Separation and relative quasiconvexity criteria for relatively geometric actions

Eduard Einstein, Daniel Groves, and Thomas Ng

Abstract. Bowditch characterized relative hyperbolicity in terms of group actions on fine hyperbolic graphs with finitely many edge orbits and finite edge stabilizers. In this paper, we define generalized fine actions on hyperbolic graphs, in which the peripheral subgroups are allowed to stabilize finite subgraphs rather than stabilizing a point. Generalized fine actions are useful for studying groups that act relatively geometrically on a CAT(0) cube complex, which were recently defined by the first two authors. Specifically, we show that any group acting relatively geometrically on a CAT(0) cube complex admits a generalized fine action on the one-skeleton of the cube complex. For generalized fine actions, we prove a criterion for relative quasiconvexity of subgroups that cocompactly stabilize quasiconvex subgroups, generalizing a result of Martínez-Pedroza and Wise in the setting of fine hyperbolic graphs. As an application, we obtain a characterization of boundary separation in generalized fine graphs and use it to prove that Bowditch boundary points in relatively geometric actions are always separated by a hyperplane stabilizer.

1. Introduction

There are many equivalent formulations of relatively hyperbolic groups, see, for example, [4, 7, 11, 13, 14, 16, 18]. Bowditch describes relative hyperbolicity in terms of an action on a fine hyperbolic graph with certain finiteness conditions (see Definition 2.1 for the definition of 'fine') [4]. A natural example of such a graph is the coned-off Cayley graph of a relatively hyperbolic pair, defined by Farb [11]. Unfortunately, fineness and other important finiteness properties of the action of $(G, \mathcal{P})$ on the coned-off Cayley graph are not quasi-isometry invariants (see, for instance, Example 3.5). The first and second authors introduced the notion of a relatively hyperbolic pair $(G, \mathcal{P})$ acting *relatively geometrically* on a CAT(0) cube complex $\widetilde{X}$ [8]. In this situation, a result of Charney and Crisp [5, Theorem 5.1] implies that both $\widetilde{X}$ and its 1-skeleton $\widetilde{X}^{(1)}$ are quasi-isometric to the coned-off Cayley graph of $(G, \mathcal{P})$. However, the edge stabilizers of $\widetilde{X}^{(1)}$ are often infinite.

In this paper, we develop tools to translate geometric features of a generalized fine action on a hyperbolic graph to the Bowditch boundary of $(G, \mathcal{P})$, a compact boundary for a relatively hyperbolic pair $(G, \mathcal{P})$ first introduced in [4]. We also apply these tools

to prove some fundamental results about relatively hyperbolic groups that act relatively geometrically on CAT(0) cube complexes. One construction of this boundary is from a fine hyperbolic graph Γ that witnesses the relative hyperbolicity of $(G, \mathcal{P})$. As a set, the Bowditch boundary is the disjoint union of the visual boundary of Γ with the vertices of Γ that have infinite stabilizer. Bowditch endows this set with the topology we describe in Definition 4.6. In our study of relatively geometric actions, we would like to take advantage of the correspondence between the CAT(0) cube complex $\widetilde{X}$ and the coned-off Cayley graph for $(G, \mathcal{P})$ to use the geometry of $\widetilde{X}$ to prove statements about the Bowditch boundary of $(G, \mathcal{P})$.

Let $(G, \mathcal{P})$ be a relatively hyperbolic group pair, where G acts combinatorially by isometries on a CAT(0) cube complex $\widetilde{X}$. The action of G is *relatively geometric* if

(1) the quotient of $G \backslash \widetilde{X}$ is compact,

(2) every infinite cell stabilizer is a finite index subgroup of P^g for some $P \in \mathcal{P}$ and $g \in G$,

(3) every $P \in \mathcal{P}$ acts elliptically.

Relatively geometric actions were defined recently by the first and second authors in [8]. There is natural motivation for identifying and studying relatively hyperbolic groups that act relatively geometrically: many of the desirable properties of hyperbolic groups that act properly and cocompactly on CAT(0) cube complexes have relatively geometric analogs when the peripheral subgroups are residually finite. For example, there is a relatively geometric version [9, Theorem 1.4] of Agol's theorem for hyperbolic cubulated groups [2, Theorem 1.1], full relatively quasiconvex subgroups are separable [9, Corollary 1.7], and there is a relatively geometric characterization of Kleinian groups for relatively hyperbolic groups with 2-sphere boundary which may provide a method for attacking the (relative) Cannon conjecture [8].

If $(G, \mathcal{P})$ acts relatively geometrically on $\widetilde{X}$, then $\widetilde{X}$ (and also its 1-skeleton $\widetilde{X}^{(1)}$) is quasi-isometric to the coned-off Cayley graph of $(G, \mathcal{P})$ by [5, Theorem 5.1]. However, in general the graph $\widetilde{X}^{(1)}$ is not a fine graph. Thus, to help study actions similar to relatively geometric actions, we introduce generalized fine actions on a hyperbolic graph.

Definition 1.1. Let G be a group that acts by isometries on a δ-hyperbolic graph Γ and let $\mathcal{P}$ be a finite and almost malnormal collection of finitely generated subgroups of G. For each $P \in \mathcal{P}$ and $g \in G$, let Γ_{Pg} be the subgraph of Γ whose cells have stabilizer commensurable to P^g. A *circuit without peripheral backtracking* is an embedded loop γ so that for all $P \in \mathcal{P}$ and $g \in G$, $\gamma \cap \Gamma_{Pg}$ is connected. We say that Γ is *generalized fine with respect to the action of* $(G, \mathcal{P})$ if

(1) the quotient $G \backslash \Gamma$ is compact,

(2) every cell with infinite stabilizer lies in Γ_{Pg} for some $P \in \mathcal{P}$ and $g \in G$,

(3) each subgraph Γ_{Pg} is compact and connected, and

(4) for every $n \in \mathbb{N}$, every edge with finite stabilizer lies in finitely many circuits without peripheral backtracking of length n.

It follows from the characterization of fineness [4, Proposition 2.1] that a cocompact action on a fine graph with finite edge stabilizers is also generalized fine with respect to a set of conjugacy representatives of the vertex stabilizers. Generalized fine actions witness relative hyperbolicity. We prove in Proposition 3.4 that $(G, \mathcal{P})$ in Definition 1.1 is a relatively hyperbolic pair when $P \in \mathcal{P}$ fixes Γ_{P^g}. Relatively geometric actions give rise to generalized fine actions: in Example 3.3 below we show that $\widetilde{X}^{(1)}$ is generalized fine with respect to the induced action of $(G, \mathcal{P})$.

We say that $H \leq G$ has a *(quasi)convex cocompact core* if H stabilizes a (quasi)convex H-cocompact connected subgraph (see Definition 4.1 for a precise definition). Our first main result shows that the existence of a quasiconvex cocompact core that interacts nicely with the subgraphs stabilized by peripheral conjugates implies relative quasiconvexity.

Theorem 1.2. *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let Γ be a hyperbolic graph with a G-action so that Γ is generalized fine with respect to the action of $(G, \mathcal{P})$. For any $P \in \mathcal{P}$ and $g \in G$, let Γ_{P^g} be the subgraph of Γ whose cell stabilizer are commensurable to P^g. If $H \leq G$ has a quasiconvex cocompact core Γ_H and one of the following holds:*

- *Γ is fine or*
- *for all $P \in \mathcal{P}$ and $g \in G$, $\Gamma_H \cap \Gamma_{P^g} \neq \emptyset$ implies $|P^g \cap H| = \infty$,*

then H is relatively quasiconvex in $(G, \mathcal{P})$.

As with the definition of relative hyperbolicity, there are many equivalent characterizations of relative quasiconvexity, see [16]. As a consequence of Theorem 1.2, we provide an alternate proof of a theorem of Martínez-Pedroza and Wise [17] characterizing relative quasiconvexity in terms of quasiconvex cores in fine hyperbolic graphs.

Corollary 1.3 ([17, Theorem 1.7]). *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair acting cocompactly on a fine hyperbolic graph so that every edge stabilizer is finite. A subgroup $H \leq G$ is relatively quasiconvex in $(G, \mathcal{P})$ if and only if H has a quasiconvex cocompact core in Γ.*

As an application, we see that if $(G, \mathcal{P})$ acts relatively geometrically on a CAT(0) cube complex $\widetilde{X}$, then hyperplane stabilizers are relatively quasiconvex in $(G, \mathcal{P})$. The first and third authors also use Theorem 1.2 to construct new families of residually finite groups using relatively geometric actions of $C'(\frac{1}{6})$-small cancellation free products with relatively geometrically cubulated factor groups [10].

Following the criterion in [3] for proper and cocompact cubulations of hyperbolic groups, the first two authors show in [8, Theorem 2.6] that if any two points in the Bowditch boundary $\partial_{\mathcal{D}} K$ of a relatively hyperbolic pair $(K, \mathcal{D})$ lie in H-distinct components of the limit set of a full relatively quasiconvex $H \leq K$, then K acts relatively geometrically on a CAT(0) cube complex. Our other main result shows that any pair of distinct points in the Bowditch boundary of a relatively geometrically cubulated group can be separated by the limit set of a hyperplane stabilizer.

Theorem 1.4. *Let $(K, \mathcal{D})$ act relatively geometrically on a* CAT(0) *cube complex $\widetilde{X}$. If $x, y \in \partial_{\mathcal{D}} K$ and $x \neq y$, then there exist a hyperplane W of $\widetilde{X}$ and a finite index subgroup $K_W \leq \mathrm{Stab}_K(W)$ so that x, y are in K_W-distinct components of $\partial_{\mathcal{D}} K \setminus \Lambda K_W$.*

More generally, if a graph Σ is generalized fine with respect to the action of a relatively hyperbolic pair $(K, \mathcal{D})$ as in Definition 1.1, we obtain a technical criterion for deciding when a subgroup of K with a cocompact core in Σ separates two points in the Bowditch boundary, see Theorem 5.7. Both Theorems 1.4 and 5.7 are essential tools in the first and third authors' work [10] on relative cubulations for small cancellation free products.

The idea behind Theorem 1.4 is to take a hyperplane W that is dual to an edge of a combinatorial geodesic between x and y. However, the Bowditch boundary is not the visual boundary of $\widetilde{X}$. Moreover, $\widetilde{X}^{(1)}$ is not proper, so it is not clear that such a combinatorial geodesic exists. In order to make statements about $\partial_{\mathcal{D}} K$, we need to *coarsely* translate the geometric features of W and $\widetilde{X}$ to a fine hyperbolic graph Σ' with a K-action that witnesses the relative hyperbolicity of $(K, \mathcal{D})$. While the image of W in Σ' will separate Σ' into two components, we need to ensure that the limit set of $\mathrm{Stab}_G(W)$ actually separates x and y into distinct complementary components of $\partial_{\mathcal{D}} K$, which is still not the visual boundary Σ'.

1.1. Outline

We introduce some background on relatively hyperbolic groups and fine hyperbolic graphs in Section 2. We also discuss a construction similar to that in [11] to relate paths in two graphs Σ, Σ', where Σ' is formed by collapsing some of the edges of Σ. Then, we recall some specific properties of relatively geometric actions in Section 2.3. In Section 3, we explore the properties of generalized fine actions and show that relatively geometric actions on a cube complex give rise to generalized fine actions on the one-skeleton of the cube complex. We then prove Theorem 1.2 and Corollary 1.3 in Section 4.

The main result of Section 5 is Theorem 5.7, a technical separation criterion for points in the Bowditch boundary of a relatively hyperbolic group in terms of a generalized fine action. Finally, in Section 6, we prove Theorem 1.4 using Theorem 5.7.

2. Obtaining fine hyperbolic graphs from relatively geometric actions

2.1. Fine hyperbolic graphs and relative hyperbolicity

First, we recall the definition of a fine graph.

Definition 2.1. Let Γ be a graph. A *circuit* is an embedded loop in Γ. The graph Γ is *fine* if for each edge e of Γ and every $n \in \mathbb{N}$, there exist finitely many circuits of length n containing e.

Here is Bowditch's definition of relative hyperbolicity in terms of fine hyperbolic graphs.

Definition 2.2 ([4, Definition 2], written as stated in [16, Definition 3.4]). Suppose G acts on a δ-hyperbolic graph Γ with finite edge stabilizers and finitely many G-orbits of edges. If Γ is fine, and $\mathcal{P}$ is a set of representatives of the conjugacy classes of infinite vertex stabilizers, then $(G, \mathcal{P})$ is a *relatively hyperbolic pair*.

In Section 4, we will also use a dynamical characterization of relative hyperbolicity and relative quasiconvexity due to Yaman [22].

As noted by Bowditch [4, p. 3], fineness is not a quasi-isometry invariant. Here is an example of quasi-isometric graphs, where one graph is fine and the other is not.

Example 2.3. Let Γ be a graph with 2 vertices and a single edge joining the two vertices. Let Σ be a graph with 2 vertices and an infinite number of edges between the two vertices. Then Γ and Σ are quasi-isometric. Any edge of Σ lies in infinitely many circuits of length 2, so Σ is not fine.

2.2. Electrification and de-electrification

Farb first introduced the notion of electrifying a space in [11], where electrification of the fundamental group of a finite volume hyperbolic 3-manifold is accomplished by collapsing the cosets of cusp subgroups to points. This idea inspires the definitions in this subsection. Similar constructions have also been performed in [1, 4, 6, 20].

Let G be a group and suppose G acts by isometries on a graph Σ. Let $\mathcal{B}$ be a collection of pairwise disjoint and connected subgraphs of Σ so that

(1) if $g \in G$ and $B \in \mathcal{B}$, then $gB \in \mathcal{B}$,

(2) every cell of $B \in \mathcal{B}$ has infinite stabilizer, and

(3) every cell with infinite stabilizer lies in some $B \in \mathcal{B}$.

Definition 2.4. The *complete electrification of Σ (with respect to $\mathcal{B}$)* is the graph Σ' that is constructed by contracting each $B \in \mathcal{B}$ to a vertex v_B of Σ'. There is a canonical quotient map $\sigma \colon \Sigma \to \Sigma'$, which we call the *electrification* map.

The *stable part of Σ under the electrification with respect to $\mathcal{B}$* is denoted by

$$\Sigma_0 = \Sigma \setminus \left(\bigcup_{B \in \mathcal{B}} B \right).$$

The stable part Σ_0 embeds naturally into both Σ and Σ'. When $\mathcal{B}$ is clear from context, we refer to Σ_0 as the *stable part of Σ*.

When the action on G induces an action on $\mathcal{B}$ (i.e., for all $g \in G$ and all $B \in \mathcal{B}$, $gB \in \mathcal{B}$), then there is a natural *induced action on the complete electrification Σ'*, where G acts on the stable part as usual, and for each $B \in \mathcal{B}$, the action of $g \in G$ takes v_B to v_{gB}. The electrification map $\sigma \colon \Sigma \to \Sigma'$ is then G-equivariant. Note that $G\Sigma_0 \subseteq \Sigma_0$ because Σ_0 consists precisely of the cells with finite stabilizers.

We are interested in how paths behave under electrification.

Definition 2.5. Let γ be a path in Σ. The path γ is *without peripheral backtracking* if for every $B \in \mathcal{B}$, $\gamma \setminus (\gamma \cap B)$ is connected when γ is a loop and has at most 2 components otherwise.

Similarly, a path γ' in Σ' is *without peripheral backtracking* if for all $B \in \mathcal{B}$, $\gamma' \setminus v_B$ is connected when γ is a loop and has at most 2 components otherwise.

We will see that peripheral backtracking is an especially important property for circuits because it characterizes circuits whose images under the complete electrification are still embedded. We can relate paths without peripheral backtracking in Σ and Σ' as follows.

Definition 2.6. Let γ be a path in Σ without peripheral backtracking. The *electrification* γ' of γ is the path in Σ' constructed by collapsing sub-segments of the form $\gamma \cap B$ to v_B.

Similarly, let ρ' be a path in Σ' without peripheral backtracking. Let $\overline{\rho}'$ be the closure of $\rho' \cap \Sigma_0$ in Σ. If $\overline{\rho}' \cap B$ fails to be connected, $\overline{\rho}' \cap B$ is exactly two points because ρ' is without peripheral backtracking. A *complete de-electrification* of ρ' is a path ρ in Σ constructed by joining any disconnected $\overline{\rho}' \cap B$ by an embedded path in B.

Proposition 2.7. *Let G act cocompactly and cellularly by isometries on a graph Σ. Let $\mathcal{B}$ be a collection of pairwise disjoint compact connected subgraphs so that for each $B \in \mathcal{B}$, $gB \in \mathcal{B}$. If Σ' is the complete electrification of Σ with respect to $\mathcal{B}$, then there exist a cocompact action of G on Σ' and a G-equivariant quasi-isometry $f : \Sigma' \to \Sigma$, whose restriction $f|_{\Sigma_0}$ to the stable part of $\Sigma_0 \subseteq \Sigma'$ is the identity.*

Furthermore, the electrification map $\sigma : \Sigma \to \Sigma'$ is a quasi-isometry.

Proof. Recall that G has a natural induced action on the complete electrification Σ' of Σ.

The quotient ${}_G \backslash^{\Sigma'}$ differs from the quotient ${}_G \backslash^{\Sigma}$ by collapsing finitely many edges of a finite graph, so ${}_G \backslash^{\Sigma'}$ is still compact.

For each $B \in \mathcal{B}$, fix a vertex $x_B \in B$. We define $f : \Sigma' \to \Sigma$ as follows:

$$f(x) = \begin{cases} x, & x \in \Sigma_0, \\ x_B, & x = v_B. \end{cases}$$

There are finitely many G-orbits of B and each B is compact, so there exists $s \geq 0$ that uniformly bounds the diameters of all $B \in \mathcal{B}$. It follows immediately that f is coarsely surjective.

Let x, y be vertices of Σ'. Choose a geodesic path γ' connecting them in Σ'. Let D be the length of γ' in Σ'. There exists a de-electrification γ of γ' so that γ' has length at most $D + Ds$ since γ' encounters at most D vertices of the form v_{Pg} in its interior. Extend γ' to a path γ'' joining $f(x)$ to $f(y)$ by adding segments of length at most s to each end. The distance in Σ between $f(x)$ and $f(y)$ is at most $D + Ds + 2s$.

On the other hand, any geodesic path between $f(x)$ and $f(y)$ in Σ can be electrified to make a shorter path in Σ' between x and y. Therefore,

$$d(x, y) \leq d(f(x), f(y)) \leq (1 + s)d(x, y) + 2s.$$

Hence f is a quasi-isometry. ∎

2.3. Obtaining an action on a fine hyperbolic graph from a relatively geometric action

For the following subsection, assume that the relatively hyperbolic pair $(G, \mathcal{P})$ acts relatively geometrically on a CAT(0) cube complex $\tilde{X}$.

The main goal of this section is to show that subgroups with convex cores in $\tilde{X}$ have quasiconvex cocompact cores in a fine hyperbolic graph that witnesses the relative hyperbolicity of $(G, \mathcal{P})$.

Proposition 2.8. *Let $H \leq G$ be a subgroup that stabilizes a convex subcomplex $\tilde{Y} \subseteq \tilde{X}$. Then $(G, \mathcal{P})$ acts on a fine hyperbolic graph Γ so that*

(1) *every $P \in \mathcal{P}$ fixes a vertex of Γ,*

(2) *every edge stabilizer is finite,*

(3) *$G \backslash \Gamma$ is compact, and*

(4) *there exists a quasiconvex H-invariant connected subgraph $\Gamma_H \subseteq \Gamma$.*

Additionally, if $\tilde{Y}$ is H-cocompact, then Γ_H is H-cocompact.

For each $P \in \mathcal{P}$ and $g \in G$, let $\tilde{X}_{P^g}$ denote the subgraph of $\tilde{X}^{(1)}$ consisting of cells whose stabilizer is commensurable to P^g. Since infinite stabilizers in $\tilde{X}$ are commensurable with a unique P^g, the following is immediate from the definitions.

Lemma 2.9. *If $P_1^{g_1} \neq P_2^{g_2}$, then $\tilde{X}_{P_1^{g_1}} \cap \tilde{X}_{P_2^{g_2}} = \emptyset$.*

Before proving Proposition 2.8, we investigate some of the properties of $\tilde{X}^{(1)}$ and its complete electrification with respect to the collection of subgraphs of the form $\tilde{X}_{P^g}$.

By [9, Proposition 3.5], the subgraph $\tilde{X}_{P^g}$ is the 1-skeleton of a compact and convex subcomplex of a CAT(0) cube complex. Then we obtain the following fact about $\tilde{X}_{P^g}$.

Proposition 2.10. *The subgraph $\tilde{X}_{P^g}$ is connected, convex and compact.*

For $\tilde{X}^{(1)}$, Proposition 2.7 implies the following.

Proposition 2.11. *Let Γ be the complete electrification of $\tilde{X}^{(1)}$ with respect to*

$$\mathcal{B} = \{\tilde{X}_{P^g} : P \in \mathcal{P}, g \in G\}$$

and let Γ_0 be the stable part of $\tilde{X}^{(1)}$. There exist a cocompact action of G on Γ and a G-equivariant quasi-isometry $f : \Gamma \to \tilde{X}^{(1)}$ so that $f|_{\Gamma_0}$ is the identity map.

Furthermore, the map $c\colon \widetilde{X}^{(1)} \to \Gamma$*, that collapses* X_{Pg} *to* v_{Pg} *and fixes* $f(\Gamma_0)$*, is a* G*-equivariant quasi-isometry that fixes the image of the stable part of* Γ_0 *in* $\widetilde{X}^{(1)}$*.*

Proposition 2.12. *The complete electrification* Γ *in Proposition* 2.11 *is a fine hyperbolic graph.*

Proof. By [5, Theorem 5.1], $\widetilde{X}^{(1)}$ is quasi-isometric to the coned-off Cayley graph for $(G, \mathcal{P})$, so $\widetilde{X}^{(1)}$ is a hyperbolic graph. By Proposition 2.11, Γ is hyperbolic.

Infinite vertex stabilizers in Γ are maximal parabolic, and each maximal parabolic stabilizes exactly one point in Γ, so Γ has finite pair stabilizers. By cocompactness, $G \backslash \Gamma$ is finite, so by [4, Lemma 4.5], Γ is a fine graph. ∎

We are ready to prove Proposition 2.8.

Proof of Proposition 2.8. By Proposition 2.11, the action of $(G, \mathcal{P})$ provides a cocompact action of G on Γ where each P^g fixes v_{Pg}. Since H stabilizes a convex subcomplex $\widetilde{Y} \subseteq \widetilde{X}$, then $\widetilde{X}^{(1)} \cap \widetilde{Y} = \widetilde{Y}^{(1)}$ is a convex subgraph of $\widetilde{X}^{(1)}$. When $\widetilde{Y}$ is H-cocompact, $\widetilde{Y}^{(1)}$ is also H-cocompact. The collapse $c\colon \widetilde{X}^{(1)} \to \Gamma$ takes $\widetilde{Y}^{(1)}$ to a subgraph of Γ_H of Γ. Since c is a quasi-isometry by Proposition 2.11, Γ_H is quasiconvex in Γ. ∎

3. Generalized fine actions

The behavior of the action of G on $\widetilde{X}^{(1)}$ in Section 2.3 is more general than relatively geometric actions on CAT(0) cube complexes and is captured by the definition of generalized fine graphs (Definition 1.1).

Hypotheses 3.1. *Let* $(K, \mathcal{D})$ *be a relatively hyperbolic pair and let* K *act on a graph* Σ*. For each* $D \in \mathcal{D}$ *and* $k \in K$*, let* Σ_{D^k} *be the subgraph of* Σ *consisting of cells whose stabilizer is commensurable to* D^k*.*

The main goal of this section is to characterize generalized fine actions as follows:

(1) If electrifying Σ with respect to the Σ_{D^k} results in a fine hyperbolic graph with an appropriate action, the action of K on Σ is generalized fine (Proposition 3.2).

(2) If the action of $(K, \mathcal{D})$ on Σ is generalized fine, then electrifying the Σ_{D^k} results in a fine hyperbolic graph (Proposition 3.4).

Proposition 3.2. *Assume Hypotheses* 3.1. *Suppose that*

(1) K *acts cocompactly on* Σ,

(2) *every cell with infinite stabilizer lies in some* Σ_{D^k},

(3) *every maximal parabolic* $D \in \mathcal{D}$ *stabilizes a vertex of* Σ, *and*

(4) *the subgraph* Σ_{D^k} *is connected and compact.*

If the complete electrification with respect to $\{\Sigma_{D^k} : D \in \mathcal{D}, k \in K\}$ *is a fine hyperbolic graph, then* Σ *is generalized fine with respect to the action of* $(K, \mathcal{D})$*.*

Proof. The first three requirements for generalized fineness follow from hypothesis.

Fix $n \geq 0$ and let e be an edge of Σ with finite stabilizer. Let γ be a length n-circuit in Σ without peripheral backtracking. Let Γ be the complete electrification of Σ. Since each of the Σ_{D^k} is compact and there are finitely many K-orbits of Σ_{D^k}, there is a constant $L(n)$ that bounds the number of length at most n embedded paths in Σ_{D^k} from above.

The map $\sigma \colon \Sigma \to \Gamma$ that collapses the subgraphs Σ_{D^k} to points v_{D^k} is injective on e because e has finite stabilizer. The electrification γ' of γ is an embedded circuit of length at most n in Γ containing $\sigma(e)$. Since Γ is a fine graph, there exists some $T \geq 0$ so that γ' is one of T circuits passing through $\sigma(e)$. Since γ intersects at most n of the v_{P^g}, γ can be obtained as a de-electrification of γ' where each of the v_{P^g} that γ intersects is replaced by an embedded path in the corresponding Σ_{D^k}. Hence there are at most $T(L(n))^n$ possibilities for γ. Thus every edge with finite stabilizer in Σ is contained in only finitely many circuits. ∎

Propositions 2.12 and 3.2 imply that a relatively geometric action gives rise to a generalized fine action.

Example 3.3. If $(G, \mathcal{P})$ acts relatively geometrically on a CAT(0) cube complex $\widetilde{X}$, then $\widetilde{X}^{(1)}$ is generalized fine with respect to the action of $(G, \mathcal{P})$.

Conversely to Proposition 3.2, generalized fine actions can always be used to electrify the underlying space into a fine hyperbolic graph witnessing relative hyperbolicity.

Proposition 3.4. *If Σ is generalized fine with respect to the action of $(K, \mathcal{D})$, then the electrification Σ' with respect to $\mathcal{B} = \{\Sigma_{D^k} : D \in \mathcal{D}, k \in K\}$ is a fine (hyperbolic) graph. Therefore, $(K, \mathcal{D}_0)$ is a relatively hyperbolic pair, where $\mathcal{D}_0 = \{\mathrm{Stab}(\Sigma_D) : D \in \mathcal{D}\}$.*

Proof. Let $\sigma \colon \Sigma \to \Sigma'$ be the complete electrification as in Proposition 2.7. Let e' be an edge in Σ' with finite stabilizer and let γ' be a circuit of length n in Σ'. The circuit γ is without peripheral backtracking because it is embedded and peripheral subgroups are only commensurable to stabilizers of vertices in Σ'. Let Σ_0 be the stable part of Σ. Let e be the unique edge of Σ whose interior is $e' = \sigma(e) \subseteq \Sigma_0$. The electrification of any de-electrification of γ' returns γ'.

Therefore, there exists a circuit γ without peripheral backtracking containing e such that γ' is a complete electrification of γ. There are finitely many K-orbits of Σ_{D^k}, so there is a uniform bound $L \geq 0$ on the diameters of the $\Sigma_{D^k} \in \mathcal{B}$. Hence the length of γ is at most $n + nL$. Since Σ is generalized fine, and γ contains e, an edge with finite stabilizer, there are only finitely many possibilities for γ. Thus there are only finitely many possibilities for γ'.

There is a natural action of K on Σ' defined as follows: let $x \in \Sigma'$ and $k \in K$. When $x \in \Sigma_0$, then $k \cdot x$ is defined according to the action of K on $\Sigma_0 \subseteq \Sigma$. Otherwise, $x = v_{D^{k_0}}$ for some $D \in \mathcal{D}$ and some $k_0 \in K$, so define $k \cdot x = v_{D^{kk_0}}$. There are finitely many K-orbits of edges because K acts cocompactly on Σ', and every edge has a finite stabilizer because electrification collapses every edge with infinite stabilizer. By Definition 2.2, the

Figure 1. The complex of $\mathcal{G}(P, Q)$ is the complex of groups with indicated local groups and a portion of the development around the edge stabilized by P. The blue circuits pass through e and v_i for $i \geq 1$, but all have peripheral backtracking because they are separated by the subcomplex $\Sigma_P = e_P$ with red endpoints.

hyperbolicity and fineness of Σ' implies that K is hyperbolic relative to any finite set of conjugacy class representatives of the infinite vertex stabilizers. The set $\mathcal{D}_0$ is a collection of conjugacy representatives of infinite vertex stabilizers for the action of K on Σ. ∎

Note that if each D fixes a point of Σ_D, then $\mathcal{D} = \mathcal{D}_0$.

We now exhibit an action that is generalized fine where the underlying graph is not fine.

Example 3.5. Let P, Q be groups and consider the complex of groups $\mathcal{G}(P, Q)$ shown in Figure 1. Let T be the Bass–Serre tree for the free product $P * Q$. The complex $\mathcal{G}(P, Q)$ is the quotient of the natural action of $P * Q$ on $T \times [0, 1]$, and this action is relatively geometric. However, the edge e_P stabilized by P lies in infinitely many circuits of length 4. Note also that even though the edge e joining vertices with stabilizers P and Q has trivial stabilizer, it is contained in infinitely many circuits. However, only finitely many of these circuits will not have peripheral backtracking.

4. Quasi-convex subgraphs of generalized fine graphs

In this section, we prove Theorem 1.2. We first formally state the definition of a quasiconvex cocompact core.

Definition 4.1. Let G act on a hyperbolic graph Γ by isometries and let $H \leq G$ be a subgroup of G. A *(quasi)convex core of H in Γ* is a connected subgraph Γ_H so that

(1) the quotient $_H\backslash^{\Gamma_H}$ is compact, and

(2) Γ_H is (quasi)convex in Γ.

Set the following hypotheses.

Hypotheses 4.2. *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and suppose that G acts on a connected hyperbolic graph Γ so that Γ is generalized fine with respect to the action of $(G, \mathcal{P})$. For $P \in \mathcal{P}$ and $g \in G$, let Γ_{P^g} be the subgraph stabilized by P^g. Let $H \leq G$ and let Γ_H be a quasiconvex cocompact core for H in Γ. If Γ is not fine (only generalized fine), we make the following additional assumption:*

$$\text{If } \Gamma_H \cap \Gamma_{P^g} \neq \emptyset, \text{ then } |H \cap P^g| = \infty.$$

Here is a rough outline of the proof of Theorem 1.2: we prove that the action of H on Γ_H implies that H is hyperbolic relative to a finite collection of vertex stabilizers $\mathcal{D}$. Then H admits a geometrically finite convergence group action on the Bowditch boundary of $(H, \mathcal{D})$. We then show that the inclusion $\Gamma_H \to \Gamma$ induces an equivariant inclusion on Bowditch boundaries whose image is the limit set of H so that the induced action of H on its limit set in the Bowditch boundary of $(G, \mathcal{P})$ is a geometrically finite convergence group action. We now recall Yaman's dynamical characterization [22] of relative hyperbolicity.

Definition 4.3 (As stated in [16, Definition 3.1]). Suppose $(G, \mathcal{P})$ has a geometrically finite convergence group action on a compact metrizable space M. Then $(G, \mathcal{P})$ is a *relatively hyperbolic pair*.

Yaman also proves (see, for example, [16, Theorem 5.2]) that the space M from Definition 4.3 is equivariantly homeomorphic to the Bowditch boundary of the pair $(G, \mathcal{P})$.

For details about geometrically finite actions, see [16, Section 3.1]. If $H \leq G$, recall that the *limit set of H in M*, denoted by ΛH, is the smallest non-empty closed H-invariant subset of M. We use Definition 4.3 in conjunction with the following definition for relative quasiconvexity.

Definition 4.4 ([16, Definition 6.2]). Let $(G, \mathcal{P})$ be a relatively hyperbolic group that acts on a compact metrizable space as a geometrically finite convergence group. A subgroup $H \leq G$ is *relatively quasiconvex* if the induced convergence action of H on the limit set $\Lambda H \subseteq M$ is geometrically finite.

Proposition 4.5. *Assume Hypotheses 4.2. Let $\mathcal{D}$ be a (finite) collection of H-conjugacy representatives (up to commensurability) of infinite vertex stabilizers for the action of H on Γ_H. Then $(H, \mathcal{D})$ is a relatively hyperbolic pair.*

Proof. If Γ is fine, then the core graph Γ_H is fine, so by Definition 2.2, $(H, \mathcal{D})$ is a relatively hyperbolic pair.

By assumption, the action of H on Γ_H is cocompact. Since $\Gamma_H \cap \Gamma_{P^g} \neq \emptyset$ implies that $H \cap P^g$ is infinite, each infinite cell stabilizer for the H-action on Γ_H is commensurable to $H \cap P^g$ for some $P \in \mathcal{P}$ and $g \in G$. Hence $\mathcal{D}$ is an almost malnormal family. The technique from [9, Proposition 3.5] can be used to show that for each $D \in \mathcal{D}$, the subgraph of Γ_H consisting of cells whose stabilizers are commensurable to D is compact and convex (and therefore connected) in Γ_H. Since the edges of Γ_H with finite

stabilizer are precisely those which have finite stabilizer with respect to the G-action on Γ, the generalized fineness of the action of $(G, \mathcal{P})$ on Γ implies that every edge of Γ_H with finite H-stabilizer lies in finitely many circuits without peripheral backtracking of length n.

Therefore, Γ_H is generalized fine with respect to the H-action. By Proposition 3.4, $(H, \mathcal{D})$ is a relatively hyperbolic pair. ∎

For clarity and completeness, we repeat Bowditch's construction of the Bowditch boundary from a fine hyperbolic graph.

Definition 4.6 ([4, Section 9]). Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and suppose that Γ is a graph that is generalized fine with respect to the action of $(G, \mathcal{P})$. Let Γ' be a complete electrification of Γ with respect to the Γ_{Pg} and for all $P \in \mathcal{P}$ and $g \in G$, let v_{Pg} be the vertex of Γ' stabilized by P^g. Let $\triangle \Gamma' = \partial \Gamma' \sqcup V(\Gamma')$ be endowed with the following topology: If A is any finite subset of the vertices of Γ' and $a \in \triangle \Gamma'$, define $N(a, A)$ to be the set of $b \in \triangle \Gamma'$ so that every geodesic from a to b avoids $A \setminus a$. A subset $U \subseteq \triangle \Gamma'$ is open if for every $a \in U$, there exists a finite set of vertices $A \subseteq V(\Gamma')$ so that $N(a, A) \subseteq U$.

Define $\Pi_{\Gamma'} = \{v_{Pg} : P \in \mathcal{P}, g \in G\}$, the *peripheral points* of the Bowditch boundary.

Let $\partial_B \Gamma' = \partial \Gamma' \cup \Pi_{\Gamma'}$, where $\partial \Gamma'$ is the visual boundary of the hyperbolic graph Γ'. We refer to the points of $\partial \Gamma'$ as the *conical limit points* of the Bowditch boundary. The topology on $\partial_B \Gamma'$ is the subspace topology induced by the topology on $\triangle \Gamma'$.

Remark 4.7. In this section, we explicitly use the notation $\partial_B \Gamma'$ to denote the construction of the Bowditch boundary of $(G, \mathcal{P})$ from the graph Γ'. Outside of Section 4, we use the notation $\partial_{\mathcal{P}} G$ to refer to the Bowditch boundary of G with respect to $\mathcal{P}$.

Assuming Hypotheses 4.2, let Γ'_H be the image of Γ_H in Γ'. Following Definition 4.6, we can define $\partial_B \Gamma'_H$, and moreover $\partial_B \Gamma'_H$ embeds in $\partial_B \Gamma'$.

Proposition 4.8. *Assuming Hypotheses* 4.2, $\partial_B \Gamma'_H$ *is closed in* $\partial_B \Gamma'$.

We prove Proposition 4.8 by showing that the complement of $\partial_B \Gamma'_H$ is open in $\partial_B \Gamma'$. Specifically, if $y \in \partial_B \Gamma' \setminus \partial_B \Gamma'_H$, we find an open neighborhood of y that does not contain any points of $\partial_B \Gamma'_H$. For the topology introduced in Definition 4.6, it suffices to prove that there exists a finite set of vertices that any geodesic from y to a point in $\partial_B \Gamma'_H$ must pass through.

Proof. Since Γ_H is convex in Γ, Γ'_H is s-quasiconvex in Γ' for some $s \geq 0$. Set $\delta > 1$ so that Γ'_H has δ-thin triangles.

Let $y \in \partial_B \Gamma' \setminus \partial_B \Gamma'_H$. To prove that $\partial_B \Gamma' \setminus \partial_B \Gamma'_H$ is open, we find a finite $A_y \subseteq V_\Gamma$ so that if $h \in \partial_B \Gamma'_H$, any geodesic from y to h passes through A_y. If so, then $N(y, A_y) \subseteq \partial_B \Gamma' \setminus \partial_B \Gamma'_H$. We now fix a geodesic γ from y to some $h \in \partial_B \Gamma'_H$ and split the proof into two cases depending on whether y is a conical limit point or a peripheral point.

Figure 2. The situation in Proposition 4.8 in case 1 where $h \notin \Pi_{\Gamma'}$.

Case 1: $y \in \partial\Gamma'$. Fix a base point $h_0 \in \Gamma'_H$. Let τ be a geodesic ray from h_0 to y and choose a vertex $y_0 \in \tau$ so that $d(y_0, \Gamma'_H) > 2s + 100\delta$. Choose a second vertex $y_1 \in V(\Gamma')$ on τ so that $9\delta < d(y_0, y_1) < 10\delta$ and y_0 lies between y_0 and h_0.

Let

$$A_y = \{v \in V(\Gamma') : v \text{ lies on an arc from } y_0 \text{ to } y_1 \text{ of length at most } 22\delta\}.$$

By [4, Proposition 2.1 (F2)] and the fineness of Γ', there are only finitely many arcs from y_0 to y_1 of length at most 22δ, so A_y is finite. We show that γ intersects A_y.

We claim that $d(y_0, \gamma), d(y_1, \gamma) < 3\delta$. First suppose $h \in \partial\Gamma'$, so we can parameterize $\gamma \colon (-\infty, \infty) \to \Gamma'$, where $\lim_{t \to \infty} \gamma(t) = y$ and $\lim_{t \to -\infty} \gamma(t) = h$. There exists M such that for $t > 0$ with t sufficiently large, we have $d(\gamma(t), \tau) < M$ and $d(\gamma(-t), \Gamma'_H) < M$. Then there are geodesics η_t whose endpoints are $x_{t,\tau} \in \tau$ and $x_{t,h} \in \Gamma'_h$ with $d(x_\tau, \gamma) < M$ and $d(x_h, \gamma) < M$. We claim that there are $x_{t,0}, x_{t,1} \in \eta_t$ so that $d(y_0, x_{t,0}) < \delta$ and $d(y_1, x_{t,1}) < \delta$. Indeed, when t is large, there is a geodesic triangle with vertices $h_0, x_{t,\tau}$, $x_{t,h}$ so that one side is η_t, one side lies in τ and the other side lies in $\mathcal{N}_s(\Gamma'_H)$. Since triangles are δ-thin and $d(y_0, \Gamma'_H), d(y_1, \Gamma'_H) > 2s + 90\delta$, then $d(y_0, \eta_t), d(y_1, \eta_t) < \delta$.

Again, for large enough $t > 0$, each of the quantities

$$d(x_{t,\tau}, x_{t,0}), \quad d(x_{t,\tau}, x_{t,1}), \quad d(x_{t,h}, x_{t,0}), \quad d(x_{t,h}, x_{t,1})$$

can be made arbitrarily large. In particular, they can all be made to exceed the constant $M + 2\delta$. A standard hyperbolic geometry argument using a 2δ-slim quadrilateral implies that $d(x_{t,0}, \gamma), d(x_{t,1}, \gamma) < 2\delta$. Then for appropriately large $T > 0$, we have points

$z_{T,0}, z_{T,1} \in \gamma$ so that $d(z_{T,0}, x_{T,0}), d(z_{T,1}, x_{T,1}) < 2\delta$ which implies that $d(y_0, z_{T,0}) < 3\delta$ and $d(y_1, z_{T,1}) < 3\delta$.

In the case that $h \in \Pi_{\Gamma'}$, a similar argument where $x_{t,h}$ is replaced by h proves the claim that $d(y_0, \gamma), d(y_1, \gamma) < 2\delta$ and that we can choose $z_{T,0}$, $z_{T,i}$ in γ so that $d(y_i, z_{T,i}) < 2\delta$ for $i = 0, 1$.

Since $9\delta < d(y_0, y_1) < 10\delta$, then $3\delta < d(z_{T,1}, z_{T,2}) < 16\delta$ by the triangle inequality. Construct an arc σ consisting of

- a path of length at most 3δ from y_0 to $z_{T,0}$,

- a subpath of γ from $z_{T,0}$ to $z_{T,1}$ whose length is more than $1 < 3\delta$ but has length at most 16δ,

- a path of length at most 3δ from y_1 to $z_{T,1}$.

The subpath of γ is long enough that it contains a vertex of Γ'. Therefore, σ is a path between y_0 and y_1 with length at most 22δ and intersects γ in a vertex.

Case 2: $y \in \Pi_{\Gamma'}$. Fix a base point $h_0 \in \Gamma'_H$ and let τ be a geodesic from h_0 to y. Let $y_0 = y$. Let $y_1 \in V(\Gamma')$ be a vertex on τ such that $\delta < d(y_0, y_1) < 2\delta$ or if no such vertex exists, set $y_1 = h_0$. Let

$$A_y = \{v \in V(\Gamma') : v \text{ lies on a path of length at most } 10\delta \text{ from } y_0 \text{ to } y_1\}.$$

As in the proof of case 1, A_y is finite by [4, Proposition 2.1 (F2)]. Let $h \in \partial_B \Gamma'_H$ and suppose that γ is a geodesic from y to h. Let $z \in \gamma$ be a vertex of Γ'_H such that $\delta < d(z, y) < 2\delta$, or if no such vertex exists, then $h \in \Pi_{\Gamma'}$ and we can set $z = h$ with $0 < d(z, y) < \delta$.

Consider a geodesic triangle with vertices z, y, h_0 and let η be the side joining z to h_0. Since $d(y, z) < 2\delta$, there must exist some point $x \in \eta$ and $y_2 \in \tau$ with $d(x, y_2) < \delta$ and $d(x, z) < 2\delta$ by δ-thinness of geodesic triangles. Then $d(y, y_2) < 5\delta$ and τ is geodesic, so $d(y_2, y_1) < 5\delta$. Therefore, there is a path from y_0 to y_1 of length at most 10δ that intersects γ in the vertex z, and so $\gamma \cap A_y \neq \emptyset$.

We have showed that for all $y \in \partial_B \Gamma' \setminus \partial_B \Gamma'_H$, there exists a finite collection of vertices A_y so that

$$N(y, A_y) \subseteq \partial_B \Gamma' \setminus \partial_B \Gamma'_H.$$

Thus $\partial_B \Gamma' \setminus \partial_B \Gamma'_H$ is open in $\partial_B \Gamma'$. ■

Proposition 4.9. *The limit set of H, ΛH, in $\partial_B \Gamma'$ is $\partial_B \Gamma'_H$.*

Proof. Immediately, $\partial_B \Gamma'_H \subseteq \Lambda H$. By Proposition 4.8, $\partial_B \Gamma'_H$ is closed and H-invariant, so $\Lambda H \subseteq \partial_B \Gamma'_H$ because the limit set of H is the smallest closed H-invariant subset of ∂M. ■

Theorem 1.2. *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let Γ be a hyperbolic graph with a G-action so that Γ is generalized fine with respect to the action of $(G, \mathcal{P})$. For any*

$P \in \mathcal{P}$ and $g \in G$, let Γ_{P^g} be the subgraph of Γ whose cell stabilizer are commensurable to P^g. If $H \leq G$ has a quasiconvex cocompact core Γ_H and one of the following holds:

- Γ is fine or
- for all $P \in \mathcal{P}$ and $g \in G$, $\Gamma_H \cap \Gamma_{P^g} \neq \emptyset$ implies $|P^g \cap H| = \infty$,

then H is relatively quasiconvex in $(G, \mathcal{P})$.

Proof. The action of H on the fine hyperbolic subgraph Γ'_H constructed by electrifying with respect to the Γ_{P^g} shows that H is hyperbolic relative to the infinite vertex stabilizers. By [16, Theorem 1.1], the induced convergence group action of H on $\partial_B \Gamma'_H$ is geometrically finite. Since $\partial_B \Gamma'_H = \Lambda H$, the subgroup H has a geometrically finite convergence group action on ΛH, and hence H satisfies Definition 4.4 for relative quasiconvexity. ∎

We can also rephrase Theorem 1.2 as a criterion for relative quasiconvexity in fine hyperbolic graphs. This special case recovers the following result of Martínez-Pedroza and Wise.

Corollary 1.3 ([17, Theorem 1.7]). *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair acting cocompactly on a fine hyperbolic graph so that every edge stabilizer is finite. A subgroup $H \leq G$ is relatively quasiconvex in $(G, \mathcal{P})$ if and only if H has a quasiconvex cocompact core in Γ.*

Proof sketch. One direction follows immediately from Theorem 1.2. If H is relatively quasiconvex, the join of ΛH in Γ provides the quasiconvex core for H, see [4, end of Section 5]. ∎

Remark 4.10. For a relatively hyperbolic pair $(G, \mathcal{P})$, we use the notation $\partial_{\mathcal{P}} G$ to denote the Bowditch boundary of G with respect to $\mathcal{P}$. When Γ' is a fine hyperbolic graph that witnesses the relative hyperbolicity of $(G, \mathcal{P})$, we henceforth conflate $\partial_{\mathcal{P}} G$ with $\partial_B \Gamma'$.

We can also prove that hyperplane stabilizers are relatively quasiconvex.

Corollary 4.11. *Let $(G, \mathcal{P})$ act relatively geometrically on a $\mathrm{CAT}(0)$ cube complex $\widetilde{X}$. Let H be the stabilizer of a hyperplane W of $\widetilde{X}$. Then H is relatively quasiconvex in $(G, \mathcal{P})$.*

Proof. Subdivide $\widetilde{X}$ cubically once to a complex $\widetilde{X}_W$ so that W is a subcomplex. The action of $(G, \mathcal{P})$ on $\widetilde{X}_W$ is still relatively geometric. Recall from Example 3.3 that $\widetilde{X}_W^{(1)}$ is generalized fine with respect to the action of $(G, \mathcal{P})$. If Γ_{P^g} intersects $W \cap \widetilde{X}_W^{(1)}$, then a finite index subgroup of P^g stabilizes an edge of $\widetilde{X}$ dual to W. Hence $P^g \leq \mathrm{Stab}_G(W)$. Since W is convex and cocompact in $\widetilde{X}_W$, $W \cap \widetilde{X}_W^{(1)}$ is an H-invariant H-cocompact connected convex subgraph of $\widetilde{X}_W^{(1)}$. The relative quasiconvexity of H now follows from Theorem 1.2. ∎

5. A separation criterion for the Bowditch boundary

A construction of Sageev [19] shows that group actions on CAT(0) cube complexes arise naturally from groups with collections of 'codimension-1' subgroups. Building on the work of Bergeron and Wise [3] for hyperbolic cubulations, the first and second authors gave a boundary criterion [8, Theorem 2.6] for relatively geometric actions that guarantees the existence of a relatively geometric action of a relatively hyperbolic pair $(G, \mathcal{P})$ on a CAT(0) cube complex whenever G contains a sufficient collection of full relatively quasiconvex subgroups that separate points in the Bowditch boundary. The main theorem of this section, Theorem 5.7, helps to show that stabilizers of quasiconvex cores in fine hyperbolic graphs exhibit 'hyperplane-like' behavior and provide a source of codimension-1 subgroups that may be used with [8, Theorem 2.6].

5.1. Hypersets and hypercarriers

Let Γ be a graph. A *hyperset L* in Γ is a collection of edge midpoints and vertices of Γ so that $\Gamma \setminus L$ has two components. A *(hyperset) carrier J for L* is a subgraph of Γ containing L. Immediately, a hyperset carrier J for L has the property that if $v_1, v_2 \in J$ and v_1, v_2 are joined by an edge in Γ whose midpoint lies in L, then J contains that edge between v_1, v_2.

Hypersets and carriers arise naturally in the one-skeleton of a CAT(0) cube complex. We will see that they are particularly helpful in the setting of relatively geometric actions.

Example 5.1. Let $\widetilde{C}$ be a CAT(0) cube complex, and let W be a hyperplane. Then $L = W \cap \widetilde{C}^{(1)}$ is a hyperset. One choice of hyperset carrier J for L is the intersection of the hyperplane carrier of W with $\widetilde{C}^{(1)}$.

In this situation, we refer to L as the *hyperset associated to W* and J as *the (hyper)carrier (of the hyperset associated to W)*.

We observe the following useful fact in the setting of generalized fine hyperbolic graphs.

Observation 5.2. *Let Σ be generalized fine with respect to the action of $(K, \mathcal{D})$ and let Σ_{D^k} be the subgraph of cells whose stabilizer is commensurable to D^k for $D \in \mathcal{D}$, $k \in K$. Let $\sigma: \Sigma \to \Gamma$ be the electrification map that collapses the Σ_{D^k}. Let $S = \bigcup\{\Sigma_{D^k} : \Sigma_{D^k} \cap L \neq \emptyset\}$. Suppose L is a hyperset and J is a connected quasiconvex hyperset carrier.*

Given an edge e of Σ, $\sigma(e)$ is either an edge or a vertex of Γ. Hence σ takes an edge midpoints to an edge midpoint or a vertex. Any path ρ' in Γ between images of the components of $\Sigma \setminus L$ has a de-electrification ρ that must pass through L an odd number of times. Since $\sigma(\rho) = \rho'$, then ρ' passes through $\sigma(L)$ an odd number of times. Therefore, $\sigma(L)$ is a hyperset in Γ.

Immediately, $\sigma(J)$ is a subgraph of Γ that contains $\sigma(L)$, so $\sigma(J)$ is a hyperset carrier for $\sigma(L)$.

5.2. The separation criterion

We set the following assumptions for the remainder of this subsection.

Hypotheses 5.3. *Let $(K, \mathcal{D})$ be a relatively hyperbolic pair and let K act on a fine δ-hyperbolic graph Γ with the following properties:*

- *the action of K is cocompact,*
- *edge stabilizers are finite, and*
- *each $D \in \mathcal{D}$ stabilizes a single vertex.*

Let L be a hyperset, and let J be a connected quasiconvex carrier for L.

Our goal is to decide whether two points in the Bowditch boundary $\partial_{\mathcal{D}} K$ lie in complementary components of the limit set of $\mathrm{Stab}_G(L)$. The hyperset L separates Γ into two complementary components, but it is not immediately apparent that the limit set $\Lambda \, \mathrm{Stab}_G(L)$ partitions the Bowditch boundary into multiple components with respect to the topology described in Definition 4.6.

Definition 5.4. With the setup in Hypotheses 5.3, J has the *two-sided carrier property* if there exist connected quasiconvex subsets J^+ and J^- of J so that

(1) $J^+ \cap J^- \subseteq L$,

(2) $J^+ \cup J^- = J$,

(3) every path in Γ between vertices in the two distinct components of $\Gamma \setminus L$ must intersect both J^+ and J^-, and

(4) if v is a vertex of $J^- \cap J^+$ with infinite stabilizer, then $v \in \Lambda \, \mathrm{Stab}_K(L)$.


Figure 3. An example of a hyperset carrier J with the two-sided carrier property in the one-skeleton of a polygonal complex. The hyperset L is the union of points indicated in black. The subset J^+ (resp. J^-) is indicated in red (resp. blue).

The two-sided carrier property arises naturally in our intended application to relatively geometric actions. See Figure 3 for an illustration of the two-sided carrier property.

Example 5.5. When $(K, \mathcal{D})$ acts relatively geometrically on a CAT(0) cube complex $\tilde{X}$, and W is a hyperplane with associated hyperset L and carrier J, there are two natural sides J^+ and J^- of the carrier J. Note that the two sides are slightly larger than the

combinatorial hyperplanes on either side of W because L needs to be contained in their union. Recall that $\widetilde{X}^{(1)}$ usually fails to be fine, but Proposition 2.12 implies that collapsing compact subgraphs of $\widetilde{X}^{(1)}$ yields a fine hyperbolic graph. The images of L and J remain a hyperset and hypercarrier, respectively, but J^+ and J^- may have overlapping images. As we will see, this may only happen at vertices that are already parabolic points in the limit set of the stabilizer of W.

Recall that under Hypotheses 5.3, the points of the Bowditch boundary are either *conical limit points* that lie in $\partial\Gamma$, the visual boundary of Γ or are *parabolic vertices* (also called *peripheral vertices*) of Γ, which are stabilized by maximal parabolics.

Definition 5.6. Assuming Hypotheses 5.3, we say that L *separates* $x, y \in \partial_{\mathcal{D}} K$ if $x, y \notin \Lambda \operatorname{Stab}_K(L)$ and one of the following holds:

- x, y are both conical limit points, and there exists some geodesic $\gamma: (-\infty, \infty) \to \Gamma$ with $\lim_{t\to\infty} \gamma(t) = x$ and $\lim_{t\to-\infty} \gamma(t) = y$ so that there exists $T > 0$ such that for all $t_- < -T < 0 < T < t_+$, $\gamma(t_+)$ and $\gamma(t_-)$ are in distinct components of $\Gamma \setminus L$,

- x is a parabolic vertex in Γ, y is a conical limit point, and there exists some geodesic $\gamma: [0, \infty) \to \Gamma$ from $x = \gamma(0)$ to y so that for sufficiently large t, $\gamma(t)$ and x are in distinct components of $\Gamma \setminus L$,

- x, y are both parabolic vertices in Γ, and x, y lie in distinct components of $\Gamma \setminus L$.

In the first two cases, we say that the geodesic γ *witnesses that L separates x and y.*

The remainder of this section is devoted to proving Theorem 5.7.

Theorem 5.7. *Assume Hypotheses 5.3 and let J satisfy the two-sided carrier property. If L separates $x, y \in \partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(L)$, then there exists a subgroup $K_L \leq \operatorname{Stab}_K(L)$ of index at most 2 so that x, y are in K_L-distinct components of $\partial_{\mathcal{D}} K \setminus \Lambda K_L$.*

We first establish some notation that we will use below.

Notation 5.8. Let Y be a subset of γ and let $r \in \mathbb{R}$. Then

$$\mathcal{N}_r(Y) = \{x \in \Gamma : \inf_{y \in Y} d_\Gamma(x, y) \leq r\}.$$

Recall from Definition 4.6 that if $x \in \partial_{\mathcal{D}} K$ and A is a set of vertices in Γ, the set $N(x, A)$ consisting of $y \in \partial_{\mathcal{D}} K$ so that some geodesic from x to y avoids $A \setminus \{x\}$ is an open neighborhood of x in $\partial_{\mathcal{D}} K$. The following lemma helps us control certain open neighborhoods of points in $\partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(L)$.

Lemma 5.9. *Assume Hypotheses 5.3. Let J_0 be a connected quasiconvex subset of J, and let $x \in \partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(J)$ so that $x \notin J_0$. There exists a finite collection of vertices V_x such that for any vertex $j \in J_0$ and any geodesic $\gamma_{x,j}$ between x and j, the intersection $\gamma_{x,j} \cap (V_x \setminus x)$ is not empty.*

Proof. Fix some $j_0 \in J_0$, and fix R so J_0 is R-quasiconvex. Consider a (possibly ideal) geodesic triangle with vertices x, j_0, j and sides $\gamma_{x,j}, \gamma_{x,j_0}, \gamma_{j,j_0}$.

Case 1: *x is a conical limit point.* Let x_0 be a vertex on γ_{x,j_0} so that $d(x_0, J_0) > R + \delta + 1$. Such a vertex exists because $x \notin \Lambda \operatorname{Stab}_K(J)$. Since $\gamma_{j,j_0} \subseteq \mathcal{N}_R(J_0)$ by quasiconvexity, $d(x_0, \gamma_{j,j_0}) > \delta$. Therefore, by hyperbolicity, there exists a vertex $y_0 \in \gamma_{x,j}$ so that $d(x_0, y_0) < \delta$. Then there exist vertices $y_1 \in \gamma_{x,j}$ and $x_1 \in \gamma_{x,j_0}$ so that y_1 lies between y_0 and j on $\gamma_{x,j}$, x_1 lies between x_0 and j_0 on γ_{x,j_0}, and $d(y_1, x_1) \le \delta$. Hence there exists a $(1, 4\delta)$-quasi-geodesic arc from j_0 to x_0 following

- from j_0 to x_1 via γ_{x,j_0},
- from x_1 to y_1 via a geodesic of length at most δ,
- from y_1 to y_0 via $\gamma_{x,j}$,
- from y_0 to x via a geodesic of length at most δ.

Then y_1 lies on a $(1, 4\delta)$-quasi-geodesic arc between x_0 and j_0. Let V_{x_0, j_0} be the subgraph of $(1, 4\delta)$-quasi-geodesic arcs between x_0 and j_0. By [4, Lemma 8.2], V_{x_0, j_0} is locally finite. The length of any such arc is uniformly bounded above by $d(x_0, j_0) + 4\delta$, so this subgraph has finite diameter and is therefore finite. Hence y_1 is one of finitely many vertices in Γ.

Case 2: *x is a parabolic point.* If $d(x, J_0) > R + \delta + 1$, carry out the same proof as in the preceding case (the choice $x_0 = x$ suffices).

Hence assume $d(x, J_0) \le R + \delta + 1$. The geodesic triangle with vertices x, j, j_0 is δ-thin. Therefore, there exist a vertex $j_1 \in J_0$ and a subpath of $\gamma_{x,j}$ of length at most $R + \delta + 2$ between x and a vertex $y_1 \in \gamma_{x,j}$ so that $y_1 \ne x$ and vertices y_1, j_1 have the following properties:

- either y_1 is the vertex on $\gamma_{x,j}$ in J_0 that is closest to x in which case we set $j_1 = y_1$,
- or $d(y_1, j_1) < R + \delta + 1$, a shortest path from y_1 to j_1 does not backtrack along $\gamma_{x,j}$ and there is an arc from x to j_1 passing through y_1 that does not contain any vertices of J_0 other than j_1.

In the first case, $y_1 \ne x$ because $x \notin J_0$. In the second case, one might worry that eliminating backtracking could force us to choose $y_1 = x$. If $d(x, j) \ge R + \delta + 1$, this is not a problem. If $d(x, j) < R + \delta + 1$, we can ensure we are in the first case by letting $y_1 = j_1$ be the vertex on $\gamma_{x,j}$ in J_0 that is closest to x. In both cases, we obtain an arc σ from x to j_1 that contains y_1 and no vertices of J_0 other than j_1.

In all of these above cases,

$$d(j_1, j_0) \le d(j_0, x) + d(j_1, y_1) + d(y_1, x) \le 3R + 3\delta + 4.$$

Since J_0 is connected and quasiconvex, and Γ is hyperbolic, there exist $\lambda \ge 1$ and $\epsilon \ge 0$ so that some (λ, ϵ)-quasi-geodesic arc ρ in J_0 connects j_0 and j_1. Thus ρ has length at most $\lambda(3R + 3\delta + 4) + \epsilon$. Note that σ cannot backtrack along ρ at j_1 because

every edge of ρ has both endpoints in J_0 while j_1 is the only vertex on σ that lies in J_0. Hence y_1 is on an arc from x_0 to j_0 of length at most

$$\lambda(3R + 3\delta + 4) + \epsilon + d(j_1, y_1) + d(y_1, x) \le \lambda(3R + 3\delta + 4) + \epsilon + 3R + 3\delta + 4$$

and there are finitely many such arcs by [4, Proposition 2.1 (F2)], since Γ is fine. Then there are finitely many possibilities for y_1. ∎

Proposition 5.10. *Assume Hypotheses* 5.3, *and suppose J has the two-sided carrier property. Suppose L separates $x, y \in \partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(L)$. If $z \in \partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(L)$ and L does not separate x, z, then L separates y, z.*

Proof. Consider a geodesic triangle with vertices x, y, z and sides $\gamma_{xy}, \gamma_{yz}, \gamma_{xz}$, where the ordered subscripts indicate the endpoints and orientation. Further, assume γ_{xy} witnesses that L separates x, y.

If z is a conical limit point, then all but a finite length of γ_{xz} lies in a single component C of $\Gamma \setminus L$ because L does not separate x from z. By hypothesis $z \notin \Lambda \operatorname{Stab}_K(L)$, so, as $t \to \infty$, the quasiconvexity of L ensures that $d(\gamma_{xz}(t), L) \to \infty$. Moreover, for all $t > 0$ large enough, $\gamma_{yz}(t)$ lies in C by hyperbolicity. If y is a parabolic point, then y lies in the other component $C' \ne C$ of $\Gamma \setminus L$, otherwise y is a conical limit point and, for all $t > 0$ sufficiently large, $\gamma_{zy}(t) = \gamma_{yz}(-t)$ lies in C'. In either case, L separates y from z.

If z is a parabolic vertex in γ, then $z \in C$ by hypothesis. As above, whether y is a parabolic or conical limit point, L separates y from z. ∎

By Proposition 5.10, there is an equivalence relation $\sim$ on $\partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(L)$ defined by $x \sim y$ if and only if $x = y$ or L does not separate x, y. There are two equivalence classes.

Proposition 5.11. *Assume Hypotheses* 5.3, *and suppose J has the two-sided carrier property. If $x, y \in \partial_{\mathcal{D}}(K) \setminus \Lambda \operatorname{Stab}_K(L)$ and $x \not\sim y$, then x, y lie in distinct components of $\partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(L)$. By passing to a subgroup K_L of index at most 2 in $\operatorname{Stab}_K(L)$, these components are K_L-distinct.*

Proof. We claim that if $z \in \partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(L)$, then there exists an open neighborhood U of $z \in \partial_{\mathcal{D}} K \setminus \Lambda \operatorname{Stab}_K(L)$ so that $u \sim z$ for all $u \in U$.

To this end, we claim that there is a finite set of vertices V_z so that any geodesic from z that crosses L must pass through V_z.

If $z \notin J$, then Lemma 5.9 with $J = J_0$ immediately provides V_z. Hence, we may assume z is a parabolic vertex in J.

The two-sided carrier property ensures that $J = J^+ \cup J^-$ and $J^+ \cap J^- \subseteq L$. Then $z \notin L$, so $z \in J^+ \setminus J^-$ or $z \in J^- \setminus J^+$. Up to relabeling, we may assume $z \in J^+ \setminus J^-$. Since every path from z to a vertex of the other component of $\Gamma \setminus L$ must pass through J^-, we can apply Lemma 5.9 to show that there exists a finite set of vertices V_z so that any geodesic γ from z to J^- has $\gamma \cap V_z \ne \emptyset$.

Hence in all cases, if L separates $z, w \in \partial_{\mathcal{D}}(K) \setminus \Lambda \operatorname{Stab}_K(L)$, any geodesic between z and w passes through V_z.

Recall that the set $N(z, V_z)$ from Definition 4.6 is an open neighborhood of z in $\partial_{\mathcal{D}}(K)$. Therefore, if $U = N(z, V_z)$, then $z \sim u$ for all $u \in U$.

Thus we conclude that $[x]$ and $[y]$ are unions of components of $\partial_{\mathcal{D}} K \setminus \operatorname{Stab}_K(L)$ because they are open and partition $\partial_{\mathcal{D}} K \setminus \operatorname{Stab}_K(L)$.

Each $k \in \operatorname{Stab}_K(L)$ permutes two components of $\Gamma \setminus L$, so $\operatorname{Stab}_K(L)$ acts on the equivalence classes of $\sim$. By passing to an index 2 subgroup K_L of $\operatorname{Stab}_K(L)$ if necessary, we can ensure that for any $k_L \in K_L$, $k_L \cdot [x] \neq [y]$, so x and y are in K_L-distinct components of $\partial_{\mathcal{D}} K \setminus \operatorname{Stab}_K(L)$. ∎

Proposition 5.11 completes the proof of Theorem 5.7.

6. Separating points in the Bowditch boundary of a group acting relatively geometrically using hyperplane stabilizers

For this section, let $(K, \mathcal{D})$ act relatively geometrically on a CAT(0) cube complex $\widetilde{X}$. For $D \in \mathcal{D}$ and $k \in K$, let Σ_{D^k} denote the subgraph induced by the vertices whose stabilizers are commensurable to D^k. Let Γ be the complete electrification of $\widetilde{X}$ with respect to

$$\{\Sigma_{D^k} : D \in \mathcal{D}, k \in K\}.$$

Recall that the electrification map $\beta \colon \widetilde{X}^{(1)} \to \Gamma$ that collapses the Σ_{D^k} to a single vertex is a continuous coarse inverse of the map in Proposition 2.11 and is K-equivariant.

Let L be the hyperset associated to a hyperplane W of $\widetilde{X}$ as in Example 5.1. Let J be the associated hyperset carrier. Then $\beta(J)$ is a quasiconvex subset of Γ. Let K_W be the stabilizer of W.

Since J is associated to a hyperplane, $J \setminus L$ has two distinct components J^+ and J^- that are connected and convex. Therefore, $\beta(J^+)$ and $\beta(J^-)$ are both quasiconvex.

Proposition 6.1. *Let W be a hyperplane of $\widetilde{X}$, and let D^k be a peripheral subgroup. Then W is dual to an edge of Σ_{D^k} if and only if D^k is commensurable to a subgroup of $\operatorname{Stab}_K(W)$.*

Proof. If W is dual to an edge e of Σ_{D^k}, then $\operatorname{Stab}_K(e)$ is commensurable to D^k. Since W is the unique hyperplane dual to e, we have $\operatorname{Stab}_K(e) \leq \operatorname{Stab}_K(W)$.

Conversely, suppose D^k is commensurable to a subgroup $D_0 \leq \operatorname{Stab}_K(W) \cap D^k$. By passing to a further finite index subgroup, we may assume that there is a vertex x in Σ_{D^k} such that $D_0 \leq \operatorname{Stab}_K(x)$. By convexity of W, there is a unique point p in W that is the nearest point projection of x. The point p is sometimes called the *gate of x in W*. When W is a hyperplane, the gate of x is a dual edge midpoint (see [15, Section 2.2]). Hence, $D_0 \leq \operatorname{Stab}_K(p)$.

Let e be the edge with midpoint p. The stabilizers $\mathrm{Stab}_K(e)$ and $\mathrm{Stab}_K(p)$ are commensurable because the action is cellular. Moreover, the peripheral D^k is commensurable to $\mathrm{Stab}_K(e)$ because cell stabilizers in a relatively geometric action are commensurable to a unique peripheral subgroup and peripheral subgroups have finite intersection. Thus, $e \subseteq \Sigma_{D^k}$ as needed. ∎

Proposition 6.1 is particularly relevant for studying hypersets in Γ, the electrified fine hyperbolic graph. In particular, peripheral points that lie in hypersets coming from images of hyperplanes in $\widetilde{X}$ are visible in the subgroup structure of the hyperplane stabilizer.

Proposition 6.2. *If $x \in \beta(J^+) \cap \beta(J^-)$ is a vertex, then x is a peripheral point whose stabilizer is commensurable to a subgroup of $\mathrm{Stab}_K(W)$.*

Proof. Since $J^+ \cap J^- \cap \widetilde{X}^{(0)} = \emptyset$, there exist distinct vertices $y_+ \in J^+ \setminus L$ and $y_- \in J^- \setminus L$ such that $\beta(y_+) = x = \beta(y_-)$. Since β is the map that collapses the Σ_{D^k}, y_+ and y_- both lie in some $\Sigma_{D_y^{k_y}}$ where $D_y \in \mathcal{D}$ and $k_y \in K$. Any combinatorial path between different sides of W contains an edge dual to W, so the connected subgraph $\Sigma_{D_y^{k_y}}$ must contain an edge e dual to W. The stabilizer of the edge e stabilizes W and must be commensurable to $D_y^{k_y}$ by Proposition 6.1. ∎

Proposition 6.3. *The hyperset carrier $\beta(J)$ has the two-sided carrier property.*

Proof. We immediately see that $\beta(J) = \beta(J^+) \cup \beta(J^-)$. If $s \in \beta(J^+) \cap \beta(J^-)$, then s has stabilizer commensurable to a subgroup of K_W. Therefore, $\beta(\Sigma_D^k)$ is a vertex in $\Lambda\,\mathrm{Stab}_K(\beta(L))$ the limit set of the hyperset stabilizer. Let ρ be a path between components of $\Gamma \setminus \beta(L)$. Let $\ell \in \beta(L) \cap \rho$. Either ℓ is the midpoint of an edge whose endpoints are in $\beta(J^+)$ and $\beta(J^-)$ or ℓ is a vertex formed by collapsing an edge dual to a hyperplane. In both cases, $\ell \in \beta(J^+) \cap \beta(J^-)$. Hence every path between vertices of $\Gamma \setminus \beta(L)$ intersects both $\beta(J^+)$ and $\beta(J^-)$. ∎

We are now ready to prove that any two points $x, y \in \partial_{\mathcal{D}} K$ can be separated by a hyperset associated to a hyperplane.

As we have seen, relatively geometric actions on CAT(0) cube complexes let us take advantage of both the cubical geometry of $\widetilde{X}$ as well as the fine graph structure of Γ. The following Definition 6.4 lets us study separating hypersets using the separating properties of hyperplanes. Note that it need not be the case that every hyperset in Γ is the image of a hyperplane in $\widetilde{X}$.

Definition 6.4. Let W be a hyperplane in $\widetilde{X}$ with associated hyperset L as in Example 5.1. We say that W separates $x, y \in \partial_{\mathcal{D}} K \setminus \Lambda\,\mathrm{Stab}_K(W)$ if $\beta(L)$ separates x, y in the sense of Definition 5.6.

We are now ready to prove that any two points $x, y \in \partial_{\mathcal{D}} K$ can be separated by a hyperset (associated to a hyperplane). Our strategy is to show that if $x, y \in \partial_{\mathcal{D}} K$, then

some hyperplane W separates x, y in the sense of Definition 6.4. Then we apply Theorem 5.7 to show that x, y lie in distinct complementary components of $\Lambda \operatorname{Stab}_K(W)$.

Lemma 6.5. *Let x be a parabolic point in $\partial_{\mathcal{D}} K$ with stabilizer K_x and let γ be a combinatorial geodesic with end-vertices v, w in $\widetilde{X}$ so that $v \in \Sigma_{K_x}$, and let γ be of minimal length among all combinatorial geodesics between w and Σ_{K_x}. Then every hyperplane dual to an edge of γ does not intersect Σ_{K_x}.*

Proof. Suppose e is the edge of γ with endpoint $v \in \Sigma_{K_x}$. By minimality, $e \not\subseteq \Sigma_{K_x}$. Let W be the hyperplane dual to e. If W intersects Σ_{K_x}, then W is dual to an edge $f \in \Sigma_{K_x}$. Note that $\operatorname{Stab}_K(f)$ and $\operatorname{Stab}_K(v)$ are commensurable, so there exists a $k \in K_x$ so that K fixes f, v and not e. Since k fixes the dual edge f, $k \cdot W = W$. Then $k \cdot e$ is adjacent to v and is dual to W. Therefore, the hyperplane W self-osculates which is impossible in a CAT(0) cube complex (see, for example, [21, pp. 20–21]).

The first paragraph shows that the first edge (counting from v) of any geodesic between v and w cannot be dual to a hyperplane that intersects Σ_{K_x}. We now assume that the first i edges of any minimal geodesic ρ between v and w are dual to a hyperplane that does not intersect Σ_{K_x} and prove that the $(i + 1)$st edge of ρ is dual to a hyperplane that does not intersect Σ_{K_x}. Now let $e_1, e_2, \ldots, e_k$ be the edges of ρ with corresponding dual hyperplanes $W_1, W_2, \ldots, W_k$. If W_{i+1} intersects Σ_{K_x}, then there is a disk diagram D enclosed by the path $e_1 \ldots e_i e_{i+1}$, a curve in the carrier of W_{i+1} and a path in K_x. Since W_i does not intersect Σ_{K_x}, it must exit D by crossing W_{i+1}.

By [21, Lemma 3.6], W_{i+1} and W_i cannot interosculate. Therefore, e_i and e_{i+1} must corner a square. Let e_i' and e_{i+1}' be the edges opposite e_i and e_{i+1}, respectively. Then let

$$\sigma' = e_1 e_2 \ldots e_{i-1} e_{i+1}' e_i' e_{i+2} \ldots e_k.$$

Now W_{i+1} is the hyperplane dual to the ith edge of σ' from v, which violates the inductive hypothesis. Hence W_{i+1} cannot intersect Σ_{K_x}. ∎

Before continuing with the proof of Theorem 1.4, we require one additional auxiliary fact. It is clear that geodesics in Γ lift to quasi-geodesics in $\widetilde{X}$ via the de-electrification map in Proposition 2.11. We need a way to show that long enough quasi-geodesics in $\widetilde{X}$ escape any finite neighborhood of some hyperplane. Recall the *Ramsey number* $\operatorname{Ram}(a, b)$ is the number of vertices such that any graph on $\operatorname{Ram}(a, b)$ vertices either contains a complete graph of size a or its complement contains a complete graph of size b (see [12] for more about Ramsey numbers). We may associate to any CAT(0) cube complex its *crossing graph* with vertex set corresponding to hyperplanes, and two vertices are adjacent if and only if their associated hyperplanes cross (see, for example, [15] for more details on crossing and related graphs).

Lemma 6.6. *Let Ω be an arbitrary CAT(0) cube complex with $d = \dim(\Omega)$. Let α be an arbitrary combinatorial geodesic in Ω. For any positive integer $N > 0$, any collection $\mathcal{H}$ consisting of at least $R = \operatorname{Ram}(d + 1, N)$ distinct hyperplanes all dual to edges of α*

contains a subset $\{\mathfrak{h}_1, \ldots, \mathfrak{h}_N\} \subseteq \mathcal{H}$ *that form a nested sequence of* N-*halfspaces* $\mathfrak{h}_1^+ \subsetneq \mathfrak{h}_2^+ \subsetneq \cdots \subsetneq \mathfrak{h}_N^+$.

In particular, if the geodesic α *has length at least* R, *then there exist a hyperplane* $\mathfrak{h}$ *dual to an edge of* α *and a vertex* $\alpha(t)$ *such that* $d(\alpha(t), \mathfrak{h}) \geq N - 1$.

Proof. Let $\mathcal{C}$ denote the crossing graph of Ω. Since $\widetilde{X}$ has finite dimension, any collection of pairwise crossing hyperplanes in Ω has cardinality at most d, so $\mathcal{C}$ does not contain a complete graph on $d + 1$ vertices. Hyperplanes dual to edges of a combinatorial geodesic are distinct by work of Sageev [19, Theorem 4.13], so the hyperplanes dual to edges of α correspond to an induced subgraph of $\mathcal{C}$.

So, any collection of at least R hyperplanes dual to α contains a subset $\{\mathfrak{h}_1, \ldots, \mathfrak{h}_N\}$ that pairwise do not cross. Each $\mathfrak{h}_i$ is dual to an edge of α, so $\{\mathfrak{h}_1, \ldots, \mathfrak{h}_N\}$ can be totally ordered by picking an orientation on α. This orientation corresponds to a choice of halfspaces $\mathfrak{h}_i^+$, which are clearly nested.

To see the last statement of Lemma 6.6, observe that if $\mathfrak{h}_1^+ \subsetneq \mathfrak{h}_2^+ \subsetneq \cdots \subsetneq \mathfrak{h}_N^+$ is a sequence of N nested halfspaces, then any points $p \in \mathfrak{h}_1$ and $q \in \mathfrak{h}_N$ are distance $d(p, q) \geq N - 1$ apart. Thus, for $p = \alpha \cap \mathfrak{h}_1$ we have $d(p, \mathfrak{h}_N) \geq N - 1$. ∎

As we saw in the proof of Lemma 6.6, (combinatorial) geodesics in CAT(0) cube complexes may only cross a given hyperplane at most once. On the other hand, (infinite) quasi-geodesics may cross a given hyperplane (infinitely) many times even when the underlying complex is locally finite. Relatively geometric actions on CAT(0) cube complexes give up local finiteness, but require that the cube complex also be δ-hyperbolic. In Lemmas 6.7 and 6.8, we will show that in δ-hyperbolic CAT(0) cube complexes there is a choice of hyperplane whose interactions with a given quasi-geodesic have many of the same useful properties of a hyperplane dual to an honest geodesic. A subspace Y of a geodesic metric space X is called *Morse* when for every $A > 0$ and $B \geq 0$ there exists a constant $D = D(A, B) \geq 0$ such that any (A, B)-quasi-geodesic joining points in Y is contained in the D-neighborhood of Y. We call D the *Morse constant* for the quasi-geodesic parameters (A, B).

Lemma 6.7. *Let* $\widetilde{X}$ *be a* δ-*hyperbolic* CAT(0) *cube complex. Let* $\gamma: (-\infty, \infty) \to \widetilde{X}^{(1)}$ *be a connected bi-infinite combinatorial* (λ, ϵ)-*quasi-geodesic. Given* $M > 1$, *there exist a hyperplane* W *of* $\widetilde{X}$ *and* $t_M > 0$ *such that for all* t *with* $|t| > t_M$,

 (1) $\gamma(\pm t)$ *lie in distinct complementary components of* W,

 (2) γ *crosses* W *an odd number of times, and*

 (3) $d(\gamma(t), W) > M$.

In particular, $\gamma(t) \in W$ *implies* $|t| \leq t_M$.

Proof. For any $t_1, t_2 \in \mathbb{R}$, we write $[\gamma(t_1), \gamma(t_2)]$ to mean any combinatorial geodesic connecting $\gamma(t_1)$ to $\gamma(t_2)$. Note also that the M-neighborhood of a convex subset of $\widetilde{X}$ is necessarily Morse.

Let $D = D(\lambda, \varepsilon) > 0$ be the Morse constant for the quasi-geodesic parameters of γ. Let R be the constant from Lemma 6.6 with $N = 2(D + M + 1) + 1$. Choose

$$t_M > \lambda(R + \varepsilon).$$

Let $\gamma_0 = [\gamma(-t_M), \gamma(t_M)]$. By choice of t_M, we have $\mathrm{diam}(\gamma_0) > R$. Lemma 6.6 guarantees that there exist hyperplanes $\{\mathfrak{h}_i : -(D + M + 1) \le i \le D + M + 1\}$ all dual to edges of γ_0 that form the following nested sequence of halfspaces:

$$\mathfrak{h}_{-(D+M+1)} \subsetneq \cdots \subsetneq \mathfrak{h}_{-1} \subsetneq \mathfrak{h}_0 \subsetneq \mathfrak{h}_1 \subsetneq \cdots \subsetneq \mathfrak{h}_{D+M+1}.$$

We will see that we may choose $W = \mathfrak{h}_0$. Since the halfspaces of the $\mathfrak{h}_i$ are nested,

$$\min\{d(\gamma(-t_M), \mathfrak{h}_0), d(\gamma(t_M), \mathfrak{h}_0)\} > D + M. \tag{$\dagger$}$$

The concatenation,

$$\Upsilon := \gamma|_{(-\infty, -t_M)} \cup \gamma_0 \cup \gamma|_{(t_M, \infty)},$$

is again a (λ, ε)-quasi-geodesic. Since γ_0 crosses $\mathfrak{h}_0$, if $\gamma(t) \in \mathcal{N}_M(\mathfrak{h}_0)$ for some t with $t > t_M$, then $\gamma(t_M) \in \mathcal{N}_{D+M}(\mathfrak{h}_0)$, contrary to ($\dagger$). We conclude that if $t > t_M$, then $\gamma(t_M) \notin \mathcal{N}_M(\mathfrak{h}_0)$. Similarly, if $t < -t_M$, then $\gamma(t) \notin \mathcal{N}_M(\mathfrak{h}_0)$. Since Υ and γ coincide for all $|t| > t_M$, we immediately have that $\gamma(t)$ and $\gamma(-t)$ lie in distinct components of $\widetilde{X} \setminus \mathcal{N}_M(\mathfrak{h}_0)$. Since γ_0 crosses $\mathfrak{h}_0$ once, γ crosses $\mathfrak{h}_0$ an odd number of times. ∎

It remains to account for the situation where γ joins a parabolic point to a conical limit point. Using a similar argument to Lemma 6.7, it is possible to prove the following.

Lemma 6.8. *Let $\gamma: [0, \infty) \to \widetilde{X}$ be an infinite combinatorial (λ, ϵ)-quasi-geodesic ray in a δ-hyperbolic* CAT(0) *cube complex $\widetilde{X}$ where $x = \gamma(0)$ is a vertex. Given $M > 1$, there exist a hyperplane W of $\widetilde{X}$ and $t_M > 0$ so that for all t with $t > t_M$,*

(1) *W separates $\gamma(0)$ and $\gamma(t)$,*

(2) *$\Sigma_{\mathrm{Stab}_K(x)} \cap W = \emptyset$,*

(3) *$d(\gamma(t), W) > M$.*

In particular, if $\gamma(t) \in W$ then $0 < t \le t_M$.

The same strategy used in Lemma 6.7 works to prove Lemma 6.8 with the following adjustments:

- The hyperplane W should be the hyperplane dual to an edge in the middle of a geodesic joining $\gamma(0)$ and $\gamma(t)$ for some suitable $t \gg 0$.

- When $t \gg \max_{D \in \mathcal{D}}\{\mathrm{diam}\,\Sigma_D\}$, then $\Sigma_{\mathrm{Stab}_K(x)} \cap W = \emptyset$.

We summarize the discussion above in the context of separating points in the Bowditch boundary as follows.

Proposition 6.9. *Let $x, y \in \partial_{\mathcal{D}} K$ be distinct. There exists a hyperplane that separates x, y in the sense of Definition 6.4 with respect to the action of K on $\widetilde{X}$.*

Proof. If x, y are both conical limit points, let γ be a bi-infinite geodesic between x and y. Then the complete de-electrification $\widehat{\gamma}$ of γ is a connected bi-infinite quasi-geodesic such that $\beta(\widehat{\gamma}) = \gamma$. Lemma 6.7 allows us to choose a hyperplane in $\widetilde{X}$ so that if $|t| \gg 0$, $\widehat{\gamma}(t)$ does not cross W, $\widehat{\gamma}(\pm t)$ are on opposite sides of W and

$$d(\widehat{\gamma}(t), W) > \max_{D \in \mathcal{D}}\{\text{diam}\, \Sigma_D + 2\delta\}.$$

Therefore, for $|t| \gg 0$, $\gamma(\pm t)$ each lie in distinct components of $\Gamma \setminus \beta(L)$ by Lemma 6.7. For all $S > 0$, $\text{diam}\{t : d(\gamma(t), W) \leq S\} < \infty$ because otherwise a standard hyperbolic geometry argument using a thin quadrilateral shows that $\text{diam}\{t : d(\gamma(t), W) < 2\delta\}$ is not bounded above, which contradicts our choice of W. Hence, $\lim_{t \to \pm\infty} d(\gamma(t), W) = \infty$, and x, y are not in the limit set of the stabilizer of W.

If one of x, y is a conical limit point, assume without loss of generality that x is the conical limit point and y is a parabolic point. Apply the argument from the previous case except use Lemma 6.8 in place of Lemma 6.7 to extract the desired hyperplane W. To see that $y \notin \Lambda \text{Stab}_K(W)$, observe that $W \cap \Sigma_{\text{Stab}_K(x)} = \emptyset$ and apply Proposition 6.1.

If x, y are both distinct parabolic vertices in Γ, let $\widehat{\gamma}$ be a minimal length geodesic in $\widetilde{X}^{(1)}$ between $\Sigma_{\text{Stab}_K(x)}$ and $\Sigma_{\text{Stab}_K(y)}$. Then γ has an edge with finite stabilizer because $\text{Stab}_K(x) \cap \text{Stab}_K(y)$ is finite. Let W be the hyperplane dual to this edge. We see that x, $y \notin \Lambda \text{Stab}_K(W)$ by Lemma 6.5 and Proposition 6.1. ∎

Finally, we prove Theorem 1.4 from the introduction.

Theorem 1.4. *Let $(K, \mathcal{D})$ act relatively geometrically on a CAT(0) cube complex $\widetilde{X}$. If $x, y \in \partial_{\mathcal{D}} K$ and $x \neq y$, then there exist a hyperplane W of $\widetilde{X}$ and a finite index subgroup $K_W \leq \text{Stab}_K(W)$ so that x, y are in K_W-distinct components of $\partial_{\mathcal{D}} K \setminus \Lambda K_W$.*

Proof. Let W be the separating hyperplane specified by Proposition 6.9 with associated hyperset L. Then $\beta(L)$ separates x, y and its carrier $\beta(J)$ is quasiconvex. Since K acts relatively geometrically, K acts cocompactly on $\widetilde{X}$ and on Γ, see Proposition 2.11. In particular, $\widetilde{X}$ is finite-dimensional. Edge stabilizers are finite because each maximal parabolic stabilizes exactly one vertex, distinct maximal parabolics have finite intersection (see, e.g., [4, p. 26]) and all cell stabilizers are parabolic. According to Proposition 6.3, $\beta(J)$ has the two-sided carrier property. Therefore, by Theorem 5.7, there exists a subgroup K_W of index at most 2 in $\text{Stab}_G(W)$ so that x, y are in K_W-distinct components of $\partial_{\mathcal{D}} K \setminus \Lambda K_W$. ∎

References

[1] C. R. Abbott and J. F. Manning, Acylindrically hyperbolic groups and their quasi-isometrically embedded subgroups. *Michigan Math. J.* (2023)

[2] I. Agol, The virtual Haken conjecture (with an appendix by Ian Agol, Daniel Groves, and Jason Manning). *Doc. Math.* **18** (2013), 1045–1087 Zbl 1286.57019 MR 3104553

[3] N. Bergeron and D. T. Wise, A boundary criterion for cubulation. *Amer. J. Math.* **134** (2012), no. 3, 843–859 Zbl 1279.20051 MR 2931226

[4] B. H. Bowditch, Relatively hyperbolic groups. *Internat. J. Algebra Comput.* **22** (2012), no. 3, article no. 1250016 Zbl 1259.20052 MR 2922380

[5] R. Charney and J. Crisp, Relative hyperbolicity and Artin groups. *Geom. Dedicata* **129** (2007), 1–13 Zbl 1161.20032 MR 2353977

[6] F. Dahmani and M. Mj, Height, graded relative hyperbolicity and quasiconvexity. *J. Éc. Polytech. Math.* **4** (2017), 515–556 Zbl 1431.20029 MR 3646028

[7] C. Druţu and M. Sapir, Tree-graded spaces and asymptotic cones of groups (with an appendix by Denis Osin and Mark Sapir). *Topology* **44** (2005), no. 5, 959–1058 Zbl 1101.20025 MR 2153979

[8] E. Einstein and D. Groves, Relative cubulations and groups with a 2-sphere boundary. *Compos. Math.* **156** (2020), no. 4, 862–867 Zbl 1481.20167 MR 4079630

[9] E. Einstein and D. Groves, Relatively geometric actions on CAT(0) cube complexes. *J. Lond. Math. Soc. (2)* **105** (2022), no. 1, 691–708 Zbl 1521.20096 MR 4411337

[10] E. Einstein and T. Ng, Relative cubulation of small cancellation free products. 2021, arXiv:2111.03008

[11] B. Farb, Relatively hyperbolic groups. *Geom. Funct. Anal.* **8** (1998), no. 5, 810–840 Zbl 0985.20027 MR 1650094

[12] R. L. Graham, B. L. Rothschild, and J. H. Spencer, *Ramsey theory*. 2nd edn., Wiley Ser. Discrete Math. Optim., John Wiley & Sons, New York, 1990 Zbl 0455.05002 MR 1044995

[13] M. Gromov, Hyperbolic groups. In *Essays in group theory*, pp. 75–263, Math. Sci. Res. Inst. Publ. 8, Springer, New York, 1987 Zbl 0634.20015 MR 919829

[14] D. Groves and J. F. Manning, Dehn filling in relatively hyperbolic groups. *Israel J. Math.* **168** (2008), 317–429 Zbl 1211.20038 MR 2448064

[15] M. F. Hagen, The simplicial boundary of a CAT(0) cube complex. *Algebr. Geom. Topol.* **13** (2013), no. 3, 1299–1367 Zbl 1267.05302 MR 3071129

[16] G. C. Hruska, Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.* **10** (2010), no. 3, 1807–1856 Zbl 1202.20046 MR 2684983

[17] E. Martínez-Pedroza and D. T. Wise, Relative quasiconvexity using fine hyperbolic graphs. *Algebr. Geom. Topol.* **11** (2011), no. 1, 477–501 Zbl 1229.20038 MR 2783235

[18] D. V. Osin, Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.* **179** (2006), no. 843, vi+100 pp. Zbl 1093.20025 MR 2182268

[19] M. Sageev, Ends of group pairs and non-positively curved cube complexes. *Proc. Lond. Math. Soc. (3)* **71** (1995), no. 3, 585–617 Zbl 0861.20041 MR 1347406

[20] D. Spriano, Hyperbolic HHS I: Factor systems and quasi-convex subgroups. 2018, arXiv:1711.10931

[21] D. T. Wise, *From riches to raags:* 3*-manifolds, right-angled Artin groups, and cubical geometry*. CBMS Reg. Conf. Ser. Math. 117, American Mathematical Society, Providence, RI, 2012 Zbl 1278.20055 MR 2986461

[22] A. Yaman, A topological characterisation of relatively hyperbolic groups. *J. Reine Angew. Math.* **566** (2004), 41–89 Zbl 1043.20020 MR 2039323

Eduard Einstein

Department of Mathematics and Statistics, Swarthmore College, 500 College Avenue, Swarthmore, PA 19081, USA; eeinste1@swarthmore.edu

Daniel Groves

Department of Mathematics, Statistics and Computer Science, University of Illinois at Chicago, 851 S. Morgan Street, 322 Science and Engineering Offices (M/C 249), Chicago, IL 60607-7045, USA; dgroves@uic.edu

Thomas Ng

Faculty of Mathematics, Technion – Israel Institute of Technology, 32000 Haifa, Israel; Department of Mathematics, Brandeis University, 415 South Street, Waltham, MA 02453, USA; thomas.ng.math@gmail.com