivision, a *vertex partition* for Λ is a partition of the

Λ -essential vertices into two sets of three vertices, such that every vertex in the first set is connected to every vertex of the second set by a branch of Λ . By a *shortest* graph with a given property, we will mean a graph having the fewest edges with that property.

Kuratowski's Theorem says that a graph is planar if and only if it contains either a K_5 subdivision or a $K_{3,3}$ subdivision. We begin with the following lemma, which will enable us to ignore the K_5 case when we are trying to establish the non-planarity of boundaries of right-angled Coxeter groups defined by non-planar graphs.

Lemma 3.1. *Let Γ be a triangle-free graph which contains a K_5 subdivision Λ . Then either Γ or the double of Γ over some vertex contains a $K_{3,3}$ subdivision.*

Proof. Choose a shortest K_5 subdivision Λ in Γ . Let a, b, c, d, e be the Λ -essential vertices. Since Λ is shortest, given any Λ -essential vertex, say a , there cannot be a Γ -edge between a and some vertex x on a branch of Λ incident to a , unless x is adjacent to a in Λ .

Next suppose there is a Γ -edge between some Λ -essential vertex and some Λ -non-essential vertex on a branch of Λ disjoint from it. (See Figure 3.1.) For definiteness, say there is an edge from a to the Λ -non-essential vertex f in the interior of $[d, e]$. Then there is a $K_{3,3}$ subdivision with vertex partition $\{a, e, d\}$ and $\{b, c, f\}$, as shown in Figure 3.1. This completes the proof of the lemma in this case.

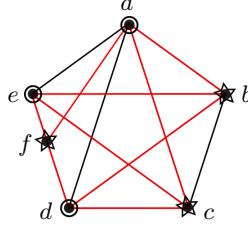


FIGURE 3.1

From the previous two paragraphs, we may assume for the remainder of the proof that if a Λ -essential vertex a is Γ -adjacent to a vertex x of Λ , then a is adjacent to x in Λ .

We claim that by relabeling the Λ -essential vertices if necessary, we may assume that a is not adjacent to vertices b and c . i.e. the branches $[a, b]$ and $[a, c]$ of Λ are subdivided. To see this, note that if there exist two Λ -essential vertices not adjacent to a , then we can simply relabel these as b and c . Otherwise, a is adjacent to at least three Λ -essential vertices, say b, c, d . Now since Γ is triangle-free, each of $[b, c]$, $[c, d]$ and $[b, d]$ is subdivided, and we can rename b to a , and d to b . This proves the claim.

We will now produce a $K_{3,3}$ subdivision in $D_a\Gamma$, assuming that b and c are not adjacent to a . By our assumption in the third paragraph, the link of a intersected with Λ consists of exactly four vertices, one on each branch incident to a . The vertices d and e could be among these. Let $\bar{\Gamma}$ and $\bar{\Gamma}'$ both denote Γ minus the open star of a , and recall that $D_a\Gamma$ is obtained by identifying $\bar{\Gamma}$ and $\bar{\Gamma}'$ along the link of a . For each vertex x of Γ , let x and x' be the corresponding vertices in $\bar{\Gamma}$ and

$\bar{\Gamma}'$ respectively. Since b and c are not adjacent to a in $\bar{\Gamma}$, we have that $b' \neq b$ and $c' \neq c$ in $D_a\bar{\Gamma}$. However, we could have $d = d'$ or $e = e'$.

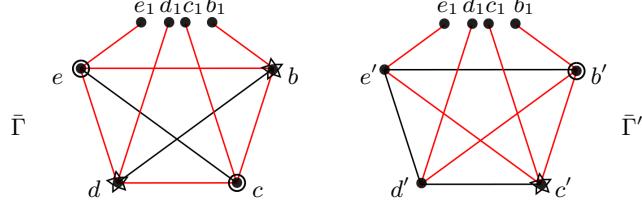


FIGURE 3.2. The graph $D_a\bar{\Gamma}$ is obtained by identifying $\bar{\Gamma}$ and $\bar{\Gamma}'$ along b_1, c_1, d_1 , and e_1 . Thus for example, the path from d to b' in $D_a\bar{\Gamma}$ consists of the branch $[d, d_1]$ in $\bar{\Gamma}$ followed by the branches $[d_1, d']$ and $[d', b']$. When a is adjacent to d (respectively e), then $d = d_1 = d'$ (respectively $e = e_1 = e'$).

We claim that there is a $K_{3,3}$ subdivision in $D_a\bar{\Gamma}$ with vertex partition $\{b, d, c'\}$ and $\{c, e, b'\}$. This is shown in Figure 3.2 in the case when $d \neq d'$ and $e \neq e'$. When $d = d'$, the path from b' to d in Figure 3.2 is replaced by the branch in $\bar{\Gamma}'$ from b' to $d' = d$. Similarly, when $e = e'$, the path from c' to e shown in Figure 3.2 is replaced by the branch in $\bar{\Gamma}'$ from c' to $e' = e$. Since the link of a in Λ consists of exactly four vertices, the subdivision constructed above is embedded in $D_a\bar{\Gamma}$. \square

The $K_{3,3}$ subdivision present in a non-planar graph Γ may not be an induced subgraph. In Proposition 3.2 we show that by successively doubling Γ along vertices finitely many times, one can find a subgraph that is either an induced $K_{3,3}$ subdivision, or one of two specific graphs.

Proposition 3.2. *Let Γ be a triangle-free graph which contains a $K_{3,3}$ subdivision. Then there exists a graph Γ' obtained from Γ by a sequence of doubling moves, and a $K_{3,3}$ subdivision $\Lambda' \subseteq \Gamma'$ such that either Λ' is induced or the subgraph of Γ' induced by Λ' is one of the graphs in Figure 3.3.*

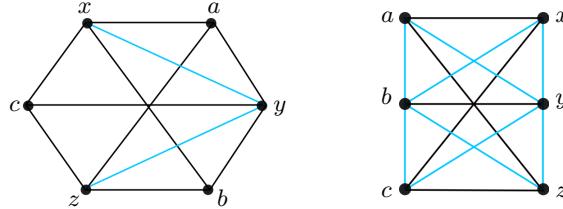


FIGURE 3.3. The figure shows the two possible graphs induced by Λ' in Γ' in Proposition 3.2. In both pictures, Λ' is the $K_{3,3}$ subdivision with vertex partition $\{a, b, c\}$ and $\{x, y, z\}$. The black paths may be subdivided, while the blue ones are edges. There are no edges of $\Gamma' \setminus \Lambda'$ connecting any pair of vertices in the graphs shown.

We begin by recalling and expanding terminology from subsection 3.1 and proving several lemmas to be used in the proof. In all of the proofs below, we assume

that Λ has Λ -essential vertex sets $\{a, b, c\}$ and $\{x, y, z\}$. Recall that a *branch* of Λ is the unique path between a pair of Λ -essential vertices which does not pass through any additional Λ -essential vertices. We assume that there are branches of Λ connecting each of a, b , and c to each of x, y , and z , and we will denote the branch between Λ -essential vertices, say a and x , by $[a, x]$. We will sometimes also use this interval notation to denote sub-paths of branches or edges. A branch or path $[s, t]$ will always include its endpoints s and t . Two branches will be called *adjacent* if they share an Λ -essential vertex endpoint.

If $\Lambda \subset \Gamma$ is not induced, we define a *bad edge* of Λ to be an edge in $E(\Gamma) \setminus E(\Lambda)$ whose endpoints are both vertices of Λ , and we define $B(\Lambda)$ to be the number of bad edges of Λ . Observe that Λ is induced if and only if $B(\Lambda) = 0$.

Idea of the proof. If the $K_{3,3}$ subdivision Λ is not induced, it has a non-trivial set of bad edges. If one doubles Γ over the endpoint v of a bad edge of Λ , then that edge disappears in $D_v\Gamma$. So the general strategy is to double over endpoints of bad edges of Λ and to find a new $K_{3,3}$ subdivision in the double which has fewer bad edges than Λ . Then after finitely many steps we end up with either an induced $K_{3,3}$ subdivision or one of the graphs in Figure 3.3.

The double of Λ in $D_v\Gamma$ typically has many more bad edges than Λ itself, and as a result, finding a $K_{3,3}$ subdivision in $D_v\Gamma$ with fewer bad edges than Λ can be a nontrivial feat. To aid this process, we do two things. Firstly, we start with a shortest subdivision Λ , and we show in Lemma 3.3 that this restricts the types of bad edges that may occur in Λ . (Types of bad edges may be differentiated based on whether they connect disjoint or adjacent branches, and whether their endpoints are Λ -essential or not.) Secondly, we choose the order of vertices to double over carefully.

Initially, we focus on reducing the number of bad edges which have at least one endpoint a Λ -non-essential vertex. In particular, we show in Lemmas 3.5 and 3.6, that if v is a Λ -non-essential vertex which is the endpoint of a bad edge of Λ , then unless the edges incident to v have a specific configuration (shown in Figure 3.8), the double $D_v\Gamma$ does contain a $K_{3,3}$ subdivision with fewer bad edges than Λ . Then, in the proof of Proposition 3.2, we show how to deal with the problematic configuration given in the statement of Lemma 3.6, Figure 3.8. This eliminates all bad edges which have at least one Λ -non-essential vertex as an endpoint. Also in the proof of Proposition 3.2, we resolve the case in which all bad edges have Λ -essential vertices as both their endpoints. It is in this case that the graphs in Figure 3.3 arise.

We begin with Lemma 3.3, which puts restrictions on the type of bad edges one could see in a shortest $K_{3,3}$ subdivision with a given number of bad edges.

Lemma 3.3. *Let Γ be a triangle-free graph containing a $K_{3,3}$ subdivision Λ . If Λ is shortest among all $K_{3,3}$ subdivisions with at most $B(\Lambda)$ bad edges, then Λ has no bad edges of the following types:*

- (1) *Bad edges with both endpoints on a single branch of Λ .*
- (2) *Bad edges connecting non- Λ -essential vertices on adjacent branches.*

Proof. Item (1) is obvious. For (2), suppose Λ has an edge connecting non- Λ -essential vertices v and w on two adjacent branches, say $[a, x]$ and $[a, y]$ (see Figure 3.4). Since Γ is triangle-free, the path from a to one of these vertices, say w ,

must be subdivided. Then there is a shorter $K_{3,3}$ subdivision with at most $B(\Lambda)$ bad edges (having vertex partition $\{x, y, z\}$ and $\{v, b, c\}$) as shown in Figure 3.4. This is a contradiction. \square

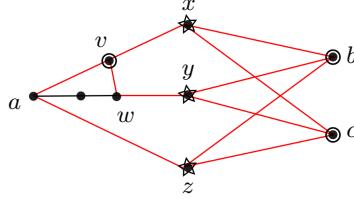


FIGURE 3.4. Consider the $K_{3,3}$ subdivision shown, with Λ -essential vertex sets $\{x, y, z\}$ and $\{v, b, c\}$. It is shorter than Λ and each of its bad edges is already a bad edge of Λ .

The next step is to begin doubling over Λ -non-essential vertices, and to find $K_{3,3}$ subdivisions with fewer bad edges in the double. We now know that bad edges between two Λ -non-essential vertices must go between two disjoint branches. Lemma 3.4 gives a useful consequence of the existence of such a bad edge. This will be used in Lemma 3.5, where we find a $K_{3,3}$ subdivision in the double over an endpoint of such an edge.

Lemma 3.4. *Let Γ be a triangle-free graph, and let Λ be shortest among all $K_{3,3}$ subdivisions with at most $B(\Lambda)$ bad edges. Suppose there is a bad edge connecting non- Λ -essential vertices on disjoint branches α and β of Λ . Then the unique branch of Λ disjoint from α and β is an edge.*

Proof. Assume without loss of generality that there is an edge connecting non- Λ -essential vertices v and w on branches $[a, x]$ and $[c, z]$ respectively. If $[b, y]$, the unique branch of Λ disjoint from $[a, x]$ and $[c, z]$, is not an edge, then one obtains a shorter $K_{3,3}$ subdivision with at most $B(\Lambda)$ bad edges, (having vertex partition $\{v, c, z\}$ and $\{w, a, x\}$) as shown in Figure 3.5. \square

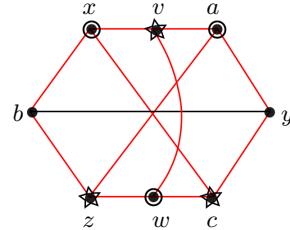


FIGURE 3.5. The $K_{3,3}$ subdivision shown, with Λ -essential vertex sets $\{v, c, z\}$ and $\{w, a, x\}$, has at most $B(\Lambda)$ bad edges. It is shorter than Λ unless $[b, y]$ is an edge.

The next two lemmas deal with finding $K_{3,3}$ subdivisions in doubles $D_v\Gamma$, where v is a Λ -non-essential vertex which is the endpoint of a bad edge. Lemma 3.5 gives a criterion on v which guarantees that $D_v\Gamma$ has a $K_{3,3}$ subdivision with fewer bad edges. This will be used in Lemma 3.6.

Lemma 3.5. *Let Γ be a triangle-free graph, and let Λ be shortest among all $K_{3,3}$ subdivisions with at most $B(\Lambda)$ bad edges. Let v be a non- Λ -essential vertex on a branch α of Λ such that v is an endpoint of a bad edge. If there is a branch β which is disjoint from α such that β does not intersect the link of v , then there exists a $K_{3,3}$ subdivision $\Lambda' \subseteq D_v(\Lambda)$ with $B(\Lambda') < B(\Lambda)$.*

Proof. Assume without loss of generality that $\alpha = [a, x]$ and that $\beta = [c, z]$ is disjoint from the link of v . See Figure 3.6. We know by Lemma 3.3(2) that the interiors of the branches $[a, z]$ and $[x, c]$ are disjoint from the link of v , and by Lemma 3.3(1) that the link of v intersects $[a, x]$ in exactly two vertices, u and w , with u possibly equal to a and w possibly equal to x . It follows that the path

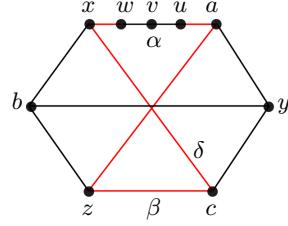


FIGURE 3.6. The link of v consists of u , w and possibly some vertices in $[z, b] \cup [b, y] \cup [y, c]$. The path δ is shown in red.

δ from u to w which consists of the concatenation of $[u, a]$, $[a, z]$, $[z, c]$, $[c, x]$, $[x, w]$ intersects the link of v only in u and w .

Now let $\bar{\Gamma}$ and $\bar{\Gamma}'$ be two copies of Γ with the open star of v removed. The double $D_v\Gamma$ is formed by identifying $\bar{\Gamma}$ and $\bar{\Gamma}'$ along the copy of the link of v in each. Define $\bar{\Lambda} \subseteq \bar{\Gamma}$ to be the copy of Λ minus the open star of v in $\bar{\Gamma}$, and note that $\bar{\Lambda}$ would be a $K_{3,3}$ subdivision if we added a path between u and w which is disjoint from $\bar{\Lambda}$. There is a copy of the path δ constructed above in $\bar{\Gamma}'$, which intersects $\bar{\Gamma}$ (and hence $\bar{\Lambda}$) only at u and w . Let γ be the shortest path in $\bar{\Gamma}'$ between u and w which intersects $\bar{\Gamma}$ only at u and w . Form Λ' by identifying $\bar{\Lambda}$ and γ along $\{u, w\}$. (This is shown in Figure 3.7 in the case $\gamma = \delta$.) Then Λ' is a $K_{3,3}$ subdivision in $D_v\Gamma$.

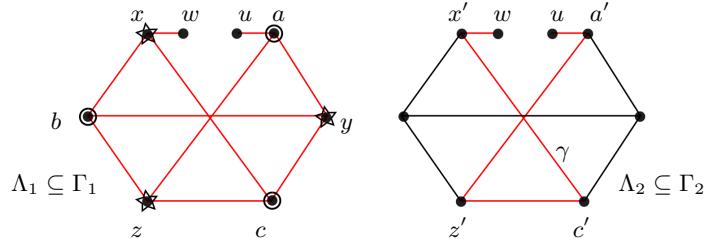


FIGURE 3.7. The two graphs shown are identified along $\text{Lk}(v)$ in $D_v\Gamma$. In particular, they are identified along u and w , and possibly some additional vertices on $[z, b] \cup [b, y] \cup [y, c]$. The red graph on the left is $\bar{\Lambda}$. The red path shown on the right is γ , in the case that $\gamma = \delta$. The choice of γ ensures that no vertex on it is identified with a vertex in $\bar{\Gamma}$.

If e is a bad edge of Λ' , our choice of γ implies that the endpoints of e are in $\bar{\Gamma}$. Thus there is a bad edge of Λ in Γ that corresponds to e . On the other hand, there is at least one bad edge of Λ incident to v , for which there is no corresponding bad edge of Λ' . Thus $B(\Lambda') < B(\Lambda)$. \square

The following lemma shows that if v is Λ -non-essential and is the endpoint of a bad edge of Λ , then $D_v\Gamma$ does contain a $K_{3,3}$ subdivision with fewer bad edges than Λ , except possibly in one particular situation. (This situation is specified in conditions (1) and (2) of the lemma).

Lemma 3.6. *Let Γ be a triangle-free graph, and let Λ be shortest among all $K_{3,3}$ subdivisions with at most $B(\Lambda)$ bad edges. Let v be a non- Λ -essential vertex on a branch α of Λ such that v is an endpoint of a bad edge. Then either there exists a $K_{3,3}$ subdivision $\Lambda' \subseteq D_v(\Lambda)$ with $B(\Lambda') < B(\Lambda)$ or both of the following statements hold. (See Figure 3.8.)*

- (1) The vertex v is adjacent to exactly one Λ -essential vertex s of Λ which is not an endpoint of α , and to at least one non- Λ -essential vertex on each of the two branches that are disjoint from both α and s .
- (2) The branch α consists of exactly the two edges incident to v .

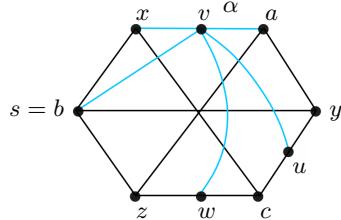


FIGURE 3.8. This illustrates the configuration from Lemma 3.6 in the case that v lies on $[a, x]$, and v is adjacent to b . Then the lemma says that v must be adjacent to at least one vertex on each of $[c, y]$ and $[c, z]$, and that $[v, x]$ and $[v, a]$ are edges.

Proof. Without loss of generality, assume v lies on $\alpha = [a, x]$. First we consider the case that v is not adjacent to any of the Λ essential vertices b, c, y , or z . In this case we show that $D_v\Gamma$ contains a $K_{3,3}$ subdivision $\Lambda' \subseteq D_v(\Lambda)$ with $B(\Lambda') < B(\Lambda)$. By hypothesis, there is a bad edge incident to v , and by Lemma 3.3, its other endpoint has to be on a branch disjoint from α , say (without loss of generality) $[c, z]$. Then by Lemma 3.4, the branch $[b, y]$ is an edge. Since by assumption v is not adjacent to b or y , we may apply Lemma 3.5 with $\beta = [b, y]$, to conclude that there exists a $K_{3,3}$ subdivision $\Lambda' \subseteq D_v(\Lambda)$ with $B(\Lambda') < B(\Lambda)$.

Thus we may assume that v is adjacent to at least one of b, c, y , and z . Now we analyze a few cases. If v is adjacent to both b and c , then one obtains a shorter $K_{3,3}$ subdivision with at most $B(\Lambda)$ bad edges, as shown on the left in Figure 3.9, which is a contradiction, since Λ was chosen to be shortest. Thus v is adjacent to at most one of b and c , and similarly, v is adjacent to at most one of y and z .

Suppose v is adjacent to exactly one from each pair, say b and y . Then since Γ is triangle-free, $[b, y]$ is not an edge. Applying Lemma 3.4, we conclude that v is not adjacent to any vertex in $[c, z]$. Then by Lemma 3.5, with $\beta = [c, z]$, there exists a $K_{3,3}$ subdivision $\Lambda' \subseteq D_v(\Lambda)$ with $B(\Lambda') < B(\Lambda)$.

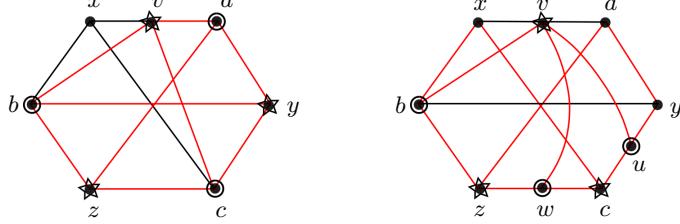


FIGURE 3.9. If v is connected to both b and c , then one obtains the red graph on the left. It is shorter than Λ and has at most $B(\Lambda)$ bad edges. If v is adjacent to b , w , and u as shown, and if one of $[a, v]$ and $[v, x]$ is not an edge, then the red graph on the right is shorter than Λ and has at most $B(\Lambda)$ bad edges.

We are left with the case that v is adjacent (via a bad edge) to exactly one Λ -essential vertex, say b . As before, if $\text{Lk}(v)$ fails to intersect one of the branches $[c, y]$ and $[c, z]$, then applying Lemma 3.5, we would find the desired Λ' in $D_v\Gamma$. If not, then v is adjacent to non- Λ -essential vertices u and w on $[y, c]$ and $[z, c]$ respectively, i.e. condition (1) in the statement of the lemma holds. Finally, if we have the configuration in (1), but one of $[x, v]$ and $[a, v]$ is not an edge, we get a shorter $K_{3,3}$ subdivision as shown on the right in Figure 3.9. The above reasoning shows that either $D_v\Gamma$ contains a $K_{3,3}$ subdivision $\Lambda' \subseteq D_v(\Lambda)$ with $B(\Lambda') < B(\Lambda)$ or (1) and (2) both hold simultaneously. \square

Using the Lemma 3.6, we can keep doubling over Λ -non-essential vertices and finding $K_{3,3}$ subdivisions with fewer bad edges until we either run out of Λ -non-essential vertices which are endpoints of bad edges, or get to the point where *every* such Λ -non-essential vertex has the configuration in Figure 3.8. In particular, by applying Lemma 3.6 to w and u (from Figure 3.8) we obtain strong restrictions on the graph, which enable us to find an induced $K_{3,3}$ subdivision. Finally, we show that when all of the bad edges have Λ -essential vertex endpoints, then either the graph is one of the graphs in Figure 3.3 or some double contains a $K_{3,3}$ subdivision with fewer bad edges. This is all put together in the following proof.

Proof of Proposition 3.2. Choose a $K_{3,3}$ subdivision $\Lambda \subseteq \Gamma$ such that Λ is shortest among all $K_{3,3}$ subdivisions with at most $B(\Lambda)$ bad edges.

Claim: If Λ is not induced, one of the following holds:

- (i) Λ induces one of the graphs in Figure 3.3.
- (ii) For some vertex v , the double $D_v\Gamma$ contains a $K_{3,3}$ subdivision Λ' such that $B(\Lambda') < B(\Lambda)$.

Before proving the claim, we explain why it is sufficient to complete the proof. Given a $K_{3,3}$ subdivision Λ , if it is induced or if (i) holds, i.e. if it induces one of the graphs in Figure 3.3, then we are done. Otherwise, (ii) holds. We take Λ_2 to be the shortest $K_{3,3}$ subdivision in $\Gamma_2 = D_v\Gamma$ with at most $B(\Lambda')$ bad edges, where $D_v\Gamma$ and Λ' are provided by (ii), and we repeat the argument with Λ_2 and Γ_2 instead of Λ and Γ . After finitely many steps we arrive at a pair $\Lambda_n \subseteq \Gamma_n$ which either satisfies (i) or such that $B(\Lambda_n) = 0$, which means Λ_n is induced. This proves the proposition.

Proof of the claim: By Lemma 3.3, Λ has no bad edges between any pair of vertices that lie on a single branch, or any pair of Λ -non-essential vertices on adjacent branches.

Case 1. Non-essential vertex on a bad edge. Suppose there exists a Λ -non-essential vertex v , say on the branch $[a, x]$, which is the endpoint of a bad edge. Then by Lemma 3.6, either (ii) in the claim above holds (in which case we are done) or we may assume that Λ has the configuration specified by conditions (1) and (2) of Lemma 3.6. In the latter case, $[v, a]$ and $[v, x]$ are edges and we can assume, by re-labeling if necessary, that the edges incident to v guaranteed by (1) are as shown in Figure 3.8. In particular, v is adjacent to b , and to Λ -non-essential vertices w on $[c, z]$ and u on $[c, y]$.

Now apply Lemma 3.6 to w . If $D_w\Gamma$ contains a $K_{3,3}$ subdivision Λ' such that $B(\Lambda') < B(\Lambda)$, we are done. If not, we conclude (from Lemma 3.6(1)) that w is adjacent to exactly one of a, b, x and y . (See the left side of Figure 3.10.) Since Γ is triangle-free, and v is already adjacent to a, x, b and w , we see that w cannot be adjacent to any of a, x , or b . Thus we conclude that w is adjacent to y . By Lemma 3.6(1), there is a Λ -non-essential vertex t on $[x, b]$ adjacent to w . Furthermore, $[w, c]$ and $[w, z]$ are edges by Lemma 3.6(2).

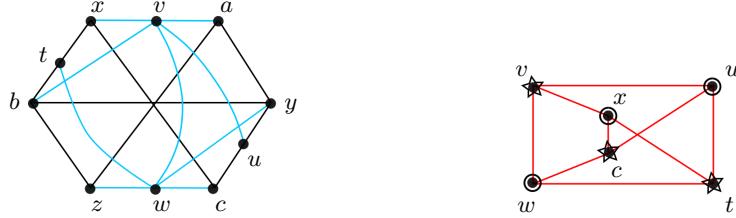


FIGURE 3.10. The picture on the left shows the configuration obtained after Lemma 3.6 has been applied to v and then to w . Blue paths are edges.

After applying Lemma 3.6 to u , we have additional edges from u to z and t . Then the resultant graph contains a $K_{3,3}$ subdivision as shown in the picture on the right.

Applying similar reasoning to u , we conclude (from Lemma 3.6(2)) that $[c, u]$, $[u, y]$ are edges and (from Lemma 3.6(1)) that u is adjacent to z and to a Λ -non-essential vertex t' on the branch $[x, b]$. Finally, applying Lemma 3.6(2) to the vertex t from the previous paragraph, we conclude that $[t, x]$ and $[b, t]$ are edges. It follows that $t = t'$, so that u is adjacent to t .

Then Γ contains a $K_{3,3}$ subdivision as shown on the right in Figure 3.10. All the branches of this $K_{3,3}$ subdivision are edges, except possibly $[c, x]$. If e is a bad edge of this graph, then Lemma 3.3(1) together with the triangle-free condition implies that e must connect one of t, u, v and w to a Λ -non-essential vertex on $[x, c]$. However, since each of t, u, v , and w lies on a branch of Λ adjacent to $[c, x]$, Lemma 3.3(2) implies that there are no bad edges of this kind. Thus the $K_{3,3}$ subdivision is induced. In particular, (ii) of the claim holds. This completes the proof of the claim in the case that there is at least one non- Λ -essential vertex of Λ which is the endpoint of a bad edge.

Case 2. Every bad edge has essential vertices. It remains to consider the case that all endpoints of bad edges are Λ -essential. Lemma 3.3 (1) implies that any bad edge has its both its endpoints in $\{x, y, z\}$ or both in $\{a, b, c\}$. Moreover, since Γ is triangle-free, for each of these sets, there can be at most two edges connecting pairs of vertices in the set.

First consider the case when one of the sides has exactly one edge (and the other side has zero, one or two edges). Assume without loss of generality, that there is a bad edge between x and y , but none between either of these and z . Then $D_x\Gamma$ contains a $K_{3,3}$ subdivision Λ' such that $B(\Lambda') < B(\Lambda)$. If there are no bad edges with endpoints among a', b', c' , then the required $K_{3,3}$ subdivision is as shown in Figure 3.11, and is induced. If there are such bad edges, then in particular, these bad edges have endpoints between non-essential vertices of adjacent branches of the $K_{3,3}$ subdivision in Figure 3.11, and they can be eliminated using the procedure in the proof of Lemma 3.3 (2). Then the resulting graph only has bad edges with endpoints among a, b, c , but we have eliminated the bad edge between x and y . So the resulting graph has fewer bad edges.

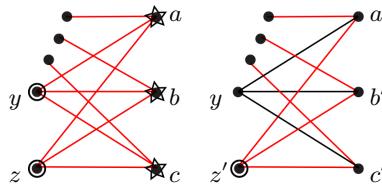


FIGURE 3.11. This illustrates the induced $K_{3,3}$ subdivision in $D_x\Gamma$, when Λ a bad edge between x and y , and no other bad edges.

Finally, we are left with the case that one of the sides has two edges (say there are edges between x and y , and between y and z) and the other side has zero or two bad edges. If the other side has zero bad edges, we have the graph on the left in Figure 3.3. Otherwise assume without loss of generality that $[a, b]$ and $[b, c]$ are the two edges on the other side. Now suppose that one of $[x, b]$, $[z, b]$, $[a, y]$, and

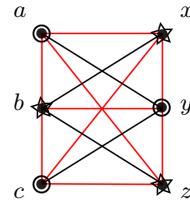


FIGURE 3.12. If one of $[x, b]$, $[z, b]$, $[a, y]$, $[c, y]$ is not an edge, then the red $K_{3,3}$ subdivision is shorter than Λ and has at most $B(\Lambda)$ bad edges.

$[c, y]$ is not an edge. Then we obtain a shorter $K_{3,3}$ subdivision with at most $B(\Lambda)$ bad edges, as shown in Figure 3.12, establishing (ii) of the claim. Otherwise, all of $[x, b]$, $[z, b]$, $[a, y]$, and $[c, y]$ are edges, and we obtain the graph on the right in Figure 3.3. \square

4. BOUNDARIES OF RIGHT-ANGLED COXETER GROUPS DEFINED BY NON-PLANAR GRAPHS.

In this section we prove Theorem 1.2. That is, we show that if Γ is a non-planar triangle-free graph, then either ∂X_Γ is non-planar for any CAT(0) space X on which W_Γ acts geometrically, or W_Γ contains a finite index subgroup whose defining graph contains an induced copy of the graph in Figure 1.2. We will show below that the graph theoretic results of the previous section together with the doubling lemma (Lemma 2.5) can be used to reduce this to proving the following two propositions.

Proposition 4.1. *If Δ is a $K_{3,3}$ subdivision and W_Δ acts geometrically on a CAT(0) space X_Δ , then ∂X_Δ is non-planar.*

Proposition 4.2. *Let Δ be the graph on the right side in Figure 3.3, subdivided enough so that it is triangle-free, and suppose that W_Δ acts geometrically on the CAT(0) space X_Δ . Then ∂X_Δ is non-planar.*

Before we prove the above propositions, we indicate how to use them to deduce Theorem 1.2.

Proof of Theorem 1.2. Let Γ be a triangle-free non-planar graph. Then Kuratowski's Theorem says that Γ contains a (not necessarily induced) $K_{3,3}$ or K_5 subdivision. If Γ contains a K_5 subdivision, then Lemma 3.1 says that Γ_1 contains a (not-necessarily induced) $K_{3,3}$ subdivision, where Γ_1 is either Γ itself, or the double of Γ over some vertex. Now we apply Proposition 3.2 to Γ_1 . We conclude that there is a graph Γ_2 which is obtained from Γ_1 by a finite sequence of doubling moves, and an induced subgraph Δ of Γ_2 , such that Δ is either a $K_{3,3}$ subdivision or one of the two graphs in Figure 3.3. Lemma 2.5 implies that W_{Γ_2} is a finite-index subgroup of W_Γ .

Now if Δ is either a $K_{3,3}$ subdivision or the graph on the right in Figure 3.3, then Propositions 4.1 and 4.2 say that every boundary of W_Δ is non-planar. The group W_Δ is a special subgroup of W_{Γ_2} , and therefore every boundary of W_{Γ_2} is non-planar. Now suppose that W_Γ acts geometrically on a CAT(0) space X_Γ . Then since W_{Γ_2} is a finite index subgroup of W_Γ , W_{Γ_2} also acts geometrically on X_Γ . Therefore the boundary of X_Γ is non-planar.

On the other hand, if Δ is the graph on the left in Figure 3.3 (which is the same as the graph in Figure 1.2), then we have produced a finite-index subgroup of W_Γ , namely W_{Γ_2} , whose defining graph Γ_2 contains an induced copy of the graph in Figure 1.2. \square

We now prove Propositions 4.1 and 4.2. We begin with a lemma about the boundaries of Θ -graph subdivisions. (By a Θ -graph we mean a graph with two essential vertices and three distinct edges between the essential vertices.) Recall that a branch in a graph is an embedded path between essential vertices of the graph. It contains its endpoints, but does not contain any other essential vertices.

Lemma 4.3. *Let Λ be a Θ -graph subdivision such that each branch has length at least 2, and at least one of the branches has length at least three. (See Figure 4.1.) If X_Λ is a CAT(0) space on which W_Λ acts geometrically, then ∂X_Λ contains an embedded Θ -graph.*

We remark that the conclusion is true even without the condition that at least one branch has length at least three, but the above lemma is sufficient for our purposes.

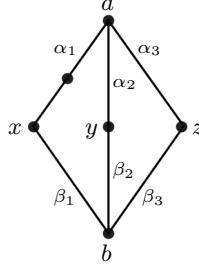


FIGURE 4.1. A Θ -graph subdivision with each branch of length at least two, and at least one branch of length at least 3.

Proof. Assume that Λ has essential vertices a and b , and non-essential vertices x, y , and z as shown in Figure 4.1. Let α_1, α_2 , and α_3 be the paths from a to x, y , and z , respectively. Define $\beta_1, \beta_2, \beta_3$ analogously with the β_i incident to b . The hypotheses imply that at least one of these paths is subdivided, so we may assume α_1 is subdivided.

Let \mathcal{D}_2 and \mathcal{D}_3 be the cycles given by $(\beta_1, \alpha_1, \alpha_2, \beta_2)$ and $(\beta_1, \alpha_1, \alpha_3, \beta_3)$ respectively, and let G_2 and G_3 be the corresponding special subgroups. Consider the quasi-isometry (coming from the orbit map) between X_Λ and the Davis complex Σ_Λ . Let X_2 denote the image of Σ_{G_2} under this quasi-isometry based at the image of the identity vertex. Note that X_2 is quasi-convex. Define X_3 analogously. We will use the cycles \mathcal{D}_2 and \mathcal{D}_3 to find a theta graph in ∂X_Λ .

Now \mathcal{D}_2 and \mathcal{D}_3 are cycles of length at least 5, so G_2 and G_3 are hyperbolic reflection groups acting geometrically on \mathbb{H}^2 . Thus $\partial X_2 \cong \partial X_3 \cong S^1$. Since \mathcal{D}_2 and \mathcal{D}_3 intersect in a path of length at least three, the corresponding special subgroup is virtually free and its boundary, which is equal to $\partial X_2 \cap \partial X_3$, is homeomorphic to either two points or a Cantor set, see Section 2.2. Furthermore, $\{(ab)^\infty, (ba)^\infty\} \subset \partial X_2 \cap \partial X_3$, and this set separates each of ∂X_2 and ∂X_3 into two components, as shown in Figure 4.2.

We now explicitly find a Θ -graph in ∂X_Λ with essential vertices $(ab)^\infty$ and $(ba)^\infty$. Let Z be the subset of ∂X_3 represented by Cayley graph geodesic rays whose first letter is the label of a vertex on the branch $\alpha_3 \cup \beta_3$. Note that Z is the arc of ∂X_3 which intersects $\partial X_2 \cap \partial X_3$ (and hence ∂X_2) only at its endpoints $(ab)^\infty$ and $(ba)^\infty$. Then $\partial X_3 \cup Z$ is an embedded Θ -graph in ∂X_Λ . \square

We now prove Proposition 4.1. In the proof we will use the embedded Θ -graph in ∂X_Γ that we constructed in the previous lemma.

Proof of Proposition 4.1. Let Δ be a $K_{3,3}$ subdivision with essential vertex sets $\{a, b, c\}$ and $\{x, y, z\}$. Let $\alpha_1, \alpha_2, \alpha_3$ be the branches from a to x, y , and z , respectively. Define $\beta_1, \beta_2, \beta_3$ and $\gamma_1, \gamma_2, \gamma_3$ analogously with the β_i incident to b , and the γ_i incident to c . (See Figure 4.3.)

If Δ is graph-isomorphic to $K_{3,3}$ (i.e. if it is the trivial $K_{3,3}$ subdivision) then ∂X_Δ is the join of two Cantor sets, and is therefore non-planar. Thus we may assume that at least one branch, say α_1 , of Δ is subdivided.

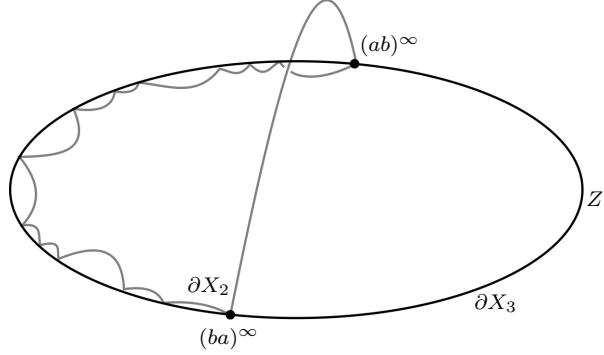


FIGURE 4.2. The figure shows ∂X_2 in grey and ∂X_3 in black. The intersection $\partial X_2 \cap \partial X_3$ is a Cantor set. The set $\{(ab)^\infty, (ba)^\infty\}$ separates the grey and black circles into two components each. The embedded Θ -graph constructed in the proof consists of the grey circle together with the arc labeled Z .

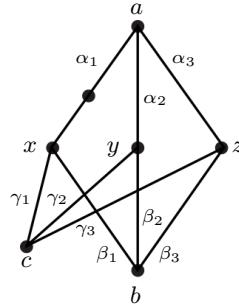


FIGURE 4.3. A $K_{3,3}$ subdivision in which at least one branch is subdivided.

Observe that the union of the α_i and the β_i is precisely the Θ -graph subdivision Λ from Lemma 4.3. (See Figure 4.1.) Retaining the notation of Lemma 4.3, we see that $\partial X_2 \cup Z$ is an embedded Θ -graph in ∂X_Δ . To complete the proof we will find another half of a Θ -graph which intersects $\partial X_2 \cup Z$ in exactly three points. (See Figure 4.4.)

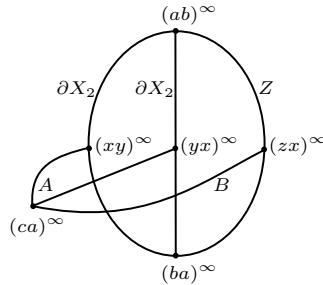


FIGURE 4.4. An embedded $K_{3,3}$ in ∂X_Δ

Let \mathcal{D}'_2 and \mathcal{D}'_3 be the cycles given by $(\gamma_1, \alpha_1, \alpha_2, \gamma_2)$ and $(\gamma_1, \alpha_1, \alpha_3, \gamma_3)$, respectively. Define X'_2 and X'_3 to be the quasiconvex subspaces of X_Δ corresponding to the special subgroups generated by \mathcal{D}'_2 and \mathcal{D}'_3 . Following the argument of Lemma 4.3 above with the cycles \mathcal{D}_2 and \mathcal{D}'_2 , we see that ∂X_2 and $\partial X'_2$ intersect in a Cantor set, and that there is a closed arc A of $\partial X'_2$ such that $A \cap \partial X_2 = \{(xy)^\infty, (yx)^\infty\}$. The interior of A contains the point $(ca)^\infty$ and is disjoint from ∂X_2 .

Similarly considering the cycles \mathcal{D}_3 and \mathcal{D}'_3 , we conclude that $\partial X_3 \cap \partial X'_3$ is a Cantor set, and that there exists closed arc C in $\partial X'_3$ which intersects ∂X_3 in $\{(xz)^\infty, (zx)^\infty\}$ and whose interior contains $(ca)^\infty$. Define B to be the subarc of C which connects $(ca)^\infty$ to $(zx)^\infty$. Notice that since z is the label of the initial edge in the Cayley graph geodesic for the ray $(zx)^\infty$ we have that $(zx)^\infty \in Z$. Moreover, points of B are represented by Cayley graph geodesic rays whose first letter is a label of a vertex on γ_3 , see Lemma 2.1, and therefore B intersects $\partial X_2 \cup Z \cup A$ exactly in the two points $(ca)^\infty$ and $(zx)^\infty$. Thus $\partial X_2 \cup Z \cup A \cup C$ is an embedded $K_{3,3}$ in ∂X_Δ , as shown in Figure 4.4. \square

Next, we prove Proposition 4.2 by showing that if Δ is the graph on the right in Figure 3.3, then ∂X_Δ is non-planar. Figure 4.5 below reproduces this graph, and also shows an alternate view of it. Note that in our application (Theorem 1.2), this graph is an induced subgraph of a triangle-free graph, which forces all of the black edges in the figure to be subdivided.

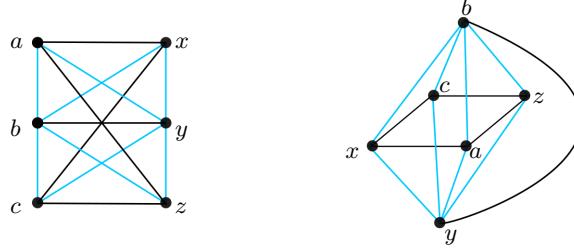


FIGURE 4.5. The figure shows two different views of the graph being considered in Proposition 4.2. The blue branches are edges. The black branches are necessarily subdivided.

Proof of Proposition 4.2. Let x, y, z, a, b , and c be the essential vertices of Δ , as shown in Figure 4.5. We now define four cycles of Δ as follows: $\mathcal{D}_1 = (x, c, z, a)$, $\mathcal{D}_2 = (x, b, z, y)$, $\mathcal{D}_3 = (c, b, a, y)$, and $\mathcal{D}_4 = (b, z, y)$. Note that the branch $[b, y]$ of \mathcal{D}_4 is subdivided by hypothesis.

For $1 \leq i \leq 4$, let X_i be the quasi-isometrically embedded copy of the Davis complex of the special subgroup generated by \mathcal{D}_i which is based at the image identity vertex in X_Δ . For each i , the cycle \mathcal{D}_i has length at least 4, so the corresponding special subgroup is virtually a surface group, and $\partial X_i \cong S^1$. The intersections of these circles are shown on the left in Figure 4.6. It follows that ∂X_Γ contains an embedded $K_{3,3}$ as shown on the right in Figure 4.6. \square

5. THE CASE OF THE BAD GRAPH

In this section we deal with the case that our defining graph contains an induced copy of the graph Π (which is subdivided along some of the black edges). As

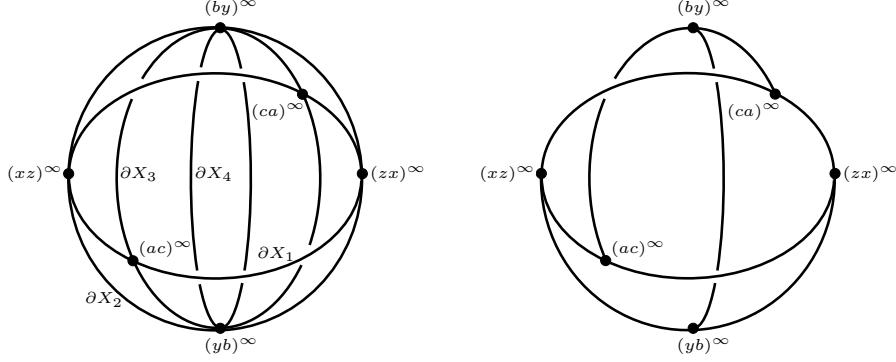


FIGURE 4.6. The picture on the left shows the pairwise intersections of the circles $\partial X_1, \partial X_2, \partial X_3$, and ∂X_4 . The picture on the right shows an embedded $K_{3,3}$ in $\partial X_1 \cup \partial X_2 \cup \partial X_3 \cup \partial X_4 \subset \partial X_\Gamma$.

discussed in the introduction, W_Π has a planar boundary. However, we will show in this section that the action of the group W_Π on any CAT(0) boundary for W_Π does not extend to the plane. Therefore, when Π is embedded in a graph Γ and W_Γ acts geometrically on a CAT(0) space X such that ∂X is connected, locally connected, and without local cut points, then ∂X cannot be planar.

Theorem 5.1. *Let Π be the graph in Figure 1.2, with the non-blue edges subdivided so that Π does not contain triangles. Let W_Π be the corresponding right-angled Coxeter group, and ∂X_Π be the boundary of any proper CAT(0) space that W_Π acts upon geometrically. Then:*

- (a) ∂X_Π contains an embedded copy of the graph in Figure 5.2.
- (b) The copies of the circles A , B , and C shown in Figure 5.2 are invariant under the induced action of y on ∂X_Π .
- (c) For any embedding of ∂X_Π in S^2 , the induced action of W_Π on ∂X_Π by homeomorphisms does not extend to S^2 .

The proof of Theorem 5.1 is delayed until the end of the section. First we discuss two corollaries. Recall that a topological space is *planar* if it embeds in S^2 .

Corollary 5.2. *Let Γ be a triangle-free finite simplicial graph which contains an induced copy of Π (with non-blue edges possibly subdivided). Suppose that ∂X_Γ is connected, locally connected, and has no local cut points. Then ∂X_Γ is not planar.*

Proof of Corollary 5.2 from Theorem 5.1. First, we claim that if ∂X_Γ satisfies all the hypotheses and is planar, then ∂X_Γ is a Sierpinski carpet. Indeed, the Sierpinski carpet is the unique 1-dimensional topological space which is compact, connected, locally connected, planar, and has no cut points or local cut points [31]. Our assumption of triangle-free ensures that the boundary is 1-dimensional by a theorem of Bestvina–Mess [1, 4]. Since ∂X_Γ is compact, and has no cut points [26, Theorem 1], this proves the claim.

Now suppose that ∂X_Γ is a Sierpinski carpet \mathcal{S} and Γ contains an induced copy of the bad graph Π . Then since W_Π is a special subgroup of W_Γ , the action of W_Π on $\partial X_\Pi \subset \partial X_\Gamma$ extends to an action by homeomorphisms on the Sierpinski carpet. Every homeomorphism of the Sierpinski carpet preserves the set of non-separating circles. Let $\mathcal{S} \cong S^2 \setminus \bigcup_i D_i$, where $\{D_i\}$ is a dense null family of open round discs

D_i in S^2 such that $\bar{D}_i \cap \bar{D}_j = \emptyset$ if $i \neq j$. Then the non-separating circles in \mathcal{S} are exactly the boundaries of the D_i . Thus every homeomorphism of \mathcal{S} extends to S^2 , so the action of W_Π on ∂W_Γ extends to an action on S^2 . But this contradicts Theorem 5.1, so ∂X_Γ must be non-planar. \square

We can now put the pieces together to prove Theorem 1.3, which we re-state for the convenience of the reader.

Theorem 1.3 *Let Γ be a triangle-free inseparable graph and let X_Γ be a CAT(0) space on which W_Γ acts geometrically. If Γ is non-planar and ∂X_Γ is locally connected and contains no local cut points, then ∂X_Γ is non-planar.*

Proof. Suppose that Γ has no triangles and is non-planar. Then by Theorem 1.2 either ∂X_Γ is non-planar, or W_Γ contains a finite index subgroup W_Λ such that Λ contains an induced copy of Π . Then since Γ is inseparable, W_Γ is one-ended. Thus ∂X_Γ ($\cong \partial X_\Lambda$) is connected. By hypothesis it is locally connected and has no local cut points. Therefore, by Corollary 5.2, ∂X_Γ is non-planar. \square

Corollary 5.3. *Suppose Γ is triangle-free and inseparable such that W_Γ is either hyperbolic or CAT(0) with isolated flats. Suppose further that Γ contains an induced (with non-blue edges possibly subdivided) copy of the bad graph Π below in Figure 5.1. Then ∂X_Γ is the Menger curve, where X_Γ is any proper CAT(0) that W_Γ acts on geometrically.*

Proof of Corollary 5.3. By a theorem of Bestvina–Mess [1, 4], the boundary ∂X_Γ is 1-dimensional since Γ is triangle-free. As W_Γ is one-ended, ∂X_Γ is connected. If W_Γ is hyperbolic, then the assumption that Γ is inseparable (so in particular W_Γ does not split over any two-ended group) implies that ∂X_Γ does not contain any local cut points by a result of Bowditch [6]. When W_Γ has isolated flats the analogous result is due to Haulmark [18]. In both cases, ∂X_Γ is locally connected; this follows by Bestvina–Mess [4] or Bowditch [6] in the hyperbolic case and by Hruska–Ruane [21] in the CAT(0) with isolated flats case. Therefore, all the hypotheses of Corollary 5.2 are satisfied, and the boundary is non-planar. Thus it must be a Menger curve [22, 18]. \square

Proof of Theorem 5.1. Note that the boundary of X_Π is possibly planar; the double cover y of Π is a planar graph, see Lemma 2.5. The point of the proof is that the action of y on any boundary of W_Π is non-planar. Consider the graph Π in Figure 5.1 below. The black edges may be subdivided; the two edges connecting x to y and y to z are not. We will also assume that our defining graph has no triangles, which will force some of the edges to be subdivided.

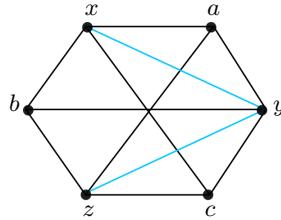


FIGURE 5.1. This is the bad graph Π .

We now define three cycles in Π as follows: $A = (x, a, y)$, $B = (x, b, y)$, and $C = (x, c, y)$, where these denote the cycles defined by the branches traversing these essential vertices (and back to the initial essential vertex).

The associated special subgroups W_A , W_B , and W_C , intersect only in the finite subgroup generated by x and y . Since every boundary of a CAT(0) space X that a virtual surface group acts upon is S^1 , we denote their boundaries using ∂W_A , etc. It follows that their boundaries ∂W_A , ∂W_B , and ∂W_C are disjoint circles, each of which is invariant by the action of y . By slight abuse of notation, we label these boundary circles in Figure 5.2 below by A , B , and C respectively.

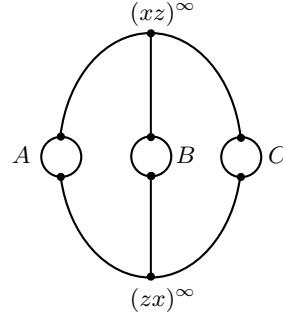


FIGURE 5.2. This figure is planar but the action of y is non-planar

Furthermore, the two fixed points of the action of y in each circle are the endpoints of the loxodromic generated by the two vertices adjacent to y . For example, if the branch $[y, a]$ of Π is not subdivided this is the limit set of the loxodromic element xa . Notice that the two points on the boundary ∂X_Π associated to the sequences $(xz)^\infty$ and $(zx)^\infty$ are also fixed by the action of y , since both x and z commute with y .

We wish to construct an embedded copy of Figure 5.2 in ∂X_Π . Throughout we will implicitly use Lemma 2.1. Consider the following two cases:

- xa is a loxodromic. In this case, use the cycle in Π defined by (x, a, z, b) . The induced special subgroup is virtually a surface group and so the boundary of this special subgroup is a circle. This circle intersects the circle A in the boundary of the subgroup defined by the branch $[x, a]$. The endpoint associated to the sequence $(xa)^\infty$ can be connected to the endpoint of the sequence $(xz)^\infty$ using the arc in this circle which avoids the boundary of the subgroup associated to the branch $[x, a]$. Note that the endpoints of the loxodromic element xz separate the endpoints of xa from the endpoints of xb , when xb is loxodromic. This follows from the order of the elements around A .
- xa is not loxodromic. In this case ya is loxodromic. In this case we will use the cycle in Π defined by (x, a, z, y) . Then we can connect the elements $(xz)^\infty$ and $(ya)^\infty$ by an arc on this circle.

We do this for each of A , B , and C using the analogously defined circles. If xb is loxodromic, use the cycle in Π defined by (x, a, z, b) , and using the other side from above, connect $(xa)^\infty$ to the endpoint of the sequence $(xz)^\infty$. If xb is not loxodromic, use (x, b, z, y) , connecting $(xz)^\infty$ to the endpoints of $(yb)^\infty$. Similarly, if xc is loxodromic, we use the cycle in Π defined by (x, a, z, c) . The endpoints of xc are separated from the endpoints of xa by the endpoints of xz . If xc is not

loxodromic, we use the cycle (x, c, z, y) , connecting $(xz)^\infty$ to $(yc)^\infty$. Moving by the action of x takes the arcs connecting the endpoint of the ray $(xz)^\infty$ to arcs connecting the endpoints of the rays $(zx)^\infty$ and another point on the circles A , B and C . Therefore, in either case, we have a copy of the Figure 5.2 in any $\text{CAT}(0)$ boundary of W_Π , and hence in the boundary ∂X_Γ , since Π is assumed to be an induced subgraph of Γ . This proves (a) in the statement of Theorem 5.1.

Now we claim that the action of the group element y acting on Figure 5.2 does not extend to the plane. Up to relabelling, any embedding of Figure 5.2 can be moved via a homeomorphism of the pair $(S^2, \text{Figure 5.2})$ to the embedding given. The circles A , B and C are all invariant under the action of y , as are the boundary points $(xz)^\infty$ and $(zx)^\infty$. The arcs connecting the points $(xz)^\infty$ and $(zx)^\infty$ to the circles A , B , and C are not invariant, nor are the points where the arcs meet the circles unless x and a are both adjacent to y in the graph Π , and the same for B and C . Assuming that ∂X_Γ is planar, y sends Figure 5.2 to a homeomorphic copy where the circles A , B and C have been flipped and the boundary points $(xz)^\infty$ and $(zx)^\infty$ are fixed. Each of the arcs connecting the points to the circles go to arcs connecting the same points to the same circles. Consider the components of $S^2 \setminus \text{Figure 5.2}$. The component containing the point at infinity must go to the bounded component between A and B , since the circle A is flipped. It must also go to the component between B and C , since the circle C is flipped. This is a contradiction so the action does not extend to the plane. This proves that the boundary ∂X_Γ is not planar. \square

REFERENCES

- [1] Mladen Bestvina. Local homology properties of boundaries of groups. *Michigan Math. J.*, 43(1):123–139, 1996.
- [2] Mladen Bestvina, Michael Kapovich, and Bruce Kleiner. Van Kampen’s embedding obstruction for discrete groups. *Invent. Math.*, 150(2):219–235, 2002.
- [3] Mladen Bestvina, Bruce Kleiner, and Michah Sageev. The asymptotic geometry of right-angled Artin groups. I. *Geom. Topol.*, 12(3):1653–1699, 2008.
- [4] Mladen Bestvina and Geoffrey Mess. The boundary of negatively curved groups. *J. Amer. Math. Soc.*, 4(3):469–481, 1991.
- [5] B. H. Bowditch. Connectedness properties of limit sets. *Trans. Amer. Math. Soc.*, 351(9):3673–3686, 1999.
- [6] Brian H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.
- [7] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [8] James W. Cannon. The theory of negatively curved spaces and groups. In *Ergodic theory, symbolic dynamics, and hyperbolic spaces (Trieste, 1989)*, Oxford Sci. Publ., pages 315–369. Oxford Univ. Press, New York, 1991.
- [9] Pierre-Emmanuel Caprace. Buildings with isolated subspaces and relatively hyperbolic Coxeter groups. *Innov. Incidence Geom.*, 10:15–31, 2009.
- [10] Pierre-Emmanuel Caprace. Erratum to “Buildings with isolated subspaces and relatively hyperbolic Coxeter groups” [MR2665193]. *Innov. Incidence Geom.*, 14:77–79, 2015.
- [11] Schieffelin Clayton. Topological immersion of Peanian continua in a spherical surface. *Ann. of Math. (2)*, 35(4):809–835, 1934.
- [12] Christopher B. Croke and Bruce Kleiner. Spaces with nonpositive curvature and their ideal boundaries. *Topology*, 39(3):549–556, 2000.
- [13] François Dahmani, Vincent Guirardel, and Piotr Przytycki. Random groups do not split. *Math. Ann.*, 349(3):657–673, 2011.

- [14] Michael W. Davis. The geometry and topology of Coxeter groups. In *Introduction to modern mathematics*, volume 33 of *Adv. Lect. Math. (ALM)*, pages 129–142. Int. Press, Somerville, MA, 2015.
- [15] Michael W. Davis and Boris Okun. Vanishing theorems and conjectures for the ℓ^2 -homology of right-angled Coxeter groups. *Geom. Topol.*, 5:7–74, 2001.
- [16] Ross Geoghegan and Pedro Ontaneda. Boundaries of cocompact proper $\text{CAT}(0)$ spaces. *Topology*, 46(2):129–137, 2007.
- [17] Peter Haïssinsky. Hyperbolic groups with planar boundaries. *Invent. Math.*, 201(1):239–307, 2015.
- [18] Matthew Haulmark. Boundary classification and two-ended splittings of groups with isolated flats. *J. Topol.*, 11(3):645–665, 2018.
- [19] Matthew Haulmark, G. Christopher Hruska, and Bakul Sathatye. Non-hyperbolic groups with menger curve boundary. preprint.
- [20] G. Christopher Hruska and Bruce Kleiner. Hadamard spaces with isolated flats. *Geom. Topol.*, 9:1501–1538, 2005. With an appendix by the authors and Mohamad Hindawi.
- [21] G. Christopher Hruska and Kim Ruane. Connectedness properties and splittings of groups with isolated flats. arXiv:1705.00784.
- [22] Michael Kapovich and Bruce Kleiner. Hyperbolic groups with low-dimensional boundary. *Ann. Sci. École Norm. Sup. (4)*, 33(5):647–669, 2000.
- [23] Michael Mihalik and Kim Ruane. $\text{CAT}(0)$ groups with non-locally connected boundary. *J. London Math. Soc. (2)*, 60(3):757–770, 1999.
- [24] Michael Mihalik and Steven Tschantz. Geodesically tracking quasi-geodesic paths for Coxeter groups. *Bull. Lond. Math. Soc.*, 45(4):700–714, 2013.
- [25] Panos Papasoglu. Quasi-isometry invariance of group splittings. *Ann. of Math. (2)*, 161(2):759–830, 2005.
- [26] Panos Papasoglu and Eric Swenson. Boundaries and JSJ decompositions of $\text{CAT}(0)$ -groups. *Geom. Funct. Anal.*, 19(2):559–590, 2009.
- [27] Kevin Schreve and Emily Stark. Non-planar graphs in boundaries of $\text{CAT}(0)$ groups. arXiv:1707.07760.
- [28] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.
- [29] G. A. Swarup. On the cut point conjecture. *Electron. Res. Announc. Amer. Math. Soc.*, 2(2):98–100, 1996.
- [30] Jacek Świątkowski. Hyperbolic Coxeter groups with Sierpiński carpet boundary. *Bull. Lond. Math. Soc.*, 48(4):708–716, 2016.
- [31] G. T. Whyburn. Topological characterization of the Sierpiński curve. *Fund. Math.*, 45:320–324, 1958.
- [32] Stephen Willard. *General topology*. Dover Publications, Inc., Mineola, NY, 2004. Reprint of the 1970 original [Addison-Wesley, Reading, MA; MR0264581].