HYPERBOLIC MODELS AND STABILITY RECOGNITION FOR COXETER GROUPS

CHRISTOPHER H. CASHEN, MICHELLE CHU, AND JING TAO

Abstract. Given a finite rank Coxeter system, we construct a hyperbolic space by coning off the wide parabolic subcomplexes of its Davis complex. The resulting space is 'stability recognizing', in the sense that stable subspaces of the Davis complex are exactly the quasigeodesically connected subspaces whose images in the coned-off space are quasiisometrically embedded. Equivalently, the coned-off space is 'Morse recognizing': a quasigeodesic in the Davis complex is Morse if and only if it is quasigeodesic in the coned-off space. Furthermore, the action of the Coxeter group on this space gives its largest acylindrically hyperbolic structure.

Contents

Date: August 16, 2026.

2020 *Mathematics Subject Classification.* Primary 20F65; Secondary 20F55, 20F67.

Key words and phrases. Coxeter group, Davis complex, Morse quasigeodesic, Morse recognition, acylindrically hyperbolic group.

This research was supported in part by the Austrian Science Fund (FWF) 10.55776/PAT7799924, NSF DMS-2304920, and NSF DMS-2441034.

1. Introduction

This paper is about the coarse geometry of finitely generated Coxeter groups, as revealed by their wall structure and parabolic subgroup structure. The CAT(0) geometry of a finite rank Coxeter system (W, S) is described by its Davis complex Σ, which admits a piecewise Euclidean, CAT(0) metric d_2. Our main result is the construction of a hyperbolic space that encodes the hyperbolic aspects of Σ:

Theorem 1.1 (Corollary 7.22 and Proposition 7.17)**.** *Let $\Sigma = \Sigma(W, S)$ be the Davis complex of a finite rank Coxeter system. There is a hyperbolic, geodesic metric space $(\widehat{\Sigma}, \hat{d})$ obtained by coning off the wide parabolic subcomplexes of Σ. Moreover, $(\widehat{\Sigma}, \hat{d})$ is "stability recognizing", in the sense that a quasigeodesically connected subspace of Σ is stable if and only if it quasiisometrically embeds into $\widehat{\Sigma}$. Equivalently, $(\widehat{\Sigma}, \hat{d})$ is "Morse recognizing", in the sense that a quasigeodesic $\gamma\colon I \to (\Sigma, d_2)$ is Morse if and only if $\mathrm{Id}_\Sigma \circ \gamma\colon I \to (\widehat{\Sigma}, \hat{d})$ is a quasigeodesic.*

The action of the Coxeter group on this coned-off space satisfies an acylindricity condition, see Section 8.

Theorem 1.2 (Theorem 8.1 and Theorem 8.5)**.** *The action of W on $\widehat{\Sigma}$ is acylindrical. In fact, this action gives the largest acylindrically hyperbolic structure on W.*

See Definition 7.15 and Definition 7.16 for definitions of stable and stability recognizing, and Definition 2.25 and Definition 7.13 for definitions of Morse and Morse recognizing. Being quasigeodesic and being Morse are invariant under quasiisometry, so the theorem is also true for quasigeodesics in the 1–skeleton $\Sigma^{(1)}$ of Σ with respect to its combinatorial metric d_1.

The "wide" condition is discussed in Section 2.7; in the context of Coxeter groups it means a parabolic subgroup of (W, S) that splits as a product of parabolic subgroups such that either both factors are infinite or one of them is irreducible higher rank affine.

The wide parabolics are the "obvious" sources of non-hyperbolicity. Indeed, Moussong's characterization of hyperbolicity in Coxeter groups [62, 61] can be rephrased: W is hyperbolic if and only if it has no wide parabolic subgroups. The main theorem says that if we collapse these known non-hyperbolic regions, the result is hyperbolic. Moreover, the fact that the coned-off space is Morse recognizing certifies that we collapsed the "correct" amount: quasigeodesics that are hyperbolic-like, in the sense of being Morse, survive as quasigeodesics in the coned-off space, while quasigeodesics that are not hyperbolic-like do not.

There are important antecedents for our main theorem for other classes of groups. The prototype is for the mapping class group of a hyperbolic surface. Minsky [60] showed that Teichmüller geodesics are Morse if and only if they avoid the "thin parts" of Teichmüller space, which have approximate product structures [59], making them the obvious sources of non-hyperbolicity. Work of Masur and Minsky [58] shows that the curve graph for the surface is a hyperbolic graph that is quasiisometric to the space obtained by coning off the thin parts of Teichmüller space and to the space obtained from the mapping class group by coning off cosets of finitely many curve stabilizers, one for each orbit of essential simple closed curves. Kent and Leininger [54] proved a result that, in an equivalent formulation established by Durham and Taylor [45], see Section 7.3.2, says that the curve graph is a stable subgroup recognizing space for the mapping class group. Bowditch [23] proved the action of the mapping class group on the curve graph is acylindrical. The fact that this action gives the largest acylindrically hyperbolic structure is one of the key motivating examples of this phenomenon, as introduced by Abbott, Balasubramanya, and Osin [1].

Other settings in which stable subgroup recognition theorems exist include groups hyperbolic relative to peripherals with linear divergence [11], right-angled Artin groups [55], some hierarchically hyperbolic groups [2], and free-by-cyclic groups [57].

Our main theorem overlaps with Abbott, Behrstock, and Durham's result for hierarchically hyperbolic groups [2] in the special case of right-angled Coxeter groups. In fact, the coned-off space $\widehat{\Sigma}$ agrees with their construction of a universal acylindrical hyperbolic space for right-angled Coxeter groups. General Coxeter groups are not known to be hierarchically hyperbolic [3], and this paper is not limited to the cocompactly cubulated setting. We allow irreducible higher rank affine parabolics. These are exactly the obstructions to W acting cocompactly on its Niblo-Reeves cube complex [33], and certain configurations[1] of affine parabolics obstruct the existence of equivariant hierarchical structures [68, 13]. We use Coxeter-specific techniques, in particular the parabolic subgroup structure, the combinatorics of the wall system, and the Tits form on the canonical root system of (W,S) to recover, for arbitrary finitely generated Coxeter groups, conclusions that in other settings are typically obtained from cocompact cubical geometry and its close relatives. This combination is part of what makes the coarse geometry of Coxeter groups interesting: wall systems and parabolic subgroup structures provide strong tools, while the class is broad enough to include examples beyond the reach of existing cocompact cubical and hierarchical methods.

The route to Theorem 1.1 includes several technical steps. First, we characterize Morse geodesics in (Σ, d_2), or, equivalently, in the combinatorial metric $(\Sigma^{(1)}, d_1)$, in terms of wall crossings and interactions with wide parabolic subcomplexes. This will be detailed in Theorem 1.4. In Section 7 we construct a non-geodesic hyperbolic metric d_f on the Niblo-Reeves cube complex X_{NR} in terms of *fine chains*. We pull this metric back to Σ via the canonical inclusion $\Sigma^{(0)} \hookrightarrow X_{\mathrm{NR}}$ and show that the resulting hyperbolic metric is Morse recognizing, Theorem 7.12, and that it is W–equivariantly quasiisometric to $(\widehat{\Sigma}, \hat{d})$, in Theorem 7.21.

A byproduct of the construction is that the metric d_f collapses the difference between the Davis complex and the Niblo-Reeves cube complex, Proposition 7.11, so (X_{NR}, d_f), (Σ, d_f) and $(\widehat{\Sigma}, \hat{d})$ are all quasiisometric to each other. This is an analogue of the Masur-Minsky result that the curve graph is quasiisometric to both the coned-off Teichmüller space and the coned-off mapping class group.

To state the Morse characterization theorem, we need some terminology. The Davis complex Σ of a finite rank Coxeter system (W,S) comes equipped with a system of walls. Each wall is the fixed set of a *reflection*, which is a conjugate of one of the generators S. Each wall is a convex subspace that cuts Σ in two.

An *admissible path* is a path in Σ that is transverse to the wall structure, see Definition 3.1. These include both the CAT(0)–geodesics and combinatorial geodesics in Σ, Lemma 3.2.

A *chain* is a set of disjoint walls $\mathcal{V} = \{V_i\}_{i \in I \subset \mathbb{Z}}$ that are linearly ordered such that $i < j < k$ in I implies V_j separates V_i and V_k in Σ, see Definition 2.6. A chain $\mathcal{V}$ and a path γ are *transverse* if γ is transverse to all the walls of $\mathcal{V}$. Chains $\mathcal{V}$ and $\mathcal{H}$ are transverse if every wall of $\mathcal{V}$ crosses every wall of $\mathcal{H}$. A *grid* is a pair $(\mathcal{V}, \mathcal{H})$ of transverse chains, each of length at least 2.

Definition 1.3. A chain is *fine* if for every pair of consecutive walls V and V' in the chain, every chain $\mathcal{H}$ transverse to both V and V' has $|\mathcal{H}| < 2\ell$, where ℓ is the constant defined in Section 2.6.

Theorem 1.4 (Characterization of the Morse property for admissible quasigeodesics)**.** *Let $\Sigma = \Sigma(W,S)$ be the Davis complex of a finite rank Coxeter system. There are constants $\mathfrak{R}$ and ℓ and a function ρ_0, depending only on (W,S), such that for all λ and ϵ, we have quantitative equivalences between the following statements that apply uniformly to every admissible (λ,ϵ)–quasigeodesic $\gamma\colon I \to \Sigma$. That is to say, the parameters of any one item effectively bound the parameters of the others, independent of the particular choice of γ.*

(1) γ is μ–Morse.
(2) There exist C and $0 < \rho \le \rho_0(\lambda,\epsilon)$ such that there exists L such that for every subinterval $I' \subset I$ of length at least L there exists a chain $\mathcal{V} \pitchfork \gamma(I')$ with $|\mathcal{V}|/|I'| \ge \rho$ such that for every pair V, V' of consecutive walls of $\mathcal{V}$, every chain transverse to $\{V, V'\}$ has length at most C.
(3) There exist $0 < \rho \le \rho_0(\lambda,\epsilon)$ such that there exists L such that for every subinterval $I' \subset I$ of length at least L there exists a fine chain $\mathcal{V} \pitchfork \gamma(I')$ with $|\mathcal{V}|/|I'| \ge \rho$.

---

[1]Specifically, it follows from [13, Corollary 7.11] that if (W,S) is a Coxeter system such that there is a $T \subset S$ with $T^{\perp}$ spherical and T is affine of type $\widetilde{A}_2$, $\widetilde{A}_n$ for $n \ge 4$, $\widetilde{E}_6$, $\widetilde{E}_7$, $\widetilde{E}_8$, $\widetilde{F}_4$, or $\widetilde{G}_2$, then W is not a hierarchically hyperbolic group. It turns out, applying work of Hagen [50], that the irreducible higher rank affine Coxeter groups of types $\tilde{A}_3$, $\tilde{B}_n$, $\tilde{C}_n$, and $\tilde{D}_n$ actually are cocompactly cubulated, even though their actions on their Niblo-Reeves cube complexes are not cocompact.

(4) There is an upper bound on the lengths of subintervals $I' \subset I$ for which there exists a wide parabolic subcomplex Υ with $\gamma(I') \subset \bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$.

(5) There is an upper bound on the lengths of chains $\mathcal{V}$ transverse to γ for which there exists a grid $(\mathcal{V}, \mathcal{H})$ with $|\mathcal{H}| = 2\ell$.

The 1–skeleton of the Davis complex is the Cayley graph of W with respect to S, and edges in Σ are labelled by elements of S. Edge paths then have a label consisting of the product of labels of their successive edges, and this label gives a word in W. Theorem 1.4 specialized to combinatorial geodesics has an additional equivalent condition in terms of labels:

Corollary 1.5 (Characterization of the Morse property for combinatorial geodesics)**.** *Let $\Sigma = \Sigma(W, S)$ be the Davis complex of a finite rank Coxeter system. There are effective quantitative equivalences between the following items, uniformly over all combinatorial geodesics γ in Σ.*

(1) γ is μ–Morse.

(2) There is an upper bound on the lengths of subsegments γ' of γ for which there exists a wide parabolic subcomplex Υ with $\gamma' \subset \Upsilon$.

(3) There is an upper bound on the lengths of subsegments γ' of γ for which there exists $T \subset S$ such that W_T is a wide special subgroup and all edges of γ' have label in T.

The proofs of Theorem 1.4 and Corollary 1.5 occupy Section 6.

Caprace and Fujiwara [29, Proposition 4.5], building on work of Caprace and Haglund [30], characterized "rank-one" elements of W. Their result is essentially the case of Theorem 1.4 in which $w \in W$ acts loxodromically on Σ with axis γ that is Morse. It is not clear how to prove Theorem 1.4 directly from that result. A rank-one geodesic in a CAT(0) space is one that does not bound a half-plane [15]. For a periodic geodesic in a proper CAT(0) space this is equivalent to being strongly contracting [19, Theorem 5.4], but this fact relies on periodicity; arbitrary geodesics, or more generally in our context, admissible paths, may encounter an irregular variety of non-hyperbolic regions, not a single half-plane.

Behrstock and Charney [16] classified rank-one elements in right-angled Artin groups. Using similar techniques, Charney and Sultan [38] gave a characterization of Morse geodesic rays in uniformly locally finite CAT(0) cube complexes that is very similar to (1) $\iff$ (2) of Theorem 1.4, except that they require the chain to be boundedly spaced along γ, not just bounded density. This would apply to Σ in the right-angled case, but not in general. An arbitrary CAT(0) cube complex does not come with a notion of parabolic subcomplexes. Section 4.2 is largely an adaptation of Charney and Sultan's construction to Davis complexes in a way that relates the Morse property to wide parabolics.

Cordes and Levcovitz proved the special case of Corollary 1.5 for Coxeter groups that are "affine-free", meaning they have no irreducible higher rank affine parabolics [40, Theorem D]. They conjectured the result should be true without the affine-free hypothesis. We confirm that conjecture. We reemphasize, however, that the affine-free hypothesis is very strong: it rules out exactly the phenomenon that prevents W from acting geometrically on its Niblo-Reeves cube complex, so affine-free Coxeter groups are still fundamentally cubical. It is precisely the affine parabolics that are usually responsible for difficulties generalizing results from the right-angled case to the general case.

Cordes and Levcovitz have essentially the same divergence construction that we use in Section 4.2, including the use of chain density instead of bounded spacing. We need some additional characterizations of Morseness in terms of the wall structure to interface with the grid constructions in Section 5. Section 4 does not really use the Coxeter structure in an essential way; it is a geometric argument about discrete wall systems in a CAT(0) space.

In Section 5 we use grids, as in [31], to make the translation between wide parabolics and chain conditions. Grids have been used before in CAT(0) cube complexes to detect non-hyperbolicity, see [51, 48], so, again, the outline of the argument will look familiar to cube complex aficionados, but the difficulty is that the tool is not as powerful in Σ as it is in CAT(0) cube complexes. The rough idea is that in a cube complex crossing walls only meet at right angles, so grids naturally correspond to product subcomplexes. That is not true for Coxeter groups. We use results of Caprace [27], recalled in Section 2.5, that are proved using the Tits symmetric bilinear form on the canonical root system for (W, S), to show that large enough grids do correspond to wide parabolics, either affine or product type. This is a key point at which Coxeter-specific technology compensates for the lack of cocompact cubical tools.

The proof of Theorem 1.4 in Section 6 is a formal combination of the results of the two preceding sections. The reinterpretation of this result in terms of edge labels in Corollary 1.5 requires additional work using Coxeter-theoretic facts.

In Section 7 we use a construction of Genevois [48] to define a non-geodesic hyperbolic *fine chain metric* on the Niblo-Reeves cube complex X_{NR} of (W,S), which we then pull back to Σ. Genevois actually defines a parameterized family of hyperbolic metrics on a cube complex, such that quasigeodesics with a common Morse gauge are eventually recognized in a hyperbolic space with large enough parameter. Notably, our result says that for a Coxeter group there is a single critical parameter for which we can already recognize all Morse quasigeodesics, so we only need one concrete hyperbolic metric, not a parameterized family or a limit of such a family.

In Section 7.3 we prove that (Σ, d_f) is Morse and stability recognizing, Theorem 7.12. A consequence is the abstract characterization that a subgroup of W is stable if and only if it is quasiisometrically embedded in (Σ, d_f), Proposition 7.18. For special subgroups we can give more concrete criteria:

Theorem 1.6 (Theorem 7.19)**.** *For $T \subset S$, let Γ_T be the subdiagram of the Coxeter diagram of the finite rank Coxeter system (W,S) spanned by the vertices T. Call a diagram 'spherical' if it defines a spherical Coxeter group.*

The special subgroup W_T is stable if and only if for every wide special subgroup W_U the diagram $\Gamma_{T\cap U}$ is spherical.

The special subgroup W_T is Morse if and only if for every wide special subgroup W_U, one of the following is true:

(1) The intersection of every connected component of Γ_U with Γ_T is spherical.
(2) Every nonspherical connected component of Γ_U is contained in Γ_T.

In the right-angled case, similar characterizations of Morse/stable special subgroups are due to Tran [76, Theorem 1.11], in the 2–dimensional case and Genevois [48, Proposition 4.9] (see also [72, Theorem 7.5]) in arbitrary dimensions.

In Section 7.4 we introduce the coned-off space and show it is equivariantly quasiisometric to Σ with the fine chain metric, thus concluding that it is hyperbolic and Morse recognizing. The payoff for this last step, passing to the coned-off space, is twofold. First, $(\widehat{\Sigma}, \hat{d})$ is a geodesic space, whereas (Σ, d_f) is not. Second, the coned-off space is Coxeter-native: it is built directly from the Coxeter structure. It interfaces with arguments based on walls and parabolic structures. We lean heavily on this interface in the acylindricity argument of Section 8, which takes place entirely in $(\widehat{\Sigma}, \hat{d})$. This perspective is complementary to that of several recent results [69, 70, 77] constructing, in a similar spirit to [48], and in great generality, abstract hyperbolic spaces that detect strongly contracting/rank-one/Morse behavior.

Finally, acylindricity of $W \curvearrowright \widehat{\Sigma}$ is proven in Section 8. The proof is guided by the corresponding argument, [17, Theorem 14.3], for hierarchically hyperbolic spaces (HHS). We do not have the full suite of HHS tools; for example, we do not have hierarchy paths or a two-sided distance formula. Nevertheless, using the geometry of walls and parabolic subgroups in Coxeter groups, we establish analogues of the particular ingredients needed to prove acylindricity, such as subsurface projection, large link condition, bounded geodesic image theorem, one-sided distance formula, and descent by induction on complexity.

2. Preliminaries

Let (W,S) be a finite rank Coxeter system, meaning W is a Coxeter group with a finite set S of fundamental generators. A Coxeter system (W,S) determines a *Coxeter diagram* that encodes a group presentation for W. The vertices of the Coxeter diagram are in bijection with S, whose elements are the generators of W, each of order 2. When s and t are distinct elements of S and st has order $m \in \{2,3,\dots\} \cup \{\infty\}$ in W, then the Coxeter diagram has no edge between s and t when $m=2$ and has an edge labelled m between s and t otherwise. (By convention, the label 3 is usually not written.)

A subset $T \subset S$ determines a *special subgroup* W_T of W generated by T that is isomorphic to the Coxeter group defined by the sub-Coxeter system involving only the generators in T. The special subgroups and their conjugates are called *parabolic subgroups*. If P is a parabolic subgroup conjugate to a special subgroup W_T then $|T|$ is the *rank* of P, and P is called *reducible* if there exists a decomposition $T = T_0 \sqcup T_1$ such that there are no edges in the Coxeter diagram between vertices of T_0 and vertices of T_1. This implies $W_T = W_{T_0} \times W_{T_1}$.

The parabolic is *irreducible* if it is not reducible. Conceivably, P could be conjugate to two different special subgroups W_T and W_U, but when this happens there is a conjugation in W taking the set T bijectively to the set U [41, Proposition 4.5.10]. Since the edge labels of the Coxeter diagram are determined by the orders of products of generators, this further implies that the subdiagram spanned by vertices of T is isomorphic as a Coxeter diagram to the subdiagram spanned by vertices of U. Thus, the rank and reducibility of P are well defined, independent of which special subgroup is taken as a representative of its conjugacy class.

The intersection of two parabolic subgroups is a parabolic subgroup [41, Lemma 5.3.6].

A parabolic subgroup is *spherical* if it is a finite group.

The *Davis complex* $\Sigma = \Sigma(W, S)$ of the Coxeter system is a cell complex whose cells correspond to cosets of spherical special subgroups. It admits a piecewise Euclidean, locally finite, CAT(0) polyhedral structure, constructed as in [41, Section 7.3 and Section 12.1]. We normalize so that edges have length 1 (by taking $\mathbf{d} \equiv 1/2$ in the construction of [41, Section 12.1]). With this normalization, the geometry of Σ is completely determined by the Coxeter system (W, S).

The reader is referred to Davis's book [41] for more background on Coxeter groups, and to Bridson and Haefliger's book [24] for the geometry of CAT(0) metric spaces.

For the remainder of the paper we regard (W, S) and Σ as fixed, so terms referred to as 'constants', e.g. in Section 2.6, are constants with respect to this fixed Coxeter system and choice of metric on Σ. Our results are true, but trivial, when W is finite, virtually cyclic, or wide, so assume not.

2.1. **Background on Coxeter groups.** A *reflection* in (W, S) is a conjugate of an element of S. A *wall* in Σ is the fixed set of a reflection. The unique reflection fixing wall M and exchanging the two components of $\Sigma \setminus M$ is denoted $\mathfrak{r}_M$. Unlike in the right-angled case, general Coxeter groups can have conjugate generators, so the reflection $\mathfrak{r}_M$ is not necessarily conjugate to a unique element of S.

A *reflection subgroup* of W is one that is generated by reflections.

A key fact in this field is that for every pair of distinct walls M and M' in Σ, the group $\langle \mathfrak{r}_M, \mathfrak{r}_{M'} \rangle$ is dihedral, and it is finite if and only if M and M' intersect transversely, $M \pitchfork M'$, and it is isomorphic to the infinite dihedral group $\mathcal{D}_\infty$ if and only if M and M' are disjoint.

Every wall M is convex and cuts Σ into two components. We say M *separates* A and B if A and B are contained in different components of $\Sigma \setminus M$. The components of $\Sigma \setminus M$ are the *open halfspaces* corresponding to M. A *closed halfspace* of M is the union of M with an open halfspace of M.

By construction, walls miss the vertices of Σ and intersect higher dimensional cells in a Euclidean codimension–1 hyperplane. This gives a natural piecewise Euclidean structure on a wall M from its intersections with the cells of Σ. In fact, by the choice $\mathbf{d} \equiv 1/2$, within each cell σ that meets M there is a metric product $\bar{\mathcal{N}}^{\sigma}_{1/2}(M \cap \sigma) := \{x \in \sigma \mid d_2(x, M \cap \sigma) \leq 1/2\} \cong (M \cap \sigma) \times [-1/2, 1/2]$ that is the convex hull in σ of the set of edges of σ that are cut by M. This structure persists globally: $\bar{\mathcal{N}}_{1/2}(M) := \{x \in \Sigma \mid d_2(x, M) \leq 1/2\}$ is isometric to $M \times [-1/2, 1/2]$, is equal to the convex hull of the set of edges cut by M, and $\bar{\mathcal{N}}_{1/2}(M) \cap \sigma = \bar{\mathcal{N}}^{\sigma}_{1/2}(M \cap \sigma)$ for each cell σ. The set $\bar{\mathcal{N}}_{1/2}(M)$ can be used in a role similar to that of a wall/hyperplane *carrier* in a cube complex. The usual definition of a wall carrier as the union of closed cells meeting M does not, in general, result in a convex set in the Davis complex of a non-right-angled Coxeter group.

Similarly, there are no tangential intersections between walls M and subcomplexes Υ of Σ: either they are disjoint or M cuts an edge of Υ. In particular, either $M \cap \Upsilon = \emptyset$ or Υ contains points from both components of $\Sigma \setminus M$.

The connected components of Σ minus the walls are the *open chambers*. Each open chamber contains a unique vertex of Σ, and every vertex is contained in a chamber. By choosing a base vertex $\mathbb{1} \in \Sigma$ we get a bijection between W and vertices of Σ by referring to the vertex $w.\mathbb{1}$ as w. The chamber containing $\mathbb{1}$ is the *fundamental chamber*, whose closure is a fundamental domain for $W \curvearrowright \Sigma$. This choice of basepoint identifies the Cayley graph of W with respect to S with the 1–skeleton of Σ.

The n–skeleton of Σ is denoted $\Sigma^{(n)}$.

It will sometimes be convenient to replace Σ by a subdivision in which walls of Σ are actual subcomplexes. The barycentric subdivision has this property, but that subdivision is finer than necessary. Instead we define an intermediate complex:

Definition 2.1. Let Σ' be the polyhedral subdivision of Σ obtained by subdividing each cell of Σ along its intersections with the walls of Σ. Equivalently, the cells of Σ' are the intersections of cells of Σ with closed chambers.

Thus, Σ' is the coarsest subdivision of Σ in which walls and their intersections are subcomplexes. Its 1–skeleton is the Hasse diagram of the face poset of Σ, i.e. one vertex for every cell of Σ and an edge connecting a cell-vertex with its codimension-1 face-vertices.

With respect to our choice of $\mathbb{1}$, the vertices of Σ corresponding to elements of a special subgroup W_T span a convex subcomplex, denoted Σ_T, isomorphic to the Davis complex of the Coxeter system (W_T, T). Similarly, a coset of a special subgroup wW_T corresponds to a *parabolic subcomplex* $w\Sigma_T$, which is the full subcomplex of Σ spanned by vertices in the coset wW_T. These subcomplexes are closely related to the *residues* used in the buildings literature. The parabolic subcomplex $w\Sigma_T$ corresponds to the residue that is the union of chambers containing the vertices of $w\Sigma_T$. The parabolic subcomplex and the residue have the same stabilizer, which is the parabolic subgroup wW_Tw^{-1}. However, a parabolic subgroup may be the stabilizer of multiple parallel parabolic subcomplexes; we return to this point in Corollary 2.14.

Lemma 2.2. *If* $a\Sigma_U \subset b\Sigma_T$ *then* $\mathrm{Stab}(a\Sigma_U) \subset \mathrm{Stab}(b\Sigma_T)$.

Proof. Since edges of $a\Sigma_U$ are labelled by elements of U and edges of $b\Sigma_T$ are labelled by elements of T, we have $U \subset T$. Passing to vertex sets, $aW_U \subset bW_T$, so $a = bc$ for some $c \in W_T$, so $cW_Uc^{-1} \subset W_T$. Thus:

$$\mathrm{Stab}(a\Sigma_U) = aW_Ua^{-1} = bcW_Uc^{-1}b^{-1} \leq bW_Tb^{-1} = \mathrm{Stab}(b\Sigma_T) \qquad \square$$

An *affine* Coxeter group is one that is virtually a Euclidean reflection group. The irreducible spherical and irreducible affine Coxeter groups are completely classified in terms of their Coxeter diagrams.

Definition 2.3. A Coxeter system is called *irreducible higher rank affine* if it is irreducible and affine of rank at least 3, in the sense that the fundamental generating set has cardinality at least 3. This corresponds to the geometric rank, the dimension of the Euclidean space on which it acts, being at least 2.

Definition 2.4. If $\mathcal{M}$ is a set of walls in Σ, let $W(\mathcal{M}) := \langle \mathfrak{r}_M \mid M \in \mathcal{M}\rangle$ be the subgroup generated by reflections in those walls.

Definition 2.5. For any subset of W there is a unique smallest parabolic subgroup containing it, known as its *parabolic closure*. In particular, when $\mathcal{M}$ is a set of walls of Σ we denote the parabolic closure of the reflections in those walls $\mathrm{Pc}(\mathcal{M}) := \mathrm{Pc}(W(\mathcal{M}))$.

Definition 2.6. A *chain* $\mathcal{K} = \{K_i\}_{i\in I}$ in Σ, indexed by $\emptyset \neq I \subset \mathbb{Z}$, is a collection of disjoint walls K_i, such that $i < j < k$ in I implies K_j separates K_i and K_k. The number of walls in the chain is denoted $|\mathcal{K}|$.

Definition 2.6 is an abuse of terminology. The complementary halfspaces of the walls $K \in \mathcal{K}$ can be labelled K^+ and K^- in such a way that for $i < j$ we have $K_i^- \subset K_j^-$ and $K_i^+ \supset K_j^+$. Then the set of halfspaces K_i^-, ordered consistently with the indices, gives a chain in the partially ordered set of halfspaces with respect to inclusion. Thus, it is really the nested sequence of halfspaces that is the chain. Since the ordered index set is part of Definition 2.6, $\mathcal{K}$ contains the same information as the nested sequence of halfspaces K_i^-, so the terminology is justified.

We say a wall M or path γ is *transverse to a chain* $\mathcal{K}$ if it is transverse to every wall in the chain, and write $M \pitchfork \mathcal{K}$ or $\gamma \pitchfork \mathcal{K}$, respectively. We say *two chains* $\mathcal{V}$ *and* $\mathcal{H}$ *are transverse*, and write $\mathcal{V} \pitchfork \mathcal{H}$, if $V \pitchfork H$ for all $V \in \mathcal{V}$ and all $H \in \mathcal{H}$.

The *extremes* of a chain are the at most 2 walls of the chain that do not separate any of the other walls of the chain, i.e., the first and last walls, if they exist.

Lemma 2.7. *Let* $\mathcal{V} := \{V_i\}_{i\in I}$ *be a chain. Then:*

- $\mathrm{Pc}(\mathcal{V})$ *is irreducible.*
- *There exist* $i_0, i_1 \in I$ *such that* $\mathrm{Pc}(\mathcal{V}) = \mathrm{Pc}(V_{i_0}, V_{i_1})$*, so* $\mathrm{Pc}(\mathcal{V})$ *is a finitely generated reflection subgroup.*

Proof. The first item follows from [31, Lemma 2.2 (ii)], since walls in a chain are disjoint so the corresponding reflections do not commute.

For the second item, [27, Lemma 17] implies that if $i < j < k$ are distinct indices in I then $\mathfrak{r}_{V_j} \in \mathrm{Pc}(\mathfrak{r}_{V_i}, \mathfrak{r}_{V_k})$, so for any finite subinterval $[i_0, i_1] \subset I$ we have $\mathrm{Pc}(\{V_i\}_{i_0\leq i\leq i_1}) = \mathrm{Pc}(\{V_{i_0}, V_{i_1})$. If I is finite we are done. If not, take an exhaustion of I by finite subintervals. This gives a nested sequence of parabolic closures of the corresponding convex subchains. A strict inclusion of parabolic subgroups requires a strict increase in rank, but the ranks of parabolic subgroups are bounded above by $|S|$, so there the sequence of parabolic closures stabilizes at some finite stage. $\square$

2.2. **Parabolic centralizers and normalizers.** There has been extensive work on describing normalizers and centralizers of parabolic subgroups of (W, S) [25, 53, 7, 26, 22, 9, 10, 64], as we now briefly recall. The main takeaway for us is that the complications arise in the reducible or spherical cases.

Lemma 2.8. *A reflection subgroup of a direct product of special subgroups splits as a direct product of reflection subgroups of the factors.*

Proof. Suppose (W, S) is a Coxeter system and r is a reflection in $W_{T_0} \times W_{T_1} \leq W$ for some $T_0, T_1 \subset S$. Then $r = wsw^{-1}$ for some $w \in W_{T_0} \times W_{T_1}$ and $s \in T_0 \sqcup T_1$ [41, Lemma 4.2.3]. Suppose $w = uv$ for $u \in W_{T_0}$ and $v \in W_{T_1}$. If $s \in T_0$ then $r = uvsv^{-1}u^{-1} = usu^{-1} \in W_{T_0}$. If $s \in T_1$ then $r = uvsv^{-1}u^{-1} = vsv^{-1} \in W_{T_1}$. Thus, each reflection is in one factor or the other, and the factors commute, by hypothesis. $\square$

Proposition 2.9. *Let (W, S) be a Coxeter system. Let $T \subset S$ be irreducible and nonspherical. Let $T^\perp := \{s \in S \setminus T \mid \forall t \in T,\ st = ts\}$. Then:*

1. *If $wTw^{-1} \subset S$ for some $w \in W$ then $wTw^{-1} = T$.*
2. *The normalizer of W_T in W is $N_W(W_T) = W_T \times W_{T^\perp}$.*
3. *The centralizer of W_T in W is $Z_W(W_T) = W_{T^\perp}$.*
4. *Every reflection in $N_W(W_T)$ is either in W_T or in $W_{T^\perp}$.*

Proof. These results are well known, see [32, Lemma 3.3], [27, Lemma 22 and the preceding line], [28, Lemma 2.2], [40, Theorem 5.7]. They are usually attributed to Deodhar [42, Proposition 5.5], with further explanation in [56, Section 3.1]. Deodhar's result, and the text prior to it, generalizes some root system facts used by Howlett in the description of parabolic normalizers for finite Coxeter groups [53], which would later be used by Brink and Howlett [26] to describe parabolic normalizers for infinite Coxeter groups.

The proposition is the irreducible nonspherical special case of the parabolic normalizer theorem obtained from Howlett's argument, together with Deodhar's refinements of the elementary conjugating elements. The final assertion follows from Lemma 2.8. $\square$

Lemma 2.10 ([41, Proposition 4.5.10])**.** *If $T, U \subset S$ and $wW_Tw^{-1} = W_U$ then there is an element $u \in W_U$ such that for $v := u^{-1}w$ we have $vTv^{-1} = U$.*

Corollary 2.11. *If P is an irreducible nonspherical parabolic subgroup of (W, S), then there is a unique $T \subset S$ such that P is conjugate to W_T.*

Proof. If P is conjugate to W_U and W_T then Lemma 2.10 says $vTv^{-1} = U$. Then since P is irreducible and nonspherical, Proposition 2.9 (1) says $T = U$. $\square$

The following corollary of Proposition 2.9 establishes the definition of the *orthogonal factor* parabolic $P^\perp$ of an irreducible nonspherical parabolic P.

Corollary 2.12. *If P is an irreducible nonspherical parabolic subgroup of (W, S), then $N_W(P)$ is parabolic and there is a canonical parabolic subgroup $P^\perp$ such that $N_W(P) = P \times P^\perp$.*

Furthermore, if $R \subset N_W(P)$ is a set of reflections then $\langle R \rangle = \langle R \cap P \rangle \times \langle R \cap P^\perp \rangle$.

Proof. By definition, there exist an irreducible, nonspherical $T \subset S$ and $w \in W$ such that $wPw^{-1} = W_T$. Proposition 2.9 says $wN_W(P)w^{-1} = N_W(W_T) = W_T \times W_{T^\perp}$ and $Z_W(W_T) = W_{T^\perp}$. Take $P^\perp := Z_W(P) = w^{-1}W_{T^\perp}w$. Then $N_W(P) = P \times P^\perp$.

The direct product splitting of $\langle R \rangle$ follows from the claim of Proposition 2.9 that reflections in $N_W(W_T)$ are contained in either W_T or $W_{T^\perp}$. $\square$

Remark. The construction of Corollary 2.12 is symmetric if $T^\perp$ is also irreducible and nonspherical, but that is not always the case. In general, $W_T \times W_{T^\perp}$ can be a proper subgroup of $N_W(W_{T^\perp})$. In the most extreme case, $T^\perp$ could be empty, with $W_{T^\perp}$ trivial and $N_W(W_{T^\perp}) = W$.

Corollary 2.13. *If P is an irreducible nonspherical parabolic that stabilizes parabolic subcomplexes Υ_0 and Υ_1 then there is a unique $T \subset S$ and there are elements $w_0, w_1 \in W$ with $w_0^{-1}w_1 \in N_W(W_T)$ such that $\Upsilon_0 = w_0\Sigma_T$ and $\Upsilon_1 = w_1\Sigma_T$. Furthermore, $\Upsilon_0^{(0)}w_0^{-1}w_1 = \Upsilon_1^{(0)}$.*

Proof. Since P is irreducible and nonspherical, Proposition 2.9 (1) implies there is a unique $T \subset S$ such that P is conjugate to W_T. Suppose $P = wW_Tw^{-1}$. Then $\Upsilon_0 = w_0\Sigma_T$ and $\Upsilon_1 = w_1\Sigma_T$ for some $w_0, w_1 \in W$. Since P stabilizes both of them, it follows that $w^{-1}w_0$ and $w^{-1}w_1$ are both in $N_W(W_T)$, so $w_0^{-1}w_1 = (w^{-1}w_0)^{-1}w^{-1}w_1$ is too. Thus:

$$\Upsilon_0^{(0)}w_0^{-1}w_1 = w_1w_1^{-1}w_0W_Tw_0^{-1}w_1 = w_1W_T = \Upsilon_1^{(0)}$$ □

Corollary 2.14. *If P is an irreducible nonspherical parabolic then there is a unique parabolic subcomplex whose stabilizer is $N_W(P)$.*

Proof. Suppose $P = xW_Tx^{-1}$ for some $x \in W$ and $T \subset S$. Proposition 2.9 (2) and Corollary 2.12 imply that $N_W(P) = x(W_T \times W_{T^\perp})x^{-1} = xW_{T\sqcup T^\perp}x^{-1}$. The parabolic subcomplex $x\Sigma_{T\sqcup T^\perp}$ has stabilizer $N_W(P)$. Suppose $w\Sigma_U$ for some $w \in W$ and $U \subset S$ is another parabolic subcomplex with stabilizer equal to $N_W(P)$. Its stabilizer is wW_Uw^{-1}, so $W_U = w^{-1}xW_{T\sqcup T^\perp}x^{-1}w$. By Lemma 2.10, there exists $u \in W_U$ such that for $v := u^{-1}w^{-1}x$ we have $v(T \sqcup T^\perp)v^{-1} = U$. Let U' be the irreducible component of U such that $vTv^{-1} = U'$. Since T is irreducible and nonspherical, Proposition 2.9 (1) implies $T = U'$, so $v \in N_W(W_T) = W_{T\sqcup T^\perp}$, which implies $W_{T\sqcup T^\perp}v^{-1} = W_{T\sqcup T^\perp}$. Then:

$$wW_U = xv^{-1}u^{-1}W_U = xv^{-1}W_U = xv^{-1}(vW_{T\sqcup T^\perp}v^{-1}) = xW_{T\sqcup T^\perp}$$

Thus, $x\Sigma_{T\sqcup T^\perp}$ and $w\Sigma_U$ are the same parabolic subcomplex. □

2.3. **Coarse geometry.** Let (X, d) be a geodesic metric space. For $x \in X$ and $Y \subset X$, $d(x, Y) := \inf_{y\in Y} d(x, y)$. For $R \in \mathbb{R}$, the *R–neighborhood* of Y is $\mathcal{N}_R(Y) := \{x \in X \mid d(x, Y) < R\}$, and the *closed R–neighborhood* is $\bar{\mathcal{N}}_R(Y) := \{x \in X \mid d(x, Y) \leq R\}$.

Subsets Y and Z of X are *coarsely equivalent* if they have finite Hausdorff distance, i.e. $\inf\{R \mid Y \subset \bar{\mathcal{N}}_R(Z) \text{ and } Z \subset \bar{\mathcal{N}}_R(Y)\}$ is finite.

For functions ϕ and ψ from X^n to $\mathbb{R}$, write $\phi \preceq \psi$ if there exist λ and ϵ such that for all $\underline{x} \in X^n$ we have $\phi(\underline{x}) \leq \lambda\psi(\underline{x}) + \epsilon$. Write $\phi \overset{+}{\preceq} \psi$ if $\phi \preceq \psi$ with $\lambda = 1$. Write $\phi \asymp \psi$ if $\phi \preceq \psi$ and $\psi \preceq \phi$, and write $\phi \overset{+}{\asymp} \psi$ if $\phi \overset{+}{\preceq} \psi$ and $\psi \overset{+}{\preceq} \phi$.

A *quasiisometric embedding* is a map $\phi\colon X \to Y$ such that for $\phi^*(d_Y)(x, x') := d_Y(\phi(x), \phi(x'))$ we have $d_X \asymp \phi^*(d_Y)$ as maps on X^2. It is a *quasiisometry* if it is a quasiisometric embedding and $\phi(X)$ is coarsely equivalent to Y. It is a *biLipschitz embedding* if it is a quasiisometric embedding with 0 additive error.

A *geodesic* is an isometric embedding of an interval of $\mathbb{R}$, which is a segment, ray, etc according to the corresponding property of the domain. Similarly, a *quasigeodesic* is a quasiisometric embedding of an interval of $\mathbb{R}$, and a *biLipschitz* path is a path that is a biLipschitz embedding of an interval of $\mathbb{R}$.

2.4. **Distances.** We have several different measures of separation in Σ:

- $d_1(x, y)$ is the length of the shortest edge path between vertices x and y.
- $d_2(x, y)$ is the CAT(0) distance between x and y.
- $d_w(x, y)$ is the number of walls separating vertices x and y (which is exactly $d_1(x, y)$).
- $d_c(x, y)$ is the length of the longest chain of walls separating x and y.

Of these, only d_2 is a metric on Σ. The combinatorial metric d_1 extends to a geodesic metric on the 1–skeleton of Σ. The wall and chain measures d_w and d_c are metrics on the 0–skeleton. There are ways to extend them to metrics on all of Σ (eg Lemma 7.8) but we will not use this fact. For us it will be enough to take coarse extensions, in the sense that if x and y are arbitrary points of Σ, then for $d \in \{d_1, d_w, d_c\}$ we let $d(x, y)$ be the minimum of $d(x', y')$ over vertices x' and y' such that x' is contained in a common closed chamber with x and y' is contained in a common closed chamber with y.

Niblo and Reeves [63] constructed a finite dimensional, locally finite CAT(0) cube complex corresponding to the Coxeter system (W, S). Its wall structure is isomorphic to that of Σ, and after choosing basepoints the orbit map gives a W–equivariant isometric embedding $(\Sigma^{(0)}, d_w) \to (X_{\mathrm{NR}}^{(0)}, d_w)$.

Lemma 2.15. *$d_2 \asymp d_1 \asymp d_w \asymp d_c$ as maps $\Sigma^2 \to \mathbb{R}$.*

Proof. There are only finitely many isometry types of cells of Σ, so every point is within uniformly bounded distance of a vertex. Moreover, any closest vertex to a point x is not separated from x by any wall. Thus, up to increasing the additive error afterwards, it suffices to prove the lemma on $\Sigma^{(0)}$.

W acts geometrically on Σ, freely and transitively on $\Sigma^{(0)}$, and $\Sigma^{(1)}$ is precisely the Cayley graph of W with respect to S, so d_1 and d_2 on $\Sigma^{(0)}$ are quasiisometric to the word metric on W with respect to S.

On $\Sigma^{(0)}$ $d_1 = d_w$; in one direction every edge is cut by a unique wall, so the number of walls separating vertices x and y is at most the length of a combinatorial geodesic between them. In the other direction, the Exchange Condition for reflection systems [41, Section 3.2] implies that a combinatorial geodesic is cut at most once by any given wall.

Since the wall structure of Σ is isomorphic to that of X_{NR}, the relation between d_w and d_c can be computed there. It is immediate from the definitions that $d_c \leq d_w$. Conversely, for any w_1 and w_2 in W, consider the set of walls separating w_1 and w_2. If H is such a wall, let H^- be the halfspace with $w_1 \in H^-$ and $\partial H^- = H$, and let H^+ be its complementary halfspace. Consider the set of halfspaces H^- partially ordered by inclusion. Since they all contain w_1, two such halfspaces H_1^- and H_2^- are incomparable if and only if H_1 and H_2 cross. Therefore, antichains in the partially ordered set correspond to collections of pairwise crossing walls. In X_{NR} such a collection has a common point of intersection, so the size of such a set is bounded by $\dim(X_{\mathrm{NR}})$. Dilworth's Theorem says that in a partially ordered set the maximum size of an antichain is equal to the minimum size of a partition of the set into chains. At least one of these has at least average size, so there is some chain separating w_1 from w_2 containing at least $d_w(w_1, w_2)/\dim(X_{\mathrm{NR}})$–many walls. Thus, $d_c \leq d_w \leq \dim(X_{\mathrm{NR}})d_c$. □

We can extend Lemma 2.15 to a point-parabolic subcomplex distance comparison:

Lemma 2.16. *The following functions on $W = \Sigma^{(0)}$ are coarsely equivalent, with constants independent of $T \subset S$:*

$$w \mapsto \begin{cases} d_1(w, \Sigma_T) \\ d_2(w, \Sigma_T) \\ d_w(w, \Sigma_T) \\ d_c(w, \Sigma_T) \end{cases}$$

Proof. Let $w \in W = \Sigma^{(0)}$. By [41, Lemma 4.3.3] there is a unique decomposition $w = vu$ where $v \in W_J$ and u is the unique minimal length representative of $W_J w$. Thus, v is the closest vertex of W_J to w. Take a d_1–geodesic γ from w to v. A wall M separating w from v cuts the interior of γ. Split γ as a concatenation $\gamma' + \gamma''$ of nontrivial subpaths meeting at $M \cap \gamma$. Then $\gamma' + \mathfrak{r}_M(\gamma'')$ tightens to an edge path from w to $\mathfrak{r}_M(v)$ of length at most $|\gamma| - 1$, with v and $\mathfrak{r}_M(v)$ on opposite sides of M. If M separates w from v but does not separate w from Σ_T then it cuts Σ_T. However, the reflections whose walls cut Σ_T are precisely the reflections in W_J, which stabilize Σ_T. This gives a contradiction, because then $\mathfrak{r}_M(v)$ is a vertex of Σ_T that is closer to w than v. Thus, the walls separating w from Σ_T are precisely the walls separating w from v. Now apply Lemma 2.15 for distances between w and v. This accounts for the wall-based distances in the statement. The CAT(0) distance follows because every point in Σ_T is within uniformly bounded distance of a vertex, so $d_2(w, \Sigma_T) \stackrel{+}{\asymp} d_2(w, \Sigma_T^{(0)})$. □

We will not use the following result in this paper, but it is worth mentioning that in addition to point-point and point-parabolic subcomplex coarse equivalences $d_1 \asymp d_2 \asymp d_c$, there are also point-wall and wall-wall versions.

Theorem 2.17 (Strong Parallel Walls Theorem, [27, Theorem C] and [34, Theorem 2.3])**.** *Let $\Sigma = \Sigma(W, S)$ be the Davis complex of a finite rank Coxeter system. Let M and M' be disjoint walls in Σ, and let x be a point in Σ.*

(1) $d_c(M, x) \asymp d_1(M, x) \asymp d_2(M, x)$
(2) $d_c(M, M') \asymp d_1(M, M') \asymp d_2(M, M')$

The constants of the equivalences are independent of M, M', and x.

We remark that the references only state that these distances are related by some function, but it is straightforward to check that the proof actually produces a linear function.

In contrast to Lemma 2.16 and Theorem 2.17, distances to arbitrary CAT(0)-convex sets are not approximated by wall distances. Here are two explicit examples:

Example 2.18. Consider the standard square tiling of $\mathbb{R}^2$.

First, let L be a line that is not vertical or horizontal. It is CAT(0)–convex, but it crosses every vertical and horizontal line, so the intersection of halfspaces containing it is an empty intersection, hence is the whole plane. There are points of the plane arbitrarily far from L, but they are not separated from L by any wall.

Second, for each $n \in \mathbb{N}$, consider $v := (n, n)$ and the closed ball B of radius n about $(0, 0)$. The intersection of halfspaces containing B is the square bounded by vertical and horizontal lines at distance $n + 1/2$ from $(0, 0)$. This set contains v, so no wall separates v and B, but $d_2(v, B) = (\sqrt{2} - 1)n$.

2.5. **Caprace toolbox.** We state some results of Caprace that we will use.

Theorem 2.19 ([27, Theorem 8])**.** *There is a constant N such that the following holds. Suppose $\mathcal{K} = \{K_i\}_{0\leq i\leq n}$ is a chain and M and M' are walls that are each transverse to $\mathcal{K}$ and such that $\emptyset \neq M \cap M' \subset K_0$. If $n > N$ then $W(\mathcal{K} \cup \{M, M'\})$ is a Euclidean triangle group that is contained in a higher rank irreducible affine parabolic.*

Lemma 2.20 ([27, Lemma 11])**.** *Let $\mathcal{M}$ be a set of walls with $W(\mathcal{M}) \cong \mathcal{D}_\infty$. If a wall M is transverse to at least 8 walls of $\mathcal{M}$ then it is transverse to all of them and either $\mathfrak{r}_M$ centralizes $W(\mathcal{M})$ or $W(\mathcal{M} \cup \{M\})$ is a Euclidean triangle group.*

Proposition 2.21 ([27, Proposition 16])**.** *If R is a higher rank irreducible affine reflection subgroup then so is its parabolic closure.*

Lemma 2.22 ([27, Lemma 17])**.** *Let $\mathcal{K}$ be a chain containing at least 2 walls. Let M be a wall. If one of the following is true then $\mathfrak{r}_M \in \mathrm{Pc}(\mathcal{K})$.*

- *M separates two walls of $\mathcal{K}$.*
- *$W(\mathcal{K} \cup M) \cong \mathcal{D}_\infty$*
- *$W(\mathcal{K} \cup M)$ is a Euclidean triangle subgroup.*

2.6. **Constants.** We fix some constants of the Coxeter system (W, S) and the choice of Σ.

Let $\ell := \max\{N + 2, 8\}$, where N is the constant of Theorem 2.19 (and 8 is the constant of Lemma 2.20).

Take $\mathfrak{R}$ large enough that for all $x \in \Sigma$ and Υ a parabolic subcomplex, Lemma 2.16 implies:

$$d_2(x, \Upsilon) \geq \mathfrak{R} \implies d_c(x, \Upsilon) \geq 4\ell - 1$$

Suppose that Lemma 2.15 gives Λ and E such that for all $x, y \in \Sigma$:

$$d_2(x, y)/\Lambda - E \leq d_c(x, y) \leq \Lambda d_2(x, y) + E$$

Given $\lambda \geq 1$ and $\epsilon \geq 0$, define $\rho_0(\lambda, \epsilon) := \frac{1}{2\lambda\Lambda}$ and $\mathfrak{L}(\lambda, \epsilon) := 2\lambda(\Lambda E + \epsilon)$.

If $\gamma \colon I \to \Sigma$ is a (λ, ϵ)–quasigeodesic and $[x, y] \subset I$ then:

$$|y - x| \geq \mathfrak{L}(\lambda, \epsilon) \implies \frac{d_c(\gamma(x), \gamma(y))}{|y - x|} \geq \rho_0(\lambda, \epsilon)$$

So if $|y - x|$ is large enough then a chain realizing $d_c(\gamma(x), \gamma(y))$ has length at least $|y - x|\rho_0(\lambda, \epsilon)$. In particular, there do exist chains separating $\gamma(x)$ and $\gamma(y)$ whose *density*, in terms of number of walls per unit length in the parameter space, is at least $\rho_0(\lambda, \epsilon)$.

Remark. All of these constants can be effectively bounded. The point is that there are only finitely many conjugacy classes of spherical subgroups in W, and they can be enumerated in terms of the known types of irreducible spherical Coxeter diagrams appearing as full subgraphs in the Coxeter diagram for (W, S). Caprace's constant from Theorem 2.19 and a bound on $\dim(X_{\mathrm{NR}})$ [63, Lemma 3] are expressed in terms of the values of the Tits form on subsystems of the canonical root system corresponding to spherical subgroups. Since the Tits form has an explicit closed form, these quantities are computable. Likewise, the cells of the Davis complex are determined by the orthogonal representations of the spherical subgroups, which can be explicitly computed, so there are only finitely many isometry types of cells in Σ and they are isometric to convex Euclidean polytopes that are described by linear algebra. Effective bounds on the distance relations in Lemma 2.15 and Lemma 2.16 can be extracted from the geometry of cells of Σ and the bound on $\dim(X_{\mathrm{NR}})$.

2.7. **Wide groups.** A group is *constricted* if all of its asymptotic cones have cut points or *unconstricted* otherwise [44]. A group is *wide* if none of its asymptotic cones have cut points [43], which is a strictly stronger condition than being unconstricted. For finitely generated groups, the existence of a biinfinite Morse geodesic implies that the group is constricted.

For Coxeter groups there is a nice wide/constricted dichotomy visible from the Coxeter diagram:

Proposition 2.23. *Let $W = P_1 \times \cdots P_n \times K$ be the canonical product decomposition of (W, S) where the P_i are special subgroups corresponding to nonspherical connected components of the Coxeter diagram, and K is the maximal spherical special subgroup. Suppose $W \neq K$. Then W is wide if and only if either $n \geq 2$ or $n = 1$ and P_1 is higher rank affine.*

Proof. If $n \geq 2$ or $n = 1$ and P_1 is higher rank affine then W is virtually a product of infinite groups, in which case it is easy to argue that it is wide. Otherwise, Caprace and Fujiwara [29, Proposition 4.5] say that W contains a rank-one element (equivalently, a Morse element), so it is constricted. □

Definition 2.24 (cf [40, Definition 3.1] and the definition $\mathbb{T}_0$ of [18, Section A.1])**.** A *wide parabolic subgroup* P of (W, S) is a nonspherical parabolic that is wide as a Coxeter group, in the sense of Proposition 2.23, which occurs exactly if either $P = B \times C$, where B and C are nonspherical parabolics of (W, S), or $P = A \times K$, where A is an irreducible higher rank affine parabolic of (W, S) and K is a spherical parabolic of (W, S).

A *wide parabolic subcomplex* is a parabolic subcomplex whose stabilizer is a wide parabolic subgroup.

2.8. **Variations of the Morse property.** Let Z be a closed subset of a proper geodesic metric space X. We define several properties that describe quantitatively how Z sits inside X in a manner similar to a quasiconvex subspace of a hyperbolic space. These properties and equivalences between them were first stated for the case of Z being a geodesic or quasigeodesic [19, 8, 75, 74], see also the survey [35].

Definition 2.25. A subset Z of a geodesic metric space X is *Morse* or has the *Morse property* if it is quasigeodesically quasiconvex, in the sense there exists a real valued function $\mu(\lambda, \epsilon) := \sup_\gamma \sup_{x \in \gamma} d(x, Z)$, where the first supremum is taken over (λ, ϵ)–quasigeodesic segments γ with both endpoints on Z. A function bounding μ from above is known as a *Morse gauge* for Z; the given μ is the *optimal Morse gauge.*

The content of the definition is whether or not μ defines a real valued function; if Z is not Morse then there is some input pair (λ, ϵ) on which μ takes value ∞.

The next definitions makes sense using the *set valued* closest point projection:

$$\pi_Z \colon X \to Z \colon x \mapsto \{z \in Z \mid d(x, z) = d(x, Z)\}$$

When X is proper and Z is closed then point preimages are nonempty, but a priori they could have arbitrarily large diameters. If Z is convex and X is CAT(0) then point images of π_Z are single points, and, in fact, $\pi_Z \colon X \to Z$ is 1–Lipschitz.

Definition 2.26. A closed subset Z of a proper geodesic metric space X is *χ–strongly contracting* if $\chi \geq \sup_{x,y} \operatorname{diam}(\pi_Z(x) \cup \pi_Z(y))$, where the supremum is taken over points x and y such that $d(x, y) \leq d(x, Z)$ and the *diameter* of $Y \subset X$ is $\operatorname{diam}(Y) := \sup_{y,y' \in Y} d(y, y')$.

Definition 2.27 (Divergence)**.** For a closed subset Z of a proper geodesic metric space X and parameters $0 < M \leq 1$ and $N > 2$, define the *optimal divergence gauge* $\delta(r; M, N) := \inf_{x,y} \bar{d}_{Mr}(x, y)$, where the infimum is taken over all pairs x and y at distance r from Z with $d(x, y) \geq Nr$. The modified distance $\bar{d}_{Mr}(x, y)$ denotes the length of the shortest path in X connecting x to y that stays outside the Mr–neighborhood of Z, or ∞ if no such path exists. Any lower bound on $\delta(r; M, N)$ is a *divergence gauge* for Z with respect to parameters M and N.

This notion of divergence appears in [35]. It is the same, up to a suitable equivalence relation on functions, to other familiar versions of divergence, and the equivalence class of the function does not depend on the choices of parameters M and N [35, Proposition 8.4]. The equivalence relation distinguishes functions of different polynomial degrees, so phrases such as "Z has at least quadratic divergence" make sense, and are established by showing there is some valid choice of M and N such that Z has a divergence gauge with respect to M and N that is bounded below by a quadratic function of r.

By convention, the infimum of the empty set of real numbers is ∞, so if for some r there are no suitable pairs (x, y) in Definition 2.27 then $\delta(r; M, N) = \infty$. In our situation, where W is infinite and not virtually

$\mathbb{Z}$, $X=\Sigma$, and Z is a quasigeodesic, suitable pairs (x,y) do exist once the domain of the quasigeodesic is long enough with respect to r and N and the quasigeodesic constants. This is the expected behavior: every closed bounded set certainly has some Morse /contraction/divergence gauge that can be defined in terms of its diameter. Sometimes those bounds are essentially best possible. For instance, a long geodesic segment in a Euclidean plane does not display any hyperbolic-like behavior on scales smaller than its length. In contrast, long geodesic segments in the hyperbolic plane have Morse/contracting/divergence behavior at small scales controlled by the ambient hyperbolicity, not just by their length.

Definition 2.28 (BGIP)**.** A closed subset Z of a proper geodesic metric space has the *Bounded Geodesic Image Property* if there exists C such that for every geodesic γ, if $\gamma\cap\mathcal{N}_C(Z)=\emptyset$ then $\operatorname{diam}(\pi_Z(\gamma))\leq C$.

Definition 2.29. A closed subset Z of a proper geodesic metric space is *strongly constricting* if there exists C such that for every geodesic segment γ with endpoints x and y, if $\operatorname{diam}(\pi_Z(x)\cup\pi_Z(y))>C$ then γ passes within distance C of both $\pi_Z(x)$ and $\pi_Z(y)$.

Proposition 2.30 ([36], [76, Proposition 3.1], [35, Proposition 8.15, Corollary 3.6])**.** *When X is CAT(0), Morse, strongly contracting, and at least quadratic divergence are equivalent.*

Proposition 2.31 ([12, Proposition 2.9])**.** *Strong contraction, strong constriction, and the Bounded Geodesic Image Property are equivalent.*

2.9. **Working in non-geodesic spaces.** To work in the setting of hyperbolic spaces that are not necessarily geodesic, Blachère, Haïssinsky, and Mathieu [21] introduced the following notion:

Definition 2.32. A map $\gamma\colon I\to X$ from a subinterval of $\mathbb{R}$ to a metric space X is a *(parameterized)* (λ,ϵ,τ)*–quasiruler* if it is a (λ,ϵ)–quasigeodesic and the triangle inequality is almost degenerate, in the sense that for all $r<s<t$ in I:

$$d(\gamma(r),\gamma(s))+d(\gamma(s),\gamma(t))-\tau\leq d(\gamma(r),\gamma(t))\leq d(\gamma(r),\gamma(s))+d(\gamma(s),\gamma(t))$$

The space X is *quasiruled* if there exist λ, ϵ, and τ such that every pair of points in X is the pair of endpoints of a (λ,ϵ,τ)–quasiruler.

Abbott and Incerti-Medici [4] make the following refinement:

Definition 2.33. A map $\gamma\colon I\to X$ from a subinterval of $\mathbb{R}$ to a metric space X is a *unparameterized* τ*–quasiruler* if it satisfies two conditions:
bounded jumps: $\forall t\in I$, $\limsup_{s\in I,\,|t-s|\to 0} d(\gamma(s),\gamma(t))\leq\tau$
almost degeneracy: $\forall r<s<t\in I$, $d(\gamma(r),\gamma(s))+d(\gamma(s),\gamma(t))-\tau\leq d(\gamma(r),\gamma(t))$

There is another definition in the literature that is related:

Definition 2.34. A $(1,\epsilon)$–quasigeodesic is also known as an ϵ*–rough geodesic*, and a space for which there exists ϵ such that every pair of points is connected by an ϵ–rough geodesic is a *rough geodesic space.*

Abbott and Incerti-Medici prove [4, Lemma A.4] that given τ there exist λ and ϵ such that every unparameterized τ–quasiruler satisfying an additional mild condition can be reparameterized to make it a (λ,ϵ,τ)–quasiruler. They also prove [4, Lemma A.6] that quasiruled spaces are roughly geodesic, although a priori the transitive family of rough geodesics might not just be reparameterizations of the quasirulers. In Section 7 we introduce a metric on Σ and X_{NR} such that combinatorial geodesics in X_{NR} can be reparameterized to be rough geodesic quasirulers, see Corollary 7.5.

2.10. **Gate projections to parabolic subcomplexes.**

Lemma 2.35. *Given a parabolic subcomplex $\Upsilon:=w\Sigma_T$ and a vertex $x\in\Sigma^{(0)}=W$ there exists a unique d_1–closest vertex $\mathfrak{p}_\Upsilon(x)$ of Υ to x. Furthermore, for every vertex $y\in\Upsilon^{(0)}$:*

$$d_1(x,y)=d_1(x,\mathfrak{p}_\Upsilon(x))+d_1(\mathfrak{p}_\Upsilon(x),y)$$

The map $\mathfrak{p}_\Upsilon\colon\Sigma^{(0)}\to\Upsilon^{(0)}$ is the gate projection *to Υ. It is Lipschitz.*

Proof. By [41, Lemma 4.3.1], there is a unique shortest element v of $W_\emptyset x^{-1}wW_T$, and there exists $z\in W_T$ such that $x^{-1}y=vz$ with $|x^{-1}y|_S=|v|_S+|z|_S$. Take $\mathfrak{p}_\Upsilon(x):=xv$. Then $d_1(x,y)=|x^{-1}y|_S=|v|_S+|z|_S=d_1(x,\mathfrak{p}_\Upsilon(x))+d_1(\mathfrak{p}_\Upsilon(x),y)$.

The fact that $\mathfrak{p}_\Upsilon$ is Lipschitz is a standard, easy computation. □

Corollary 2.36. *If $U \subset T \subset S$ then $\mathfrak{p}_{\Sigma_U} = \mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}$.*

Proof.

$$\begin{aligned}
&d_1(x, \mathfrak{p}_{\Sigma_U}(x)) \\
&\quad = d_1(x, \mathfrak{p}_{\Sigma_T}(x)) + d_1(\mathfrak{p}_{\Sigma_T}(x), \mathfrak{p}_{\Sigma_U}(x)) \\
&\quad = d_1(x, \mathfrak{p}_{\Sigma_T}(x)) + d_1(\mathfrak{p}_{\Sigma_T}(x), \mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}(x)) + d_1(\mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}(x), \mathfrak{p}_{\Sigma_U}(x)) \\
&\quad \geq d_1(x, \mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}(x)) + d_1(\mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}(x), \mathfrak{p}_{\Sigma_U}(x)) \\
&\quad \geq d_1(x, \mathfrak{p}_{\Sigma_U}(x)) + d_1(\mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}(x), \mathfrak{p}_{\Sigma_U}(x))
\end{aligned}$$

Thus, $d_1(\mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}(x), \mathfrak{p}_{\Sigma_U}(x)) = 0$. ◻

Corollary 2.37. *A wall M separates x from $\mathfrak{p}_\Upsilon(x)$ if and only if M separates x from Υ.*

Proof. A wall that separates x from all of Υ certainly separates x from $\mathfrak{p}_\Upsilon(x) \in \Upsilon$. For the converse, suppose M is a wall that separates x from $y := \mathfrak{p}_\Upsilon(x)$ but not from Υ, so there is a vertex $z \in \Upsilon$ such that x and z are on one side of M, and y is on the other. Then the concatenation of the d_1–geodesic from x to y with the d_1–geodesic from y to z crosses M twice, so it is not a combinatorial geodesic, so $d_1(x, z) < d_1(x, y) + d_1(y, z)$. This contradicts the gate property $d_1(x, z) = d_1(x, y) + d_1(y, z)$. ◻

The following interaction between gate maps and the group action follows directly from the definition.

Corollary 2.38. *For $w \in W$, Υ a parabolic subcomplex, and $x \in \Sigma^{(0)}$:*

$$\mathfrak{p}_{w\Upsilon}(x) = w\mathfrak{p}_\Upsilon(w^{-1}x)$$

Lemma 2.39. *For any $U \subset S$ and $g \in U^\perp$, for all $x \in \Sigma^{(0)}$:*

$$\mathfrak{p}_{\Sigma_U}(gx) = \mathfrak{p}_{\Sigma_U}(x)$$

Proof. Let $T := U \sqcup U^\perp$. Write $\mathfrak{p}_{\Sigma_T}(x) = ab$ for $a \in W_U$ and $b \in W_{U^\perp}$. Using that $g\Sigma_T = \Sigma_T$ and Corollary 2.38, $\mathfrak{p}_{\Sigma_T}(gx) = \mathfrak{p}_{g\Sigma_T}(gx) = g\mathfrak{p}_{\Sigma_T}(g^{-1}gx) = g\mathfrak{p}_{\Sigma_T}(x)$. This plus Corollary 2.36 gives:

$$\mathfrak{p}_{\Sigma_U}(g\mathfrak{p}_{\Sigma_T}(x)) = \mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}(gx) = \mathfrak{p}_{\Sigma_U}(gx)$$

Thus, $\mathfrak{p}_{\Sigma_U}(gx) = \mathfrak{p}_{\Sigma_U}(gab) = \mathfrak{p}_{\Sigma_U}(agb)$. The latter is agb times the shortest element of $(agb)^{-1}W_U = b^{-1}g^{-1}W_U$. Since $b^{-1}g^{-1} \in W_{U^\perp}$, there is no cancellation between $b^{-1}g^{-1}$ and words in W_U, so the shortest element of $b^{-1}g^{-1} \in W_{U^\perp}$ is $b^{-1}g^{-1}$, which gives $\mathfrak{p}_{\Sigma_U}(agb) = agbb^{-1}g^{-1} = a$. The same argument gives $\mathfrak{p}_{\Sigma_U}(ab) = a$. Thus:

$$\mathfrak{p}_{\Sigma_U}(x) = \mathfrak{p}_{\Sigma_U} \circ \mathfrak{p}_{\Sigma_T}(x) = \mathfrak{p}_{\Sigma_U}(ab) = a = \mathfrak{p}_{\Sigma_U}(gab) = \mathfrak{p}_{\Sigma_U}(gx)$$ ◻

Lemma 2.40. *For $i \in \{0, 1\}$, let $\Upsilon_i := w_i\Sigma_{T_i}$ be parabolic subcomplexes. Let u be the unique shortest element of $W_{T_0}w_0^{-1}w_1W_{T_1}$. Take $v_0 \in W_{T_0}$ and $v_1 \in W_{T_1}$ such that $w_0^{-1}w_1 = v_0uv_1^{-1}$. Let $U_0 := \{s \in T_0 \mid u^{-1}su \in T_1\}$ and $U_1 := \{s \in T_1 \mid usu^{-1} \in T_0\}$. Then $\mathfrak{p}_{\Upsilon_0}(\Upsilon_1) = w_0v_0W_{U_0}$ is the vertex set of a parabolic subcomplex $w_0v_0\Sigma_{U_0}$. Furthermore, right-multiplication by u defines an isomorphism $w_0v_0\Sigma_{U_0} \to w_1v_1\Sigma_{U_1}$.*

Proof. The element u exists by [41, Lemma 4.3.1], and $d_1(\Upsilon_0, \Upsilon_1) = |u|_S$. For every vertex $w_0v_0x \in w_0v_0W_{U_0} \subset \Upsilon_0^{(0)}$ we have $d_1(w_0v_0x, w_0v_0xu) = |u|_S$ and $w_0v_0xu = w_0v_0uu^{-1}xu \in w_0v_0uW_{U_1} = w_1v_1W_{U_1} \subset \Upsilon_1^{(0)}$. Thus $w_0v_0W_{U_0} \subset \mathfrak{p}_{\Upsilon_0}(\Upsilon_1)$.

Conversely, suppose $z \in W_{T_0}$ such that $w_0v_0z \in \mathfrak{p}_{\Upsilon_0}(\Upsilon_1)$. Then there is some $y \in W_{T_1}$ such that $w_0v_0z = \mathfrak{p}_{\Upsilon_0}(w_1v_1y)$. By construction, $\mathfrak{p}_{\Upsilon_0}(w_1v_1y)$ is w_1v_1yx, where x is the unique shortest element of $(w_1v_1y)^{-1}w_0v_0W_{T_0} = y^{-1}u^{-1}W_{T_0}$. Let $a := x^{-1}y^{-1}u^{-1} \in W_{T_0}$. We have $w_0v_0z = w_1v_1 \cdot yx = w_0v_0u \cdot u^{-1}a^{-1}$, so $z = a^{-1}$ and $x = y^{-1}u^{-1}z$. Since u^{-1} is the shortest element of $W_{T_1}u^{-1}W_{T_0} \ni x$, there are $b \in W_{T_0}$ and $c \in W_{T_1}$ such that $x = cu^{-1}b$ and $|x|_S = |c|_S + |u^{-1}|_S + |b|_S$. But since x is shortest in $y^{-1}u^{-1}W_{T_0} \subset W_{T_1}u^{-1}W_{T_0}$, the element b is trivial. Thus, $x = cu^{-1}$, which gives $W_{T_0} \ni z = uycu^{-1} \in uW_{T_1}u^{-1}$. By the remark following [41, Lemma 5.3.6], $uW_{T_1}u^{-1} \cap W_{T_0} = W_{U_0}$, so $z \in W_{U_0}$. Thus, $w_0v_0W_{U_0} \supset \mathfrak{p}_{\Upsilon_0}(\Upsilon_1)$.

An edge in the full subcomplex spanned by $\mathfrak{p}_{\Upsilon_0}(\Upsilon_1)$ is of the form $w_0v_0p - w_0v_0pq$ for some $p \in W_{U_0}$ and $q \in U_0$. Right-multiplication by u sends these vertices to $w_0v_0pu = w_0v_0uu^{-1}pu = w_1v_1u^{-1}pu$ and $w_0v_0pqu = w_1v_1u^{-1}pu \cdot u^{-1}qu$, respectively. These are two elements of $w_1v_1W_{U_1}$ that differ by $u^{-1}qu \in U_1$, so they are adjacent vertices of $w_1v_1W_{U_1}$. Thus, right-multiplication by u gives a bijection from $w_0v_0W_{U_0}$ to $w_1v_1W_{U_1}$ that preserves adjacency, so extends to an isomorphism of the full subcomplexes spanned by those vertex sets. □

Corollary 2.41. $\mathrm{Stab}(\mathfrak{p}_{\Upsilon_0}(\Upsilon_1)) = \mathrm{Stab}(\mathfrak{p}_{\Upsilon_1}(\Upsilon_0))$

The next result is a partial converse of Corollary 2.41.

Lemma 2.42. *Let P be a nonspherical irreducible parabolic subgroup of (W, S). Let $T \subset S$ be the unique subset such that P is conjugate to W_T. Suppose $\Upsilon_0 = w_0\Sigma_{T_0}$ and $\Upsilon_1 = w_1\Sigma_{T_1}$ are parabolic subcomplexes with $P = \mathrm{Stab}(\Upsilon_0) = \mathrm{Stab}(\Upsilon_1)$. Then $T = T_0 = T_1$ and the unique shortest element u of $W_Tw_0^{-1}w_1W_T$ belongs to $W_{T^\perp}$ and satisfies:*

- $\mathfrak{p}_{\Upsilon_1}(\Upsilon_0) = \Upsilon_0^{(0)}u = \Upsilon_1^{(0)}$.
- *Left-multiplication by $w_0uw_0^{-1} \in P^\perp$ and right-multiplication by $u \in W_T^\perp$ induce the same isomorphism $\Upsilon_0 \to \Upsilon_1$.*

Proof. The set T exists by Corollary 2.11, and Proposition 2.9 (1) further implies $T = T_0 = T_1$. Proposition 2.9 (2) gives $w_0^{-1}w_1 \in N_W(W_T) = W_T \times W_{T^\perp}$, so we can write $w_0^{-1}w_1 = vu$ where $v \in W_T$ and $u \in W_{T^\perp}$ is the shortest element of $W_Tw_0^{-1}w_1W_T$. The projection formula in Lemma 2.40 gives $U_0 = U_1 = T$, and that right-multiplication by u is the gate map $\mathfrak{p}_{\Upsilon_1}(\Upsilon_0) = \Upsilon_0^{(0)}u = \Upsilon_1^{(0)}$. Finally, for any $x \in W_T$ we have $w_0uw_0^{-1} \cdot w_0x = w_0ux = w_0x \cdot u$, so on $\Upsilon_0^{(0)}$ left-multiplication by $w_0uw_0^{-1}$, right-multiplication by u, and $\mathfrak{p}_{\Upsilon_1}$ all define the same map. □

Lemma 2.43. *If $x \in \Sigma^{(0)}$ and Υ is a parabolic subcomplex then $\mathfrak{p}_\Upsilon(x)$ is the unique d_2–closest vertex of Υ to x. Furthermore, there exists a closed chamber containing both $\pi_\Upsilon(x)$ and $\mathfrak{p}_\Upsilon(x)$.*

Proof. Let $y \in \Upsilon^{(0)} \setminus \{\mathfrak{p}_\Upsilon(x)\}$. Let z be the first vertex after y on a d_1–geodesic from y to $\mathfrak{p}_\Upsilon(x)$. Since Υ is d_1–convex, $z \in \Upsilon$. Let M be the wall dual to the edge between z and y. Then M cuts Υ, so by Corollary 2.37 it does not separate x and $\mathfrak{p}_\Upsilon(x)$. Thus, vertices x, $\mathfrak{p}_\Upsilon(x)$, and z are in one open halfspace M^- of M and y is in the other. Let p be the point at which the unique d_2–geodesic from x to y crosses M. The reflection through M is an isometry that fixes p and exchanges y and z, so:

$$d_2(x, y) = d_2(x, p) + d_2(p, y) = d_2(x, p) + d_2(p, z) \geq d_2(x, z)$$

In fact, $d_2(x, p) + d_2(p, z) > d_2(x, z)$, because x and z are contained in the d_2–convex open halfspace M^- and $p \in M$ is not, so p does not lie on the unique d_2–geodesic from x to z. Iterating along the d_2–geodesic, conclude $d_2(x, \mathfrak{p}_\Upsilon(x)) < d_2(x, y)$.

The proof of the further statement is similar: suppose M is a wall with $\mathfrak{p}_\Upsilon(x) \in M^-$ and $\pi_\Upsilon(x) \in M^+$. Then M cuts Υ, so M does not separate x and $\mathfrak{p}_\Upsilon(x)$, so $x \in M^-$. Using the reflect and straighten trick as in the previous argument, this produces a point $\mathfrak{r}_M(x)$ of Υ that is strictly closer to x than $\pi_\Upsilon(x)$ is. That is a contradiction, so no wall separates $\mathfrak{p}_\Upsilon(x)$ and $\pi_\Upsilon(x)$. □

3. Admissible paths

We work with a family of paths that is a common generalization of CAT(0) and combinatorial geodesics:

Definition 3.1. An *admissible path* $\gamma\colon I \to (\Sigma, d_2)$ is a continuous map from an interval $I \subset \mathbb{R}$ into Σ such that for every wall M exactly one of the following is true:

- $\gamma \cap M = \emptyset$
- $\gamma \subset M$
- There is a single time $t \in I$ such that $\gamma(t) \in M$, which is either:
  - an endpoint of I, or
  - an interior point of I such that for every sufficiently small neighborhood $(t - \epsilon, t + \epsilon) \subset I$ of t the points $\gamma(t - \epsilon)$ and $\gamma(t + \epsilon)$ are on opposite sides of M. In this case we say M *cuts* γ or M is *transverse to* γ, and write $\gamma \pitchfork M$.

An *admissible quasigeodesic* is a quasigeodesic that is an admissible path.

Lemma 3.2. *Combinatorial geodesics and CAT(0) geodesics are admissible and uniformly quasigeodesic.*

Proof. Let γ be a CAT(0) geodesic. It is 1–biLipschitz. Let M be a wall. Walls are convex, so if γ is not contained in M or disjoint from M then it intersects M in a single subsegment. Suppose M contains an interior point $\gamma(t)$ of γ. Consider the component of $\gamma \cap \bar{\mathcal{N}}_{1/2}(M)$ containing $\gamma(t)$. It is a geodesic of positive length in a metric product $M \times [-1/2, 1/2]$, so it is either contained in or transverse to $M \times \{0\}$.

Now suppose γ is a combinatorial geodesic. Transversality to the wall structure is immediate, since walls only meet $\Sigma^{(1)}$ in the midpoints of edges, and, as in the preceding proof, a wall cuts a combinatorial geodesic at most once. Lemma 2.15 implies that combinatorial geodesics are uniformly quasigeodesic. □

Lemma 3.3. *The combinatorial geodesics are precisely the admissible edge paths.*

Proof. Every edge is transverse to a unique wall. An edge path that is admissible does not contain distinct edges transverse to the same wall, so every edge of the path is transverse to a wall separating the ends of the path. Such an edge path is a combinatorial geodesic. □

Lemma 3.4 (Admissible paths are close to combinatorial geodesics)**.** *There exists D such that for every admissible path $\gamma\colon I \to \Sigma$ there exists a combinatorial geodesic at Hausdorff distance at most D from γ.*

Proof. If γ is contained in some wall M then it can be moved an arbitrarily small distance in the second coordinate of $\bar{\mathcal{N}}_{1/2}(M) = M \times [-1/2, 1/2]$ to be disjoint from M. Also, since walls have positive codimension, by an arbitrarily small perturbation we may assume γ does not meet any wall intersections. Thus, there is an admissible path γ' arbitrarily close to γ such that there is a sequence of times $\dots, t_0, t_1, \dots$ such that γ' meets a single wall at each time t_i and meets no wall in between. The closure of a component of the complement of the walls in Σ is a W–translate of the fundamental chamber. Each contains a unique vertex. Let v_i be the vertex in the same chamber as $\gamma'(t_i, t_{i+1})$. Since the chambers containing v_i and v_{i+1} are adjacent across the wall through $\gamma'(t_{i+1})$, there is a corresponding edge e_i of $\Sigma^{(1)}$ dual to that wall connecting v_i and v_{i+1}. Let α be the edge path with vertices v_i and edges e_i. The chambers of a fixed finite rank Coxeter system have finite diameter, so there is a D' such that α and γ' are D'–Hausdorff equivalent. Let $D := D' + 1$ to account for passing from γ to γ'. Then we have an edge path α at Hausdorff distance at most D from γ. Since α crosses the same sequence of walls at γ', it is also an admissible path. Lemma 3.3 says α is a combinatorial geodesic. □

Lemma 2.15 enables the following sort of computation, which we use frequently:

Lemma 3.5 (4–gon trick)**.** *Let α, β, γ, and δ be admissible paths such that the concatenation $\alpha + \beta + \gamma$ is defined and has the same endpoints as δ. Let α^- and α^+ denote the endpoints of α, and similarly for the other paths. Let $\mathcal{V}$ be a chain separating δ^- and δ^+. Let $\mathcal{V}_\alpha$ be the walls of $\mathcal{V}$ that meet α. Let $\mathcal{V}_\gamma$ be the walls of $\mathcal{V}$ that meet γ. Let $\mathcal{V}' := \mathcal{V} \setminus (\mathcal{V}_\alpha \cup \mathcal{V}_\gamma)$. Then $\mathcal{V}_\alpha$ is an initial subchain of $\mathcal{V}$, $\mathcal{V}_\gamma$ is a terminal subchain of $\mathcal{V}$, and $\mathcal{V}'$ is a convex subchain transverse to β. See Figure 1.*

If $\mathcal{V}'$ is nonempty then $\mathcal{V}_\alpha$ and $\mathcal{V}_\gamma$ are disjoint. Furthermore:

$$|\mathcal{V}| - |\mathcal{V}'| \leq d_c(\alpha^-, \alpha^+) + d_c(\gamma^-, \gamma^+) + 2 \asymp d_2(\alpha^-, \alpha^+) + d_2(\gamma^-, \gamma^+)$$

In particular, when β and δ are much longer than α and γ then $\mathcal{V}'$ is "most of" $\mathcal{V}$.

Proof. Assume $\alpha + \beta + \gamma$, δ, and $\mathcal{V} := \{V_i\}_{i_0 \leq i \leq i_1}$ are oriented from $x := \delta^-$ to $y := \delta^+$.

Since every wall $V \in \mathcal{V}$ separates x from y and δ is a path from x to y, V cuts δ an odd number of times. Since δ is admissible, that number is 1. Similarly, every wall $V \in \mathcal{V}$ cuts $\alpha + \beta + \gamma$ an odd number of times, but meets each subsegment at most once, so if V does not meet α or γ then it must cut β.

Suppose $i < j$ and $V_j \pitchfork \alpha$. Then V_i separates V_j from x, but α gives a path from x to V_j, so V_i cuts α. Conversely, if $i < j$ and V_i does not cut α then every point of α is on the same side of V_i as x, and V_j is on the other side of V_i, so V_j does not cut α. This gives the desired monotonicity—$\mathcal{V}$ is composed of a (possibly empty) initial subchain consisting of walls that cut α, a (possibly empty) terminal subchain consisting of walls that do not meet α, and possibly a single wall in between that meets α at its endpoint.

Every wall of $\mathcal{V}_\alpha$ separates α^- from α^+ except possibly that the last wall contains α^+. Thus, $|\mathcal{V}_\alpha| - 1 \leq d_c(\alpha^-, \alpha^+)$. Lemma 2.15 implies $d_c(\alpha^-, \alpha^+) \asymp d_2(\alpha^-, \alpha^+)$. The same is true for γ. □

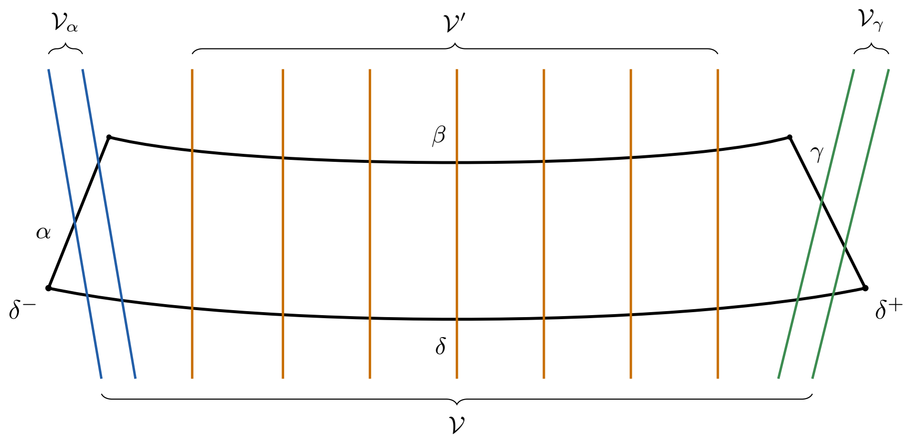

FIGURE 1. 4-gon trick: If α and γ are short then most of the walls cutting δ also cut β.

4. Bridges and divergence

Section 4.1 recalls the CAT(0) bridge machinery of Bestvina, Kleiner, and Sageev [20], and specializes the general results to walls in Davis complexes. Section 4.2 states and proves the main result of this section, Theorem 4.6, which characterizes the Morse property for admissible quasigeodesics in terms of dense, transverse chains.

4.1. **CAT(0) bridges.**

Lemma 4.1 (CAT(0) bridge lemma (cf [20, Lemma 2.3]))**.** *Let X be a complete CAT(0) space. Let A and B be nonempty closed convex subsets of X such that $d_2(A,B)$ is realized. Define $A_0 := \{a \in A \mid d_2(a,B) = d_2(A,B)\}$, $B_0 := \{b \in B \mid d_2(b,A) = d_2(A,B)\}$, and*

$$\mathrm{Bridge}(A,B) := \bigcup_{a \in A_0} [a, \pi_B(a)]$$

(1) A_0 and B_0 are closed and convex.
(2) $\pi_B|_{A_0} \colon A_0 \to B_0$ and $\pi_A|_{B_0} \colon B_0 \to A_0$ are inverse isometries.
(3) $\mathrm{Bridge}(A,B)$ is closed and convex and is isometric to a product $A_0 \times [0, d_2(A,B)]$.
(4) Suppose there exists $\epsilon > 0$ such that:

$$a \in A,\ d_2(a, A_0) = 1 \implies d_2(a,B) \geq d_2(A,B) + \epsilon \qquad (\prec)$$

Then for every $a \in A$ with $d_2(a, A_0) \geq 1$ we have:

$$d_2(a,B) \geq \ d_2(A,B) + \epsilon d_2(a, A_0)$$

(5) If X is a locally finite complex with a cocompact isometry group then for every D there exists ϵ such that if A and B are convex subcomplexes with $d_2(A,B) \leq D$ then ϵ is a flaring bound *for A and B satisfying* ($\prec$).

As an example, consider two lines in $\mathbb{E}^2$ intersecting at angle $0 < \theta < \pi/2$. Their bridge is the single point of intersection. Projection of one line to the other is surjective, so projections between convex sets need not be coarsely equivalent to the bridge (unlike in the cubical case). The optimal flaring bound satisfying ($\prec$) is $\epsilon = \sin(\theta)$, so it depends on the pair of convex sets.

Lemma 4.1 is just [20, Lemma 2.3] stated slightly differently, with the additional hypothesis that X is complete, and with the final item of [20, Lemma 2.3] split into two separate statements.

Corollary 4.2. *For all D there exists ϵ such that if M and M' are walls of Σ with $d_2(M,M') \leq D$ then ϵ is a flaring bound for M and M' satisfying* ($\prec$).

Proof. There is a subdivision of Σ, as in Definition 2.1, in which the walls of Σ are subcomplexes. $\square$

Lemma 4.3. *For all D there exists $k > 0$ such that for all walls M, M' of Σ with $d_2(M, M') \leq D$, all paths γ from M to M', and all R we have:*

$$\gamma \not\subset \mathcal{N}_R(\mathrm{Bridge}(M, M')) \implies |\gamma| \geq kR - 2$$

Proof. Let $D_0 := d_2(M, M') \leq D$. By Corollary 4.2, given D there exists a flaring bound ϵ for M and M'. Define $k := \frac{\epsilon}{1+\epsilon}$.

Let $M_0 := \{m \in M \mid d_2(m, M') = d_2(M, M')\}$ and $M_0' := \{m \in M' \mid d_2(m, M) = d_2(M, M')\}$. Suppose there exists $x \in \gamma$ at least R–far from the bridge. Let $m := \pi_M(x)$ and $m' := \pi_{M'}(x)$.

$$|\gamma| \geq d_2(x, m) + d_2(x, m') \geq d_2(m, m') \geq d_2(m, M') \tag{1}$$

If $d_2(x, m) \geq R$ or $d_2(x, m') \geq R$ then $|\gamma| \geq R > kR - 2$.

If $d_2(m, M_0) < 1$ and $d_2(m', M_0') < 1$ then $|\gamma| > 2(R - 1) > kR - 2$.

Up to swapping the roles of M and M', it remains to consider the case that $d_2(m, M_0) \geq 1$ and $d_2(x, m) < R$. In this case Lemma 4.1 gives:

$$d_2(m, M') \geq D_0 + \epsilon d_2(m, M_0) \geq D_0 + \epsilon(d_2(x, M_0) - d_2(x, m))$$

Since $M_0 \subset \mathrm{Bridge}(M, M')$, we have $d_2(x, M_0) \geq d_2(x, \mathrm{Bridge}(M, M')) \geq R$, so:

$$d_2(m, M') \geq D_0 + \epsilon(R - d_2(x, m)) \tag{2}$$

Combining equations (1) and (2):

$$\begin{aligned} |\gamma| &\geq \max\{D_0 + \epsilon(R - d_2(x, m)), d_2(x, m)\} \\ &\geq \frac{\epsilon}{1+\epsilon} R + \frac{D_0}{1+\epsilon} \\ &> kR - 2 \end{aligned}$$

□

Corollary 4.4. *For all D there exists $k > 0$, as in Lemma 4.3, such that for all walls M, M' of Σ with $d_2(M, M') \leq D$, all (λ, ϵ)–quasigeodesics $\gamma\colon [0, L] \to \Sigma$ from M to M', and all R we have:*

$$\gamma \not\subset \mathcal{N}_R(\mathrm{Bridge}(M, M')) \implies L \geq (k/\lambda)R - (2/\lambda)(\epsilon + 1)$$

Proof. Suppose $t \in [0, L]$ is such that $\gamma(t) \notin \mathcal{N}_R(\mathrm{Bridge}(M, M'))$. By Lemma 4.3, the length of the concatenation of the geodesic from $\gamma(0)$ to $\gamma(t)$ with the geodesic from $\gamma(t)$ to $\gamma(L)$ has length at least $kR - 2$, so $\lambda t + \epsilon + \lambda(L - t) + \epsilon \geq kR - 2$, which implies $L \geq (k/\lambda)R - (2/\lambda)(\epsilon + 1)$. □

Lemma 4.5. *For all C and D there exists B such that for all walls M, M' of Σ such that $d_2(M, M') \leq D$ and such that the longest chain transverse to both M and M' contains at most C–many walls, the diameter of* $\mathrm{Bridge}(M, M')$ *is at most B.*

Proof. Let $D_0 := d_2(M, M')$ and $M_0 := \{m \in M \mid d_2(m, M') = D_0\}$. Since $\mathrm{Bridge}(M, M')$ is isometric to a product $M_0 \times [0, D_0]$ and $D_0 \leq D$, it suffices to bound the diameter of M_0. For any $x, y \in M_0$, consider the geodesic 4–gon whose sides are $[x, y]$, $[y, \pi_{M'}(y)]$, $[\pi_{M'}(y), \pi_{M'}(x)]$, and $[\pi_{M'}(x), x]$. Let $\mathcal{K}$ be a chain realizing $d_c(x, y)$ and $\mathcal{K}'$ the subchain of $\mathcal{K}$ consisting of walls separating $\pi_{M'}(x)$ and $\pi_{M'}(y)$. By hypothesis $|\mathcal{K}'| \leq C$, so Lemma 2.15 and Lemma 3.5 give:

$$d_2(x, y) \asymp d_c(x, y) \stackrel{+}{\asymp} d_c(x, y) - C \leq |\mathcal{K}| - |\mathcal{K}'| \preceq 2D_0$$

□

Remark. The converse of Lemma 4.5 is not true: having a small bridge does not limit the number of common crossers. For example, the Davis complex of the 3,3,3 triangle groups is the regular hexagonal tiling of the plane. The walls are three families of parallel lines. Any line M from the first family crosses any line M' from the second family, so their bridge is the single point $M \cap M'$, but the entire infinite third family consists of lines that cross both M and M'.

4.2. **Wall characterization of the Morse property.**

Theorem 4.6. *For all λ and ϵ the following are equivalent for all admissible (λ, ϵ)–quasigeodesics $\gamma\colon I \to \Sigma$. Moreover, there are explicit bounds between the constraints of each item that are independent of γ.*

(a) There exist $K > 2$, r_0, and a quadratic polynomial $q(r)$ with positive leading coefficient such that γ has divergence gauge $\delta(r; 1, K)$ satisfying $\delta(r; 1, K) \geq q(r)$ for all $r \geq r_0$.
(b) γ is μ–Morse.
(c) γ is χ–strongly contracting.
(d) There exist C and L_0 such that if $|I| \geq L_0$ then there exists a chain $\mathcal{V} \pitchfork \gamma$ such that the components of $I \setminus \gamma^{-1}(\gamma \cap \mathcal{V})$ have length at most $5L_0$ and for every pair V, V' of consecutive walls of $\mathcal{V}$ there are at most C–many walls transverse to $\{V, V'\}$.
(e) There exist C and $0 < \rho \leq \rho_0(\lambda, \epsilon)$ such that there exists L_1 such that for every subinterval $I' \subset I$ of length at least L_1 there exists a chain $\mathcal{V} \pitchfork \gamma(I')$ with $|\mathcal{V}|/|I'| \geq \rho$ such that for every pair V, V' of consecutive walls of $\mathcal{V}$ the longest chain $\mathcal{H}$ transverse to $\{V, V'\}$ has $|\mathcal{H}| \leq C$.

The proof occupies this subsection. The equivalence of (a), (b), and (c) is a standard result about subsets of CAT(0) spaces, as discussed in Section 2.8. Lemma 4.7 and Lemma 4.8 use strong contraction to produce walls transverse to γ with small projection diameter to γ. Proposition 4.9 collects these into a chain to show (c) $\implies$ (d). The implication (d) $\implies$ (e) follows by taking $\rho := \min\{\frac{1}{10L_0}, \rho_0(\lambda, \epsilon)\}$ and $L_1 := 20L_0$, choosing the chain $\mathcal{V}$ for all of γ, and then restricting it to the walls that cut γ'. Proposition 4.10 shows (e) $\implies$ (a).

Lemma 4.7. *For all χ there exists χ' such that for all λ and ϵ there exists χ'' such that for every χ–strongly contracting admissible (λ, ϵ)-quasigeodesic $\gamma\colon I \to \Sigma$, and for every wall M such that* $\operatorname{diam}(\pi_\gamma(M)) > 2\chi'$, *there is a subinterval $J_M \subset I$ such that:*

$$\gamma^{-1}(\pi_\gamma(M)) \subset J_M \quad \textit{and} \quad \gamma(J_M) \subset \bar{\mathcal{N}}_{\chi''}(M)$$

Proof. By Proposition 2.31, strong contraction is equivalent to strong constriction, so there exists χ' depending only on χ such that for any geodesic $[x, y]$, $\operatorname{diam}(\pi_\gamma([x, y])) > \chi'$ implies $[x, y]$ passes within distance χ' of $\pi_\gamma(x)$ and $\pi_\gamma(y)$. Suppose M is a wall such that $\operatorname{diam}(\pi_\gamma(M)) > 2\chi'$. For every $x_0 \in \pi_\gamma(M)$ there exist $x_1 \in \pi_\gamma(M)$ with $d_2(x_0, x_1) > \chi'$, else $\operatorname{diam}(\pi_\gamma(M))$ would be at most $2\chi'$. Choose $m_0, m_1 \in M$ such that $x_i \in \pi_\gamma(m_i)$. Since M is convex, it contains the geodesic $[m_0, m_1]$, which, by the result above, passes within distance χ' of $\pi_\gamma(m_0)$ and $\pi_\gamma(m_1)$. Since $\operatorname{diam}(\pi_\gamma(m_0)) \leq \chi$, the geodesic $[m_0, m_1] \subset M$ passes within distance $\chi + \chi'$ of x_0. Since x_0 was an arbitrary point of $\pi_\gamma(M)$, this gives $\pi_\gamma(M) \subset \bar{\mathcal{N}}_{\chi+\chi'}(M)$.

Take J_M to be the smallest closed subinterval of I containing $\gamma^{-1}(\pi_\gamma(M))$. Let I'' be a maximal open subinterval of J_M with $\gamma(I'') \subset \bar{\mathcal{N}}^c_{\chi+\chi'}(M)$. Convexity of distance to M implies the geodesic between the endpoints of $\gamma(\bar{I}'')$ is contained in $\bar{\mathcal{N}}_{\chi+\chi'}(M)$. But γ is χ–strongly contracting, so $\gamma(\bar{I}'')$ and the geodesic between its endpoints are at Hausdorff distance at most D from each other, where D can be computed in terms of χ, λ, and ϵ. Take $\chi'' := \chi + \chi' + D$. □

Lemma 4.8. *For all χ, λ, and ϵ there exist L_0 such that for every χ–strongly contracting admissible (λ, ϵ)–quasigeodesic $\gamma\colon I \to \Sigma$, for every subinterval $I' \subset I$ of length at least L_0 there exists a wall M transverse to $\gamma(I')$ and a subinterval $J_M \subset I$ such that $\gamma^{-1}(\pi_\gamma(M)) \subset J_M$ and $|J_M| \leq L_0$.*

Proof. Take χ' and χ'' of Lemma 4.7. By local finiteness of the wall structure, there exists C such that there are at most C–many walls intersecting any ball of radius χ''. By Lemma 2.15 we can choose $L_0 \geq \lambda(2\chi' + \epsilon)$ large enough that $|[x, y]| \geq L_0$ implies $d_w(\gamma(x), \gamma(y)) \geq 2C + 1$. Choose any subinterval $I' \subset I$ with $|I'| = L_0$, and suppose $\bar{I}' = [x, y]$. Suppose that for every wall M transverse $\gamma(I')$, the smallest closed interval J_M such that $\gamma^{-1}(\pi_\gamma(M)) \subset J_M$ has length strictly greater than L_0. Since J_M is smallest, each endpoint of J_M is either in, or is an accumulation point of, $\gamma^{-1}(\pi_\gamma(M))$, so $\operatorname{diam}(\pi_\gamma(M))$ is at least the d_2–distance between the endpoints of $\gamma(J_M)$, which is at least $|J_M|/\lambda - \epsilon > L_0/\lambda - \epsilon \geq 2\chi'$. Lemma 4.7 says $\gamma(J_M) \subset \bar{\mathcal{N}}_{\chi''}(M)$. Note that $[x, y]$ and J_M are intersecting closed intervals, since both contain the time t such that $\gamma(t) \in M$, so $|J_M| > L_0 = |[x, y]|$ implies J_M contains at least one of x or y. Thus, M enters $\bar{\mathcal{N}}_{\chi''}(\{\gamma(x), \gamma(y)\})$.

By the choice of C there are at most $2C$–many walls that enter $\bar{\mathcal{N}}_{\chi''}(\{\gamma(x), \gamma(y)\})$, but by the choice of L_0 there are at least $2C + 1$–many walls transverse to $\gamma(I')$, so there is at least one wall M transverse to $\gamma(I')$ that does not enter $\bar{\mathcal{N}}_{\chi''}(\{\gamma(x), \gamma(y)\})$. This wall must have $|J_M| \leq L_0$. □

Proposition 4.9. *For all χ, λ, and ϵ there exist C and L_0 such that for every χ–strongly contracting admissible (λ, ϵ)–quasigeodesic $\gamma\colon I \to \Sigma$ with I an interval of length at least L_0, there exists a chain $\mathcal{V}$ transverse to γ such that the components of $I \setminus \gamma^{-1}(\gamma \cap \mathcal{V})$ have length at most $5L_0$ and for every pair V, V' of consecutive walls of $\mathcal{V}$ the number of walls transverse to $\{V, V'\}$ is at most C.*

Proof. Given χ, λ, and ϵ take χ', χ'', L_0, and C as in Lemma 4.8. For $\gamma\colon I \to \Sigma$ such that $|I| < L_0$ there is nothing to prove, so we may restrict our attention to cases with $|I| \geq L_0$, and by shifting the parameterization of γ, if necessary, we may assume $(0, L_0) \subset I$. Let $s_i := iL_0$ for $i \in \mathbb{Z}$. If $(s_i, s_{i+1}) \subset I$ then by Lemma 4.8 there exists a wall $V_i \pitchfork \gamma(s_i, s_{i+1})$ and an interval $J_i \subset I$ of length at most L_0 such that $\gamma^{-1}(\pi_\gamma(V_i)) \subset J_i$. Let t_i be the unique time such that $\gamma(t_i) \in V_i$. Since J_i has length at most L_0 and $t_i \in J_i \cap (s_i, s_{i+1})$, we have $J_i \subset ((i-1)L_0, (i+2)L_0) \cap I$. In particular, if $j \geq i+3$ and $(s_i, s_{j+1}) \subset I$ then J_i is disjoint from $J_j \subset ((j-1)L_0, (j+2)L_0)$, which implies that V_i and V_j are disjoint, since a point in $V_i \cap V_j$ would project via π_γ to a point in $\pi_\gamma(V_i) \cap \pi_\gamma(V_j)$, but $\gamma^{-1}(\pi_\gamma(V_i) \cap \pi_\gamma(V_j) \subset J_i \cap J_j = \emptyset$.

Suppose for $i + 1 < j < k - 1$ in $\mathbb{Z}$ that $(s_i, s_{k+1}) \subset I$ and M is a wall transverse to both V_i and V_k. For $z \in M \cap V_i$, we have $\pi_\gamma(z) \subset \pi_\gamma(M) \cap \pi_\gamma(V_i)$, so $\gamma^{-1}(\pi_\gamma(M))$ contains a point in J_i. Similarly, $\gamma^{-1}(\pi_\gamma(M))$ contains a point in J_k. Now, $J_i \subset ((i-1)L_0, (i+2)L_0)$ and $J_k \subset ((k-1)L_0, (k+2)L_0)$ so $[jL_0, (j+1)L_0]$ separates J_i and J_k in I. It follows that $\operatorname{diam}(\pi_\gamma(M)) > L_0/\lambda - \epsilon > 2\chi'$, so Lemma 4.7 implies $\gamma([jL_0, (j+1)L_0]) \subset \bar{\mathcal{N}}_{\chi''}(M)$.

Thus, every point of $\gamma([jL_0, (j+1)L_0])$ is χ''–close to every wall that is transverse to both V_i and V_k. The number of walls entering a χ''–ball around any point is at most C, so there are at most C–many walls transverse to both V_i and V_k when $k \geq i+4$ and $(s_i, s_{k+1}) \subset I$.

Take $\mathcal{V} := \{V_i\}_{i \in 4\mathbb{Z} \wedge (iL_0, (i+1)L_0) \subset I}$. It contains at least one wall, V_0, since $0 \in 4\mathbb{Z} \wedge (0, L_0) \subset I$. Since wall indices in $\mathcal{V}$ differ by at least 4, the argument above gives that they are disjoint and there are at most C walls transverse to any distinct pair of them. Furthermore, since they are all transverse to γ, by construction, $\mathcal{V}$ contains no facing triples, so it is a chain.

Finally, we compute the length of components of $I \setminus \gamma^{-1}(\gamma \cap \mathcal{V})$. Let t_i be the time at which $\gamma(t_i) \in V_i$. If there is a first wall V_i then $(iL_0, (i+1)L_0) \subset I$ but $(i-4)L_0 \notin I$. Since $t_i \in (iL_0, (i+1)L_0)$, the initial segment of I ending at t_i has length at most $5L_0$. If there is a last wall V_i then similarly $(iL_0, (i+1)L_0) \subset I$ but $(i+5)L_0 \notin I$, so the terminal segment of I starting at t_i has length at most $5L_0$. If both V_{4i} and $V_{4(i+1)}$ are consecutive walls in $\mathcal{V}$ then $t_{4i} \in (4iL_0, (4i+1)L_0)$ and $t_{4(i+1)} \in (4(i+1)L_0, (4(i+1)+1)L_0)$, so $0 < t_{4(i+1)} - t_{4i} \leq 5L_0$. ☐

Proposition 4.10. *Given λ, ϵ, C, $0 < \rho \leq \rho_0(\lambda, \epsilon)$, and L, suppose $\gamma\colon I \to \Sigma$ is an admissible (λ, ϵ)–quasigeodesic with the property that for every subinterval $I' \subset I$ of length at least L there exists a chain $\mathcal{V} \pitchfork \gamma(I')$ with $|\mathcal{V}|/|I'| \geq \rho$ such that for every pair of consecutive walls V, V' of $\mathcal{V}$, every chain $\mathcal{H} \pitchfork \{V, V'\}$ has $|\mathcal{H}| \leq C$. Then there are $K = K(\lambda, \rho)$, $r_0 = r_0(\lambda, \epsilon, \rho, C, L)$, and a quadratic polynomial $q(r)$ with positive leading coefficient whose coefficients depend on λ, ϵ, ρ, and C, such that γ has divergence gauge $\delta(r; 1, K)$ satisfying $\delta(r; 1, K) \geq q(r)$ for $r \geq r_0$.*

Proof. We make a refinement of the 4–gon trick of Lemma 3.5. Suppose λ_0 and ϵ_0 are constants such that Lemma 2.15 implies that for all $x, y \in \Sigma$:

$$d_2(x, y)/\lambda_0 - \epsilon_0 \leq d_c(x, y) \leq \lambda_0 d_2(x, y) + \epsilon_0$$

Take $K := 2 + 2\lambda(1 + 2\lambda_0)/\rho$. Consider $r \geq r_0 := (\lambda L + \epsilon)/(K-2)$.

Suppose there exist points x and y such that $d_2(x, \gamma) = d_2(y, \gamma) = r$ and $d_2(x, y) \geq Kr$, so that $d_2(\pi_\gamma(x), \pi_\gamma(y)) \geq (K-2)r$.

Pick $x' \in \pi_\gamma(x)$ and $y' \in \pi_\gamma(y)$ and let I' be the smallest closed interval whose endpoints map to x' and y'. This implies $|I'| \geq ((K-2)r - \epsilon)/\lambda \geq L$, so there exists a chain $\mathcal{V}$ as in the statement.

Since γ is admissible and $\mathcal{V}$ is a chain, for each $V_i \in \mathcal{V}$ there is a unique $t_i \in I$ such that $\gamma(t_i) \in V_i$, and $i < j$ implies $t_i < t_j$.

Suppose there exists a path α from x to y in $N_r^c(\gamma)$. Let $\mathcal{V}'$ be the convex subchain of $\mathcal{V}$ consisting of the walls that cut α. At most the first $1 + d_c(x, x') \leq \lambda_0 d_2(x, x') + \epsilon_0 + 1 \leq \lambda_0 r + \epsilon_0 + 1$ walls of $\mathcal{V}$ intersect $[x, x']$, and, similarly, at most the last $d_c(y, y') \leq \lambda_0 r + \epsilon_0 + 1$ many intersect $[y, y']$. See Figure 2. Now we count the number of subsegments of $I' \setminus \gamma^{-1}(\gamma \cap \mathcal{V})$. The number of these from the origin of I' to the time corresponding to the first wall of $\mathcal{V}'$ is at most $\lambda_0 r + \epsilon_0 + 2$, and the number from the time of the last wall of

$\mathcal{V}'$ to the terminus of I' is at most $\lambda_0 r + \epsilon_0 + 2$. There are $|\mathcal{V}|+1$–many intervals in I', with average length:

$$\frac{|I'|}{|\mathcal{V}|+1} \leq \frac{|I'|}{\rho|I'|+1} < \frac{1}{\rho}$$

The proportion of them of length greater than $2/\rho$ is less than $1/2$, so the number of length at most $2/\rho$ is at least:

$$(|\mathcal{V}|+1)/2 \geq \rho|I'|/2 \geq (\rho/2)((K-2)r-\epsilon)/\lambda$$

The number of components of $I' \setminus \gamma^{-1}(\mathcal{V})$ that occur between two walls of $\mathcal{V}$ and have length at most $2/\rho$ is therefore bounded below by:

$$\begin{aligned}\frac{\rho(K-2)r-\rho\epsilon}{2\lambda} - 2(\lambda_0 r + \epsilon_0 + 2) &= \left(\frac{\rho(K-2)}{2\lambda} - 2\lambda_0\right) r - \frac{\rho\epsilon}{2\lambda} - 2(\epsilon_0+2) \\ &= r - \left(\frac{\rho\epsilon}{2\lambda} + 2(\epsilon_0+2)\right)\end{aligned} \tag{3}$$

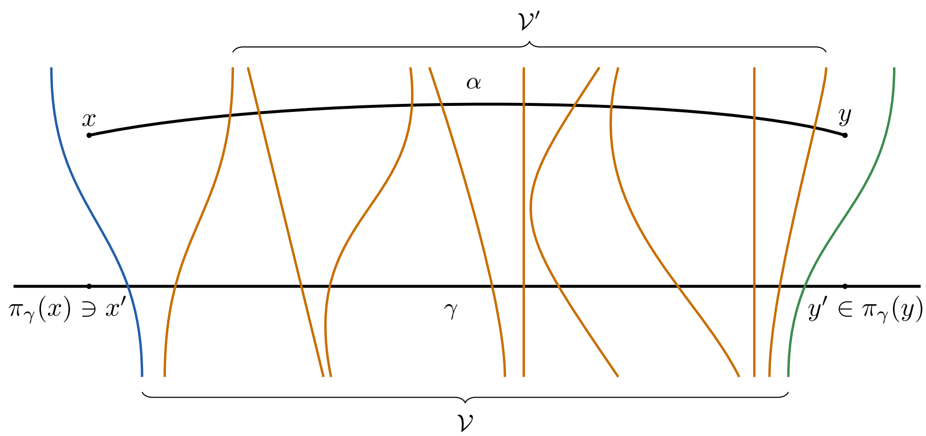

FIGURE 2. Density control on the chain transverse to γ forces linearly many of the interwall subsegments of γ to be short and the corresponding interwall subsegments of α to be long.

Let V_i and V_{i+1} be a pair of consecutive walls of $\mathcal{V}'$ such that $0 < t_{i+1} - t_i \leq 2/\rho$. This implies $d_2(V_i, V_{i+1}) \leq 2\lambda/\rho + \epsilon$.

Corollary 4.4 says the length of a quasigeodesic with fixed constants connecting V_i and V_{i+1} not contained in $\mathcal{N}_R(\text{Bridge}(V_i, V_{i+1}))$ has parameter interval length growing linearly in R, with constants depending on the quasigeodesic constants and $d_2(V_i, V_{i+1})$ but not on V_i and V_{i+1}. Since $t_{i+1} - t_i \leq 2/\rho$, the corresponding segment of γ is contained in a uniform neighborhood of $\text{Bridge}(V_i, V_{i+1})$. Lemma 4.5 says $\text{diam}(\text{Bridge}(V_i, V_{i+1}))$ is bounded in terms of $2\lambda/\rho + \epsilon$ and C. Thus, $\text{Bridge}(V_i, V_{i+1})$ and $\gamma([t_i, t_{i+1}])$ are coarsely equivalent, with constants depending on ρ, λ, ϵ, and C.

On the other hand, since α is required to stay far from γ, this implies α stays outside $\mathcal{N}_R(\text{Bridge}(V_i, V_{i+1}))$ for some $R \overset{+}{\asymp} r$, so Lemma 4.3 says the length of the subsegment of α from V_i to V_{i+1} is bounded below by a linear function of r, with positive slope, whose constants depend on $d_2(V_i, V_{i+1})$, which is bounded above by $2\lambda/\rho + \epsilon$. The bound from (3) says the number of such subsegments of α is linear in r with slope 1, so the total length of α is bounded below by a quadratic polynomial $q(r)$ with positive leading coefficient whose coefficients depend on ρ, λ, ϵ, and C.

We have shown that $\delta(r; 1, K) \geq q(r)$ when $r \geq r_0$ as desired. Notice that if for some r there is no such pair of points (x, y) or there is no such path α, then $\delta(r; 1, K) = \infty$, so it is still true that $\delta(r; 1, K) \geq q(r)$. $\square$

5. GRIDS AND WIDE PARABOLICS

Section 5.1 introduces grids and uses them to produce wide parabolics. Section 5.2 states and proves the main result of this section, Theorem 5.10, which relates an admissible quasigeodesic's interactions with wide parabolic subcomplexes and dense, transverse chains.

5.1. **Grids.**

Definition 5.1. A *grid* is a pair $(\mathcal{V},\mathcal{H})$ (vertical/horizontal) of transverse chains that each contain at least two walls.

Definition 5.2. Let ℓ be as in Section 2.6. If $\mathcal{V}$ is a chain, define its *interior* $\mathring{\mathcal{V}}$ to be all the walls V of $\mathcal{V}$ such that there are at least $\ell-1$ walls of $\mathcal{V}$ that come before V in the chain and at least $\ell-1$ walls that come after V.

The goal of this subsection is to prove[2] the following proposition, which is essentially the Grid Lemma attributed to Caprace and Przytycki by Caprace and Marquis [31, Lemma 2.8].

Proposition 5.3. *Let $(\mathcal{V},\mathcal{H})$ be a grid in Σ such that $|\mathcal{V}|\geq 2\ell-1$ and $|\mathcal{H}|\geq 2\ell$. Then $\mathrm{Pc}(\mathcal{V})\times\mathrm{Pc}(\mathcal{V})^\perp$ is a wide parabolic containing $\mathfrak{r}_H$ for all $H\in\mathring{\mathcal{H}}$.*

Lemma 5.4 (cf [28, Lemma 2.3][3])**.** *Let R_1 and R_2 be finite sets of reflections in (W,S) such that $\langle R_1\rangle$ and $\langle R_2\rangle$ are irreducible, and such that $r_1r_2=r_2r_1$ for all $r_1\in R_1$ and $r_2\in R_2$. Suppose $\langle R_1\rangle$ is infinite. Then one of the following is true:*

- $\mathrm{Pc}(R_1\cup R_2)=\mathrm{Pc}(R_1)\times\mathrm{Pc}(R_2)$
- $R_2\leq\mathrm{Pc}(R_1)$

If, in addition, $\langle R_2\rangle$ is infinite, then $R_2\leq\mathrm{Pc}(R_1)$ implies $\mathrm{Pc}(R_1)=\mathrm{Pc}(R_2)$ is higher rank affine.

Example 5.5. Consider the Coxeter system (W,S) defined by the Coxeter diagram $a\overset{4}{-}b\overset{4}{-}c\overset{4}{-}d$, which is not affine. Consider the finite sets of reflections $R_1:=\{bab,c,d\}$ and $R_2:=\{a\}$. Then $aca=c$, $ada=d$ and $a(bab)a=(abab)a=(baba)a=bab$, so R_1 and R_2 commute. In fact, from Allcock's description of reflection centralizers [10], one finds $\langle R_1\rangle$ is the entire reflection part of $Z_W(\langle R_2\rangle)$. The group $\mathrm{Pc}(R_2)\cong\mathbb{Z}/2\mathbb{Z}$ is irreducible. The group $\langle R_1\rangle$ is infinite and irreducible; to see this, consider $P:=\langle a,b,c\rangle$, which is a Euclidean triangle parabolic subgroup of W. In its usual $\mathrm{Isom}(\mathbb{E}^2)$ representation the elements $bab,c\in R_1\cap P$ are reflections through parallel lines, so $\langle R_1\rangle$ contains an infinite dihedral subgroup. Furthermore, P is 2–spherical and rank 3, so every proper parabolic subgroup is finite, so $\mathrm{Pc}(\langle bab,c\rangle)=P$. But $d\in R_1\leq\mathrm{Pc}(R_1)\geq\mathrm{Pc}(\langle bab,c\rangle)$, so $\mathrm{Pc}(R_1)=W$. Thus, $\mathrm{Pc}(R_2)\lneq\mathrm{Pc}(R_1)$ and $\mathrm{Pc}(R_1)$ is not affine. This contradicts both conclusions in the $R_2\leq\mathrm{Pc}(R_1)$ branch of [28, Lemma 2.3].

Lemma 5.6. *Let $(\mathcal{V},\mathcal{H})$ be a grid in Σ. Suppose one of the following is true:*

- *$W(\mathcal{V})$ and $W(\mathcal{H})$ commute.*
- *$W(\mathcal{V})$ is dihedral and $\mathcal{V}$ contains at least 8 walls.*

Then one of the following is true:

- $\mathrm{Pc}(\mathcal{H})\leq\mathrm{Pc}(\mathcal{V})^\perp$
- $\mathrm{Pc}(\mathcal{H})=\mathrm{Pc}(\mathcal{V})$ *is an irreducible higher rank affine parabolic.*

Proof. By Lemma 2.7, $W(\mathcal{V})$ and $W(\mathcal{H})$ are irreducible and nonspherical. Furthermore, since it does not change $\mathrm{Pc}(\mathcal{V})$ and $\mathrm{Pc}(\mathcal{H})$, we may assume $\mathcal{V}$ and $\mathcal{H}$ are finite chains, so that $W(\mathcal{V})$ and $W(\mathcal{H})$ are finitely generated reflection groups.

By Lemma 5.4 we can partition the set $R:=\{\mathfrak{r}_H\mid H\in\mathcal{H}\}$ of reflections through walls of $\mathcal{H}$ into disjoint subsets:

- $R_0:=\{\mathfrak{r}_H\mid H\in\mathcal{H},\,\mathfrak{r}_H\notin Z_W(W(\mathcal{V}))\}$
- $R_1:=\{\mathfrak{r}_H\mid H\in\mathcal{H},\,\mathfrak{r}_H\in Z_W(W(\mathcal{V}))\cap\mathrm{Pc}(\mathcal{V})\}$
- $R_2:=\{\mathfrak{r}_H\mid H\in\mathcal{H},\,\mathfrak{r}_H\in\mathrm{Pc}(\mathcal{V})^\perp\}$

[2]The proof of the Grid Lemma presented in [31] has a gap. In our language, that proof supposes that there is a grid $(\mathcal{V},\mathcal{H})$ with $\mathcal{V}=\{V_i\}_{0\leq i\leq k}$, and supposes there is *any* $0\leq i_0\leq k$ such that there is an $H_j\in\mathring{\mathcal{H}}$ such that $\mathfrak{r}_{V_{i_0}}$ and $\mathfrak{r}_{H_j}$ do not commute. It then applies [31, Lemma 2.9], which is Theorem 2.19, to $\{V_i\}_{0\leq i\leq i_0}$. However, Theorem 2.19 requires the chain to be sufficiently long, which is not true if i_0 is small. Proposition 5.8 below argues that given the noncommuting pair of walls there are four possible configurations to which one might try to apply Theorem 2.19, and at least one of them satisfies that theorem's hypotheses. This fills the gap.

[3]Lemma 5.4 is a corrected version of [28, Lemma 2.3] stating what is actually proven there. The original omits the hypothesis that $\langle R_2\rangle$ is infinite in the final sentence, but still concludes $\mathrm{Pc}(R_1)=\mathrm{Pc}(R_2)$ is higher rank affine. This is an error, as demonstrated by Example 5.5. The specific error in the proof of [28, Lemma 2.3] is the application of [28, Lemma 2.2] in the case $R_2\leq\mathrm{Pc}(R_1)$ without explicitly requiring $\langle R_2\rangle$ to be nonspherical.

Furthermore, if $|R_1| \geq 2$ then $\mathrm{Pc}(R_1) = \mathrm{Pc}(\mathcal{V})$ is higher rank affine.

Reflections in distinct walls of $\mathcal{H}$ do not commute, so at least one of $R \cap \mathrm{Pc}(\mathcal{V})$ or R_2 is empty.

$W(\mathcal{V})$ and $W(\mathcal{H})$ commute if and only if $R_0 = \emptyset$, in which case we are done, since then either $R = R_1$ or $R = R_2$. The former gives $|R_1| \geq 2$, so $\mathrm{Pc}(\mathcal{H}) = \mathrm{Pc}(R_1) = \mathrm{Pc}(\mathcal{V})$ is higher rank affine. The latter gives $\mathrm{Pc}(\mathcal{H}) = \mathrm{Pc}(R_2) \leq \mathrm{Pc}(\mathcal{V})^{\perp}$.

If $R_0 \neq \emptyset$ then, by hypothesis, $W(\mathcal{V})$ is dihedral and $\mathcal{V}$ contains at least 8 walls. For $\mathfrak{r}_H \in R_0$, Lemma 2.20 says $W(H, \mathcal{V})$ is a Euclidean triangle group, so Lemma 2.22 gives $\mathfrak{r}_H \in \mathrm{Pc}(\mathcal{V})$, and Proposition 2.21 says $\mathrm{Pc}(\mathcal{V})$ is higher rank affine. Since $\emptyset \neq R_0 \subset R \cap \mathrm{Pc}(\mathcal{V})$, $R_2 = \emptyset$, so $R = R_0 \cup R_1$, which gives $\mathrm{Pc}(\mathcal{H}) \leq \mathrm{Pc}(\mathcal{V})$. Since irreducible affine parabolics do not contain proper nonspherical parabolic subgroups and $\mathrm{Pc}(\mathcal{H})$ is a nonspherical parabolic, $\mathrm{Pc}(\mathcal{H}) = \mathrm{Pc}(\mathcal{V})$. □

Lemma 5.7. *Let $(\mathcal{V}, \mathcal{H})$ be a grid in Σ, with $\mathcal{V} = \{V_i\}_{i \in I}$ and $\mathcal{H} = \{H_j\}_{j \in J}$. Suppose, for some $i_0 \in I$ and $j_0 \in J$ there is a wall M that contains $V_{i_0} \cap H_{j_0}$. Then M is transverse to at least two of the following chains:*

$$\mathcal{V}^- := \{V_i\}_{\{i \in I | i \leq i_0\}} \qquad \mathcal{V}^+ := \{V_i\}_{\{i \in I | i \geq i_0\}}$$

$$\mathcal{H}^- := \{H_j\}_{\{j \in J | j \leq j_0\}} \qquad \mathcal{H}^+ := \{H_j\}_{\{j \in J | j \geq j_0\}}$$

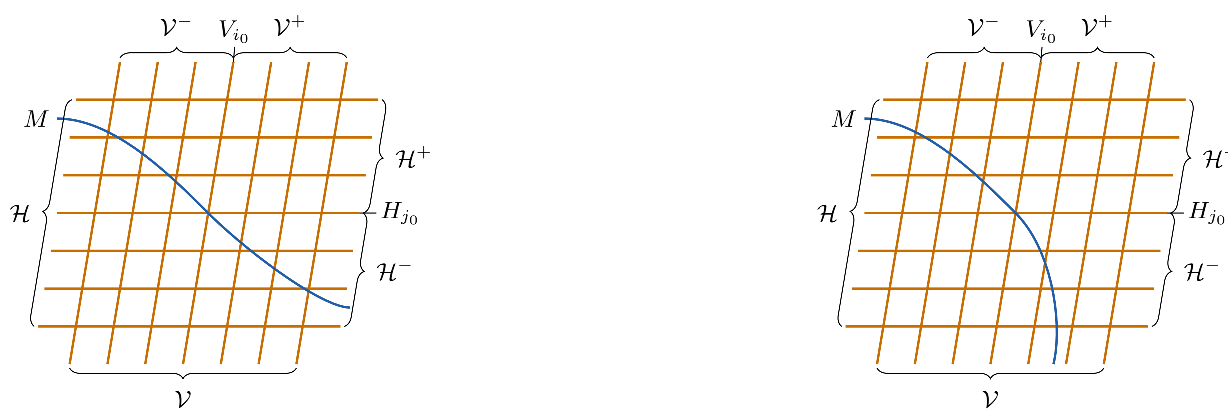

FIGURE 3. Two possibilities for a wall M containing $V_{i_0} \cap H_{j_0}$, separating $V_{i_0}^+ \cap H_{j_0}^+$ from $V_{i_0}^- \cap H_{j_0}^-$ and $M \not\pitchfork \mathcal{H}^+$. We have $M \pitchfork \mathcal{V}^-$ and at least one of $M \pitchfork \mathcal{V}^+$ or $M \pitchfork \mathcal{H}^-$.

Proof. If M is equal to V_{i_0} then it is transverse to all of $\mathcal{H}$, so it is transverse to both of $\mathcal{H}^-$ and $\mathcal{H}^+$. If M is equal to H_{j_0} then it is transverse to both $\mathcal{V}^-$ and $\mathcal{V}^+$. Now assume M is not one of V_{i_0} or H_{j_0}. Then $V_{i_0} \cup H_{j_0}$ divides Σ into four regions, $V_{i_0}^- \cap H_{j_0}^-$, $V_{i_0}^+ \cap H_{j_0}^-$, $V_{i_0}^+ \cap H_{j_0}^+$, and $V_{i_0}^- \cap H_{j_0}^+$. The wall M meets two of these, which are opposites, either $V_{i_0}^- \cap H_{j_0}^-$ and $V_{i_0}^+ \cap H_{j_0}^+$ or $V_{i_0}^- \cap H_{j_0}^+$ and $V_{i_0}^+ \cap H_{j_0}^-$, and separates the other two. Suppose, without loss of generality, that M separates $V_{i_0}^- \cap H_{j_0}^-$ from $V_{i_0}^+ \cap H_{j_0}^+$.

Suppose i_0 and j_0 are not extremes. Suppose M is not transverse to $\mathcal{H}^+$. Then there is some first $j_1 > j_0$ in J such that M is disjoint from H_{j_1}. For any $i < i_0$ there is a path from $V_{i_0}^- \cap H_{j_0}^-$ to $V_{i_0}^+ \cap H_{j_0}^+$ by travelling through V_i and H_{j_1}. Specifically, choose a point $a \in V_i \subset V_{i_0}^-$ in the interior of $H_{j_0}^-$, which exists, since $V_i \pitchfork H_{j_0}$, a point $b \in V_i \cap H_{j_1}$, which exists, since $V_i \pitchfork H_{j_1}$, and a point $c \in H_{j_1} \subset H_{j_0}^+$ in the interior of $V_{i_0}^+$, which exists, since $V_{i_0} \pitchfork H_{j_1}$. By convexity of walls, $[a, b] \subset V_i \subset V_{i_0}^-$ and $[b, c] \subset H_{j_1} \subset H_{j_0}^+$. Now, since M separates $V_{i_0}^- \cap H_{j_0}^-$ from $V_{i_0}^+ \cap H_{j_0}^+$, it cuts either $[a, b]$ or $[b, c]$, but if M is disjoint from $H_{j_1} \supset [b, c]$ then it must be that M cuts $[a, b] \subset V_i$, so $M \pitchfork V_i$. Since this was true for every $i < i_0$, we conclude $M \not\pitchfork \mathcal{H}^+$ implies $M \pitchfork \mathcal{V}^-$.

An analogous argument yields that M is transverse to at least one of $\mathcal{H}^-$ or $\mathcal{V}^+$.

When one or both of i_0 and j_0 are extremes then the argument is similar, with the added observation that if, say, i_0 is the least element of I, then $\mathcal{V}^-$ is the single wall V_{i_0}, which is transverse to M by hypothesis. □

Proposition 5.8. *Let $(\mathcal{V}, \mathcal{H})$ be a grid in Σ such that $\mathcal{V}$ and $\mathcal{H}$ each contain at least $(2\ell - 1)$–many walls. If $W(\mathring{\mathcal{H}}) \not\subset Z_W(W(\mathcal{V}))$ or $W(\mathring{\mathcal{V}}) \not\subset Z_W(W(\mathcal{H}))$ then $\mathrm{Pc}(\mathcal{V}) = \mathrm{Pc}(\mathcal{H})$ is higher rank affine.*

Proof. The bound on the chain size implies $\mathring{\mathcal{V}}$ and $\mathring{\mathcal{H}}$ are nonempty. The cases are symmetric, so suppose $W(\mathring{\mathcal{H}}) \not\subset Z_W(W(\mathcal{V}))$. For $\mathcal{V} = \{V_i\}_{i\in I}$ and $\mathcal{H} = \{H_j\}_{j\in J}$ there are $i_0 \in I$ and $j_0 \in J$ such that $\mathfrak{r}_{V_{i_0}}$ does not commute with $\mathfrak{r}_{H_{j_0}}$, and, furthermore, $H_{j_0} \in \mathring{\mathcal{H}}$. The wall $\mathfrak{r}_{V_{i_0}}(H_{j_0})$ is distinct from V_{i_0} and H_{j_0}. The fixed point set of $\mathfrak{r}_{V_{i_0}}$ is precisely V_{i_0}, so $V_{i_0} \cap H_{j_0} \subset \mathfrak{r}_{V_{i_0}}(H_{j_0})$ and $V_{i_0} \cap \mathfrak{r}_{V_{i_0}}(H_{j_0}) \subset H_{j_0}$ and $H_{j_0} \cap \mathfrak{r}_{V_{i_0}}(H_{j_0}) \subset V_{i_0}$. As described by Lemma 5.7, there are four subchains $\mathcal{V}^+$, $\mathcal{V}^-$, $\mathcal{H}^+$, and $\mathcal{H}^-$ with respect to i_0 and j_0, and $\mathfrak{r}_{V_{i_0}}(H_{j_0})$ is transverse to at least two of them. Since j_0 is an interior index, both of $\mathcal{H}^+$ and $\mathcal{H}^-$ contain at least ℓ–many walls, and at least one of $\mathcal{V}^+$ or $\mathcal{V}^-$ does, since $|\mathcal{V}^+| + |\mathcal{V}^-| = |\mathcal{V}| + 1 \geq 2\ell$. Thus, there is at least one of $\mathcal{V}^+$, $\mathcal{V}^-$, $\mathcal{H}^+$, or $\mathcal{H}^-$ that contains at least ℓ–many walls and is transverse to $\mathfrak{r}_{V_{i_0}}(H_{j_0})$.

Suppose that $\mathcal{V}^+$ is a chain transverse to $\mathfrak{r}_{V_{i_0}}(H_{j_0})$ with at least ℓ–many walls. Then Theorem 2.19 for $\mu := \mathfrak{r}_{V_{i_0}}(H_{j_0})$ and $\mu' := H_{j_0}$ and the chain $\mathcal{V}^+$ gives that $W(H_{j_0}, \mathfrak{r}_{V_{i_0}}(H_{j_0}), \mathcal{V}^+) = W(H_{j_0}, \mathcal{V}^+)$ is a Euclidean triangle group that is contained in an irreducible affine parabolic $\mathcal{A}$, which has rank at least 3 since it contains the triangle group. Then Lemma 2.22 says $\mathfrak{r}_{H_{j_0}} \in \mathrm{Pc}(\mathcal{V}^+) \leq \mathcal{A}$. Actually, the proof of Lemma 2.22 uses the fact that a proper parabolic subgroup of an irreducible affine parabolic is finite, but $\mathrm{Pc}(\mathcal{V}^+)$ is infinite, so $\mathrm{Pc}(\mathcal{V}^+) = \mathcal{A}$ is an irreducible affine parabolic of rank at least 3 and $W(\mathcal{V}^+)$ is infinite dihedral.

Since $\mathfrak{r}_{H_{j_0}} \in \mathrm{Pc}(\mathcal{V}^+)$, we do not have $\mathrm{Pc}(\mathcal{H}) \leq \mathrm{Pc}(\mathcal{V})^\perp$, so Lemma 5.6 implies $\mathrm{Pc}(\mathcal{H}) = \mathrm{Pc}(\mathcal{V}^+)$.

The proof in the case that $\mathcal{V}^-$ is the long chain transverse to $\mathfrak{r}_{V_{i_0}}(H_{j_0})$ is exactly the same. If instead we use one of the chains $\mathcal{H}^+$ or $\mathcal{H}^-$ the only change is to take $\mu' := V_{i_0}$ in the application of Theorem 2.19. ☐

Corollary 5.9. *Let $(\mathcal{V}, \mathcal{H})$ be a grid in Σ such that $|\mathcal{V}| \geq 2\ell - 1$ and $|\mathcal{H}| \geq 2\ell$. If $W(\mathcal{V}) \ncong \mathcal{D}_\infty$ then $\mathrm{Pc}(\mathring{\mathcal{H}}) \leq \mathrm{Pc}(\mathcal{V})^\perp$.*

Proof. If $\mathrm{Pc}(\mathcal{V})$ is affine then $W(\mathcal{V})$ is dihedral, contrary to hypothesis, so Proposition 5.8 forces $W(\mathring{\mathcal{H}})$ to centralize $W(\mathcal{V})$. Apply Lemma 5.4 to $W(\mathcal{V})$ and $W(\mathring{\mathcal{H}})$, both of which are nonspherical since $|\mathcal{V}| \geq 2\ell - 1 > 2$ and $|\mathring{\mathcal{H}}| \geq 2$. Again, since $\mathrm{Pc}(\mathcal{V})$ cannot be affine, the remaining alternative is $\mathrm{Pc}(\mathring{\mathcal{H}}) \leq \mathrm{Pc}(\mathcal{V})^\perp$. ☐

Proof of Proposition 5.3. The cardinality lower bounds give that $\mathrm{Pc}(\mathcal{V})$ and $\mathrm{Pc}(\mathring{\mathcal{H}})$ are irreducible and nonspherical, by Lemma 2.7. If $W(\mathcal{V}) \ncong \mathcal{D}_\infty$ then Corollary 5.9 says $\mathrm{Pc}(\mathring{\mathcal{H}}) \leq \mathrm{Pc}(\mathcal{V})^\perp$. If $W(\mathcal{V}) \cong \mathcal{D}_\infty$ then Lemma 5.6 says either $\mathrm{Pc}(\mathcal{H})$ is contained in $\mathrm{Pc}(\mathcal{V})^\perp$ or $\mathrm{Pc}(\mathcal{V}) = \mathrm{Pc}(\mathcal{H})$ is irreducible higher rank affine. Thus, either $\mathrm{Pc}(\mathcal{V}) \times \mathrm{Pc}(\mathcal{V})^\perp$ is a product of two infinite factors, or the first factor is irreducible higher rank affine. ☐

5.2. **Avoiding wide parabolic subcomplexes.** Recall we defined constants $\mathfrak{R}$ and ℓ and functions $\rho_0(\lambda, \epsilon)$ and $\mathfrak{L}(\lambda, \epsilon)$ in Section 2.6.

Theorem 5.10. *For all λ and ϵ the following are equivalent for all admissible (λ, ϵ)–quasigeodesics $\gamma\colon I \to \Sigma$. Moreover, there are explicit bounds between the constraints of each item that are independent of γ.*

(i) For all L there exists a wide parabolic subcomplex Υ and subinterval $I' \subset I$ of length L such that $\gamma(I') \subset \bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$.

(ii) There exists R such that for all L there exists a wide parabolic subcomplex Υ and subinterval $I' \subset I$ of length L such that $\gamma(I') \subset \bar{\mathcal{N}}_R(\Upsilon)$.

(iii) For all $0 < \rho \leq \rho_0(\lambda, \epsilon)$ there exists L_0 such that for all $L_1 \geq L_0$ there exists a subinterval $I' \subset I$ of length at least L_1 such that every chain $\mathcal{V}$ transverse to $\gamma(I')$ with $|\mathcal{V}|/|I'| \geq \rho$ contains a pair of consecutive walls mutually transverse to an infinite chain.

(iv) For all $0 < \rho \leq \rho_0(\lambda, \epsilon)$ there exists L_0 such that for all $L_1 \geq L_0$ there exists a subinterval $I' \subset I$ of length at least L_1 such that every chain $\mathcal{V}$ transverse to $\gamma(I')$ with $|\mathcal{V}|/|I'| \geq \rho$ contains a pair of consecutive walls mutually transverse to a chain of length 2ℓ.

(v) For all L' there exists a grid $(\mathcal{V}, \mathcal{H})$ with $\mathcal{V} \pitchfork \gamma$, $|\mathcal{V}| \geq L'$, and $|\mathcal{H}| = 2\ell$.

The remainder of the subsection proves the theorem in steps. The implication (i) $\Longrightarrow$ (ii) follows by choosing $R = \mathfrak{R}$. Lemma 5.11 gives (ii) $\Longrightarrow$ (iii). The implication (iii) $\Longrightarrow$ (iv) follows by passing to a subchain. Lemma 5.12 gives (iv) $\Longrightarrow$ (v). The proof for (v) $\Longrightarrow$ (i) is broken into three pieces: Lemma 5.13 is the inductive step for Proposition 5.14, which is then applied in Lemma 5.15.

We will see in the proof of Theorem 1.4 that Theorem 5.10 characterizes when γ fails to be Morse. Notice that the items of Theorem 5.10 are impossible to satisfy if I is bounded, because each of them asserts the existence of arbitrarily long subintervals with some property. This is the expected behavior; bounded quasigeodesic segments are always Morse.

Lemma 5.11. *Suppose $\gamma\colon I \to \Sigma$ is an admissible (λ,ϵ)–quasigeodesic such that there exists R such that for all L there exists a wide parabolic subcomplex Υ and a subinterval $I' \subset I$ of length L such that $\gamma(I') \subset \bar{\mathcal{N}}_R(\Upsilon)$. Then for all $0 < \rho \leq \rho_0(\lambda,\epsilon)$ there exists L_0 such that for all $L_1 \geq L_0$ there exists a subinterval $I' \subset I$ of length at least L_1 such that every chain $\mathcal{V}$ transverse to $\gamma(I')$ with $|\mathcal{V}|/|I'| \geq \rho$ contains a pair of consecutive walls mutually transverse to a biinfinite chain.*

Proof. Fix R as in the hypothesis and take any $0 < \rho \leq \rho_0(\lambda,\epsilon)$. By Lemma 2.15 there exists $L_0 \geq \mathfrak{L}(\lambda,\epsilon)$ such that for all $x, y \in \Sigma$:

$$d_2(x,y) \leq R \implies d_c(x,y) \leq \rho L_0/4 - 2$$

By hypothesis, given $L_1 \geq L_0$ there exists a wide parabolic subcomplex Υ and $I' \subset I$ of length at least L_1 such that $\gamma(I') \subset \bar{\mathcal{N}}_R(\Upsilon)$. Suppose that $\mathcal{V}$ is a chain transverse to $\gamma(I')$ with $|\mathcal{V}|/|I'| \geq \rho$.

Let $\bar{I}' = [z_0, z_1]$. Consider the concatenation of d_2–geodesics $[\gamma(z_0), \pi_\Upsilon(\gamma(z_0))] + [\pi_\Upsilon(\gamma(z_0)), \pi_\Upsilon(\gamma(z_1))] + [\pi_\Upsilon(\gamma(z_1)), \gamma(z_1)]$ and apply the 4–gon trick, Lemma 3.5, to find $\mathcal{V}'$ consisting of walls of $\mathcal{V}$ that separate $\pi_\Upsilon(z_0)$ and $\pi_\Upsilon(z_1)$:

$$|\mathcal{V}| - |\mathcal{V}'| \leq d_c(\gamma(z_0), \pi_\Upsilon(\gamma(z_0))) + d_c(\gamma(z_1), \pi_\Upsilon(\gamma(z_1))) + 2 \leq 2(\rho L_0/4 - 2) + 2$$

But $|\mathcal{V}| \geq \rho L_1$, so $|\mathcal{V}'| \geq |\mathcal{V}|/2 + 2$. In particular, there is at least one consecutive pair V,V' of walls of $\mathcal{V}$ that cut $[\pi_\Upsilon(z_0), \pi_\Upsilon(z_1)] \subset \Upsilon$.

If Υ is a product of two infinite factors then up to translation by W we may assume $\Upsilon = \Sigma_{T \sqcup T'}$ where T and T' are subsets of S with no edges connecting them in the Coxeter diagram. A wall M cutting Υ has reflection $\mathfrak{r}_M \in W_T \times W_{T'}$. Two walls M and M' cutting Υ give a reflection subgroup $\langle \mathfrak{r}_M, \mathfrak{r}_{M'} \rangle$ that is finite if and only if $M \pitchfork M'$. Lemma 2.8 says $\langle \mathfrak{r}_M, \mathfrak{r}_{M'} \rangle$ splits as a direct product of reflection subgroups of the two factors. If $\mathfrak{r}_M$ and $\mathfrak{r}_{M'}$ live in opposite factors then $\langle \mathfrak{r}_M, \mathfrak{r}_{M'} \rangle \cong (\mathbb{Z}/2\mathbb{Z})^2$. If they live in the same factor, say W_T, then $\langle \mathfrak{r}_M, \mathfrak{r}_{M'} \rangle$ is a subgroup of W_T, so M and M' cross in Σ if and only if they cross in Σ_T. Thus, walls from $\mathcal{V}$, being disjoint in Σ, either all have reflections in W_T or all have reflections in $W_{T'}$. Without loss of generality, assume the former. By hypothesis $W_{T'}$ is infinite, so there exist walls H and H' of Σ such that $\langle \mathfrak{r}_H, \mathfrak{r}_{H'} \rangle \cong \mathcal{D}_\infty \leq W_{T'}$. Then the $\langle \mathfrak{r}_H, \mathfrak{r}_{H'} \rangle$ orbits of H and H' give a chain in Σ consisting of walls with reflections in $W_{T'}$. These are transverse to $\{V, V'\}$, since every reflection in W_T commutes with every reflection in $W_{T'}$.

The other option is that up to W–translation $\Upsilon = \Sigma_{T \sqcup T'}$ where T is irreducible higher rank affine and T' is spherical. Walls of $\mathcal{V}$ are disjoint, so they have reflections in W_T. Since W_T is affine type, Σ_T is a polyhedral tiling of $\mathbb{E}^n$ for some $n > 1$, and walls of Σ_T are affine codimension–1 hyperplanes that come in finitely many parallel families, with every wall from one family crossing every wall from every other family. Since the walls of $\mathcal{V}$ are disjoint in Σ, their intersections with Σ_T all belong to the same parallel family. Any one of the other parallel families gives a chain transverse to $\Sigma_T \cap \mathcal{V}$, so taking the corresponding walls of Σ gives an infinite chain transverse to $\{V, V'\}$. □

Lemma 5.12. *Suppose $\gamma\colon I \to \Sigma$ is an admissible (λ,ϵ)–quasigeodesic such that for all $0 < \rho \leq \rho_0(\lambda,\epsilon)$ there exists L_0 such that for all $L_1 \geq L_0$ there exists a subinterval $I' \subset I$ of length at least L_1 such that every chain $\mathcal{V}$ transverse to $\gamma(I')$ with $|\mathcal{V}|/|I'| \geq \rho$ contains a pair of consecutive walls mutually transverse to a chain of length 2ℓ. Then for all L' there exists a grid $(\mathcal{V}, \mathcal{H})$ with $\mathcal{V} \pitchfork \gamma$, $|\mathcal{V}| \geq L'$, and $|\mathcal{H}| = 2\ell$.*

Proof. Assume $L' \geq 2$, and set $\rho := \rho_0(\lambda,\epsilon)/(L'-1)$. Take $L_1 := \max\{\mathfrak{L}(\lambda,\epsilon), L_0(\rho)\}$. The definitions of $\mathfrak{L}(\lambda,\epsilon)$ and $\rho_0(\lambda,\epsilon)$ imply that for any subinterval $I' \subset I$ of length at least L_1, there does exist a chain $\mathcal{V} = \{V_i\}_{i_0 \leq i \leq i_1}$ transverse to $\gamma(I')$ with $|\mathcal{V}|/|I'| \geq \rho_0(\lambda,\epsilon)$. Consider the subchain $\{V_{j_k}\}_{0 \leq k \leq (i_1 - i_0)/(L'-1)}$ where $j_k = i_0 + (L'-1)k$, which contains at least $\rho|I'|$–many walls, so, by hypothesis, contains a pair of consecutive walls V and V', transverse to a chain $\mathcal{H}$ of length 2ℓ. Let $\mathcal{V}'$ be the subchain of $\mathcal{V}$ consisting of all walls from V to V'. By construction, $|\mathcal{V}'| = L'$, and the two extremes are transverse to $\mathcal{H}$, so $\mathcal{V}' \pitchfork \mathcal{H}$. □

Lemma 5.13. *Let γ be an admissible path, and let $\mathcal{V}$ be a chain such that $\mathcal{V} \pitchfork \gamma$ and $|\mathcal{V}| \geq 2\ell$. Let Υ be the unique parabolic subcomplex with stabilizer $\mathrm{Pc}(\mathcal{V}) \times \mathrm{Pc}(\mathcal{V})^\perp$ (recall Corollary 2.14). Suppose $\mathcal{V} = \{V_i\}_{i_0 \leq i \leq i_2}$. For any $i_0 < i_1 < i_2$, let $\mathcal{V}^- := \{V_i\}_{i_0 \leq i \leq i_1}$ and $\mathcal{V}^+ := \{V_i\}_{i_1+1 \leq i \leq i_2}$. Suppose that there is a point x on the closed subsegment of γ between V_{i_1} and V_{i_1+1} such that $d_c(x, \Upsilon) \geq 4\ell - 1$. Then at least one of $\mathcal{V}^+$ or $\mathcal{V}^-$ is transverse to a chain $\mathcal{H}$ with $|\mathcal{H}| \geq 2\ell$, $\mathcal{H} \pitchfork \gamma$, and $\mathfrak{r}_H \notin \mathrm{Pc}(\mathcal{V}) \times \mathrm{Pc}(\mathcal{V})^\perp$ for all $H \in \mathcal{H}$.*

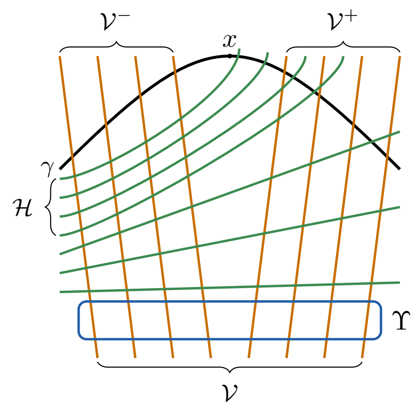

FIGURE 4. Inductive step: A point $x \in \gamma$ between the extremes of $\mathcal{V}$ and far from Υ gives a splitting of $\mathcal{V}$ into subchains $\mathcal{V}^+$ and $\mathcal{V}^-$ such that one of them is part of a grid with a chain $\mathcal{H}$ consisting of 'most' of the walls separating x from Υ; Only at most $2\ell - 1$ of the walls separating x from Υ can be transverse to all of $\mathcal{V}$ without cutting Υ.

Proof. By hypothesis, there is a chain $\mathcal{M}$ with $|\mathcal{M}| \geq 4\ell - 1$ separating x from Υ. Every wall $M \in \mathcal{M}$ is disjoint from $\{x\} \cup \Upsilon$. Every reflection in $\mathrm{Pc}(\mathcal{V}) \times \mathrm{Pc}(\mathcal{V})^\perp$ cuts Υ, so $\mathfrak{r}_M \notin \mathrm{Pc}(\mathcal{V}) \times \mathrm{Pc}(\mathcal{V})^\perp$ for all $M \in \mathcal{M}$.

Let the subsegment of γ between the extreme walls of $\mathcal{V}$ be $\gamma^- + \gamma^+$ with concatenation point at x and labelled so that $\mathcal{V}^- \pitchfork \gamma^-$, $\mathcal{V}^+ \pitchfork \gamma^+$, and $\{x\} = \gamma^- \cap \gamma^+$.

For every $i \leq i_1$ there is a path δ_i from Υ to x that travels from a point in $V_i \cap \Upsilon$ through V_i to $V_i \cap \gamma \in \gamma^-$, and then along γ^- to x. The same is true for $i \geq i_1 + 1$, with γ^+ replacing γ^-. Since $M \in \mathcal{M}$ separates x from Υ, it cuts every such path. If M does not cut γ^- then for $i_0 \leq i \leq i_1$, M cuts the V_i segment of δ_i, so $M \pitchfork \mathcal{V}^-$. Similarly, if M does not cut γ^+ then $M \pitchfork \mathcal{V}^+$. Since M cuts γ at most once, and does not do so at x, at least one of these is true. Moreover, suppose some wall $M \in \mathcal{M}$ cuts γ^+. If there is a wall $M' \in \mathcal{M}$ separating M from x then it must cut the subsegment of γ^+ between x and $M \cap \gamma^+$, so M' cuts γ^+ and does not cut γ^-. Suppose there is a wall $M' \in \mathcal{M}$ that is separated from x by M. Since M cuts γ^+ it does not cut γ^-, so M' also does not cut γ^-, else we would have a path from M' to x through γ^- avoiding M. Thus, if $\mathcal{M}$ is ordered from Υ toward x, then $\mathcal{M}$ splits into an initial subchain $\mathcal{M}'$ consisting of walls disjoint from γ and a terminal subchain $\mathcal{M}''$ consisting of walls that either all cut γ^+ or all cut γ^-.

Suppose $|\mathcal{M}'| \geq 2\ell$. Since $\mathcal{M}'$ is disjoint from γ it is transverse to $\mathcal{V}$. Apply Proposition 5.3 to $(\mathcal{V}, \mathcal{M}')$. Then for $M \in \mathring{\mathcal{M}}'$ we have $\mathfrak{r}_M \in \mathrm{Pc}(\mathcal{V}) \times \mathrm{Pc}(\mathcal{V})^\perp$, which is a contradiction. Thus, $|\mathcal{M}'| \leq 2\ell - 1$, which implies $|\mathcal{M}''| \geq 2\ell$. Take $\mathcal{H} := \mathcal{M}''$. If $\mathcal{M}'' \pitchfork \gamma^+$ then $\mathcal{H} \pitchfork \mathcal{V}^-$, and if $\mathcal{M}'' \pitchfork \gamma^-$ then $\mathcal{H} \pitchfork \mathcal{V}^+$. $\square$

Proposition 5.14. *Let $C := 3^{|S|-3}$. For every admissible path γ, for every grid $(\mathcal{V}, \mathcal{H})$ with $\mathcal{V} \pitchfork \gamma$, $|\mathcal{V}| \geq C(2\ell - 1)$, and $|\mathcal{H}| \geq 2\ell$, there exists a wide parabolic subcomplex Υ and a convex subchain $\mathcal{V}'$ of $\mathcal{V}$ such that $|\mathcal{V}'| \geq |\mathcal{V}|/C$ and the subsegment of γ between the extreme walls of $\mathcal{V}'$ is contained in $\bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$.*

Proof. Suppose there is a grid $(\mathcal{V}, \mathcal{H})$ satisfying the hypotheses. The proof is by induction on rank: we guess a wide product parabolic subcomplex and subsegment of γ, and if the lemma is not satisfied we change the guesses in such a way that the rank of the first factor of the parabolic decreases.

Let $\mathcal{V}_0 := \mathcal{V} = \{V_i\}_{i_0 \leq i \leq i_1}$ and $\mathcal{H}_0 := \mathcal{H}$. Let $j_0 := i_0$, $j_1 := i_0 + \lfloor \frac{i_1 - i_0}{3} \rfloor$, $j_2 := j_1 + \lfloor \frac{i_1 - i_0}{3} \rfloor$, and $j_3 := i_1$. Let $\mathcal{V}_{01} := \{V_j\}_{j_0 \leq j \leq j_1}$, $\mathcal{V}_{02} := \{V_j\}_{j_1 \leq j \leq j_2}$, and $\mathcal{V}_{03} := \{V_j\}_{j_2 \leq j \leq j_3}$; each of these is a convex subchain of $\mathcal{V}_0$ containing at least $\frac{|\mathcal{V}_0|}{3}$-many walls. Define Υ_0 to be the unique parabolic subcomplex with stabilizer $\mathrm{Pc}(\mathcal{V}_0) \times \mathrm{Pc}(\mathcal{V}_0)^\perp$, which is wide by Proposition 5.3 applied to $(\mathcal{V}, \mathcal{H})$. Consider the subsegment γ' of γ between the extremes of $\mathcal{V}_{02}$. If $\gamma' \subset \bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon_0)$ then we are done, since γ' is transverse to $\mathcal{V}_{02}$ and $|\mathcal{V}_{02}| \geq |\mathcal{V}|/3 \geq |\mathcal{V}|/C$. Otherwise, let $x \in \gamma'$ be a point maximizing the distance to Υ_0. By hypothesis $d_2(x, \Upsilon_0) > \mathfrak{R}$, so, by definition of $\mathfrak{R}$, there is a chain $\mathcal{M}$ separating x from Υ_0 with $|\mathcal{M}| \geq 4\ell - 1$. By definition of γ', there exists $j_1 \leq k < j_2$ such that x occurs on the closed subsegment of γ' between V_k and V_{k+1}. Define $\mathcal{V}_0^- := \{V_i\}_{j_0 \leq i \leq k}$ and $\mathcal{V}_0^+ := \{V_i\}_{k+1 \leq i \leq j_3}$. Apply Lemma 5.13 to produce a chain $\mathcal{M}' \subset \mathcal{M}$ with $|\mathcal{M}'| \geq 2\ell$ that is transverse to either $\mathcal{V}_0^-$ or $\mathcal{V}_0^+$. Define $\mathcal{H}_1' := \mathcal{M}'$, and let $\mathcal{V}_1$ be whichever of $\mathcal{V}_0^-$ or

$\mathcal{V}_0^+$ is transverse to $\mathcal{H}_1'$. Observe that $\mathcal{V}_{01} \subset \mathcal{V}_0^-$ and $\mathcal{V}_{03} \subset \mathcal{V}_0^+$, so in either case $|\mathcal{V}_1| \geq |\mathcal{V}_0|/3 \geq 2\ell - 1$. Also, Lemma 5.13 says reflections in walls of $\mathcal{H}_1'$ are not in $\mathrm{Pc}(\mathcal{V}_0) \times \mathrm{Pc}(\mathcal{V}_0)^\perp$.

If $W(\mathcal{V}_0) \cong \mathcal{D}_\infty$ we get a contradiction, because $\mathcal{V}_1$ contains more than 8 walls of $\mathcal{V}_0$, so Lemma 2.20 upgrades $\mathcal{H}_1' \pitchfork \mathcal{V}_1$ to $\mathcal{H}_1' \pitchfork \mathcal{V}_0$, but then Lemma 5.6 says $\mathrm{Pc}(\mathcal{H}_1')$ is either contained in $\mathrm{Pc}(\mathcal{V}_0)$ or $\mathrm{Pc}(\mathcal{V}_0)^\perp$. Thus, if $W(\mathcal{V}_0) \cong \mathcal{D}_\infty$ then $\gamma' \subset \bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon_0)$, and we are done.

Otherwise, Proposition 5.3 says $\mathrm{Pc}(\mathcal{V}_1) \times \mathrm{Pc}(\mathcal{V}_1)^\perp$ is wide and contains reflections through all walls of $\mathcal{H}_1 := \mathring{\mathcal{H}}_1'$. Moreover, if $W(\mathcal{V}_1) \not\cong \mathcal{D}_\infty$ then $\mathrm{Pc}(\mathcal{V}_1)$ is not higher rank irreducible affine, so $\mathrm{Pc}(\mathcal{H}_1) \leq \mathrm{Pc}(\mathcal{V}_1)^\perp$. Now, $\mathcal{V}_1 \subset \mathcal{V}_0$ implies $\mathrm{Pc}(\mathcal{V}_1) \leq \mathrm{Pc}(\mathcal{V}_0)$, which implies $\mathrm{Pc}(\mathcal{V}_0)^\perp \leq \mathrm{Pc}(\mathcal{V}_1)^\perp$. In the $W(\mathcal{V}_1) \not\cong \mathcal{D}_\infty$ case these cannot be equalities, because in that case we know that for all $H \in \mathcal{H}_1$ we have $\mathfrak{r}_H \in \mathrm{Pc}(\mathcal{V}_1)^\perp \smallsetminus \mathrm{Pc}(\mathcal{V}_0)^\perp$. Thus $W(\mathcal{V}_1) \not\cong \mathcal{D}_\infty \implies \mathrm{Pc}(\mathcal{V}_1) \lneq \mathrm{Pc}(\mathcal{V}_0)$.

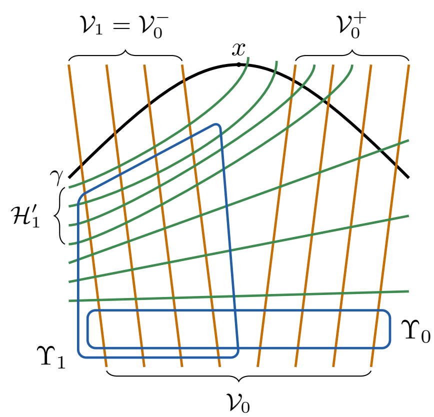

FIGURE 5. Using the induction step to produce a new candidate wide parabolic subcomplex Υ_1 corresponding to $\mathcal{V}_1 \subset \mathcal{V}_0$.

Define Υ_1 to be the unique parabolic subcomplex with stabilizer $\mathrm{Pc}(\mathcal{V}_1) \times \mathrm{Pc}(\mathcal{V}_1)^\perp$, and repeat the construction. See Figure 5. Either the induction will stop with the segment of γ corresponding to the middle third of $\mathcal{V}_1$ contained in $\bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon_1)$, which is guaranteed to be the case if $W(\mathcal{V}_1) \cong \mathcal{D}_\infty$, or $\mathrm{Pc}(\mathcal{V}_1) \lneq \mathrm{Pc}(\mathcal{V}_0)$ and we produce a further $\mathcal{V}_2$, $\mathcal{H}_2$, Υ_2, and repeat.

The ranks of parabolics in a strictly decreasing sequence are strictly decreasing, so this iteration terminates in at most $|S|$ steps. In fact, it terminates in at most $|S| - 4$ steps, since $\mathrm{Pc}(\mathcal{V}_0)^\perp$ has rank at least 2, so $\mathrm{Pc}(\mathcal{V}_0)$ has rank at most $|S| - 2$, while each $\mathrm{Pc}(\mathcal{V}_i)$ is nonspherical, so also has rank at least 2. Thus, in $0 \leq i \leq |S| - 4$ steps we produce a wide parabolic subcomplex Υ_i such that $\bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon_i)$ contains a subsegment of γ transverse to the walls of $\mathcal{V}_{i2}$, where:

$$|\mathcal{V}_{i2}| \geq \frac{|\mathcal{V}_i|}{3} \geq \frac{|\mathcal{V}|}{3^{i+1}} \geq \frac{|\mathcal{V}|}{3^{|S|-3}} = \frac{|\mathcal{V}|}{C}$$ □

Lemma 5.15. *Suppose $\gamma\colon I \to \Sigma$ is an admissible (λ, ϵ)–quasigeodesic such that for all L' there exists a grid $(\mathcal{V}, \mathcal{H})$ with $\mathcal{V} \pitchfork \gamma$, $|\mathcal{V}| \geq L'$, and $|\mathcal{H}| = 2\ell$. Then for all L there exists a wide parabolic subcomplex Υ and subinterval $I' \subset I$ of length at least L such that $\gamma(I') \subset \bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$.*

Proof. Fix arbitrary L. Let C be the constant of Proposition 5.14. By Lemma 2.15, we can choose L' sufficiently large that for all $x, y \in \Sigma$:

$$d_c(x, y) \geq L'/C \implies d_2(x, y) \geq \lambda L + \epsilon$$

Set $L'' := \max\{C(2\ell - 1), L'\}$. By hypothesis, there exists a grid $(\mathcal{V}, \mathcal{H})$ with $\mathcal{V} \pitchfork \gamma$, $|\mathcal{V}| \geq L''$, and $|\mathcal{H}| = 2\ell$. Since $|\mathcal{V}| \geq C(2\ell - 1)$, we can apply Proposition 5.14 to $(\mathcal{V}, \mathcal{H})$ to get a wide parabolic subcomplex Υ and subinterval $I' \subset I$ such that $\gamma(I') \subset \bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$ is transverse to $|\mathcal{V}|/C \geq L'/C$-many walls of $\mathcal{V}$. Then the choice of L' guarantees $|I'| \geq L$. □

6. Characterization of the Morse property for admissible quasigeodesics

In this section we prove Theorem 1.4 and Corollary 1.5 from Section 1.

6.1. **Proof of Theorem 1.4.** Recall that the goal is to characterize the Morse property for admissible quasigeodesics by items (1)-(5) of Theorem 1.4.

Theorem 4.6 characterized the Morse property in terms of equivalent items (a)-(e). In particular, (1)=(b) is the Morse property, and (2)=(e) describes the interaction of the quasigeodesic with dense, transverse chains.

Theorem 5.10 gave equivalent items (i)-(v) describing the interaction of the quasigeodesic with wide parabolics, grids, and dense, transverse chains. Now we claim that these conditions characterize being non-Morse. In particular, (3)= ¬(iv), (4)= ¬(i), (5) = ¬(v).

To finish the proof of Theorem 1.4 we relate the three transverse, dense chain statements: The negation of (2) is that for every choice of C and ρ there exists L_0 such that for all $L_1 \geq L_0$ there exists a subinterval $I' \subset I$ of length at least L_1 such that every chain $\mathcal{V} \pitchfork \gamma(I')$ with density at least ρ has some pair of consecutive walls that are mutually transverse to a chain $\mathcal{H}$ with $|\mathcal{H}| > C$. Item (iii) says the same thing with $\mathcal{H}$ infinite, so (iii) $\implies$ ¬(2). Item (iv) says the same thing specifically for $C = 2\ell - 1$, so ¬(2) $\implies$ (iv). This sandwiches ¬(2) between the equivalent statements (iii) and (iv), so ¬(2) is equivalent to both of them, hence to all of (i)-(v).

Conversely, (1)=(b)=(e)=(2) are all equivalent to the negation of any of (i)-(v), so (1) and (2) are equivalent to (3)= ¬(iv), (4)= ¬(i), and (5) = ¬(v). This completes the proof of Theorem 1.4.

6.2. **Proof of Corollary 1.5.** We call a subset of S wide/affine/spherical according to the corresponding property of the special subgroup it generates.

The implications (1) $\implies$ (2) and (2) $\iff$ (3) are easy. We will prove (3) $\implies$ (1) by showing ¬(1) $\implies$ ¬(3). There are two parts to the proof. The first is to go from a combinatorial geodesic having a long subsegment contained in the $\mathfrak{R}$–neighborhood of a wide parabolic subcomplex $\Upsilon = w\Sigma_T$ to having a long subsegment such that every edge e_i is dual to a wall that cuts Υ. The second part is to argue that the edge labels of the e_i are all contained in a common wide subset of S. The complication in this step is that in non-right-angled Coxeter groups the edges dual to a fixed wall might not all have the same label; the labels might only be conjugate generators. This means the subset of S containing the labels of the edges e_i might not be T—it might be some subset of S related to T by a conjugation.

For the first step, we start with a counting argument. Suppose that $\mathcal{M}$ is the set of walls dual to edges of a combinatorial geodesic segment α, and suppose $\mathcal{M}'$ is some subset of those walls. Partition α into a concatenation of subpaths that are either single edges dual to a wall in $\mathcal{M} \setminus \mathcal{M}'$ or a maximal subpath such that every edge is dual to a wall in $\mathcal{M}'$. Suppose that the lengths of the latter are bounded above by N. There are at most $|\mathcal{M}| - |\mathcal{M}'| + 1$ such subsegments, so $|\mathcal{M}| = |\alpha| \leq N(|\mathcal{M}| - |\mathcal{M}'| + 1) + |\mathcal{M}| - |\mathcal{M}'|$. This implies:

$$(\clubsuit) \qquad \frac{|\mathcal{M}'|}{|\mathcal{M}| - |\mathcal{M}'| + 1} \leq N$$

Suppose $\gamma\colon I \to \Sigma$ is a combinatorial geodesic. As in Lemma 3.2, there are λ and ϵ such that γ is an admissible (λ, ϵ)–quasigeodesic. Furthermore, γ is injective and parameterized by arc length, and the quasigeodesic constants are uniform over all combinatorial geodesics.

Suppose ¬(1), so γ is not Morse. It follows from Theorem 1.4 that for every L there exists a wide parabolic subcomplex Υ and subinterval $I' \subset I$ of length L such that $\gamma' := \gamma(I')$ is an edge path contained in $\bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$. Let P be the parabolic subgroup stabilizing Υ, so that walls that cut Υ have reflection in P. Let x and y be the end vertices of γ', and consider the set of all walls $\mathcal{M}$ separating x and y, of which there are exactly L–many, since they are the walls cutting the edges of γ'. Let $\mathcal{M}'$ be the subset of $\mathcal{M}$ consisting of walls that cut Υ. A d_w–version of the 4–gon trick, Lemma 3.5, gives:

$$|\mathcal{M}| - |\mathcal{M}'| \preceq d_w(x, \pi_\Upsilon(x)) + d_w(y, \pi_\Upsilon(y)) \preceq 2\mathfrak{R}$$

Since the right-hand side is a fixed constant, this gives $|\mathcal{M}'| \stackrel{+}{\asymp} |\mathcal{M}| = L$, so:

$$(\spadesuit) \qquad \frac{|\mathcal{M}'|}{|\mathcal{M}| - |\mathcal{M}'| + 1} \succeq L$$

The coarse comparison constants of ($\spadesuit$) are independent of the choices of L and γ, so, given N we can choose L large enough that ($\spadesuit$) implies ¬($\clubsuit$), which guarantees that γ' contains a subsegment γ'' consisting of at least N–many consecutive edges dual to walls of $\mathcal{M}'$. Thus, γ contains arbitrarily long subsegments $\alpha := \gamma''$

such that there exists a wide parabolic subcomplex Υ such that $\alpha \subset \bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$ is an edge path, every edge of which is dual to a wall that cuts Υ.

For the second step, suppose a and b are the first and last vertices of α, respectively. Pick a vertex $c \in \Upsilon^{(0)}$ that is d_1–closest to a. Replace Υ, α, a, and b by their c^{-1}–translates. Then $\Upsilon^{(0)} = W_T$ for some wide $T \subset S$. Suppose the vertices of α are $a = x_0, x_1, \dots, x_n = b$ and e_i is the edge from x_{i-1} to x_i. Suppose the label of e_i is $s_i \in S$. Let $U := \{s_1, \dots, s_n\}$, so that the vertices of α are contained in the parabolic subcomplex aW_U.

Since the wall dual to e_i cuts Υ, its reflection $r_i := x_i s_i x_i^{-1}$ is in W_T.

Now, $x_1 = r_1 x_0 \in W_T x_0$. By induction, $x_j = r_j x_{j-1} = r_j \cdots r_1 x_0 \in W_T x_0$. Plugging this expression for x_j into the definition of r_j gives $r_j = r_j \cdots r_1 x_0 s_j x_0^{-1} r_1 \cdots r_j$. Rearranging gives $s_j = x_0^{-1} r_1 \cdots r_{j-1} r_j r_{j-1} \cdots r_1 x_0 \in x_0^{-1} W_T x_0 = a^{-1} W_T a$.

The choice of c makes a the minimal length element of $W_T a$, so for all $u \in W_T$ we have $|ua| = |u| + |a|$ [41, Lemma 4.3.3]. Since $W_T \ni w_j := a s_j a^{-1}$, we have $w_j a = a s_j$. The left-hand side is an element of $W_T a$, so the minimal length words in S representing this element have length $|w_j| + |a|$. The right-hand side is the same group element, expressed as a word of length $|a| + 1$. Thus $|w_j| = 1$, which means $w_j \in T$. We have shown $U \subset S \cap a^{-1} T a$.

Our goal is to show that U is contained in a wide subset of S, establishing $\neg(3)$. This is true, in particular, if U itself is wide, so suppose not. Once α is longer than the maximal d_1–diameter of a spherical subgroup of (W, S) then U cannot be spherical. Consider the canonical decomposition of W_U into irreducible components. If U is nonspherical and not wide then there is a unique nonspherical irreducible component U_0 of U, which, moreover, is not higher rank affine. Let $U_1 := U \setminus U_0$. This implies $U_1 \subset U_0^\perp$, so $W_U \leq W_{U_0} \times W_{U_0^\perp}$. Thus, it suffices to show $U_0^\perp$ is nonspherical.

By Proposition 2.9 (1), since $U_0 \subset S$ is irreducible and nonspherical, $aU_0a^{-1} \subset T \subset S$ implies $U_0 = aU_0a^{-1} \subset T$. Let T_0 be the irreducible component of T containing U_0. If T_0 is higher rank affine then $T_0 = U_0$, since irreducible affine parabolics have no proper nonspherical subparabolics. This is impossible, since we already said that U_0 is not higher rank affine. Since T is wide, the alternative is that $T_1 := T \setminus T_0$ is nonspherical. But then $U_0 \subset T_0$ implies $T_1 \subset T_0^\perp \subset U_0^\perp$, so $U_0^\perp$ is nonspherical. This completes the proof of Corollary 1.5.

7. Hyperbolic models

In Section 7.1 we use fine chains, as in Theorem 1.4, to define a metric d_f on X_{NR}, which we show is hyperbolic. Since it is defined in terms of walls, it also gives a hyperbolic metric on Σ.

A hyperbolic metric on Σ must somehow eliminate the known sources of non-hyperbolicity, the wide parabolics. In fact, in Section 7.2 we observe that it collapses wide parabolics to diameter 1 subsets. This has the side effect, noted in Proposition 7.11, that it collapses the Niblo-Reeves cube complex to within bounded distance of the canonical copy of $\Sigma^{(0)}$ in $X_{\mathrm{NR}}^{(0)}$.

Using Theorem 1.4 to relate Morse geodesics with fine chains along them, we deduce that (Σ, d_f) is a *Morse recognizing space*, in Section 7.3.

A common strategy for producing hyperbolic spaces is to cone off known non-hyperbolic regions to make them have bounded diameter. In Section 7.4 we construct such a *coned-off space* $\widehat{\Sigma}$ by coning off the wide parabolic subcomplexes of Σ and show that it is W–equivariantly quasiisometric to (Σ, d_f). This gives a more concrete construction of a hyperbolic Morse recognizing space for W.

7.1. **The fine chain metric.** Recall, from Definition 1.3, that a chain $\mathcal{V} := \{V_i\}_{i \in I}$ is *fine* if for all $i \neq j$, or, equivalently, for all consecutive i, j, every chain transverse to $\{V_i, V_j\}$ has length less than 2ℓ.

Definition 7.1. Define $d_f \colon \Sigma^{(0)} \times \Sigma^{(0)} \to \mathbb{R}$ by:

$$d_f(x, y) := \max\{|\mathcal{V}| \mid \mathcal{V} \text{ is a fine chain separating } x \text{ and } y\}$$

Make the same definition for $d_f \colon X_{\mathrm{NR}}^{(0)} \times X_{\mathrm{NR}}^{(0)} \to \mathbb{R}$.

Recall that a vertex of $\Sigma^{(0)}$ defines a principal ultrafilter on the halfspaces system that for each wall picks the complementary halfspace containing that vertex. This gives a canonical inclusion $\Sigma^{(0)} \hookrightarrow X_{\mathrm{NR}}^{(0)}$ with the additional property that a wall separates two vertices of $\Sigma^{(0)}$ if and only if the corresponding wall in X_{NR} separates the $X_{\mathrm{NR}}^{(0)}$ images of those vertices. Thus, for $x, y \in \Sigma^{(0)}$ the distance $d_f(x, y)$ is the same measured in Σ as in X_{NR}.

In this section we will show that d_f defines a hyperbolic metric on $(X^{(0)}_{\mathrm{NR}}, d_f)$. This is a variation of a construction of Genevois. We further show that combinatorial geodesics are unparameterized quasirulers that can be reparameterized to be rough geodesic quasirulers, and that d_f extends to a hyperbolic metric on Σ.

Lemma 7.2. *d_f is a metric on $\Sigma^{(0)}$ and on $X^{(0)}_{\mathrm{NR}}$.*

Proof. The proof is the same, so we do it in X_{NR}. A single wall is a fine chain, so $x \neq y \in X^{(0)}_{\mathrm{NR}}$ implies $d_f(x, y) \geq 1$.

Vertices of X_{NR} do not lie on walls, so for all x, y, z in $X^{(0)}_{\mathrm{NR}}$, if $\mathcal{V}$ is a chain realizing $d_f(x, z)$ then $\mathcal{V} = \mathcal{V}' \sqcup \mathcal{V}''$ where $\mathcal{V}'$ is an initial subchain separating x from y and $\mathcal{V}''$ is a terminal subchain separating y from z. Both $\mathcal{V}'$ and $\mathcal{V}''$ are fine, so $d_f(x, z) = |\mathcal{V}| = |\mathcal{V}'| + |\mathcal{V}''| \leq d_f(x, y) + d_f(y, z)$. □

Lemma 7.3 (cf [48, Lemma 6.54])**.** *Combinatorial geodesics in X_{NR} are unparameterized $(4\ell - 1)$–quasirulers in $(X^{(0)}_{\mathrm{NR}}, d_f)$.*

Proof. Let x, y, and z be vertices on a combinatorial geodesic of X_{NR} such that y occurs between x and z. Let $\mathcal{V} := \{V_0, V_1, \ldots, V_m\}$ be a longest fine chain separating x and y and let $\mathcal{H} := \{H_0, H_1, \ldots, H_n\}$ be a longest fine chain separating y and z. Assume that indices and halfspaces are chosen so that for all i and j, $y \in V_i^- \subset V_{i+1}^-$ and $y \in H_j^- \subset H_{j+1}^-$.

Suppose for some i_0 and j_0 we have that V_{i_0} is disjoint from H_{j_0}. Since $y \in V_{i_0}^- \cap H_{j_0}^-$, we have $V_{i_0}^+ \cap H_{j_0}^+ = \emptyset$. Thus, for all $i > i_0$ and $j > j_0$ we have V_i and H_j are disjoint, since $V_i \subset V_{i_0}^+$ and $H_j \subset H_{j_0}^+$. Conversely, if $V_{i_0} \pitchfork H_{j_0}$ then $V_i \pitchfork H_j$ for all $i \leq i_0$ and all $j \leq j_0$. In particular, if $\max\{i_0, j_0\} \geq 2\ell - 1$ this contradicts fineness of $\mathcal{V}$ or $\mathcal{H}$, so V_i is disjoint from H_j for all $i, j \geq 2\ell - 1$. Then $\mathcal{K} := \{V_m, V_{m-1}, \ldots, V_{2\ell-1}, H_{2\ell}, H_{2\ell+1}, \ldots, H_n\}$ is a chain. It is fine because $H_{2\ell-1}$ separates $V_{2\ell-1}$ from $H_{2\ell}$, so any chain transverse to $V_{2\ell-1}$ and $H_{2\ell}$ is transverse to $H_{2\ell-1}$ and $H_{2\ell}$, so has length less than 2ℓ, by fineness of $\mathcal{H}$.

The hypothesis that y is between x and z in X_{NR} implies that any wall separating y from one of x or z in X_{NR} also separates x and z, so $\mathcal{K}$ is a fine chain separating x and z, which gives a bound:

$$d_f(x, z) \geq |\mathcal{K}| = |\mathcal{V}| + |\mathcal{H}| - 4\ell + 1 = d_f(x, y) + d_f(y, z) - (4\ell - 1)$$ □

The next result says the vertex sequence of a combinatorial geodesic can be parameterized as a rough geodesic in $(X^{(0)}_{\mathrm{NR}}, d_f)$. The point is that along a combinatorial geodesic the d_f–distance from the initial vertex is nondecreasing, but it can be constant over long subintervals. However, since it is discrete valued, we can reparameterize so that those constant valued intervals are each traversed in time 1.

Lemma 7.4. *Let $\alpha\colon [0, n] \to X_{\mathrm{NR}}$ be a combinatorial geodesic, let $\tau := 4\ell - 1$, and let $T := d_f(\alpha(0), \alpha(n))$. There exist numbers $0 = b_0 < b_1 < \cdots < b_n = T$ such that the map $\beta\colon [0, T] \to X^{(0)}_{\mathrm{NR}}$ defined by sending t to $\alpha(i)$ when $b_i \leq t < b_{i+1}$ is a $(\tau + 2)$–rough geodesic in $(X^{(0)}_{\mathrm{NR}}, d_f)$ that traverses the vertices of the image of α in the same order as α.*

Proof. For each $i \in [0, n] \cap \mathbb{Z}$, define $a_i := d_f(\alpha(0), \alpha(i))$. Since α is a combinatorial geodesic, for all $i \leq j$ in $[0, n] \cap \mathbb{Z}$, every wall separating $\alpha(0)$ and $\alpha(i)$ also separates $\alpha(0)$ and $\alpha(j)$, and the only wall separating $\alpha(0)$ and $\alpha(i + 1)$ that does not also separate $\alpha(0)$ and $\alpha(i)$ is the wall transverse to the edge between $\alpha(i)$ and $\alpha(i + 1)$. Thus, $0 \leq a_{i+1} - a_i \leq 1$.

Lemma 7.3 says α is an unparameterized τ–quasiruler. Rearranging the almost degeneracy condition gives that for integers $i \leq j$:

$$a_j - a_i \leq d_f(\alpha(i), \alpha(j)) \leq a_j - a_i + \tau \tag{4}$$

Define $b_0 := 0$. For each integer D in $[1, T]$ consider the indices $\{i \mid a_i = D\}$, which is an integer interval $[j, j + k]$ for some j and k, by monotonicity of $i \mapsto a_i$. For $i \in [j, j + k]$ define $b_i := D + \frac{i-j-k}{k+1}$, so that $b_j = D - 1 + \frac{1}{k+1}$ and $b_{j+k} = D$, with the other values b_i linearly interpolating between these two values. By construction, $i \mapsto b_i$ is strictly increasing and for all i we have:

$$0 \leq a_i - b_i < 1 \tag{5}$$

For integers $i < j$ in $[0, n]$, relations (4) and (5) give:

$$b_j - b_i - 1 \leq d_f(\beta(b_i), \beta(b_j)) \leq b_j - b_i + \tau + 1$$

Since $0 < b_{i+1} - b_i \leq 1$, if $x \leq y$ with $b_i \leq x < b_{i+1}$ and $b_j \leq y < b_{j+1}$ then $d_f(\beta(x), \beta(y)) = d_f(\beta(b_i), \beta(b_j))$, by definition of β, and $y - x \in b_j - b_i \pm 1$, so:

$$y - x - 2 \leq b_j - b_i - 1 \leq d_f(\beta(x), \beta(y)) \leq b_j - b_i + \tau + 1 \leq y - x + \tau + 2 \qquad \square$$

Corollary 7.5. *With $\tau := 4\ell - 1$, for every combinatorial geodesic α in X_{NR} there is a parameterized $(1, \tau+2, \tau)$–quasiruler in $(X^{(0)}_{\mathrm{NR}}, d_f)$ that has the same vertex sequence as α. The same is true for combinatorial geodesics in Σ.*

More generally, every admissible path has image that is uniformly Hausdorff close to the image of a $(1, \tau + 2, \tau)$–quasiruler.

Proof. In X_{NR}, the combination of Lemma 7.3 and Lemma 7.4 says combinatorial geodesics admit reparameterization as $(1, 4\ell + 1, 4\ell - 1)$–quasirulers, so the choice of τ gives the desired result.

The same is true in Σ because the canonical map $\Sigma^{(0)} \to X^{(0)}_{\mathrm{NR}}$ takes combinatorial geodesics to combinatorial geodesics, and d_f gives the same distances on $\Sigma^{(0)}$ measured in Σ as in X_{NR}.

Finally, Lemma 3.4 says that admissible paths are uniformly Hausdorff close to combinatorial geodesics. $\square$

The following is essentially Genevois's proof that his δ_L metric on a CAT(0) cube complex defines a non-geodesic hyperbolic metric, [48, Proposition 6.53]. The difference is that his metric is defined in terms of two disjoint walls being L–well-separated if every *set of walls* transverse to both of them and containing no facing triple has cardinality at most L, whereas we are saying two disjoint walls make a fine chain if every *chain* transverse to both of them contains at most $2\ell - 1$ walls. A pair of $(2\ell - 1)$–well-separated walls makes a fine chain, and a two-wall fine chain gives a $(2\ell - 1) \cdot \dim X_{\mathrm{NR}}$–well-separated pair, so these two notions are related, but it turns out to be fairly easy to just rerun the hyperbolicity argument using d_f, as opposed to trying to compare d_f to $\delta_{2\ell-1}$ and $\delta_{(2\ell-1)\cdot \dim X_{\mathrm{NR}}}$.

Proposition 7.6 (cf [48, Lemma 2.9 and Proposition 6.53])**.** *$(X^{(0)}_{\mathrm{NR}}, d_f)$ is $(18\ell - 5)$–hyperbolic.*

Proof. Let x_0, x_1, x_2, and x_3 be vertices of X_{NR}. Since X_{NR} is a CAT(0) cube complex, its 1–skeleton is a median graph. For each i, let y_i be the median of the triple $(x_{i+1}, x_{i+2}, x_{i+3})$, and let z_i be the median of the triple $(y_{i+1}, y_{i+2}, y_{i+3})$, with subscripts mod 4. Figure 6 shows the idealized version when all of these points are distinct. A key fact from median algebras is that sides of the median cube in Figure 6 have a parallelism

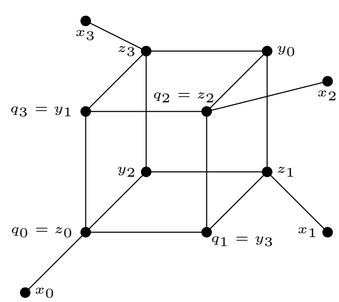

FIGURE 6. Setup for Proposition 7.6 pictured as the free median algebra on $\{x_0, x_1, x_2, x_3\}$.

relation such that a wall that separates the endpoints of one edge in the cube also separates the endpoints of the 3 edges parallel to it.

The intuition for the hyperbolicity proof is to show that in any configuration as in Figure 6 at least two of the dimensions of the cube are small, so the figure is uniformly tree-like.

Formally, there is a four-point δ–hyperbolicity condition [24, Definition III.H.1.20]:

$$(Q(\delta)) \qquad d_f(x_0, x_2) + d_f(x_1, x_3) \leq \max\{d_f(x_0, x_1) + d_f(x_2, x_3), d_f(x_0, x_3) + d_f(x_1, x_2)\} + 2\delta$$

We will establish condition $Q(18\ell - 5)$.

It will be enough in the argument to focus on one face of the cube. We arbitrarily choose the front face in the figure, and define its corners to be $q_0 := z_0$, $q_1 := y_3$, $q_2 := z_2$, and $q_3 := y_1$. The parallelism relation implies $a := d_f(q_0, q_1) = d_f(q_2, q_3)$ and $b := d_f(q_0, q_3) = d_f(q_1, q_2)$. Furthermore, every wall separating q_0 and q_1 crosses every wall separating q_0 and q_3. In particular, if $\mathcal{V}$ is a fine chain realizing $a = d_f(q_0, q_1)$ and $\mathcal{H}$ is a fine chain realizing $b = d_f(q_0, q_3)$ then $\mathcal{V} \pitchfork \mathcal{H}$, so fineness demands:

$$\min\{a, b\} = \min\{|\mathcal{V}|, |\mathcal{H}|\} \leq 1 \quad \text{or} \quad \max\{a, b\} = \max\{|\mathcal{V}|, |\mathcal{H}|\} \leq 2\ell - 1 \tag{6}$$

Let $[p_0, p_1]$ denote the *interval* between p_0 and p_1, which is all vertices that lie on combinatorial geodesics between p_0 and p_1. From the median point definitions we have, for each i:

$$[x_i, q_i] \cup [q_i, q_{i+1}] \cup [q_{i+1}, x_{i+1}] \subset [x_i, x_{i+1}]$$

Applying Lemma 7.3 twice:

$$\begin{aligned} d_f(x_3, x_0) &\geq d_f(x_3, q_3) + d_f(q_3, x_0) - 4\ell + 1 \\ &\geq d_f(x_3, q_3) + d_f(q_3, q_0) + d_f(q_0, x_0) - 8\ell + 2 \\ &= d_f(x_0, q_0) + d_f(x_3, q_3) + b - 8\ell + 2 \end{aligned}$$

A similar estimate holds for $d_f(x_1, x_2)$. Combining them gives:

$$d_f(x_3, x_0) + d_f(x_1, x_2) \geq 2b - 16\ell + 4 + \sum_{i=0}^{3} d_f(x_i, q_i) \tag{7}$$

A similar argument gives:

$$d_f(x_0, x_1) + d_f(x_2, x_3) \geq 2a - 16\ell + 4 + \sum_{i=0}^{3} d_f(x_i, q_i) \tag{8}$$

On the other hand, the triangle inequality gives:

$$d_f(x_0, x_1) + d_f(x_2, x_3) \leq 2a + \sum_{i=0}^{3} d_f(x_i, q_i) \tag{9}$$

Suppose, without loss of generality, that:

$$d_f(x_0, x_3) + d_f(x_1, x_2) \leq d_f(x_0, x_1) + d_f(x_2, x_3) \tag{10}$$

Then (7), (9), and (10) combine to give:

$$b - a \leq 8\ell - 2 \tag{11}$$

Now, from the triangle inequality:

$$d_f(x_0, x_2) + d_f(x_1, x_3) \leq 2a + 2b + \sum_{i=0}^{3} d_f(x_i, q_i)$$

Combined with (7):

$$d_f(x_0, x_2) + d_f(x_1, x_3) \leq d_f(x_0, x_3) + d_f(x_1, x_2) + 16\ell - 4 + 2a \tag{12}$$

Similarly, using (8):

$$d_f(x_0, x_2) + d_f(x_1, x_3) \leq d_f(x_0, x_1) + d_f(x_2, x_3) + 16\ell - 4 + 2b \tag{13}$$

Recall from (6) that if $b \geq 2$ then $a \leq 2\ell - 1$, so (11) says $b \leq 10\ell - 3$, whereas if $b < 2$ then $b \leq 10\ell - 3$ is still true, so (13) gives:

$$d_f(x_0, x_2) + d_f(x_1, x_3) \leq d_f(x_0, x_1) + d_f(x_2, x_3) + 36\ell - 10 \tag{14}$$

The assumption (10) says that (14) realizes the maximum appearing in $Q(\delta)$, so (14) establishes $Q(\delta)$ with $\delta = 18\ell - 5$. $\square$

We can extend d_f to a metric on all of Σ by adding some padding for non-vertices.

Definition 7.7. Extend d_f to $d_f\colon \Sigma\times\Sigma\to\mathbb{R}$ symmetrically as follows. Let $a\odot b$ be 0 if $a=b$ and 1 if $a\neq b$.

$$d_f(x,y) := \begin{cases} 0 & \text{if } x=y \\ \max_{x',y'} d_f(x',y') + \frac{x\odot x' + y\odot y'}{2} & \text{else} \end{cases}$$

The maximum is taken over vertices x' and y' such that x and x' are contained in a common closed cell and no wall separates x and x', and similarly for y and y'.

Lemma 7.8. *Definition 7.7 defines a metric on Σ extending d_f on $\Sigma^{(0)}$.*

For $x,y\in\Sigma$ we have:

$$d_f(x,y)-3 \leq \max\{|\mathcal{V}| \mid \mathcal{V} \textit{ is a fine chain separating } x \textit{ and } y\} \leq d_f(x,y)$$

Proof. It is reflexive, symmetric, and extends d_f, by construction. Suppose $x,y,z\in\Sigma$. We check the triangle inequality. If $x=z$ there is nothing to prove, so assume $x\neq z$. Let x' and z' be as in Definition 7.7 such that $d_f(x,z)=d_f(x',z')+\frac{x\odot x'+z\odot z'}{2}$. Let y' and y'' be any two vertices contained in the same smallest closed cell with y and not separated from y by any wall. Then estimate:

$$\begin{aligned} d_f(x,z) &= d_f(x',z') + \frac{x\odot x' + z\odot z'}{2} \\ &\leq d_f(x',y') + d_f(y',y'') + d_f(y'',z') + \frac{x\odot x' + z\odot z'}{2} \\ &= d_f(x',y') + \frac{x\odot x' + y\odot y'}{2} + d_f(y'',z') + \frac{y\odot y'' + z\odot z'}{2} \\ &\quad + d_f(y',y'') - \frac{y\odot y' + y\odot y''}{2} \\ &\leq d_f(x,y) + d_f(y,z) + d_f(y',y'') - \frac{y\odot y' + y\odot y''}{2} \end{aligned}$$

The vertices y' and y'' were chosen in the same smallest closed cell containing y, and not separated from y by any wall. Every wall separating y' and y'' contains y, so they all intersect, so maximal chains separating y' and y'' have length 1. Thus, $d_f(y',y'')\in\{0,1\}$. Furthermore, for $d_f(y',y'')=1$ to occur it must have been that y was not already a vertex, so y, y', and y'' are distinct. It follows that in any case $d_f(y',y'')-\frac{y\odot y'+y\odot y''}{2}$ is nonpositive, so $d_f(x,z)\leq d_f(x,y)+d_f(y,z)$. This shows that d_f is a metric on Σ.

When x and y are distinct we have, by definition, that $d_f(x,y)=d_f(x',y')+\frac{x\odot x'+y\odot y'}{2}$, and

$$d_f(x',y') = \max\{|\mathcal{V}| \mid \mathcal{V} \text{ is a fine chain separating } x' \text{ and } y'\}$$

No walls separate x and x' or y and y', so a fine chain separating x and y also separates x' and y'. Consider a fine chain separating x' and y' with $x'\in V_0^-\subset V_1^-\subset\cdots\subset V_n^-$ and $y'\in V_n^+$. No wall separates x' from x, so $x\in V_0^-$, but x might sit on V_0. Similarly, $y\in V_n^+$, but possibly $y\in V_n$. However, $V_1,\dots,V_{n-1}$ is a fine chain separating x and y. Now let $E:=\frac{x\odot x'+y\odot y'}{2}$ and estimate:

$$\begin{aligned} d_f(x,y) &= d_f(x',y') + E \\ &= \max\{|\mathcal{V}| \mid \mathcal{V} \text{ is a fine chain separating } x' \text{ and } y'\} + E \\ &\geq \max\{|\mathcal{V}| \mid \mathcal{V} \text{ is a fine chain separating } x \text{ and } y\} + E \\ &\geq \max\{|\mathcal{V}| \mid \mathcal{V} \text{ is a fine chain separating } x' \text{ and } y'\} - 2 + E \\ &= d_f(x',y') - 2 + E \\ &= d_f(x,y) - 2 \end{aligned}$$

Conclude that:

$$d_f(x,y) - 2 - E \leq \max\{|\mathcal{V}| \mid \mathcal{V} \text{ is a fine chain separating } x \text{ and } y\} \leq d_f(x,y)$$

Since $E\leq 1$, this concludes the proof. ∎

Corollary 7.9. *(Σ, d_f) is hyperbolic.*

Proof. This follows from Proposition 7.6 since $\Sigma^{(0)}$ is coarsely dense in Σ, so the four-point hyperbolicity condition $Q(\delta)$ is satisfied, after increasing δ to account for the density constant. □

7.2. **The fine chain metric collapses wide parabolics.**

Lemma 7.10. *If Υ is a wide parabolic subcomplex of Σ then $\Upsilon^{(0)}$ has diameter 1 in $(\Sigma^{(0)}, d_f)$.*

Proof. Recall from the analysis in the proof of Lemma 5.11 that fine chains in a wide Coxeter group are single walls, since every pair of disjoint walls is mutually transverse to an infinite set of walls. It follows that fine chains that cut a wide parabolic subcomplex of Σ consist of single walls. □

Recall that when W contains higher rank affine parabolics the W–action on the Niblo-Reeves cube complex X_{NR} is not cocompact. It turns out that since the d_f metric collapses such parabolics, it collapses the difference between $\Sigma^{(0)}$ and $X^{(0)}_{\mathrm{NR}}$:

Proposition 7.11. *The vertices of $\Sigma^{(0)}$ are coarsely dense in $(X^{(0)}_{\mathrm{NR}}, d_f)$.*

Proof. We identify $\Sigma^{(0)}$ with its canonical image in $X^{(0)}_{\mathrm{NR}}$. For $x \in X^{(0)}_{\mathrm{NR}}$, choose $v_0 \in \Sigma^{(0)}$ minimizing the combinatorial distance in $X^{(1)}_{\mathrm{NR}}$ from x to $\Sigma^{(0)}$.

We first recall that x lies in the cubical chamber containing v_0, as in [27, Section 7.5]: let $\Psi(v_0)$ be the set of halfspaces of X_{NR} that contain v_0 but not one of its neighbors in $\Sigma^{(0)}$. If x were not contained in some $\psi \in \Psi(v_0)$, then a combinatorial geodesic in X_{NR} from x to v_0 would cross the wall $\partial\psi$ in its last edge. Let v_1 be the neighbor of v_0 in X_0 not contained in ψ. Then the same geodesic, with its last edge replaced by the edge to v_1, would show that v_1 is strictly closer to x than v_0 is, contradicting the choice of v_0. Thus, x belongs to the cubical chamber of v_0.

Let $\mathcal{F}$ be a longest fine chain separating v_0 from x. The walls of $\mathcal{F}$ are pairwise disjoint and contained in the set of walls $\mathcal{M}(v_0, x)$ separating v_0 from x.

By [27, Proposition 24] there is a constant $K = K(W, S)$ such that for any set of pairwise disjoint walls $\mathcal{M}$ contained in $\mathcal{M}(v_0, x)$ with $|\mathcal{M}| > K$, we have $W(\mathcal{M}) \cong \mathcal{D}_\infty$ and $\mathrm{Pc}(\mathcal{M})$ is irreducible higher rank affine.

Thus, $|\mathcal{F}| > K$ implies $\mathrm{Pc}(\mathcal{F})$ is irreducible higher rank affine, which implies $|\mathcal{F}| = 1$, which is a contradiction, since K is clearly at least 1.

Conclude that $d_f(x, v_0) = |\mathcal{F}| \leq K$. □

7.3. **Recognizing Morse quasigeodesics and parabolics.** In Section 7.3.1 we show that quasigeodesics in (Σ, d_2) are Morse if and only if they are (parameterized) quasigeodesics in (Σ, d_f). In Section 7.3.2 we discuss related notions in the literature. In Section 7.3.3 we characterize Morse and stable parabolic subgroups.

7.3.1. *The fine chain metric is Morse recognizing.*

Theorem 7.12. *A quasigeodesic γ is Morse in (Σ, d_2) if and only if it is quasigeodesic in (Σ, d_f). More precisely:*

(1) *For all λ, ϵ, χ there exist λ' and ϵ' such that for every χ–strongly contracting (λ, ϵ)–quasigeodesic path $\gamma\colon I \to (\Sigma, d_2)$ the path $\gamma\colon I \to (\Sigma, d_f)$ is a (λ', ϵ')–quasigeodesic.*

(2) *For all λ, ϵ, λ', ϵ' there exists χ such that for every (λ, ϵ)–quasigeodesic path $\gamma\colon I \to (\Sigma, d_2)$ such that $\gamma\colon I \to (\Sigma, d_f)$ is (λ', ϵ')–quasigeodesic, γ is χ–strongly contracting in (Σ, d_2).*

Proof. It suffices to assume I is compact, because if compact subintervals of a path are uniformly contracting or uniformly quasigeodesic then so is the whole path.

First we prove Item (1). It is an exercise to argue that if $\gamma\colon I \to (\Sigma, d_2)$ is a χ–strongly contracting (λ, ϵ)–quasigeodesic path then there is D depending on λ, ϵ, and χ such that the Hausdorff distance between γ and the geodesic α between its endpoints is at most D. It also follows that α is χ'–strongly contracting for χ' depending on χ and D.

Suppose we can show that α is (λ'', ϵ'')–quasigeodesic in (Σ, d_f). Let ϕ be a map from the domain of γ to the domain of α such that $d_2(\gamma(t), \alpha(\phi(t))) \leq D$ for all $t \in I$. Since $d_f \leq d_2$, the same is true for d_f.

Since γ is a d_2–quasigeodesic, for $s < t$ in I:

$$d_f(\gamma(s), \gamma(t)) \leq d_2(\gamma(s), \gamma(t)) \leq \lambda(t - s) + \epsilon$$

Conversely, we need the d_f–quasigeodesic lower bound for γ. We have the d_2–quasigeodesic lower bound:

$$\frac{1}{\lambda}(t-s)-\epsilon \leq d_2(\gamma(s),\gamma(t)) \leq d_2(\alpha(\phi(s)),\alpha(\phi(t)))+2D = |\phi(t)-\phi(s)|+2D$$

Thus:

$$|\phi(t)-\phi(s)| \geq \frac{1}{\lambda}(t-s)-\epsilon-2D \tag{15}$$

Since we know α is a d_f quasigeodesic:

$$d_f(\alpha(\phi(s)),\alpha(\phi(t))) \geq \frac{1}{\lambda''}|\phi(t)-\phi(s)|-\epsilon'' \tag{16}$$

Combining (15) and (16) gives:

$$\begin{aligned} d_f(\gamma(s),\gamma(t)) &\geq d_f(\alpha(\phi(s)),\alpha(\phi(t)))-2D \\ &\geq \frac{1}{\lambda''}|\phi(t)-\phi(s)|-\epsilon''-2D \\ &\geq \frac{1}{\lambda\lambda''}(t-s)-\frac{2D+\epsilon}{\lambda''}-(2D+\epsilon'') \end{aligned}$$

This shows that γ is a d_f–quasigeodesic with constants depending only on λ, ϵ, and χ. Thus, it suffices to prove Item (1) in the case that γ is a d_2–geodesic. Since d_2–geodesics are admissible, Theorem 1.4 (3) says there exists a $0<\rho\leq\rho_0(1,0)$ and L such that for every subinterval $I'\subset I$ of length at least L there is a fine chain $\mathcal{V}$ transverse to $\gamma(I')$ with $|\mathcal{V}|/|I'|\geq\rho$. Thus, for every $t_0,t_1\in I$, if $|t_1-t_0|\geq L$ then Lemma 7.8 gives $d_f(\gamma(t_0),\gamma(t_1))\geq\rho|t_1-t_0|$. Thus, for arbitrary t_0 and t_1 in I we have (since $\rho\leq 1$) that $d_f(\gamma(t_0),\gamma(t_1))\geq\rho|t_1-t_0|-L$. In the other direction, $d_f(\gamma(t_0),\gamma(t_1))\leq d_2(\gamma(t_0),\gamma(t_1))\leq|t_1-t_0|$. Thus, when γ is a d_2–geodesic it is a $(1/\rho,L)$–quasigeodesic in (Σ,d_f).

Now we prove Item (2). The proof is to notice that γ being d_f–quasigeodesic essentially gives the hypotheses of Proposition 4.10, and then follow that proof. Suppose γ is a (λ,ϵ)–quasigeodesic path $I\to(\Sigma,d_2)$ that is (λ',ϵ')–quasigeodesic in (Σ,d_f). Let $J:=[a,b]\subset I$. Let $\mathcal{V}=\{V_i\}_{1\leq i\leq n}$ be a longest fine chain separating $\gamma(a)$ and $\gamma(b)$, oriented from $\gamma(a)$ to $\gamma(b)$. By Lemma 7.8, $|\mathcal{V}|\geq d_f(\gamma(a),\gamma(b))-3$. Combined with the hypothesis that γ is d_f–quasigeodesic:

$$|\mathcal{V}|/|J| \geq \frac{d_f(\gamma(a),\gamma(b))-3}{|J|} \geq \frac{|J|/\lambda'-\epsilon'-3}{|J|} = \frac{1}{2\lambda'}+\frac{1}{2\lambda'}-\frac{\epsilon'+3}{|J|}$$

Thus, there exists a fine chain separating $\gamma(a)$ and $\gamma(b)$ of density at least $\rho:=\frac{1}{2\lambda'}$, provided:

$$|J| \geq 2\lambda'(\epsilon'+3) \tag{17}$$

Since γ is continuous it crosses every wall of $\mathcal{V}$.

Supposing we have chosen a long enough subinterval J as above, define t_i to be the first time in J such that $\gamma(t_i)\in\bar{V}_i^+$. Since $\mathcal{V}$ is a chain separating $\gamma(a)$ and $\gamma(b)$, these times satisfy $a<t_1<t_2<\cdots<t_n<b$. We have $\sum_{i=1}^{n-1}t_{i+1}-t_i<b-a=|J|\leq|\mathcal{V}|/\rho=n/\rho$. Thus, there are at least $(n-2)/2$ intervals $[t_i,t_{i+1}]$ of length at most $2/\rho$. For any such 'short' interval $[t_i,t_{i+1}]$, we have:

$$D := 2\lambda/\rho+\epsilon \geq d_2(\gamma(t_i),\gamma(t_{i+1})) \geq d_2(V_i,V_{i+1})$$

Now for any large $K>2$ and large r, suppose there exist points x and y satisfying $d_2(\gamma,x)=d_2(\gamma,y)=r$ and $d_2(x,y)\geq Kr$. Let $x'=\gamma(a)$ be some closest point of γ to x, let $y'=\gamma(b)$ be some closest point of γ to y, and let $J:=[a,b]$. Then $d_2(x',y')\geq(K-2)r$, so:

$$|J| \geq ((K-2)r-\epsilon)/\lambda \tag{18}$$

In particular, (17) is satisfied for all sufficiently large r.

Suppose there exists a path α from x to y that stays outside $\mathcal{N}_r(\gamma)$. Let $\mathcal{V}'$ be the convex subchain of $\mathcal{V}$ consisting of walls that separate x from y. By the usual 4–gon trick, $|\mathcal{V}|-|\mathcal{V}'|\leq d_c(x,x')+d_c(y,y')+2$, which is bounded above by a linear function of r, with constants from Lemma 2.15. By choosing K large enough with respect to these constants we can arrange that the number of short intervals between consecutive walls of $\mathcal{V}'$ is bounded below by a linear function of r.

Proceed as in the proof of Proposition 4.10. Since $\mathcal{V}'$ is a fine chain, if $[t_i,t_{i+1}]$ is a short interval then Lemma 4.5 says the diameter of Bridge(V_i,V_{i+1}) is bounded in terms of D and 2ℓ. Since γ is quasigeodesic and the segment $\gamma|_{[t_i,t_{i+1}]}$ is short, it is contained in a bounded neighborhood of Bridge(V_i, V_{i+1}), by Lemma 4.3. Conversely, the subsegment of α between V_i and V_{i+1} stays at least r-far from γ, so it stays linearly far, in terms of r, from Bridge(V_i, V_{i+1}), so Lemma 4.3 says its length is bounded below by a linear function of r. Since α contains a linear in r number of disjoint subsegments whose lengths are uniformly bounded below by a linear in r function, the total length of α is quadratic in r. This gives a quadratic lower bound on the divergence gauge $\delta(r; 1, K)$ of γ, so γ is strongly contracting.

Recall that if no suitable pair (x, y) exists, or if no suitable path α exists, then $\delta(r; 1, K) = \infty$, so these are not problematic cases—$\delta(r; 1, K)$ is still greater than the computed quadratic lower bound at r. □

7.3.2. *Morse recognizing is equivalent to stability recognizing.* The purpose of this subsection is to give definitions of "Morse recognizing" and "stability recognizing", compare them to related notions that appear in the literature, and show that these two recognition properties are equivalent.

Definition 7.13. $\phi\colon X \to Y$ is *Morse recognizing* if:

- For all λ, ϵ, μ there exist λ' and ϵ' such that if $\gamma\colon I \to X$ is a μ–Morse (λ, ϵ)–quasigeodesic then $\phi \circ \gamma\colon I \to Y$ is a (λ', ϵ')–quasigeodesic.
- For all λ, ϵ, λ', ϵ' there exists μ such that if $\gamma\colon I \to X$ is a (λ, ϵ)–quasigeodesic and $\phi \circ \gamma\colon I \to Y$ is a (λ', ϵ')–quasigeodesic then $\gamma\colon I \to X$ is μ–Morse.

Definition 7.14. A subset $Z \subset X$ is (λ, ϵ)*–quasigeodesically connected* if for every pair of points z_0 and z_1 in Z there exists a (λ, ϵ)–quasigeodesic $\gamma\colon I \to X$ whose image is in Z and whose endpoints are z_0 and z_1.

Definition 7.15. (cf [45, 39, 2][4]) A subset $Z \subset X$ is (λ, ϵ, μ)*–stable* if for every pair of points z_0 and z_1 in Z there exists a μ–Morse (λ, ϵ)–quasigeodesic $\gamma\colon I \to X$ whose image is in Z and has endpoints z_0 and z_1.

Definition 7.16. $\phi\colon X \to Y$ is *stability recognizing* if:

- For all λ, ϵ, μ there exist λ' and ϵ' such that for every (λ, ϵ, μ)–stable $Z \subset X$ the map $\phi|_Z$ is a (λ', ϵ')–quasiisometric embedding of $(Z, d_X|_Z)$ into Y.
- For all λ, ϵ, λ', ϵ' there exists μ such that for every (λ, ϵ)–quasigeodesically connected set $Z \subset X$ if $\phi|_Z$ is a (λ', ϵ')–quasiisometric embedding of $(Z, d_X|_Z)$ into Y then Z is (λ, ϵ, μ)–stable in X.

Proposition 7.17. *$\phi\colon X \to Y$ is Morse recognizing if and only if it is stability recognizing.*

Usually one also requires that the space Y is hyperbolic.[5] In that case the Morse property, stability, and quasiconvexity are equivalent, so the recognition properties can be restated. When Y is hyperbolic, '$\phi\colon X \to Y$ is Morse recognizing' is equivalent to 'a quasigeodesic γ in X is Morse if and only if $\phi \circ \gamma$ is a Morse quasigeodesic in Y, with quantitative control on the quasigeodesic constants and Morse gauges in both directions.' Similarly, '$\phi\colon X \to Y$ is stability recognizing' is equivalent to 'a quasigeodesically connected subspace Z is stable in X if and only if $\phi|_Z$ is a quasiisometric embedding with quasiconvex image, with quantitative control between the stability parameters and the quasiisometric embedding/quasiconvexity parameters, in both directions.' We do not need Y to be hyperbolic for Proposition 7.17.

Russell, Spriano, and Tran [71, Definition 4.17] defined a *Morse detectable space* X to be one such that there exists a Morse recognizing map $\phi\colon X \to Y$ as above, with the additional conditions that Y is hyperbolic and ϕ is coarse Lipschitz. This gave a name to the property characterizing contracting quasigeodesics in hierarchically hyperbolic spaces described by Abbott, Behrstock, and Durham [2]. This terminology has been picked up in subsequent work related to the Morse-local-global property [5, 67].

Durham and Taylor [45] defined a notion of *stable subgroup* abstracting the properties of what were called convex cocompact subgroups of mapping class groups of hyperbolic surfaces [47]. A subgroup H of a finitely generated group G is stable[6] if it is stable as a subspace of a Cayley graph of G. Kent and Leininger [54] (see also [52]) proved what we would now call a stable subgroup recognition theorem, showing that convex cocompact subgroups of the mapping class group are those for which the orbit map into the curve graph is a quasiisometric embedding of the subgroup. Analogous stable subgroup recognition theorems exist for groups

[4]We are treating stability as a property of a subspace Z of X with the subspace metric. It is not hard to show this is equivalent to the standard definition in which Z is considered as a separate metric space and stability is a property of a quasiisometric embedding of Z into X.

[5]A notable early example of a recognition-type theorem where the target space is not hyperbolic is the characterization of convex cocompact subgroups of the mapping class group in terms of having quasiconvex orbits in Teichmüller space [47].

[6]Quasigeodesic connectivity of H already forces it to be finitely generated and undistorted in G.

hyperbolic relative to peripherals with linear divergence [11], right-angle Artin groups [55], hierarchically hyperbolic groups [2], etc [69, 70, 77, 14].

Definition 7.16 for stable *subspace* recognition is formally stronger than stable *subgroup* recognition, since if $\phi\colon G \to Y$ recognizes stable subgroups of G it is conceivable that there are stable subspaces of G that are not close to any finitely generated subgroup and for which ϕ does not give a quasiisometric embedding. We do not know any examples.

Proof of Proposition 7.17. Suppose $\phi\colon X \to Y$ is Morse recognizing.

Suppose $Z \subset X$ is (λ, ϵ, μ)–stable. For all $z_0, z_1 \in Z$ there exists a μ–Morse (λ, ϵ) quasigeodesic $\gamma\colon [t_0, t_1] \to X$ with image in Z and $\gamma(t_i) = z_i$ for $i \in \{0, 1\}$. Since $\phi\colon X \to Y$ is Morse recognizing, $\phi \circ \gamma\colon I \to Y$ is (λ', ϵ')–quasigeodesic, so:

$$d_Y(\phi(z_0), \phi(z_1)) = d_Y(\phi(\gamma(t_0)), \phi(\gamma(t_1))) \overset{\lambda',\epsilon'}{\asymp} |t_1 - t_0| \overset{\lambda,\epsilon}{\asymp} d_X(z_0, z_1)$$

Thus, $\phi|_Z$ is a quasiisometric embedding with constants determined by λ, ϵ, λ', ϵ', and λ' and ϵ' were determined by λ, ϵ, and μ.

Conversely, suppose $Z \subset X$ is (λ, ϵ)–quasigeodesically connected and $\phi|_Z$ is a (λ', ϵ')–quasiisometric embedding of Z into Y. For all $z_0, z_1 \in Z$ there exists a (λ, ϵ)–quasigeodesic $\gamma\colon [t_0, t_1] \to X$ with image in Z and $\gamma(t_i) = z_i$ for $i \in \{0, 1\}$. Since $\phi|_Z$ is a (λ', ϵ')–quasiisometric embedding, $\phi \circ \gamma$ is a quasigeodesic with constants determined by λ, ϵ, λ', and ϵ'. Since $\phi\colon X \to Y$ is Morse recognizing, there exists μ such that γ is μ–Morse, for μ depending on λ, ϵ and the quasigeodesic constants of $\phi \circ \gamma$, hence on λ, ϵ, λ', and ϵ'. Thus, the transitive family of uniform quasigeodesics verifying quasigeodesic connectivity of Z is also uniformly Morse, so Z is stable.

We have shown that $\phi\colon X \to Y$ is Morse recognizing implies that it is stability recognizing. Now suppose $\phi\colon X \to Y$ is stability recognizing.

Suppose $\gamma\colon I \to X$ is a μ–Morse (λ, ϵ)–quasigeodesic. Then $Z := \gamma(I)$ is a (λ, ϵ, μ)–stable subset of X. Since $\phi\colon X \to Y$ is stability recognizing, $\phi|_Z$ is a (λ', ϵ')–quasiisometric embedding. Thus, $\phi \circ \gamma$ is a composition of quasiisometric embeddings, so $\phi \circ \gamma\colon I \to Y$ is a quasigeodesic with constants determined by λ, ϵ, λ', and ϵ', hence by λ, ϵ, and μ.

Conversely, suppose $\gamma\colon I \to X$ is (λ, ϵ)–quasigeodesic and $\phi \circ \gamma\colon I \to Y$ is (λ', ϵ')–quasigeodesic. Then $Z := \gamma(I)$ is (λ, ϵ)–quasigeodesically connected in X. Furthermore, for any $z_0, z_1 \in Z$ there exist $t_0, t_1 \in I$ such that $\gamma(t_i) = z_i$, and:

$$d_X(z_0, z_1) = d_Z(\gamma(t_0), \gamma(t_1)) \overset{\lambda,\epsilon}{\asymp} |t_1 - t_0| \overset{\lambda',\epsilon'}{\asymp} d_Y(\phi(\gamma(t_0)), \phi(\gamma(t_1))) = d_Y(\phi(z_0), \phi(z_1))$$

This says $\phi|_Z$ is a quasiisometric embedding of Z into Y, with constants depending on λ, ϵ, λ', and ϵ'. Since $\phi\colon X \to Y$ is stability recognizing, there exists μ depending on λ, ϵ, λ', and ϵ' such that Z is (λ, ϵ, μ)–stable. Consider an interval $[t_0, t_1] \subset I$. Since Z is stable, there exists a μ–Morse (λ, ϵ)–quasigeodesic $\delta\colon J \to X$ with image in Z and endpoints at $\gamma(t_0)$ and $\gamma(t_1)$. So $\gamma([t_0, t_1])$ and δ are (λ, ϵ)–quasigeodesics with the same endpoints such that δ is μ–Morse. It follows that the Hausdorff distance between δ and $\gamma([t_0, t_1])$ can be bounded in terms of λ, ϵ, and μ. Then it follows that there is μ' depending on λ, ϵ, and μ such that $\gamma([t_0, t_1])$ is μ'–Morse. Since this is true for every subsegment of γ, conclude that γ is μ'–Morse. □

7.3.3. *Recognizing Morse and stable parabolics.* The combination of Theorem 7.12 and Proposition 7.17 gives:

Proposition 7.18. *A subgroup $H \leq W$ is stable if and only if it is finitely generated and the orbit map is a quasiisometric embedding of H with respect to a word metric into (Σ, d_f).*

It is not immediately clear from Theorem 7.12 how to recognize Morse subgroups of W. However, for parabolic subgroups there are nice criteria for being stable and Morse. Given a finite rank Coxeter system (W, S) and $T \subset S$, let Γ_T denote the subdiagram of the Coxeter diagram of (W, S) spanned by the vertices T. Call the subdiagram 'spherical' if it defines a spherical Coxeter group.

Theorem 7.19. *Let (W, S) be a finite rank Coxeter system, and let $T \subset S$.*

W_T is stable if and only if $\Gamma_{T \cap U}$ is spherical for every wide $U \subset S$.

W_T is Morse if and only if for every wide $U \subset S$, one of the following is true:

1. *The intersection of every connected component of Γ_U with Γ_T is spherical.*
2. *Every nonspherical connected component of Γ_U is contained in Γ_T.*

Proof. The stable case follows immediately from Corollary 1.5: if U is wide and $T \cap U$ is nonspherical then there is a biinfinite d_1–geodesic γ contained in $\Sigma_{T\cap U} = \Sigma_T \cap \Sigma_U$, so the entire γ has edges with labels in the wide subset U, so γ is not Morse. On the other hand, if $T \cap U$ is spherical for every wide U, then, since there are only finitely many spherical subsets of S, there is a uniform diameter bound on $\Sigma_{T\cap U}$, so for every d_1–geodesic γ in Σ_T there is a uniform bound on the lengths of its subsegments having label in some wide subset of S. Corollary 1.5 then says that geodesics in Σ_T are uniformly Morse.

Now consider Morseness of Σ_T. Condition (1) is equivalent to: $W_{T\cap U} = W_T \cap W_U$ is spherical. Since nonspherical irreducible Coxeter groups do not have finite index proper parabolic subgroups, condition (2) is equivalent to saying that the index of $W_{T\cap U} = W_T \cap W_U$ in W_U is finite. The two conditions together say that every wide special subgroup is either virtually contained in W_T or virtually has trivial intersection with W_T.

Suppose $U \subset S$ is wide and neither condition (1) nor (2) is true. Let $U = U_1 \sqcup \cdots \sqcup U_k \sqcup K$ be the decomposition of U into irreducible nonspherical components U_i and spherical K. Since U is wide, if $k = 1$ then U_1 is irreducible higher rank affine, so it has no proper nonspherical subparabolics. Thus $T \cap U_1$ is either spherical or all of U_1. This contradicts the hypothesis, so we must have $k \geq 2$, and there is some i such that $W_{T\cap U_i}$ has infinite index in W_{U_i}. If $T \cap U_i$ is spherical then there is some $j \neq i$ such that $T \cap U_j$ is nonspherical. If $T \cap U_i$ is nonspherical then consider any $j \neq i$. If $T \cap U_j$ is spherical then $W_{T\cap U_j}$ has infinite index in W_{U_j}; in this case swap i and j. Thus, we may assume we have indices $i \neq j$ such that $W_{T\cap U_i}$ has infinite index in W_{U_i} and $T \cap U_j$ is nonspherical. For $n \in \mathbb{N}$, choose $b_n \in W_{U_i}$ such that $d_1(b_n, \Sigma_{T\cap U_i}) \xrightarrow{n\to\infty} \infty$, which is possible since $W_{T\cap U_i}$ has infinite index in W_{U_i}. Then Lemma 2.40 implies that $d_1(b_n, \Sigma_T) \to \infty$. Pick $a_n \in W_{T\cap U_j}$ with $|a_n|_S \geq |b_n|_S$. Since W_{U_i} and W_{U_j} commute, the geodesics $[\mathbb{1}, a_n] \subset \Sigma_T$ have endpoints that are connected by uniform quasigeodesics $[\mathbb{1}, b_n] + [b_n, a_n b_n] + [a_n b_n, a_n]$ that contain points b_n arbitrarily far from Σ_T. Thus, Σ_T is not Morse.

Now suppose that for every wide $U \subset S$ either (1) or (2) is true. We derive a contradiction by supposing that Σ_T is not Morse. Let D be the diameter of a chamber of Σ. Since (Σ, d_2) is CAT(0), Proposition 2.30 and Proposition 2.31 combine to say that if Σ_T is not Morse then it does not have the Bounded Geodesic Image Property. Thus, for every $n \in \mathbb{N}$ there exists a d_2–geodesic $[a_n, b_n]$ that does not come n–close to Σ_T, but $\operatorname{diam}(\pi_{\Sigma_T}([a_n, b_n])) > n$. By passing to a subsegment, we may assume that the diameter of $\pi_{\Sigma_T}([a_n, b_n])$ is realized by $d_2(\pi_{\Sigma_T}(a_n), \pi_{\Sigma_T}(b_n))$. Let x_n be a vertex closest to a_n, and let y_n be a vertex closest to b_n. Then $d_2(a_n, x_n) \leq D$, $d_2(b_n, y_n) \leq D$, and the Hausdorff distance between $[a_n, b_n]$ and $[x_n, y_n]$ is at most D, so $d_2(\pi_{\Sigma_T}(x_n), \pi_{\Sigma_T}(y_n)) \geq d_2(\pi_{\Sigma_T}(a_n), \pi_{\Sigma_T}(b_n)) - 2D \geq n - 2D$ and $d_2(\Sigma_T, [x_n, y_n]) \geq n - D$. Recalling the definition of the gate map from Section 2.10, set $x'_n := \mathfrak{p}_{\Sigma_T}(x_n)$ and $y'_n := \mathfrak{p}_{\Sigma_T}(y_n)$. Lemma 2.43 implies $d_2(\pi_{\Sigma_T}(x_n), x'_n) \leq D$ and $d_2(\pi_{\Sigma_T}(y_n), y'_n) \leq D$, so $d_2(x'_n, y'_n) \geq n - 4D$.

Fix n large enough that d_2–distance at least $n - 4D$ implies d_c distance at least 4ℓ in Lemma 2.15 and Lemma 2.16. The former implies $d_c(x'_n, y'_n) \geq 4\ell$, so there is a chain $\mathcal{V}^- \sqcup \mathcal{V}^+$ oriented from x'_n to y'_n such that the walls of $\mathcal{V}^-$ come first, the walls of $\mathcal{V}^+$ come second, and each of the two subchains contains at least 2ℓ–many walls.

Every wall of $\mathcal{V}^- \sqcup \mathcal{V}^+$ cuts Σ_T, so $\operatorname{Pc}(\mathcal{V}^-)$ and $\operatorname{Pc}(\mathcal{V}^+)$ are both irreducible nonspherical subgroups of W_T. By Corollary 2.37, no wall of $\mathcal{V}^- \sqcup \mathcal{V}^+$ separates x_n from x'_n or separates y_n from y'_n, so they all cut $\gamma := [x_n, y_n]$. Let p be a point on γ that occurs between the last wall of $\mathcal{V}^-$ and the first wall of $\mathcal{V}^+$ and is not contained in any wall. Let γ^- be the subsegment of γ from x_n to p, and let γ^+ be the subsegment of γ from p to y_n. Let p' be a vertex closest to p. Then $d_2(p, p') \leq D$, no wall separates p and p', and $d_2(p', \Sigma_T) \geq n - 2D$. Then the choice of n with respect to Lemma 2.16 gives a chain $\mathcal{H}_0$ of length at least 4ℓ separating p' and Σ_T. Since no walls separate p and p', the chain $\mathcal{H}_0$ also separates p and Σ_T. The walls of $\mathcal{H}_0$ do not cut Σ_T, but all of them must cut the paths between Σ_T and p given by $[x'_n, x_n] + \gamma^-$ and $\gamma^+ + [y_n, y'_n]$. Since any wall cuts γ at most once, each wall of $\mathcal{H}_0$ cuts at least one of $[x'_n, x_n]$ and $[y_n, y'_n]$. Let $\mathcal{H}^-$ be the subchain of $\mathcal{H}_0$ consisting of walls that cut $[x_n, x'_n]$, and let $\mathcal{H}^+ := \mathcal{H}_0 \smallsetminus \mathcal{H}^-$. Without loss of generality, assume $|\mathcal{H}^-| \geq 2\ell$, and set $\mathcal{V} := \mathcal{V}^-$ and $\mathcal{H} := \mathcal{H}^-$. This is possible because otherwise $|\mathcal{H}^+| \geq 2\ell$, so then we would take $\mathcal{V} := \mathcal{V}^+$ and $\mathcal{H} := \mathcal{H}^+$ and swap x_n and x'_n with y_n and y'_n in what follows.

Each wall $H \in \mathcal{H}$ cuts $[x_n, x'_n]$ and separates p and Σ_T. Since H separates p from Σ_T and separates x_n from $x'_n \in \Sigma_T$, we have that x_n and p are on the same side of H. Since the open halfspaces of H are convex, γ^- lies in the same open halfspace of H as x_n and p. For each $V \in \mathcal{V}$ there is a path from p to Σ_T that follows γ^- from p to $V \cap \gamma$ and then follows V to Σ_T. Since H cuts this path but does not cut γ^-, we must have $H \pitchfork V$.

Since this is true for all $H \in \mathcal{H}$ and $V \in \mathcal{V}$, we have a grid $(\mathcal{V}, \mathcal{H})$ with $|\mathcal{V}| \geq 2\ell$ and $|\mathcal{H}| \geq 2\ell$. Proposition 5.3 says that $N_W(\mathrm{Pc}(\mathcal{V})) = \mathrm{Pc}(\mathcal{V}) \times \mathrm{Pc}(\mathcal{V})^\perp$ is a wide parabolic, and the proof shows that either $\mathrm{Pc}(\mathring{\mathcal{H}}) \leq \mathrm{Pc}(\mathcal{V})^\perp$ or $\mathrm{Pc}(\mathcal{V}) = \mathrm{Pc}(\mathcal{H})$ is higher rank affine.

Since the walls of $\mathcal{H}$ do not cut Σ_T, their reflections are not in W_T.

The affine case is impossible because it yields $W_T \geq \mathrm{Pc}(\mathcal{V}) = \mathrm{Pc}(\mathcal{H}) \not\leq W_T$.

The commuting case is also impossible. Up to the W_T–action, we may assume $\mathrm{Pc}(\mathcal{V}) = W_U$ for some $U \subset T$, and it follows that $N_W(\mathrm{Pc}(\mathcal{V})) = W_U \times W_{U^\perp}$. Now, U is a nonspherical irreducible component of the wide set $U \sqcup U^\perp$ that is contained in T, so condition (1) is not satisfied for $U \sqcup U^\perp$. By hypothesis, condition (2) must be satisfied, meaning that T contains every nonspherical irreducible component of $U \sqcup U^\perp$, hence of $U^\perp$. However, this is not true, because $\mathrm{Pc}(\mathring{\mathcal{H}})$ is a nonspherical, irreducible reflection subgroup of $W_{U^\perp}$, so by Lemma 2.8 it is contained in $W_{\mathcal{C}}$ for some connected component of $\Gamma_{U^\perp}$. The component $\mathcal{C}$ is necessarily nonspherical, since $\mathrm{Pc}(\mathring{\mathcal{H}})$ is. We get the contradiction $W_T \not\geq \mathrm{Pc}(\mathring{\mathcal{H}}) \leq W_{\mathcal{C}} \leq W_T$.

Since both cases lead to contradictions, Σ_T is Morse. □

7.4. The coned-off space.

Definition 7.20. The *coned-off space* $(\widehat{\Sigma}, \hat{d})$ is the metric graph $\widehat{\Sigma}$ and geodesic metric $\hat{d}$ obtained from $\Sigma^{(1)}$ by adding a new vertex v_Υ for each wide parabolic subcomplex Υ of Σ, and connecting v_Υ to each vertex of Υ by an edge of length $1/2$.

In the case that (W, S) is a right-angled Coxeter group, $\widehat{\Sigma}$ is W–equivariantly quasiisometric to the universal acylindrical hyperbolic space for W constructed by Abbott, Behrstock, and Durham [2]. In the 2-dimensional right-angled case this also coincides, again up to W–equivariant quasiisometry, with the intersection complex of Oh [65], see [37, Section 6.2]. We also note that Genevois [49] has results producing hyperbolic spaces by coning off certain convex subcomplexes of CAT(0) cube complexes, but Σ is not a cube complex outside of the right-angled case.

Theorem 7.21. *The identity map on $\Sigma^{(0)}$ is a quasiisometry between $(\Sigma^{(0)}, d_f)$ and $(\widehat{\Sigma}, \hat{d})$.*

Corollary 7.22. *$(\widehat{\Sigma}, \hat{d})$ is a geodesic hyperbolic space and $(\Sigma^{(1)}, d_1) \xrightarrow{\mathrm{Id}_{\Sigma^{(1)}}} (\widehat{\Sigma}, \hat{d})$ is Morse recognizing.*

Proof. Hyperbolicity is preserved by quasiisometries between quasiruled spaces [21, Section 2.1]. We have that $(\Sigma^{(1)}, d_f)$ is quasiruled, Corollary 7.5, hyperbolic, Corollary 7.9, and quasiisometric to the geodesic space $(\widehat{\Sigma}, \hat{d})$, Theorem 7.21, so $(\widehat{\Sigma}, \hat{d})$ is hyperbolic. The quasiisometry transfers the Morse recognizing property of $(\Sigma^{(1)}, d_f)$, Theorem 7.12, to $(\widehat{\Sigma}, \hat{d})$. □

Since every vertex of $\widehat{\Sigma}$ is within distance $1/2$ of a vertex of $\Sigma^{(0)}$, it suffices to show that $\hat{d}$ and d_f are comparable on $\Sigma^{(0)}$. One direction is easy: Lemma 7.23 says $d_f \leq \hat{d}$, because d_f gives wide parabolic subcomplexes diameter 1, so it collapses as much as $\hat{d}$. In the other direction we need to argue that d_f does not collapse too much more than $\hat{d}$. That is Proposition 7.28. The rough idea of that proof is that we know that combinatorial geodesics in Σ are unparameterized d_f quasirulers, so they are useful test paths for relating d_f and $\hat{d}$. We define the *shortcut graph* $\hat{\gamma}$ of a combinatorial geodesic γ to model its length in $\widehat{\Sigma}$, and show that the d_f length of γ cannot be too much shorter than the length of $\hat{\gamma}$.

By construction $(\widehat{\Sigma}, \hat{d})$ is a geodesic metric space. Let γ be a $\hat{d}$–geodesic. There are two kinds of edges in $\widehat{\Sigma}$, those coming from $\Sigma^{(1)}$ and those that connect a vertex $a \in \Upsilon^{(0)}$, for some wide parabolic subcomplex Υ, to the cone vertex v_Υ. Furthermore, v_Υ is only connected by an edge to vertices from $\Upsilon^{(0)}$, so a $\hat{d}$–geodesic γ between x and y in $\Sigma^{(0)}$ can be decomposed as a concatenation $\beta_0 + c_1 + \beta_1 + \cdots + c_m + \beta_m$ where the β_i are disjoint (possibly singleton) d_1–geodesics and the c_i are cone paths of the form $a_i - v_{\Upsilon_i} - b_i$ for some $a_i \neq b_i \in \Upsilon_i^{(0)}$. Moreover, two vertices of any β_i are contained in a common wide parabolic subcomplex only if they are adjacent in β_i, since otherwise β_i would not be $\hat{d}$–geodesic.

Compare with the 'electric paths' and 'electric path length' in the relatively hyperbolic case [46].

Lemma 7.23. *On $\Sigma^{(0)}$, $d_f \leq \hat{d}$.*

Proof. Let $x, y \in \Sigma^{(0)}$, and let γ be a $\hat{d}$–geodesic connecting them, decomposed as above. Since γ is $\hat{d}$–geodesic, so are its subpaths. Thus, if β_i is a d_1–geodesic subsegment from β_i^- to β_i^+ then $|\beta_i| = \hat{d}(\beta_i^-, \beta_i^+)$. Since it is a combinatorial geodesic, we also have $|\beta_i| = d_w(\beta_i^-, \beta_i^+) \geq d_f(\beta_i^-, \beta_i^+)$. For cone paths $c_i = a_i - v_{\Upsilon_i} - b_i$ Lemma 7.10 says $d_f(a_i, b_i) = 1$. Thus:

$$d_f(x, y) \leq \sum_{i=0}^{m} d_f(\beta_i^-, \beta_i^+) + \sum_{i=1}^{m} d_f(a_i, b_i) \leq \sum_{i=0}^{m} \hat{d}(\beta_i^-, \beta_i^+) + \sum_{i=1}^{m} \hat{d}(a_i, b_i) = \hat{d}(x, y) \qquad \square$$

Given a combinatorial geodesic γ, define its *shortcut graph* $\hat{\gamma}$ to be the graph whose vertices are the vertices of γ, with an edge between two vertices if they are either adjacent in $\Sigma^{(1)}$ or contained in a common wide parabolic subcomplex.

Observe that the distance between the endpoints of γ in $\hat{\gamma}$ is an upper bound on their $\hat{d}$ distance, because edges in $\hat{\gamma}$ correspond to steps of length one in $\widehat{\Sigma}$, but there could be some shorter path in $\widehat{\Sigma}$ with the same endpoints.

If γ is a combinatorial geodesic segment whose vertex sequence is $z_0, z_1, \dots, z_n$ define the *greedy path* in $\hat{\gamma}$ to be the path with vertices $p_i = z_{\phi(i)}$ where ϕ is defined inductively by $\phi(0) := 0$ and:

$$\phi(i+1) := \max\{m \mid z_{\phi(i)} \text{ and } z_m \text{ are adjacent in } \hat{\gamma}\}$$

Lemma 7.24. *The greedy path is a shortest path from z_0 to z_n in $\hat{\gamma}$.*

Proof. Let ϕ be the indexing function of the greedy path, and inductively define a function ϕ' with $\phi'(0) := 0$ and:

$$\phi'(i+1) := \max\{m \mid \exists k \leq \phi'(i),\, z_k \text{ and } z_m \text{ are adjacent in } \hat{\gamma}\}$$

We claim that $\phi(i) = \phi'(i)$ for all i. By definition, $\phi(0) = \phi'(0) = 0$. Suppose inductively that $\phi = \phi'$ on $[0, i]$, and consider $i + 1$. By definition, the maximum defining $\phi'(i + 1)$ is taken over a superset of the maximum defining $\phi(i + 1)$, so $\phi'(i + 1) \geq \phi(i + 1)$. To derive the opposite bound, suppose $j = \phi'(i + 1)$. Then there is some $k \leq \phi'(i)$ such that z_k and z_j are adjacent in $\hat{\gamma}$. If they are adjacent via a $\Sigma^{(1)}$ edge then $j = k + 1$, so $\phi'(i + 1) - 1 = k \leq \phi'(i) = \phi(i) < \phi(i + 1)$ implies $\phi'(i + 1) \leq \phi(i + 1)$. If they are adjacent via a shortcut through a wide parabolic subcomplex Υ then since parabolic subcomplexes are d_1–convex and $k \leq \phi'(i) = \phi(i) < \phi(i + 1) \leq \phi'(i + 1) = j$, the vertex $z_{\phi(i)}$ is also in Υ. Since there is an edge from $z_{\phi(i)}$ to z_j in $\hat{\gamma}$, it follows that $\phi(i + 1) \geq j = \phi'(i + 1)$. Thus $\phi \equiv \phi'$.

Now take any injective edge path from z_0 to z_n in $\hat{\gamma}$, and let $q_0, q_1, \dots, q_m$ be its vertex sequence, and define ψ by $q_i = z_{\psi(i)}$. It follows immediately that $\psi(i) \leq \phi'(i)$ for all i, so by our previous argument $\psi(i) \leq \phi(i)$. Thus, at every step the greedy path is at least as far along γ as the q path. In particular, since $\psi(m) = n$, the greedy path takes at most m steps. Applying this argument to a $\hat{\gamma}$ geodesic implies the greedy path is a $\hat{\gamma}$ geodesic. $\square$

For $N \in \mathbb{N}$, define:

$$K(N) := \sup\{d_{\hat{\gamma}}(x, y) \mid d_f(x, y) \leq N \text{ and } \gamma \text{ is a } d_1\text{–geodesic from } x \text{ to } y\}.$$

The key estimate, in Lemma 7.25, is that $K(4\ell - 1)$ is finite. The proof uses our previous method of producing wide parabolic subcomplexes from grids to show that if d_f collapses the length of a combinatorial geodesic enough then there must be a long subsegment contained in a wide parabolic subcomplex. Recall that $4\ell - 1$ is the quasiruler constant of Lemma 7.3. Lemma 7.26 uses this fact to promote $K(4\ell - 1) < \infty$ to $K(N) < \infty$ for all N.

Lemma 7.25. $K(4\ell - 1) < \infty$

Proof. Suppose not. Then for all $n \in \mathbb{N}$ there exists a combinatorial geodesic $\gamma_n = [x_n, y_n]$ such that $d_f(x_n, y_n) \leq 4\ell - 1$ and $d_{\hat{\gamma}_n}(x_n, y_n) \geq n$. Fix $n \gg 0$, and let $p_0, p_1, \dots, p_q$ be the "greedy vertices", that is, the vertices of the greedy path in $\hat{\gamma}_n$. By Lemma 7.24, the greedy path is a $\hat{\gamma}_n$-geodesic, so $q = d_{\hat{\gamma}_n}(x_n, y_n) \geq n$.

Since γ_n is a combinatorial geodesic, the walls dual to the edges of γ_n are distinct and separate x_n from y_n. In particular, consider the set of "greedy walls" dual to the first edge of γ_n after vertex p_i, for each $0 \leq i < q$. Dilworth's Theorem implies there is a subset of these walls forming a chain $\mathcal{V} = \{V_i\}_{0 \leq i \leq I}$ such that $I + 1 = |\mathcal{V}| \geq q/\dim(X_{\mathrm{NR}}) \geq n/\dim(X_{\mathrm{NR}})$, where the indexing and orientation are as usual: $x_n \in V_i^- \subset V_{i+1}^-$ for all $i < I$.

Consider the subchain $\mathcal{V}' := \{V_0, V_E, V_{2E}, \dots, V_{(4\ell-1)E}\}$, for $E := \lfloor \frac{|\mathcal{V}|-1}{4\ell-1} \rfloor$, assuming $n > 4\ell \dim(X_{\mathrm{NR}})$ so that $E \geq 1$, making the walls of $\mathcal{V}'$ distinct. Then $|\mathcal{V}'| = 4\ell$, so the hypothesis $d_f(x_n, y_n) \leq 4\ell - 1$ implies $\mathcal{V}'$ is not fine. Thus, there is some consecutive pair of walls V_{Ej}, $V_{E(j+1)}$ of $\mathcal{V}'$ that are transverse to a chain $\mathcal{H}$ with $|\mathcal{H}| = 2\ell$. But then $\mathcal{H}$ is transverse to the convex subchain $\mathcal{V}''$ of $\mathcal{V}$ with extreme walls V_{Ej} and $V_{E(j+1)}$, which contains $E + 1 \geq |\mathcal{V}|/4\ell$ many walls.

By Proposition 5.14 there exists a constant C, (depending only on the size of the fundamental generating set) such that $|\mathcal{V}''| \geq C(2\ell - 1)$ implies there exists a wide parabolic subcomplex Υ and a convex subchain $\mathcal{V}'''$ of $\mathcal{V}''$ such that $|\mathcal{V}'''| \geq |\mathcal{V}''|/C$ and the subsegment β of γ_n between the extreme walls of $\mathcal{V}'''$ is contained in $\bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$. We have $|\mathcal{V}''| \geq n/4\ell \dim(X_{\mathrm{NR}})$, so when $n \geq 8C\ell^2 \dim(X_{\mathrm{NR}})$ this result applies and gives $|\mathcal{V}'''| \geq n/4C\ell \dim(X_{\mathrm{NR}})$.

Now proceed as in Section 6.2: Since $\beta \subset \bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon)$ and Υ is convex, the d_w analogue of the 4-gon trick (Lemma 3.5) implies that, for β^- and β^+ the endpoints of β, if $\mathcal{M}$ is the set of all walls transverse to β and $\mathcal{M}'$ is the subset of those that cut Υ, we have:

$$D := |\mathcal{M}| - |\mathcal{M}'| \overset{+}{\preceq} d_w(\beta^-, \Upsilon) + d_w(\beta^+, \Upsilon) \asymp d_2(\beta^-, \Upsilon) + d_2(\beta^+, \Upsilon) \leq 2\mathfrak{R}$$

The first coarse inequality accounts for the potentially uniformly finitely many walls of $\mathcal{M}$ that go through points $\pi_\Upsilon(\beta^+)$ or $\pi_\Upsilon(\beta^-)$ but do not cut Υ, and the second coarse inequality is the relation between d_w and d_2. Both of these are uniform, independent of n and choices of γ_n or β, so there exists D_0, depending only on (W, S), such that $D \leq D_0$ for all n.

The count $D = |\mathcal{M}| - |\mathcal{M}'|$ gives us a decomposition of β as D–many single edges whose dual walls do not cut Υ separating $(D+1)$–many (possibly trivial) subpaths consisting of edges all of whose dual walls cut Υ. By the same argument as in Section 6.2, each of the latter subpaths is contained in a wide parabolic subcomplex. Thus, the diameter of β in $\hat{\gamma}_n$ is at most $2D + 1 \leq 2D_0 + 1$.

For all but the first wall of $\mathcal{V}'''$, the edge path β contains the vertex of γ_n immediately prior to that wall. Since the walls of $\mathcal{V}'''$ are greedy walls, this means that β contains $|\mathcal{V}'''| - 1$ many greedy vertices. Since the greedy path is a $\hat{\gamma}_n$ geodesic, this implies that:

$$\frac{n}{4C\ell \dim(X_{\mathrm{NR}})} - 2 \leq |\mathcal{V}'''| - 2 \leq \mathrm{diam}_{\hat{\gamma}_n}(\beta) \leq 2D_0 + 1$$

This is a contradiction for $n \geq 8C\ell \dim(X_{\mathrm{NR}})(D_0 + 2)$, so we conclude:

$$K(4\ell - 1) \leq 8C\ell \dim(X_{\mathrm{NR}}) \cdot \max\{\ell, D_0 + 2\}$$ □

Lemma 7.26. $\forall N \in \mathbb{N}$, $K(N) \leq NK(4\ell - 1) + N < \infty$

Proof. Let γ be a combinatorial geodesic segment between vertices x and y with vertex sequence $x = z_0, \dots, z_m = y$, such that $d_f(x, y) \leq N$. Define $f(i) := d_f(x, z_i)$. As in the proof of Lemma 7.4, $0 \leq f(i+1) - f(i) \leq 1$.

Let $p_i = z_{\phi(i)}$ be the vertex sequence of the greedy path from z_0 to z_m in $\hat{\gamma}$. Then $0 \leq f(\phi(i+1)) - f(\phi(i)) \leq 1$, since f is nondecreasing and $z_{\phi(i)}$ and $z_{\phi(i+1)}$ are either adjacent in $\Sigma^{(1)}$ or contained in a common wide parabolic subcomplex, so in either case $d_f(z_{\phi(i)}, z_{\phi(i+1)}) = 1$.

The graph of the function f has horizontal 'plateaus', along which f is constant, separated by 'step-ups' with $f(i+1) - f(i) = 1$. Let $P_q := \{i \in [0, m] \cap \mathbb{Z} \mid f(i) = q\}$ be the plateau at height q. The greedy path decomposes into plateau runs, maximal subpaths contained in a single plateau, and single edge plateau step-ups moving from some P_q to P_{q+1}. There are at most N step-up edges. The first edge is a step-up edge, since $P_0 = \{0\}$. Thus, there are at most N non-singleton plateau runs. Suppose $z_a, z_b \in P_q$. By Lemma 7.3:

$$d_f(z_0, z_a) + d_f(z_a, z_b) - (4\ell - 1) \leq d_f(z_0, z_b)$$

Since $q = d_f(z_0, z_a) = d_f(z_0, z_b)$, this gives $d_f(z_a, z_b) \leq 4\ell - 1$, which implies $d_{\hat{\gamma}}(z_a, z_b) \leq K(4\ell - 1)$, which is finite, by Lemma 7.25. Conclude that the plateau runs of the greedy path each contribute at most $K(4\ell - 1)$ to its $\hat{\gamma}$ length, so $d_{\hat{\gamma}}(x, y) \leq N + NK(4\ell - 1)$, which gives $K(N) \leq NK(4\ell - 1) + N$. □

Lemma 7.27. *Let γ be a combinatorial geodesic segment in Σ. Let $\mathcal{V}$ be a longest fine chain separating the endpoints of γ. The d_f-distance between the end vertices of maximal edge-subpaths of $\gamma \setminus \mathcal{V}$ is at most $8\ell - 2$.*

Proof. Let $\gamma\colon [0, T] \to \Sigma^{(1)}$ be the combinatorial geodesic parameterized by arc length. Let $\mathcal{V} = \{V_i\}_{1 \leq i \leq n}$ with $\gamma(0) \subset V_i^- \subset V_{i+1}^-$. Let t_i be the times such that $\gamma(t_i) \in V_i$. These times are half-integral. Let

$t_0 := -1/2$ and $t_{n+1} := T + 1/2$. Suppose for some $0 \le k \le n$ that $\mathcal{H} = \{H_j\}_{0 \le j \le N}$ is a fine chain separating vertices $\gamma(t_k + 1/2)$ and $\gamma(t_{k+1} - 1/2)$, with $N + 1 > 8\ell - 2$.

If $2\ell - 1 \le k \le n - (2\ell - 1)$, then by the argument of Lemma 7.3, there is a fine chain:

$$\{V_i\}_{1 \le i \le k-(2\ell-1)} + \{H_j\}_{2\ell \le j \le N-2\ell} + \{V_i\}_{k+1+2\ell-1 \le i \le n}$$

That chain has length:

$$(k - 2\ell + 1) + (N - 4\ell + 1) + (n - (k + 2\ell) + 1) = n + N - 8\ell + 3$$

That is a contradiction, because since $N > 8\ell - 3$, the new chain is a fine chain separating $\gamma(0)$ and $\gamma(T)$ that is longer than $\mathcal{V}$.

If $k < 2\ell - 1$ and $k \le n - (2\ell - 1)$ then we get a similar contradiction from the chain:

$$\{H_j\}_{0 \le j \le N-2\ell} + \{V_i\}_{k+1+2\ell-1 \le i \le n}$$

This is already longer than $\mathcal{V}$ when $N \ge 6\ell - 3$. A similar argument holds when $k \ge 2\ell - 1$ and $k > n - (2\ell - 1)$.

If both $k < 2\ell - 1$ and $k > n - (2\ell - 1)$, then $n < 4\ell - 2$, so we get a contradiction since $|\mathcal{H}| > |\mathcal{V}|$. □

Proposition 7.28. *There exists C such that for every combinatorial geodesic segment γ in $\Sigma^{(1)}$ with endpoints $x, y \in \Sigma^{(0)}$:*

$$d_{\hat{\gamma}}(x, y) \le C d_f(x, y) + C$$

In particular:

$$\hat{d}(x, y) \le C d_f(x, y) + C$$

Proof. Let $M := 8\ell - 2$. By Lemma 7.26:

$$C := K(M) + 1 \le MK(4\ell - 1) + M + 1 < \infty$$

Let $\gamma\colon [0, T] \to \Sigma^{(1)}$ be a combinatorial geodesic segment parameterized by arc length with endpoints $x := \gamma(0)$ and $y := \gamma(T)$ in $\Sigma^{(0)}$. Let $\mathcal{V} := \{V_i\}_{1 \le i \le n}$ be a longest fine chain separating x and y, with $x \in V_i^- \subset V_{i+1}^-$ for all i. Let t_i be the half-integral time such that $\gamma(t_i) \in V_i$, and set $t_0 := -1/2$ and $t_{n+1} := T + 1/2$.

For each $0 \le i \le n$, let γ_i be the maximal edge-subpath of $\gamma \setminus \mathcal{V}$ with endpoints $a_i := \gamma(t_i + 1/2)$ and $b_i := \gamma(t_{i+1} - 1/2)$. By Lemma 7.27, $d_f(a_i, b_i) \le M$, so $d_{\hat{\gamma}_i}(a_i, b_i) \le K(M)$.

Since $\hat{\gamma}_i$ is a subgraph of $\hat{\gamma}$, this implies $d_{\hat{\gamma}}(a_i, b_i) \le K(M)$ for every i.

Also, for each $1 \le i \le n$ the vertices b_{i-1} and a_i are the endpoints of the edge of γ crossing V_i, so they are adjacent in $\Sigma^{(1)}$, so they are adjacent in $\hat{\gamma}$.

Concatenating these edges with shortest paths in each $\hat{\gamma}_i$, we obtain a path in $\hat{\gamma}$ from x to y of length at most $(n + 1)K(M) + n$. Thus:

$$\hat{d}(x, y) \le d_{\hat{\gamma}}(x, y) \le (n + 1)K(M) + n \le Cn + C = Cd_f(x, y) + C$$ □

8. Acylindricity

Bowditch [23] defines a group action $G \curvearrowright X$ on a metric space to be *acylindrical* if for all $\epsilon \ge 0$ there exist R_ϵ and N_ϵ such that for all points x and y with $d(x, y) \ge R_\epsilon$, the set $\{g \in G \mid d(x, gx) \le \epsilon \text{ and } d(y, gy) \le \epsilon\}$ has cardinality at most N_ϵ.

Theorem 8.1 (Acylindricity)**.** *The action $W \curvearrowright (\widehat{\Sigma}, \hat{d})$ is acylindrical. That is, for all $\epsilon \ge 0$ there are constants R_ϵ and N_ϵ such that $\hat{d}(x, y) \ge R_\epsilon$ implies $|A_\epsilon(x, y)| \le N_\epsilon$, where:*

$$A_\epsilon(x, y) := \{w \in W \mid \hat{d}(x, wx) \le \epsilon \text{ and } \hat{d}(y, wy) \le \epsilon\}$$

Most of this section is dedicated to proving Theorem 8.1, but before diving in we draw some further conclusions, assuming that Theorem 8.1 holds.

Osin [66] formalized the notion of an *acylindrically hyperbolic group* as a group G that admits an acylindrical, non-elementary action on a hyperbolic space. In that work, he also defined a *generalized loxodromic* as an element $g \in G$ such that there exists some acylindrical action of G on a hyperbolic space X_g for which the element g acts by a loxodromic isometry. A *universal acylindrical action* $G \curvearrowright X$ is an acylindrical action of G on a hyperbolic space X such that every generalized loxodromic element $g \in G$ acts loxodromically on X. See also [6].

Generalized loxodromic elements $g \in G$ are Morse [73]; that is, $\mathbb{Z} \to G : z \mapsto g^z$ is a Morse quasigeodesic with respect to any word metric on G. The converse is not true in general; Morse elements need not be generalized loxodromic [6].

Proposition 8.2. *The following are equivalent, for each $w \in W$:*

- *w is a Morse element of W.*
- *w is loxodromic in the action $W \curvearrowright (\widehat{\Sigma}, \hat{d})$.*
- *w is a generalized loxodromic element of W.*

Proof. An element $w \in W$ is Morse if and only if $n \mapsto w^n$ is a Morse quasigeodesic in $(\Sigma^{(0)}, d_1)$. By Corollary 7.22, this is true if and only if $n \mapsto w^n$ is a quasigeodesic in $(\widehat{\Sigma}, \hat{d})$. This is equivalent to w being loxodromic for $W \curvearrowright (\widehat{\Sigma}, \hat{d})$. Since this particular action is an acylindrical action on a hyperbolic space, every element that acts loxodromically here is a generalized loxodromic element of W. All generalized loxodromic elements of W are Morse elements of W [73]. □

Proposition 8.3. *The following are equivalent:*

- *$W \curvearrowright (\widehat{\Sigma}, \hat{d})$ is non-elementary.*
- *(W, S) is nonspherical, not wide, and not virtually infinite cyclic.*
- *The canonical decomposition $P_1 \times \cdots \times P_n \times F$ of (W, S) into a product of irreducible nonspherical factors P_i and a spherical factor F has $n = 1$ and P_1 is not affine.*

Proof. The action is acylindrical, by Theorem 8.1, so by [66, Theorem 1.1], there are three possibilities:

- The action is non-elementary.
- The action has bounded orbits.
- W is virtually infinite cyclic and contains a loxodromic element.

If (W, S) is spherical or wide then $\widehat{\Sigma}$ is bounded. Otherwise, W contains a Morse element [29], and such an element acts loxodromically on $\widehat{\Sigma}$, by Proposition 8.2. In particular, $\widehat{\Sigma}$ is not bounded. Therefore, outside the spherical and wide cases, the action contains a loxodromic element and the action does not have bounded orbits, so the only possibilities are W is virtually infinite cyclic or the action is non-elementary.

To see the equivalence with the product decomposition description, observe that (W, S) is spherical if and only if $n = 0$. It is wide if and only if either $n > 1$ or $n = 1$ and P_1 is higher rank affine. The remaining case to rule out is that $n = 1$ and P_1 is infinite dihedral, which is precisely the case that W is virtually infinite cyclic. □

Corollary 8.4. *The action $W \curvearrowright (\widehat{\Sigma}, \hat{d})$ is a universal acylindrical action. Moreover, if (W, S) is nonspherical, not wide, and not virtually infinite cyclic, then the action is non-elementary.*

Proof. It is an acylindrical action, by Theorem 8.1. The action is universal, by Proposition 8.2. With the additional conditions, the action is non-elementary, by Proposition 8.3. □

Abbott, Balasubramanya, and Osin [1] define a poset of hyperbolic structures on a group: if X and Y are (possibly infinite) generating sets of a group G then Y *dominates* X if the identity map on G induces a Lipschitz map from the Cayley graph $\Gamma(G, Y)$ of G with respect to Y to $\Gamma(G, X)$. Two generating sets are equivalent if they dominate each other. A *hyperbolic structure* is an equivalence class of generating sets of G for which the corresponding Cayley graphs are hyperbolic. An *acylindrically hyperbolic structure* is a hyperbolic structure with the additional property that the actions of G on the corresponding Cayley graphs are acylindrical. The *largest* acylindrically hyperbolic structure, if it exists, is the one that dominates all the others.

Theorem 8.5. *$\widehat{\Sigma}$ defines the largest acylindrically hyperbolic structure for W.*

It is noted in [1] that if a largest acylindrically hyperbolic structure exists then it is universal, so Theorem 8.5 also implies Corollary 8.4.

Proof. The 1–skeleton of the Davis complex is the Cayley graph $\Gamma(W, S)$ of W with respect to S. The coned-off space $\widehat{\Sigma}$ contains the graph $\Sigma^{(1)}$ and adds additional length $1/2$ edges and cone vertices whose effect is to make any two distinct vertices contained in a wide parabolic subcomplex have distance exactly 1. Restricted to $\Sigma^{(0)}$, which is coarsely dense in $\widehat{\Sigma}$, this is exactly the same metric as the word metric defined by the generating set $S' := S \cup \bigcup_{\text{wide } T\subset S} W_T$. Therefore, $\Gamma(W,S')$ is equivariantly quasiisometric to $\widehat{\Sigma}$, so it is hyperbolic and the W–action is acylindrical, by Corollary 7.22 and Theorem 8.1.

The remaining claim is that S' dominates every other acylindrically hyperbolic structure on W. This part of the argument is the same as in [1, Theorem 7.9]. Suppose that X is some other acylindrically hyperbolic structure on W. The restriction of $W \curvearrowright \Gamma(W,X)$ to W_T is also acylindrical, for any $T \subset S$.

We claim that when T is wide every infinite order element $w \in W_T$ has a power with non-virtually cyclic centralizer. If W_T is virtually irreducible higher rank affine then it has a finite index translation subgroup isomorphic to $\mathbb{Z}^n$ for $n \geq 2$, so some power of w is contained in that subgroup, hence has centralizer containing $\mathbb{Z}^2$. If $W_T = A \times B$ with A and B infinite then write $w = (a,b)$. If both a and b have infinite order then $\langle (a,1),(1,b)\rangle \cong \mathbb{Z}^2 \leq Z_W(w)$. If one of them, say b, has finite order n then $w^n = (a^n,1)$ has $\langle w^n\rangle \times B \leq Z(w^n)$.

Powers of loxodromic elements are loxodromic, and centralizers of loxodromic elements in acylindrical actions are virtually cyclic [66, Corollary 6.9], so this shows that $W_T \curvearrowright \Gamma(W,X)$ contains no loxodromic elements, and then Osin's trichotomy implies it has bounded orbits. Since there are only finitely many wide subsets of S, we can take such a bound uniformly. Since S is also finite, $\sup_{s\in S'} |s|_X = C < \infty$, which implies that the identity map on W from $\Gamma(W,S')$ to $\Gamma(W,X)$ is C–Lipschitz. Hence, S' dominates X. □

The remainder of this section builds towards the proof of acylindricity. The proof is strongly inspired by that of acylindricity for the action of a group on a hierarchically hyperbolic space [17, Theorem 14.3]. While we do not know if Coxeter groups are hierarchically hyperbolic, we can prove analogues of some of the HHS tools: one-sided distance formula, large link condition, bounded geodesic image, descent by induction on a notion of 'complexity', which for us is rank of a parabolic subgroup. We do not have analogues for hierarchy paths and the two-sided distance formula, but it turns out that for acylindricity we do not need the full power of those tools, and what we do need we can recover with arguments specific to walls and parabolics in Coxeter groups.

Note that acylindricity is trivial if $\widehat{\Sigma}$ is bounded—just take all R_ϵ bigger then the diameter of $\widehat{\Sigma}$—so we may assume throughout that (W,S) is not wide and nonspherical.

The setup for Theorem 8.1 is written assuming x and y are vertices of the Davis complex. Since every point of $\widehat{\Sigma}$ is distance at most $1/2$ from $\Sigma^{(0)}$, the extension to arbitrary $x,y \in \widehat{\Sigma}$ is an easy step at the very end of the proof.

8.1. **Projections.** Recall from Section 2.10 that given a parabolic subcomplex Υ there is a gate projection $\mathfrak{p}_\Upsilon \colon \Sigma^{(0)} \to \Upsilon^{(0)}$, and it agrees with d_1–closest point projection.

We will now abuse the gate map notation to define a map $\mathfrak{p}_P$, where P is a nonspherical irreducible parabolic subgroup. Given such a P, recall from Proposition 2.9 and Corollary 2.12 that there exists a parabolic subgroup $P^\perp$ such that $Z_W(P) = P^\perp$ and $N_W(P) = P \times P^\perp$. Let $\mathcal{R}(P)$ be the set of parabolic subcomplexes with stabilizer equal to P. By Lemma 2.42, for any two Ξ_0 and Ξ_1 in $\mathcal{R}(P)$ there is an element $g \in P^\perp$ such that left-multiplication by g induces an isomorphism $\Xi_0 \to \Xi_1$. On the other hand, suppose $P = wW_Tw^{-1}$, $\Xi = w\Sigma_T \in \mathcal{R}(P)$, and suppose there are $x,y \in \Sigma_T^{(0)}$ such that there is an $h \in P^\perp$ with $hwx = wy$. Then $h' := w^{-1}hw \in W_T^\perp$, so $wy = hwx = wh'w^{-1}wx = wxh'$, which implies $W_{T^\perp} \ni h' = x^{-1}y \in W_T$. Since $W_T \cap W_{T^\perp}$ is trivial, so are h and h'. Thus, $P^\perp$–orbits in $\coprod_{\Xi\in\mathcal{R}(P)} \Xi^{(0)}$ meet each $\Xi \in \mathcal{R}(P)$ in exactly one vertex. Notice that $\coprod_{\Xi\in\mathcal{R}(P)} \Xi^{(0)} = \Upsilon^{(0)}$, where Υ is the unique parabolic subcomplex whose stabilizer is $N_W(P)$, as in Corollary 2.14. We define $\mathfrak{p}_P$ by fixing any $\Xi \in \mathcal{R}(P)$ and setting:

$$\mathfrak{p}_P \colon \Sigma^{(0)} \to P^\perp \backslash \Upsilon^{(0)} : x \mapsto P^\perp \mathfrak{p}_\Xi(x)$$

By Corollary 2.38 and Lemma 2.39, if $g \in P^\perp$ and $\Xi \in \mathcal{R}(P)$ then $\mathfrak{p}_{g\Xi} = g\mathfrak{p}_\Xi$. Thus, for all $x \in \Sigma^{(0)}$ and $\Xi' \in \mathcal{R}(P)$, we have $\mathfrak{p}_P(x) = P^\perp\mathfrak{p}_\Xi(x) = P^\perp\mathfrak{p}_{\Xi'}(x)$, so $\mathfrak{p}_P(x)$ is a well defined $P^\perp$–orbit of vertices in $\Upsilon^{(0)}$, independent of the choice of $\Xi \in \mathcal{R}(P)$.

Definition 8.6. Given a parabolic subcomplex Υ, define its projection pseudometric on $\Sigma^{(0)}$:

$$d_\Upsilon(x,y) := d_1(\mathfrak{p}_\Upsilon(x), \mathfrak{p}_\Upsilon(y))$$

Given a nonspherical irreducible parabolic subgroup P, define its projection pseudometric on $\Sigma^{(0)}$:

$$d_P(x,y) := d_1(\mathfrak{p}_P(x), \mathfrak{p}_P(y))$$

The definition $d_P(x,y) := d_1(\mathfrak{p}_P(x), \mathfrak{p}_P(y))$ is in terms of the infimal d_1–distance between the $P^\perp$–orbits $\mathfrak{p}_P(x)$ and $\mathfrak{p}_P(y)$. Since the parabolic subcomplex with stabilizer $N_W(P)$ is a product and the $P^\perp$–action induces isomorphisms between P–fibers, an equivalent formulation would be to choose $\Xi \in \mathcal{R}(P)$, set $d_P = d_\Xi$, and observe that the resulting pseudometric is independent of the choice of Ξ.

Lemma 8.7. *For any parabolic subcomplex Υ and $x, y \in \Sigma^{(0)}$:*

$$d_\Upsilon(x,y) = \#\{\text{walls } M \mid M \text{ separates } x \text{ and } y \text{ and cuts } \Upsilon\}$$

Proof. The number of walls separating $\mathfrak{p}_\Upsilon(x)$ and $\mathfrak{p}_\Upsilon(y)$ is equal to $d_1(\mathfrak{p}_\Upsilon(x), \mathfrak{p}_\Upsilon(y)) = d_\Upsilon(x,y)$. It follows from Corollary 2.37 and convexity of Υ that M separates x from y and cuts Υ if and only if M separates $\mathfrak{p}_\Upsilon(x)$ and $\mathfrak{p}_\Upsilon(y)$. ☐

8.2. Coned-off graphs of parabolic subgroups.

Definition 8.8 (Parabolic coned-off graph)**.** Given a parabolic subcomplex Υ, let $C(\Upsilon)$ be the graph obtained from $\Upsilon^{(1)}$ by coning off every proper parabolic subcomplex $\Xi \subsetneq \Upsilon$ by adding a vertex v_Ξ and an edge of length $1/2$ from v_Ξ to every vertex of $\Xi^{(0)}$.

Given a nonspherical irreducible parabolic subgroup P, define:

$$\widetilde{C(P)} := \coprod_{\Upsilon \in \mathcal{R}(P)} C(\Upsilon)$$

Then $P^\perp$ acts on $\widetilde{C(P)}$ by left-multiplication, with $z \in P^\perp$ taking $v \in C(\Upsilon)$ to $zv \in zC(\Upsilon) = C(z\Upsilon)$. Define $C(P)$ to be the quotient of $\widetilde{C(P)}$ by the left $P^\perp$–action.

Remark. In contrast to the curve complex/HHS setting, the graphs $C(\Upsilon)$ and $C(P)$ need not be hyperbolic. Indeed, suppose that $\Upsilon = \Sigma_T$ for $T \subset S$ irreducible higher rank affine. Then W_T has no proper nonspherical parabolics. Since special spherical subgroups have uniformly bounded diameter, $\Upsilon^{(1)} \hookrightarrow C(\Upsilon)$ is a quasiisometry. Thus, $C(\Upsilon)$ is quasiisometric to $\mathbb{Z}^{|T|-1}$.

Lemma 8.9. *If P is a nonspherical irreducible parabolic subgroup and $\Upsilon \in \mathcal{R}(P)$, then the composition $C(\Upsilon) \hookrightarrow \widetilde{C(P)} \to C(P)$ is a graph isomorphism. Moreover, for any two such parabolic subcomplexes Υ and Υ', the resulting isomorphism $C(\Upsilon) \to C(P) \to C(\Upsilon')$ is left-multiplication by the unique $z \in P^\perp$ such that $z\Upsilon = \Upsilon'$.*

There is a canonical assignment of a parabolic subgroup of P to each cone vertex of $C(P)$ by taking the common stabilizer of corresponding $P^\perp$–orbit of cone vertices of the $C(\Upsilon)$.

Proof. This is essentially the same argument as in the construction of $\mathfrak{p}_P$ and d_P in the previous section. The 'ordinary' vertices of $C(P)$ can be identified with $P^\perp$–orbits of pairs (Υ, v), where $\Upsilon \in \mathcal{R}(P)$ and $v \in \Upsilon^{(0)}$. The 'cone' vertices of $C(P)$, which are distinguished by the fact that all of their incident edges have length $1/2$, can be identified with $P^\perp$–orbits of pairs (Υ, Ξ), where $\Upsilon \in \mathcal{R}(P)$ and Ξ is a sub-(parabolic subcomplex) of Υ.

Note that if $z \in P^\perp$ and (Υ, Ξ) is a parabolic subcomplex pair with $\Upsilon \in \mathcal{R}(P)$ then $\mathrm{Stab}(z\Xi) = z\,\mathrm{Stab}(\Xi)z^{-1} = \mathrm{Stab}(\Xi)$, so we can associate the subgroup $\mathrm{Stab}(\Xi)$ to the cone vertex of $C(P)$ corresponding to the $P^\perp$–orbit of (Υ, Ξ), independent of the choice of representative pair. ☐

Just as we defined projection pseudometrics d_Υ and d_P in Definition 8.6, we now define their $C(\Upsilon)$ and $C(P)$ analogues:

Definition 8.10. If Υ is a parabolic subcomplex, let $d_{C(\Upsilon)}$ be the combinatorial distance in the graph $C(\Upsilon)$. Define a projection pseudometric on $\Sigma^{(0)}$:

$$\delta_\Upsilon(x,y) := d_{C(\Upsilon)}(\mathfrak{p}_\Upsilon(x), \mathfrak{p}_\Upsilon(y))$$

Here we implicitly compose the gate map $\mathfrak{p}_\Upsilon \colon \Sigma^{(0)} \to \Upsilon^{(0)}$ with the inclusion $\Upsilon^{(0)} \to C(\Upsilon)^{(0)}$.

Similarly, if P is a nonspherical irreducible parabolic subgroup, define:

$$\delta_P(x,y) := d_{C(P)}(\mathfrak{p}_P(x), \mathfrak{p}_P(y))$$

An equivalent formulation of δ_P is to choose $\Upsilon \in \mathcal{R}(P)$, take $\delta_P := \delta_\Upsilon$, and observe that the resulting pseudometric is independent of the choice of Υ.

Definition 8.11. An irreducible nonspherical parabolic subgroup P is *pre-wide* if $N_W(P)$ is wide.

Lemma 8.12. *If Υ is a wide parabolic subcomplex, then every irreducible nonspherical parabolic $P \leq \mathrm{Stab}(\Upsilon)$ is pre-wide.*

Proof. Let $Q_1 \times \cdots \times Q_n \times K$ be the canonical decomposition of $\mathrm{Stab}(\Upsilon)$ into irreducible nonspherical factors Q_i and a spherical factor K. Since the parabolic subgroup P is a reflection subgroup, Lemma 2.8 says it splits as a direct product of its intersections with the factors of $\mathrm{Stab}(\Upsilon)$. Since it is irreducible, this implies that it is contained in one Q_i, which we may assume is Q_1. Since $\mathrm{Stab}(\Upsilon)$ is wide, either $n > 1$ or $n = 1$ and Q_1 is higher rank affine. In the former case, $Q_2 \leq Q_1^\perp \leq P^\perp$, so $P^\perp$ is nonspherical, so $N_W(P) = P \times P^\perp$ is wide. In the latter case, since proper parabolic subgroups of an irreducible affine Coxeter group are spherical, $P = Q_1$, so P is irreducible higher rank affine, which also implies $N_W(P)$ is wide. □

Lemma 8.13. *Let P be a pre-wide subgroup and let Ξ be any parabolic subcomplex. Then either $P \leq \mathrm{Stab}(\Xi)$ or the δ_P–diameter of Ξ is at most 1.*

Proof. Choose any parabolic subcomplex Υ with $\mathrm{Stab}(\Upsilon) = P$. The projection $\mathfrak{p}_\Upsilon(\Xi)$ is the vertex set of a parabolic subcomplex of Υ. If $\mathfrak{p}_\Upsilon(\Xi) \subsetneq \Upsilon$, then $\mathfrak{p}_\Upsilon(\Xi)$ is coned off in $C(\Upsilon)$, so its image has diameter at most 1. If $\mathfrak{p}_\Upsilon(\Xi) = \Upsilon$, then Υ is parallel to $\mathfrak{p}_\Xi(\Upsilon) \subset \Xi$, and Corollary 2.41 and Lemma 2.2 give:

$$P = \mathrm{Stab}(\Upsilon) = \mathrm{Stab}(\mathfrak{p}_\Upsilon(\Xi)) = \mathrm{Stab}(\mathfrak{p}_\Xi(\Upsilon)) \leq \mathrm{Stab}(\Xi)$$ □

8.3. **Intersection large links and bounded geodesic image.**

Lemma 8.14. *There exists a constant B_{-1} depending only on (W, S) such that for any wide special subgroups P and Q and element $w \in W$, if $|P \cap wQw^{-1}| = \infty$, then $\hat{d}(\mathbb{1}, w) \leq B_{-1}$.*

Proof. We will show it suffices to take $B_{-1} := B_{-2} + 2$, where B_{-2} is the maximal diameter of a spherical special subgroup in $(\widehat{\Sigma}, \hat{d})$. Single generators give spherical special subgroups isomorphic to $\mathbb{Z}/2\mathbb{Z}$, so $B_{-2} \geq 1$. Since there are only finitely many spherical special subgroups, B_{-2} is finite.

Suppose $P \cap wQw^{-1}$ is infinite. Intersections of parabolic subgroups are parabolic, and an infinite parabolic has at least one nonspherical factor in its canonical decomposition as a product of irreducible factors. Thus there are nonspherical irreducible subsets $T, U \subset S$ and elements $p \in P$ and $q \in Q$ such that $W_T \leq P$, $W_U \leq Q$, and $pW_Tp^{-1} = wqW_Uq^{-1}w^{-1}$. By Corollary 2.11, $T = U$, so $h := p^{-1}wq \in N_W(W_T)$. By Proposition 2.9 (2), $N_W(W_T) = W_T \times W_{T^\perp}$, so $h = ab$ for some $a \in W_T$ and $b \in W_{T^\perp}$, and then $w = pabq^{-1}$.

We now bound the $\hat{d}$–length of these factors uniformly. Since $W_T \subseteq P$ and P is wide, every vertex of Σ_{W_T} has $\hat{d}$–distance at most 1 from $\mathbb{1}$, so $\hat{d}(\mathbb{1}, pa) \leq 1$. Similarly, $\hat{d}(\mathbb{1}, q) \leq 1$. If $W_{T^\perp}$ is finite then $\hat{d}(\mathbb{1}, b) \leq B_{-2}$, so $\hat{d}(\mathbb{1}, pabq^{-1}) \leq B_{-2} + 2 = B_{-1}$. If $W_{T^\perp}$ is not finite then $W_T \times W_{T^\perp}$ is wide, so $\hat{d}(\mathbb{1}, ab) \leq 1$ and $\hat{d}(\mathbb{1}, pabq^{-1}) \leq 3 \leq B_{-1}$. □

Corollary 8.15. *If Υ and Ξ are wide parabolic subcomplexes of Σ with corresponding cone vertices v_Υ and v_Ξ in $\widehat{\Sigma}$, and $|\mathrm{Stab}(\Upsilon) \cap \mathrm{Stab}(\Xi)| = \infty$, then $\hat{d}(v_\Upsilon, v_\Xi) \leq B_{-1} + 1$, where B_{-1} is the constant of Lemma 8.14.*

Proof. Write $\Upsilon = a\Sigma_T$ and $\Xi = b\Sigma_U$, where T and U are wide subsets of S. We have $\mathrm{Stab}(\Upsilon) = aW_Ta^{-1}$ and $\mathrm{Stab}(\Xi) = bW_Ub^{-1}$. Thus:

$$|\mathrm{Stab}(\Upsilon) \cap \mathrm{Stab}(\Xi)| = \infty \implies |W_T \cap a^{-1}bW_Ub^{-1}a| = \infty$$

By Lemma 8.14 $\hat{d}(a, b) = \hat{d}(\mathbb{1}, a^{-1}b) \leq B_{-1}$. Then, by the triangle inequality:

$$\hat{d}(v_\Upsilon, v_\Xi) \leq \hat{d}(v_\Upsilon, a) + \hat{d}(a, b) + \hat{d}(b, v_\Xi) \leq B_{-1} + 1$$ □

Lemma 8.16. *Let Υ be a wide parabolic subcomplex, and let $\mathcal{V}$ be a chain of walls cutting Υ with $|\mathcal{V}| \geq 2$. There exists a biinfinite chain $\mathcal{H}$ cutting Υ such that $(\mathcal{V}, \mathcal{H})$ is a grid.*

Proof. See the proof of Lemma 5.11. □

Lemma 8.17. *Let $x, y \in \Sigma^{(0)}$, and let γ be a geodesic in $\widehat{\Sigma}$ from x to y. Denote by $\mathcal{P}(\gamma)$ the collection of wide parabolic subcomplexes whose cone vertices occur on γ. For any wide parabolic subcomplex Υ satisfying $d_\Upsilon(x, y) > \dim(X_{\mathrm{NR}}) \cdot \hat{d}(x, y)$, there exists $\Xi \in \mathcal{P}(\gamma)$ such that $|\mathrm{Stab}(\Upsilon) \cap \mathrm{Stab}(\Xi)| = \infty$.*

Proof. Let $x = z_0, z_1, \ldots, z_m = y$ be the sequence of vertices of $\Sigma^{(0)}$ occurring in γ, so that at each step $z_{i-1} - z_i$ either z_{i-1} and z_i are adjacent in γ and in $\Sigma^{(1)}$, or γ contains a cone shortcut $z_{i-1} - v_\Xi - z_i$, where Ξ is a wide parabolic subcomplex containing z_{i-1} and z_i.

Since ordinary edges have length 1 and cone edges have length 1/2, we have:

$$\big|\mathcal{P}(\gamma)\big| \leq \hat{d}(x, y) = m$$

Now let:

$$\mathcal{W}_\Upsilon(x, y) := \big\{\text{walls } M \mid M \text{ separates } x \text{ and } y \text{ and cuts } \Upsilon\big\}$$

By Lemma 8.7, $|\mathcal{W}_\Upsilon(x, y)| = d_\Upsilon(x, y)$. Since the dimension of the Niblo-Reeves cube complex is the size of the largest collection of pairwise crossing walls, and since the fact that walls of $\mathcal{W}_\Upsilon(x, y)$ separate x and y implies the collection contains no facing triples, Dilworth's theorem implies $\mathcal{W}_\Upsilon(x, y)$ contains a chain $\mathcal{V}$ with:

$$|\mathcal{V}| \geq \frac{d_\Upsilon(x, y)}{\dim(X_{\mathrm{NR}})}$$

Thus, supposing $d_\Upsilon(x, y) > \dim(X_{\mathrm{NR}}) \cdot \hat{d}(x, y)$, we have $|\mathcal{V}| > m$.

Every wall $M \in \mathcal{V}$ separates x from y, so some step of γ has endpoints in opposite halfspaces of M. Assign M to one such step. An ordinary Davis edge crosses exactly one wall, so the 'excess' ($|\mathcal{V}| - m$)–many walls come from cone shortcut steps. In particular, there is some cone shortcut step $z_{i-1} - v_\Xi - z_i$ such that z_{i-1} and z_i are separated by distinct walls $M, M' \in \mathcal{V}$. Because $z_{i-1}, z_i \in \Xi$, both walls cut Ξ. By construction, both walls also cut Υ. Furthermore, they are disjoint, since they came from the chain $\mathcal{V}$. Thus:

$$\mathcal{D}_\infty \cong \langle \mathfrak{r}_M, \mathfrak{r}_{M'} \rangle \leq \mathrm{Stab}(\Upsilon) \cap \mathrm{Stab}(\Xi)$$ □

Proposition 8.18 (Intersection large links)**.** *There are constants λ, D_0, and K_0 such that for every pair of vertices $x, y \in \Sigma^{(0)}$ there exists a collection of wide parabolic subcomplexes $\mathcal{P}(x, y)$ satisfying the following conditions:*

- $|\mathcal{P}(x, y)| \leq \lambda \hat{d}(x, y) + \lambda^2$.
- *For all $\Xi \in \mathcal{P}(x, y)$ and any geodesic γ in $\widehat{\Sigma}$ from x to y, $\hat{d}(v_\Xi, \gamma) \leq D_0$.*
- *If Υ is a wide parabolic subcomplex and $d_\Upsilon(x, y) > K_0$, then there exists $\Xi \in \mathcal{P}(x, y)$ such that $|\,\mathrm{Stab}(\Upsilon) \cap \mathrm{Stab}(\Xi)| = \infty$.*

Compare [17, Proposition 9.4 "large link lemma"], which is similar but concludes that Υ is parallel into Ξ, whereas we only show that the intersection of their stabilizers is nonspherical.

Proof. Let α be a d_1–geodesic in $\Sigma^{(1)}$ from x to y, and let β be the greedy path in the shortcut graph $\hat{\alpha}$ of α. Then β is an edge path that starts at $x = \beta(0)$, ends at $y = \beta(|\beta|)$, and all of whose vertices are vertices of $\alpha^{(0)} \subset \Sigma^{(0)}$. Recall for consecutive vertices $\beta(i)$ and $\beta(i+1)$ of β, either $\beta(i)$ and $\beta(i+1)$ are adjacent in $\Sigma^{(1)}$ or they belong to a common wide parabolic subcomplex. The latter type is a 'shortcut edge'. In both cases, $\hat{d}(\beta(i), \beta(i+1)) = 1$. We also note that it follows from the construction in the proof of Lemma 7.24 that the order of vertices in β is the same as their order in α.

For each shortcut edge $\beta(i) - \beta(i+1)$ that appears along β, choose a wide parabolic subcomplex containing $\beta(i)$ and $\beta(i+1)$. Let $\mathcal{P}(x, y) = \mathcal{P}(\beta)$ be the collection of wide parabolic subcomplexes so chosen. By Corollary 7.5 and Theorem 7.21, there is a λ, which we may assume is greater than $4\ell - 1$, such that β is a uniform (λ, λ)–quasigeodesic from x to y in $\widehat{\Sigma}$, so:

$$|\mathcal{P}(x, y)| \leq |\beta| \leq \lambda \hat{d}(x, y) + \lambda^2$$

Since $(\widehat{\Sigma}, \hat{d})$ is δ–hyperbolic, the quasigeodesic β is contained in a fixed radius D_0 neighborhood of any $\hat{d}$–geodesic γ from x to y in $\widehat{\Sigma}$, where the radius D_0 depends only on λ and δ.

Let C and $\mathfrak{R}$ be the constants of Proposition 5.14. Let $K = K(\mathfrak{R})$ be such that whenever $d_2(u, v) \leq \mathfrak{R} + 1/2 + \mathrm{diam}(\text{chamber})$, then $d_1(u, v) \leq K$. Since $\hat{d} \leq d_1$, this also gives $\hat{d}(u, v) \leq K$. We further set $D := \lambda(2K + 3 + \lambda)$. Finally, let:

$$K_0 := CD \dim(X_{\mathrm{NR}})^2$$

Now let Υ be a wide parabolic subcomplex such that $d_\Upsilon(x, y) > K_0$. By Lemma 8.7:

$$d_\Upsilon(x, y) = |\mathcal{W}_\Upsilon(x, y)| = |\big\{\text{walls } M \mid M \text{ separates } x \text{ and } y \text{ and cuts } \Upsilon\big\}|$$

By Dilworth's theorem, $\mathcal{W}_\Upsilon(x,y)$ contains a chain $\mathcal{V}$ with:

$$|\mathcal{V}| \geq \frac{d_\Upsilon(x,y)}{\dim(X_{\mathrm{NR}})} > \frac{K_0}{\dim(X_{\mathrm{NR}})} = CD\dim(X_{\mathrm{NR}})$$

By Lemma 8.16 there is a chain $\mathcal{H}$ of 2ℓ–many walls cutting Υ such that $(\mathcal{V},\mathcal{H})$ is a grid. By Proposition 5.14, there exists a wide parabolic subcomplex Υ' and a convex subchain $\mathcal{V}'\subset\mathcal{V}$ with $|\mathcal{V}'|\geq\frac{|\mathcal{V}|}{C}$, and a subpath α' of α between the extremal walls of $\mathcal{V}'$ lying within $\bar{\mathcal{N}}_{\mathfrak{R}}(\Upsilon')$. Let a be the last vertex of α before α', and let b be the first vertex of α after α', so that $d_2(a,\Upsilon')\leq\mathfrak{R}+1/2$, $d_2(b,\Upsilon')\leq\mathfrak{R}+1/2$, and every wall of $\mathcal{V}'$ separates a and b. Let a' and b' be vertices of Υ' contained in a common chamber with $\pi_{\Upsilon'}(a)$ and $\pi_{\Upsilon'}(b)$, respectively. Then:

$$\hat{d}(a,b)\leq\hat{d}(a,a')+\hat{d}(a',v_{\Upsilon'})+\hat{d}(v_{\Upsilon'},b')+\hat{d}(b',b)\leq 2K+1$$

Let $\beta(l)$ be the vertex of β closest to a in $\hat{\alpha}$ lying on the left side of a in α. Similarly, let $\beta(r)$ be the vertex of β closest to b in $\hat{\alpha}$ lying to the right of b along α. We have $\hat{d}(a,\beta(l))\leq 1$ and $\hat{d}(b,\beta(r))\leq 1$, so $\hat{d}(\beta(l),\beta(r))\leq 2K+3$. In particular, the quasigeodesic condition gives:

$$r-l\leq\lambda(2K+3)+\lambda^2=D$$

Thus, the segment of β from $\beta(l)$ to $\beta(r)$ has $\hat{d}$–length at most D.

Since every wall of $\mathcal{V}'$ separates a and b, our choices of $\beta(l)$ and $\beta(r)$ imply that every wall of $\mathcal{V}'$ separates $\beta(l)$ and $\beta(r)$. Since $|\mathcal{V}'|>D\dim(X_{\mathrm{NR}})$ and the $\hat{\alpha}$–distance from $\beta(l)$ to $\beta(r)$ is $r-l\leq D$, there is some $s\in\{l,\dots r-1\}$ such that there are more than $\dim(X_{\mathrm{NR}})$–many walls of $\mathcal{V}'$ separating $\beta(s)$ and $\beta(s+1)$. Since $\mathcal{V}'\subset\mathcal{V}$, those walls cut Υ, so Lemma 8.7 yields $d_\Upsilon(\beta(s),\beta(s+1))>\dim(X_{\mathrm{NR}})$.

Note that since $\beta(s)$ and $\beta(s+1)$ are separated by multiple walls, they are not adjacent in $\Sigma^{(1)}$, so their adjacency in $\hat{\alpha}$ is due to a shortcut edge corresponding one of our chosen wide parabolic subcomplexes $\Xi\in\mathcal{P}(x,y)$. Thus:

$$d_\Upsilon(\beta(s),\beta(s+1))>\dim(X_{\mathrm{NR}})=\dim(X_{\mathrm{NR}})\cdot 1=\dim(X_{\mathrm{NR}})\cdot\hat{d}(\beta(s),\beta(s+1))$$

By Lemma 8.17, $\mathrm{Stab}(\Upsilon)\cap\mathrm{Stab}(\Xi)$ is infinite. □

Theorem 8.19 (Bounded Geodesic Image)**.** *Let D_0 and K_0 be the constants of Proposition 8.18. Let $B_0:=D_0+B_{-1}+1$, where B_{-1} is the constant of Lemma 8.14. Let $x,y\in\Sigma^{(0)}$, and let γ be a $\hat{d}$–geodesic in $\widehat{\Sigma}$ from x to y. For any wide parabolic subcomplex Υ, if $\hat{d}(v_\Upsilon,\gamma)>B_0$, then $d_\Upsilon(x,y)\leq K_0$.*

Proof. By Proposition 8.18, if $d_\Upsilon(x,y)>K_0$, then there exist $\Xi\in\mathcal{P}(x,y)$ with $|\mathrm{Stab}(\Upsilon)\cap\mathrm{Stab}(\Xi)|=\infty$. Corollary 8.15 together with the same proposition yields:

$$\hat{d}(v_\Upsilon,\gamma)\leq\hat{d}(v_\Upsilon,v_\Xi)+\hat{d}(v_\Xi,\gamma)\leq B_{-1}+1+D_0=B_0$$ □

Lemma 8.20 (One-sided distance formula)**.** *There is a constant C such that for all $x,y\in\Sigma^{(0)}$:*

$$d_1(x,y)\leq C(1+\hat{d}(x,y))\cdot\max\{1,\sup_\Upsilon d_\Upsilon(x,y)\}$$

The supremum is taken over wide parabolic subcomplexes Υ.

Proof. It suffices to take C to be the constant of Proposition 7.28.

Let α be a d_1–geodesic in $\Sigma^{(1)}$ from x to y, and let β be the greedy path in the shortcut graph $\hat{\alpha}$. Let $x=z_0,\dots,z_{|\beta|}=y$ be the vertex sequence of β. By Proposition 7.28 and Lemma 7.23, $|\beta|\leq C\hat{d}(x,y)+C$.

If the edge from z_i to z_{i+1} in β is an edge of Σ then $d_1(z_i,z_{i+1})=1$. Otherwise it is a shortcut edge, so z_i and z_{i+1} are contained in a common wide parabolic subcomplex Υ. By d_1–convexity, the subsegment α_i of α between z_i and z_{i+1} is a d_1–geodesic in Υ. By definition, $d_\Upsilon(x,y)$ is the d_1–distance between the gate projections $\mathfrak{p}_\Upsilon(x)$ and $\mathfrak{p}_\Upsilon(y)$ of x and y to Υ. Now:

$$\begin{aligned}d_1(x,z_i)+&|\alpha_i|+d_1(z_{i+1},y)=|\alpha|=d_1(x,y)\\&\leq d_1(x,\mathfrak{p}_\Upsilon(x))+d_1(\mathfrak{p}_\Upsilon(x),\mathfrak{p}_\Upsilon(y))+d_1(\mathfrak{p}_\Upsilon(y),y)\end{aligned}$$

Since the gate map agrees with d_1–closest point projection, this gives:

$$0\leq d_1(x,z_i)-d_1(x,\mathfrak{p}_\Upsilon(x))+d_1(z_{i+1},y)-d_1(\mathfrak{p}_\Upsilon(y),y)\leq d_1(\mathfrak{p}_\Upsilon(x),\mathfrak{p}_\Upsilon(y))-|\alpha_i|$$

Thus, $|\alpha_i|\leq d_1(\mathfrak{p}_\Upsilon(x),\mathfrak{p}_\Upsilon(y))=d_\Upsilon(x,y)$. The claim follows. □

8.4. **Factoring paths in the coned-off space.**

Lemma 8.21. *Let $P = W_T$ be a nonspherical irreducible special subgroup that is the stabilizer of a parabolic subcomplex Υ. Given $u, v \in \Upsilon^{(0)}$, let $m := \delta_P(u, v)$. The element $w = u^{-1}v$ admits a factorization $w = h_1 \cdots h_m$, where $h_i \in W_{J_i}$ for $J_i \subsetneq T$. Moreover, this factorization can be chosen so that each h_i is a reduced word in the corresponding W_{J_i} and $h_1 \cdots h_m$ is a reduced word in W.*

Proof. Let $u = z_0, z_1, \ldots, z_m = v$ be a length-m path in $C(\Upsilon)$, where each z_i is a vertex in $\Upsilon^{(0)}$. Each unit step joins two vertices lying in a common proper parabolic subcomplex of Υ, say $z_{i-1}, z_i \in aW_{J_i}$ for some $J_i \subsetneq T$. Thus $h_i = z_{i-1}^{-1} z_i \in W_{J_i}$. Therefore, a length-$m$ path in $C(\Upsilon)$ gives a factorization of w.

Among such paths choose one that minimizes $\sum_{i=1}^{m} |h_i|$. Choose a reduced word for each h_i. If their concatenation was not reduced, then the Coxeter deletion condition gives two letters whose deletion leaves a word representing the same element w. Recall that a reduced word representing an element of W_J only uses letters in J; therefore, the remaining letters in the i-th block still belong to J_i, and hence represent an element of W_{J_i}. Thus, after further reducing each remaining W_{J_i}–word, either some block becomes trivial, contradicting the definition of m, or all m factors remain nontrivial, but the total word length decreases, contradicting minimality of the sums of the word lengths. □

Proposition 8.22 (Witness)**.** *Let $c_2, \ldots, c_{|S|}$ be the positive integers defined in Section 8.6. There exists a constant $M_\epsilon = M_\epsilon(c_2, \ldots, c_{|S|})$ such that for all $x, y \in \Sigma^{(0)}$ and all wide parabolic subcomplexes Υ, if $d_\Upsilon(x, y) > M_\epsilon$, then there exists a pre-wide $P \leq \mathrm{Stab}(\Upsilon)$ such that $\delta_P(x, y) > c_{\mathrm{rank}(P)}$.*

Proof. Let $\Delta \geq 1$ be the maximal d_1–diameter of a spherical special subgroup of W. We will choose $M'_2 < M'_3 < \ldots < M'_{|S|}$ recursively. Set $M'_1 := 0$. Then, for $2 \leq r \leq |S|$, choose:

$$L_r := \Delta + \left\lfloor \frac{r-1}{2} \right\rfloor M'_{r-1} \qquad M'_r := c_r L_r \tag{19}$$

We first prove an irreducible statement by induction on rank.

Claim 8.22.1. Suppose that $P = \mathrm{Stab}(\Upsilon)$ is irreducible and nonspherical with $d_P(x, y) > M'_{\mathrm{rank}(P)}$. Then either $\delta_P(x, y) > c_{\mathrm{rank}(P)}$ or there exists an irreducible nonspherical $Q \subsetneq P$ such that $d_Q(x, y) > M'_{\mathrm{rank}(Q)}$.

Proof of Claim: Up to the W–action, we may assume $P = W_T$ and $\Upsilon = \Sigma_T$ for some $T \subset S$ with $|T| = r := \mathrm{rank}(P)$. If $\delta_P(x, y) > c_r$ there is nothing to prove, so suppose $m := \delta_P(x, y) \leq c_r$. Let $u := \mathfrak{p}_\Upsilon(x)$ and $v := \mathfrak{p}_\Upsilon(y)$. By Lemma 8.21, the element $w = u^{-1}v$ admits a factorization $w = h_1 \cdots h_m$ satisfying $h_i \in W_{J_i}$ for $J_i \subsetneq T$ and:

$$d_P(x, y) = d_P(u, v) = \sum_{i=1}^{m} |h_i|$$

For at least one i the word h_i has at least average length:

$$|h_i| \geq \frac{d_P(x, y)}{m} > \frac{M'_r}{m} = \frac{c_r L_r}{m} \geq \frac{m L_r}{m} = L_r \tag{20}$$

For that i, first observe that if J_i is spherical then $\Delta \geq |h_i| > L_r \geq L_2 = \Delta$, which is a contradiction. Therefore, the canonical decomposition of W_{J_i} as a product of irreducible nonspherical factors $W_{J_{i,j}}$ and a spherical factor F_i contains at least one nonspherical factor:

$$W_{J_i} = W_{J_{i,1}} \times W_{J_{i,2}} \times \cdots \times W_{J_{i,k}} \times F_i$$

Since $2 \leq |J_i| \subsetneq |T| = r$, and since the irreducible nonspherical components of J_i are disjoint sets containing at least two vertices each, the number of non-spherical components is $0 < k \leq \lfloor \frac{r-1}{2} \rfloor$.

Since different factors commute, we can further factorize:

$$h_i = h_{i,1} h_{i,2} \cdots h_{i,k} \cdot f_i$$

Where:

$$h_{i,j} \in W_{J_{i,j}} \quad \text{and} \quad f_i \in F_i \quad \text{and} \quad |h_i| = \sum_{j=1}^{k} |h_{i,j}| + |f_i|$$

Thus, by (20), there is some $1 \leq j \leq k$ such that:

$$|h_{i,j}| \geq \frac{|h_i| - |f_i|}{k} > \frac{L_r - \Delta}{\lfloor \frac{r-1}{2} \rfloor} = M'_{r-1}$$

Because the full concatenated expression for w is reduced, we may commute the component words within h_i and obtain a reduced word for w in which a reduced expression for $h_{i,j}$ occurs as a contiguous subword. The corresponding subpath of the combinatorial geodesic lies in a parabolic subcomplex Ξ whose stabilizer Q is conjugate to $W_{J_{i,j}}$. Moreover, the prefix in w preceding this subword is a word in T, so $\Xi \subsetneq \Upsilon$ and $Q \lneq P$. Every wall crossed by this subpath separates the gates of x and y in Υ, hence separates x and y and cuts Ξ. Thus:

$$d_Q(x, y) \geq |h_{i,j}| > M'_{r-1} \geq M'_{\mathrm{rank}(Q)}$$

This completes the proof of the claim. ◊

Now let Υ be wide. Take the canonical decomposition of $\mathrm{Stab}(\Upsilon)$ as a product of nonspherical irreducible factors and a spherical factor:

$$\mathrm{Stab}(\Upsilon) = P_1 \times \cdots \times P_k \times F$$

The gate distance splits as the sum of the component distances, while the spherical contribution is at most Δ. Therefore, setting $M_\epsilon = \Delta + \lfloor \frac{|S|}{2} \rfloor M'_{|S|}$ guarantees that whenever $d_\Upsilon(x, y) > M_\epsilon$, then at least one of the $k \leq \lfloor \frac{|S|}{2} \rfloor$–many P_i will have $d_{P_i}(x, y) > M'_{\mathrm{rank}(P_i)}$. Applying the claim by inducting on $\mathrm{rank}(P_i)$, we see that there is an irreducible nonspherical $P \leq P_i$ such that $\delta_P(x, y) > c_{\mathrm{rank}(P)}$. Finally, P is pre-wide by Lemma 8.12. □

8.5. **Large links and stability.**

Lemma 8.23. *There is a map ρ from pre-wide parabolics to cone vertices of $\widehat{\Sigma}$ defined by sending a pre-wide parabolic P to the cone vertex corresponding to the unique parabolic subcomplex $\Upsilon_{N(P)}$ whose stabilizer is $N_W(P)$, as in Corollary 2.14. The map has the following properties:*

1. *$g\rho(P) = \rho(gPg^{-1})$*
2. *If $P \leq Q$ are pre-wide parabolic subgroups, then $\hat{d}(\rho(P), \rho(Q)) \leq 1$.*
3. *For all $x, y \in \Sigma^{(0)}$, $\delta_P(x, y) \leq d_{\Upsilon_{N(P)}}(x, y)$.*
4. *If $P \leq Q$ are pre-wide parabolic subgroups, then for all $x, y \in \Sigma^{(0)}$, $\delta_P(x, y) = \delta_P(\mathfrak{p}_Q(x), \mathfrak{p}_Q(y))$.*

Proof. Item (1) is clear.

For Item (2), let Υ_P, Υ_Q, $\Upsilon_{N(P)}$, and $\Upsilon_{N(Q)}$ be parabolic subcomplexes with stabilizers P, Q, $N_W(P)$, and $N_W(Q)$, respectively. The parabolic subcomplexes $\Upsilon_{N(P)}$ and $\Upsilon_{N(Q)}$ are unique, by Corollary 2.14. For Υ_P and Υ_Q there are, potentially, choices to make; choose them in such a way that $\Upsilon_P \subset \Upsilon_Q$. Then $\Upsilon_P \subset \Upsilon_{N(P)} \cap \Upsilon_{N(Q)}$. In particular, every vertex of Υ_P is a $\Sigma^{(0)}$ vertex at distance 1/2 from both the cone vertex for $\Upsilon_{N(P)}$, which is $\rho(P)$, and the cone vertex for $\Upsilon_{N(Q)}$, which is $\rho(Q)$. Thus, $\hat{d}(\rho(P), \rho(Q)) \leq 1$.

Similarly, for Item (4), choose $\Upsilon_P \subset \Upsilon_Q$ and apply Corollary 2.36:

$$\begin{aligned}
\delta_P(x, y) &= d_{C(\Upsilon_P)}(\mathfrak{p}_{\Upsilon_P}(x), \mathfrak{p}_{\Upsilon_P}(y)) \\
&= d_{C(\Upsilon_P)}(\mathfrak{p}_{\Upsilon_P} \circ \mathfrak{p}_{\Upsilon_Q}(x), \mathfrak{p}_{\Upsilon_P} \circ \mathfrak{p}_{\Upsilon_Q}(y)) \\
&= \delta_P(\mathfrak{p}_{\Upsilon_Q}(x), \mathfrak{p}_{\Upsilon_Q}(y)) \\
&= \delta_P(\mathfrak{p}_Q(x), \mathfrak{p}_Q(y))
\end{aligned}$$

Finally, for Item (3), it follows from Corollary 2.36 and the Lipschitz property of gate maps that:

$$\begin{aligned}
\delta_P(x, y) &= d_{C(\Upsilon_P)}(\mathfrak{p}_{\Upsilon_P}(x), \mathfrak{p}_{\Upsilon_P}(y)) \\
&\leq d_1(\mathfrak{p}_{\Upsilon_P}(x), \mathfrak{p}_{\Upsilon_P}(y)) \\
&= d_1(\mathfrak{p}_{\Upsilon_P} \circ \mathfrak{p}_{\Upsilon_{N(P)}}(x), \mathfrak{p}_{\Upsilon_P} \circ \mathfrak{p}_{\Upsilon_{N(P)}}(y)) \\
&\leq d_1(\mathfrak{p}_{\Upsilon_{N(P)}}(x), \mathfrak{p}_{\Upsilon_{N(P)}}(y)) \\
&= d_{\Upsilon_{N(P)}}(x, y)
\end{aligned}$$

□

Corollary 8.24. *If P is a pre-wide parabolic and γ is a $\hat{d}$–geodesic segment in $\widehat{\Sigma}$ between $x, y \in \Sigma^{(0)}$ then $\hat{d}(\gamma, \rho(P)) > B_0$ implies $\delta_P(x, y) \leq K_0$, where B_0 and K_0 are as in Theorem 8.19.*

Proof. Since $\rho(P)$ is the cone vertex corresponding to the unique parabolic subcomplex $\Upsilon_{N(P)}$ with stabilizer $N_W(P)$, we have $\delta_P(x,y) \le d_{\Upsilon_{N(P)}}(x,y) \le K_0$, where the first inequality is Lemma 8.23 (3) and the second is from Theorem 8.19. ☐

8.6. **Some parallel geodesic estimates and constant roundup.**

Lemma 8.25. *Let X be a δ–hyperbolic metric space. Given ϵ there exists H_ϵ such that for any geodesic segments $\alpha\colon [0,L] \to X$ and $\beta\colon [0,L] \to X$, if $d(\alpha(0),\beta(0)) \le \epsilon$ and $d(\alpha(L),\beta(L)) \le \epsilon$, then $d(\alpha(t),\beta(t)) \le H_\epsilon$ for every $t \in [0,L]$.*

Proof. This is an elementary δ–hyperbolicity calculation. We claim it suffices to take $H_\epsilon := 4\delta + 3\epsilon$. ☐

Lemma 8.26. *There exists J such that if α is a $\hat{d}$–geodesic between $x,y \in \Sigma^{(0)} \subset \widehat{\Sigma}$ and $m \in \Sigma^{(0)}$ is a vertex of α then for every d_1–geodesic γ from x to y there exists a vertex $p \in \gamma$ such that $\hat{d}(m,p) \le J$.*

Proof. Let β be the greedy geodesic in the shortcut graph $\hat{\gamma}$ of γ. The vertices of β are vertices of γ, and the edges of β are either edges of γ or they correspond to cone shortcuts in $\widehat{\Sigma}$ between pairs of vertices of γ. There is a δ such that $\widehat{\Sigma}$ is δ–hyperbolic, by Corollary 7.22. As in the proof of Proposition 8.18, there is a path $\bar{\beta}$ in $\widehat{\Sigma}$ with the same $\Sigma^{(0)}$ vertex sequence as β, such that edges of β are replaced by cone shortcuts, and $\bar{\beta}$ is uniformly quasigeodesic in $\widehat{\Sigma}$. Since $\widehat{\Sigma}$ is hyperbolic, quasigeodesics are Morse, so $\bar{\beta}$ is μ–Morse for a Morse gauge μ depending on δ and the quasigeodesic constants, but not on $\bar{\beta}$. Since α is a geodesic with the same endpoints as $\bar{\beta}$, every point of α is within $\hat{d}$–distance $\mu(1,0)$ of some point of $\bar{\beta}$. Since every point of $\bar{\beta}$ is within distance $1/2$ of a $\Sigma^{(0)}$ vertex, which is then a vertex of γ, it suffices to take $J := \mu(1,0) + 1/2$. ☐

We now collect the constants that have appeared so far and define some more.

- δ is the hyperbolicity constant of $\widehat{\Sigma}$.
- ϵ is the acylindricity displacement parameter.
- H_ϵ is the fellow-traveler constant from Lemma 8.25.
- J is the $\hat{d}$–quasiconvexity constant of Lemma 8.26.
- B_{-1} is the constant of Lemma 8.14
- D_0 and K_0 are the constants of Proposition 8.18.
- $B_0 = D_0 + B_{-1} + 1$, as in Theorem 8.19.

Note that B_{-1}, D_0 can be taken to be integers, so B_0 is also an integer. For our later convenience we replace H_ϵ, J, and K_0 by their integer ceilings. Then we make the following proleptic choices of additional constants, which are also integers.

- $C_0 := H_\epsilon + B_0 + 3J$
- $C_1 := C_0 + 2$
- $C_2 := C_1 + H_\epsilon$
- $C_3 := C_2 + B_0 + 1$

Also for each $2 \le r \le |S|$, we choose positive integers $a_r < c_r < b_r$ such that:

- $a_r > 2K_0 + 2C_3$
- $c_r > a_r + 2K_0$
- $b_r > c_r + 2K_0$
- $a_s > b_r + 4K_0$ when $s < r$

Such choices exist, for example by choosing $a_{|S|} > 2K_0 + 2C_3$ first, then choosing $c_{|S|}$, $b_{|S|}$, and $a_{|S|-1}$ satisfying the required conditions, then continuing inductively.

8.7. **Parabolic carriers.** Let $u,v \in \Sigma^{(0)}$. For $i = 1,\dots,|S|-1$, we inductively define three collections $\mathcal{P}_i$, $\mathcal{D}_i$, and $\mathcal{F}_i$ of parabolic subgroups. Elements of $\mathcal{P}_i$ are called the *parent carriers* at level i; those of $\mathcal{D}_i$ are the *descendent carriers* at level i; and those of $\mathcal{F}_i$ are the *factor carriers* at level i.

Set $\mathcal{P}_1 := \{W\}$.

Choose a $\hat{d}$–geodesic γ from u to v in $\widehat{\Sigma}$.

Let $\mathcal{D}_1$ be the set of stabilizers of wide parabolic subcomplexes associated to the cone vertices that occur on γ.

Define $\mathcal{F}_1$ to be the set of irreducible nonspherical factors occurring in the canonical product decomposition of some $Q \in \mathcal{D}_1$.

Now assume for $i \geq 1$, the sets $\mathcal{P}_i$, $\mathcal{D}_i$, and $\mathcal{F}_i$ have been defined. Let:

$$\mathcal{P}_{i+1} := \{P \in \mathcal{F}_i \mid \delta_P(u,v) \leq b_{\mathrm{rank}(P)} + 2K_0\}$$

For each $P \in \mathcal{P}_{i+1}$, choose a geodesic $\hat{\gamma}_P$ in $C(P)$ from $\mathfrak{p}_P(u)$ to $\mathfrak{p}_P(v)$, where $C(P)$ and $\mathfrak{p}_P$ are as defined in Section 8.1 and Definition 8.8.

Let $\mathcal{D}_{i+1}$ be the set of parabolic subgroups associated to the cone vertices on $\hat{\gamma}_P$ for each $P \in \mathcal{P}_{i+1}$.

Let $\mathcal{F}_{i+1}$ be the set of irreducible nonspherical factors occurring in the canonical product decomposition of some $Q \in \mathcal{D}_{i+1}$.

We set $\mathrm{rank}(\mathcal{P}_i) := \sup\{\mathrm{rank}(P) \mid P \in \mathcal{P}_i\}$, which is a maximum unless $\mathcal{P}_i = \emptyset$, and define $\mathrm{rank}(\mathcal{D}_i)$ and $\mathrm{rank}(\mathcal{F}_i)$ similarly. By construction, if $\mathcal{P}_i \neq \emptyset$ then:

$$\mathrm{rank}(\mathcal{P}_i) > \mathrm{rank}(\mathcal{D}_i) \geq \mathrm{rank}(\mathcal{F}_i) \geq \mathrm{rank}(\mathcal{P}_{i+1})$$

In particular, $\mathrm{rank}(\mathcal{P}_i) \leq |S| + 1 - i$, so $\mathrm{rank}(\mathcal{P}_{|S|-1}) \leq 2$, which implies $\mathcal{P}_{|S|} \subset \mathcal{F}_{|S|-1} = \emptyset$.

Lemma 8.27. *For all i, each $F \in \mathcal{F}_i$ is pre-wide.*

Proof. Each $F \in \mathcal{F}_1$ is a nonspherical irreducible parabolic subgroup of a wide parabolic subgroup, by definition of $\mathcal{F}_1$ and $\mathcal{D}_1$, so F is pre-wide, by Lemma 8.12.

Suppose the lemma is true for all $F \in \mathcal{F}_i$ for all $i \leq i_0$. Consider $F \in \mathcal{F}_{i_0+1}$. By construction, F is a nonspherical irreducible factor in the canonical decomposition of some $Q \in \mathcal{D}_{i_0+1}$. In turn, Q is the parabolic subgroup associated to some cone vertex of $C(P)$ for some $P \in \mathcal{P}_{i_0+1} \subset \mathcal{F}_{i_0}$. By the induction hypothesis, P is pre-wide, so $N_W(P)$ is a wide parabolic subgroup. Thus, $F \leq Q \leq P \leq N_W(P)$. Since F is a nonspherical irreducible parabolic contained in the wide parabolic $N_W(P)$, it is pre-wide, again by Lemma 8.12. □

Lemma 8.28. *There is a function $F\colon \mathbb{N} \to \mathbb{N}$ defined by:*

$$F(L) := \sup_{\hat{d}(u,v)=2L} \sup \sum_{i=1}^{|S|-2} |\mathcal{F}_i|$$

The first supremum is taken over pairs $u, v \in \Sigma^{(0)}$ such that $\hat{d}(u,v) = 2L$. The second supremum is taken over all possible choices of carrier systems for the given u and v.

Proof. Take any $u, v \in \Sigma^{(0)}$ such that $\hat{d}(u,v) = 2L$. Construct parent, descendant, and factor carriers for each level, including the choice of $\hat{d}$–geodesic γ from u to v in $\widehat{\Sigma}$. We must show $\sum_{i=1}^{|S|-2} |\mathcal{F}_i|$ is bounded above, depending only on L.

By construction, elements of $\mathcal{F}_i$ are irreducible nonspherical factors of the canonical product decomposition of elements of $\mathcal{D}_i$. If $\mathcal{D}_i = \emptyset$ then $|\mathcal{F}_i| = 0$. Otherwise, since such a factor has rank at least 2, the number of such factors of $Q \in \mathcal{D}_i$ is at most $\left\lfloor \frac{\mathrm{rank}(Q)}{2} \right\rfloor$. Therefore:

$$|\mathcal{F}_i| \leq \left\lfloor \frac{\mathrm{rank}(\mathcal{D}_i)}{2} \right\rfloor \cdot |\mathcal{D}_i| < \frac{|S| \cdot |\mathcal{D}_i|}{2}$$

Since $\hat{d}(u,v) = 2L$, there can be at most $2L$ cone vertices on γ, so $|\mathcal{D}_1| \leq 2L$ and $|\mathcal{F}_1| < L|S|$. For $i \geq 1$, recall the collection of parent carriers:

$$\mathcal{P}_{i+1} := \{P \in \mathcal{F}_i \mid \delta_P(u,v) \leq b_{\mathrm{rank}(P)} + 2K_0\}$$

Each element $P \in \mathcal{P}_{i+1}$ has a corresponding path $\hat{\gamma}_P$ in $C(P)$ of length at most $b_{\mathrm{rank}(P)} + 2K_0$, therefore contributing at most that number of cone vertex parabolics to the descent carriers $\mathcal{D}_{i+1}$. Since $b_2 > b_3 > \cdots$, we have for $i \geq 2$:

$$|\mathcal{D}_i| \leq (b_2 + 2K_0)\, |\mathcal{P}_i| \leq (b_2 + 2K_0)\, |\mathcal{F}_{i-1}|$$

Therefore, $|\mathcal{F}_i| \leq \frac{|S|}{2}(b_2 + 2K_0)\, |\mathcal{F}_{i-1}|$. Recursively, this yields:

$$|\mathcal{F}_i| \leq L|S| \left(\frac{|S|(b_2 + 2K_0)}{2} \right)^{i-1}$$

This gives a uniform upper bound on $\sum_{i=1}^{|S|-2} |\mathcal{F}_i|$, depending only on L and constants:

$$\sum_{i=1}^{|S|-2} |\mathcal{F}_i| \leq L|S| \sum_{j=0}^{|S|-3} \left(\frac{|S|(b_2+2K_0)}{2}\right)^j \qquad \square$$

8.8. **The strong and weak sets for trisection vertices.** Let $x, y \in \Sigma^{(0)}$, and let γ be a $\hat{d}$–geodesic from x to y in $\widehat{\Sigma}$. Choose two vertices m_- and m_+ in $\Sigma^{(0)}$ on γ that are nearest to the points 1/3 and 2/3 of the way along γ. For all $g \in A_\epsilon(x,y)$ and $m \in \{m_\pm\}$, we have $\hat{d}(m, gm) \leq H_\epsilon$, by definition of H_ϵ.

Definition 8.29 (The strong set L_1 and weak set L_2)**.** For fixed x, y, and γ as above, for each $m \in \{m_\pm\}$, let $L_1(m)$ be the collection of pre-wide parabolics P such that:

- $\delta_P(x,y) > c_{\mathrm{rank}(P)}$
- $\hat{d}(m, \rho(P)) \leq C_1$
- There is no pre-wide parabolic $Q \gneq P$ with $\delta_Q(x,y) > c_{\mathrm{rank}(Q)}$.

Similarly, let $L_2(m)$ be the collection of pre-wide parabolics P such that:

- $\delta_P(x,y) > a_{\mathrm{rank}(P)}$
- $\hat{d}(m, \rho(P)) \leq C_2$
- There is no pre-wide parabolic $Q \gneq P$ with $\delta_Q(x,y) > b_{\mathrm{rank}(Q)}$.

In this section we will show that $|L_2(m)|$ is uniformly finite (Proposition 8.30), that $A_\epsilon(x,y)$ conjugates $L_1(m)$ into $L_2(m)$ (Proposition 8.31), and that if $L_1(m) = \emptyset$ then there is a vertex of $\Sigma^{(0)}$ uniformly $\hat{d}$–close to m that has bounded d_1–displacement under the $A_\epsilon(x,y)$–action (Proposition 8.33).

Proposition 8.30 (Uniform finiteness of $L_2(m)$)**.** *Let C_3 be as in Section 8.6. Let F be the function of Lemma 8.28. Let $x, y \in \Sigma^{(0)}$ with $\hat{d}(x,y) \geq R_\epsilon > 3(C_3+1)$, Let γ be any $\hat{d}$–geodesic from x to y, and let $m_\pm$ be the trisection vertices of γ. For $m \in \{m_\pm\}$ we have $|L_2(m)| \leq F(C_3)$.*

Proof. The lower bound on R_ϵ guarantees that the subsegment γ' of γ centered about m of radius C_3 exists and does not contain x or y. Let u and v be the end vertices of γ'. Let $\mathcal{P}_i$, $\mathcal{D}_i$, and $\mathcal{F}_i$ be parent, descendent, and factor carriers for u and v at levels $i = 1, \ldots, |S|-1$, with the assumption that γ' is the geodesic used to define $\mathcal{D}_1$. Our goal is to show that $L_2(m) \subset \bigcup \mathcal{F}_i$. By Lemma 8.28, this would imply the desired bound $|L_2(m)| \leq F(C_3)$.

Let $P \in L_2(m)$. Since $\hat{d}(m, \rho(P)) \leq C_2$ and $\hat{d}(m,u) = \hat{d}(m,v) = C_3 = C_2 + B_0 + 1$, both segments $[x,u]$ and $[v,y]$ are more than B_0 away from $\rho(P)$. Therefore, by Corollary 8.24:

$$\delta_P(x,u),\, \delta_P(v,y) \leq K_0$$

It follows that:

$$\delta_P(u,v) \geq \delta_P(x,y) - 2K_0 > a_{\mathrm{rank}(P)} - 2K_0 \tag{21}$$

We now make the following claim.

Claim 8.30.1. Let $P \in L_2(m)$. If $P \lneq Q$ for some $Q \in \mathcal{P}_i$, then either $P \in \mathcal{F}_i$ or $P \lneq Q'$ for some $Q' \in \mathcal{P}_{i+1}$.

Proof of Claim: We deal with the case of $i = 1$ and $i > 1$ separately. They are similar, but the definition of $\mathcal{D}_1$ is slightly different than $\mathcal{D}_i$ for $i > 1$.

For $i = 1$, recall that $\mathcal{D}_1$ is the collection of wide parabolic subgroups $\mathrm{Stab}(\Upsilon)$, where v_Υ is a cone vertex on γ'. We show that $P \leq \mathrm{Stab}(\Upsilon)$ for some $\mathrm{Stab}(\Upsilon) \in \mathcal{D}_1$. If not, then by Lemma 8.13, the δ_P–diameter of Υ is at most 1, for every such Υ. Thus, every cone shortcut step in γ' has δ_P–diameter at most 1. Every $\Sigma^{(1)}$ step in γ' also has δ_P–diameter at most 1, so $\delta_P(u,v) \leq 2C_3$. Since $a_{\mathrm{rank}(P)} - 2K_0 > 2C_3$, this contradicts (21). Therefore, there exists a cone vertex v_Υ on γ' such that $P \leq \mathrm{Stab}(\Upsilon)$. Since P is irreducible and nonspherical, $P \leq Q$ for one of the canonical nonspherical irreducible factors Q of $\mathrm{Stab}(\Upsilon)$. By construction, such a Q belongs to $\mathcal{F}_1$, so if $P = Q$ we are done. Otherwise, if $P \lneq Q$ then the definition of $L_2(m)$ requires that $\delta_Q(x,y) \leq b_{\mathrm{rank}(Q)}$. Since $P \leq Q$, Lemma 8.23 (2) gives:

$$\hat{d}(m, \rho(Q)) \leq 1 + \hat{d}(m, \rho(P)) \leq 1 + C_2$$

Thus, the subsegments of γ from x to u and from v to y are both more than $\hat{d}$–distance B_0 away from $\rho(Q)$, so Corollary 8.24 gives $\delta_Q(x,u), \delta_Q(y,v) \le K_0$. Then the triangle inequality yields:

$$\delta_Q(u,v) \le \delta_Q(x,y) + 2K_0 \le b_{\mathrm{rank}(Q)} + 2K_0$$

This shows $Q \in \mathcal{P}_2$.

For $i > 1$, suppose $P \lneq Q \in \mathcal{P}_i$. By definition of $\mathcal{P}_i$, this implies $\delta_Q(u,v) \le b_{\mathrm{rank}(Q)} + 2K_0$. Let $\hat{\gamma}_Q$ be the geodesic in $C(Q)$ connecting $\mathfrak{p}_Q(u)$ to $\mathfrak{p}_Q(v)$ used to define $\mathcal{D}_i$. It has length at most $b_{\mathrm{rank}(Q)} + 2K_0$. The same argument as in the $i = 1$ case gives that if $P \not\le \mathrm{Stab}(\Upsilon)$ for every cone vertex v_Υ of $\hat{\gamma}_Q$, then $\delta_P(\mathfrak{p}_Q(u), \mathfrak{p}_Q(v)) \le b_{\mathrm{rank}(Q)} + 2K_0$. Lemma 8.23 (4) says $\delta_P(u,v) = \delta_P(\mathfrak{p}_Q(u), \mathfrak{p}_Q(v))$. However, since $\mathrm{rank}(P) < \mathrm{rank}(Q)$ and $a_{\mathrm{rank}(P)} > b_{\mathrm{rank}(Q)} + 4K_0$, this contradicts (21). Thus, there is some cone vertex v_Υ on $\hat{\gamma}_Q$ with $P \le \mathrm{Stab}(\Upsilon)$, and, by construction, $\mathrm{Stab}(\Upsilon) \in \mathcal{D}_i$. Let Q' be the irreducible nonspherical factor of $\mathrm{Stab}(\Upsilon)$ containing P. Then $Q' \in \mathcal{F}_i$. A similar argument to the $i = 1$ case gives that either $P = Q' \in \mathcal{F}_i$, or $Q' \in \mathcal{P}_{i+1}$. ◊

Now we will apply the claim to see that eventually $P \in \mathcal{F}_i$. At $i = 1$ we have $\mathcal{P}_1 = \{W\}$, so $P \le Q_1 := W \in \mathcal{P}_1$, but P is certainly not all of W since that would mean W is wide and $\widehat{\Sigma}$ is bounded. Thus, $P \lneq Q_1 \in \mathcal{P}_1$, and the assumption of the claim holds at level 1, so the claim says that either $P \in \mathcal{F}_1$, in which case we are done, or $P \lneq Q_2$ for some $Q_2 \in \mathcal{P}_2$. Continue inductively. Since $\mathrm{rank}(\mathcal{P}_j) \le |S| + 1 - j$, there is some first $j_0 \le |S|$ such that $\mathcal{P}_{j_0} = \emptyset$. Thus, $P \in \mathcal{F}_i$ for some $i < j_0 \le |S|$. □

Proposition 8.31 (Stability)**.** *Let $x, y \in \Sigma^{(0)}$ with $\hat{d}(x,y) \ge R_\epsilon > 3(C_3 + \epsilon + 1)$. For any $\hat{d}$–geodesic from x to y, let $m_\pm$ be its trisection vertices. For all $m \in \{m_\pm\}$, $g \in A_\epsilon(x,y)$, and $P \in L_1(m)$, we have $gPg^{-1} \in L_2(m)$.*

Proof. Let $P' := gPg^{-1}$. The definition of $L_1(m)$ gives $\hat{d}(m, \rho(P)) \le C_1$. Since $g\rho(P) = \rho(P')$ and $\hat{d}(m, gm) \le H_\epsilon$, we have $\hat{d}(m, \rho(P')) \le H_\epsilon + C_1 = C_2$. This shows the second condition of $P' \in L_2(m)$.

By our assumption on R_ϵ, the $\hat{d}$–distance from $\rho(P')$ to any geodesic from x to gx is at least:

$$\hat{d}(m,x) - \hat{d}(m,\rho(P')) - \hat{d}(x,gx) \ge R_\epsilon/3 - 1 - C_2 - \epsilon > C_3 - C_2 = B_0 + 1 \tag{22}$$

Therefore, by Corollary 8.24, $\delta_{P'}(x, gx) \le K_0$. Similarly, $\delta_{P'}(y, gy) \le K_0$. Hence:

$$\begin{aligned}
\delta_{P'}(x,y) &\ge \delta_{P'}(gx,gy) - \delta_{P'}(gx,x) - \delta_{P'}(y,gy) \\
&\ge \delta_P(x,y) - 2K_0 \\
&> c_{\mathrm{rank}(P)} - 2K_0 \\
&> a_{\mathrm{rank}(P)} \\
&= a_{\mathrm{rank}(P')}
\end{aligned}$$

This shows the first condition of $P' \in L_2(m)$.

Lastly, suppose that $P' \lneq Q$, for some pre-wide Q. Suppose that $\delta_Q(x,y) > b_{\mathrm{rank}(Q)}$. Then, by Lemma 8.23 (2), $\hat{d}(\rho(P'), \rho(Q)) \le 1$. The extra $+1$ slack in (22) accommodates an extra $\hat{d}(\rho(P'), \rho(Q))$ term in a similar computation, so again Corollary 8.24 implies:

$$\delta_Q(x, gx), \delta_Q(y, gy) \le K_0$$

By the triangle inequality, we have:

$$\begin{aligned}
\delta_{g^{-1}Qg}(x,y) &= \delta_Q(gx,gy) \\
&\ge \delta_Q(x,y) - \delta_Q(x,gx) - \delta_Q(y,gy) \\
&> b_{\mathrm{rank}(Q)} - 2K_0 > c_{\mathrm{rank}(Q)}
\end{aligned}$$

This contradicts $P \in L_1(m)$, so we must have had $\delta_Q(x,y) \le b_{\mathrm{rank}(Q)}$, which gives the final condition for $P' \in L_2(m)$. □

Lemma 8.32. *Fix $x, y \in \Sigma^{(0)}$ and a $\hat{d}$–geodesic γ from x to y in $\widehat{\Sigma}$. Let $m_\pm$ be the trisection vertices of γ. If $L_1(m) = \emptyset$ for some $m \in \{m_\pm\}$ and Υ is a wide parabolic subcomplex with $\hat{d}(m, v_\Upsilon) \le C_0$, then $d_\Upsilon(x,y) \le M_\epsilon$, where M_ϵ is the constant of Proposition 8.22.*

Proof. Suppose $d_\Upsilon(x,y) > M_\epsilon$. Then Proposition 8.22 says there exists some pre-wide witness $P \leq \mathrm{Stab}(\Upsilon)$ with $\delta_P(x,y) > c_{\mathrm{rank}(P)}$. If $\Upsilon = w\Sigma_T$ then $P \leq \mathrm{Stab}(\Upsilon)$ implies $P = wvW_Uv^{-1}w^{-1}$ for some $v \in W_T$ and $U \subset T$. Then $\Xi := wv\Sigma_U$ has $\mathrm{Stab}(\Xi) = P$ and $\Xi \subset \Upsilon$. On the other hand, every parabolic subcomplex with stabilizer P is a subcomplex of the unique parabolic subcomplex with stabilizer $N_W(P)$. Hence, the cone vertices $\rho(P)$ and v_Υ have every vertex of $\Xi^{(0)}$ as mutual neighbors, so $\hat{d}(v_\Upsilon, \rho(P)) \leq 1$.

Take a maximal Q among the set of pre-wide parabolics that contain P and satisfy $\delta_Q(x,y) > c_{\mathrm{rank}(Q)}$. The set is not empty, since it contains P, and chains are uniformly bounded in length by $|S| - 2$, since $\mathrm{rank}(P) \geq 2$ and strict containment of parabolic subgroups forces strict increase in rank. Thus, a maximal element exists.

Since P and Q are nested, Lemma 8.23 (2) says $\hat{d}(\rho(P), \rho(Q)) \leq 1$. Therefore:

$$\hat{d}(m, \rho(Q)) \leq \hat{d}(m, v_\Upsilon) + \hat{d}(v_\Upsilon, \rho(P)) + \hat{d}(\rho(P), \rho(Q)) \leq C_0 + 2 = C_1$$

We have confirmed the three membership conditions for $Q \in L_1(m)$, but $L_1(m) = \emptyset$, so this is a contradiction and we must have had $d_\Upsilon(x,y) \leq M_\epsilon$. □

Proposition 8.33. *Suppose $\hat{d}(x,y) \geq R_\epsilon > 3(B_0 + C_0 + \epsilon + 1)$. Fix a $\hat{d}$–geodesic γ from x to y, and let $m_\pm$ be its trisection vertices. There exists D_ϵ such that if $L_1(m) = \emptyset$ for some $m \in \{m_\pm\}$, then for any choice of a d_1–geodesic from x to y and a vertex p on that geodesic $\hat{d}$–closest to m, we have $d_1(p, gp) \leq D_\epsilon$ for all $g \in A_\epsilon(x,y)$.*

Proof. We show it suffices to take $D_\epsilon := C(1 + H_\epsilon + 2J)(2M_\epsilon + K_0)$, where C is the constant of Lemma 8.20, M_ϵ is as in Lemma 8.32, and H_ϵ, J, and K_0 are as in Section 8.6.

Suppose $L_1(m) = \emptyset$. For any choice of d_1–geodesic from x to y, take a vertex p that is $\hat{d}$–closest to m. Then $\hat{d}(m,p) \leq J$, by Lemma 8.26.

For all $g \in A_\epsilon(x,y)$, we have $\hat{d}(m, gm) \leq H_\epsilon$, by Lemma 8.25, so $\hat{d}(p, gp) \leq H_\epsilon + 2J$.

We claim $d_\Upsilon(p, gp) \leq 2M_\epsilon + K_0$, for all wide parabolic subcomplexes Υ. If $d_\Upsilon(p, gp) \leq K_0$ we are happy, so suppose not. Then Theorem 8.19 implies that the $\hat{d}$–distance from v_Υ to any $\hat{d}$–geodesic from p to gp is at most B_0. Since the length of any such geodesic is at most $H_\epsilon + 2J$, it follows that:

$$\hat{d}(v_\Upsilon, m) \leq \hat{d}(v_\Upsilon, p) + \hat{d}(p, m) \leq B_0 + H_\epsilon + 3J = C_0$$

The same estimate holds using gp and gm, so:

$$\hat{d}(m, v_{g^{-1}\Upsilon}) = \hat{d}(gm, v_\Upsilon) \leq \hat{d}(v_\Upsilon, gp) + \hat{d}(gp, gm) \leq C_0$$

Lemma 8.32 gives $d_\Upsilon(x,y) \leq M_\epsilon$ and $d_{g^{-1}\Upsilon}(x,y) \leq M_\epsilon$.

By Lemma 8.7, $d_\Upsilon(x,y)$ is equal to the number of walls that separate x from y and cut Υ. Since p is a vertex on a d_1–geodesic between x and y:

$$d_\Upsilon(x,p) \leq d_\Upsilon(x,p) + d_\Upsilon(p,y) = d_\Upsilon(x,y) \leq M_\epsilon$$

Similarly:

$$d_\Upsilon(gx, gp) = d_{g^{-1}\Upsilon}(x,p) \leq d_{g^{-1}\Upsilon}(x,y) \leq M_\epsilon$$

Any $\hat{d}$–geodesic from x to gx has $\hat{d}$–distance from v_Υ at least:

$$\hat{d}(m,x) - \hat{d}(m, v_\Upsilon) - \hat{d}(x, gx) \geq R_\epsilon/3 - 1 - C_0 - \epsilon > B_0$$

Then Theorem 8.19 implies $d_\Upsilon(x, gx) \leq K_0$.

Putting the last three bounds together:

$$\begin{aligned} d_\Upsilon(p, gp) &\leq d_\Upsilon(p,x) + d_\Upsilon(x,gx) + d_\Upsilon(gx,gp) \\ &\leq M_\epsilon + K_0 + M_\epsilon \end{aligned}$$

This proves the claim.

Now apply Lemma 8.20:

$$\begin{aligned} d_1(p, gp) &\leq C(1 + \hat{d}(p, gp)) \cdot \max\{1, \sup_{\text{wide } \Upsilon} d_\Upsilon(p, gp)\} \\ &\leq C(1 + H_\epsilon + 2J)(2M_\epsilon + K_0) \end{aligned}$$

□

8.9. **Proof of acylindricity.** Choose R_ϵ large enough so that:

$$R_\epsilon/3 - 1 > \max\{C_3 + \epsilon,\, B_{-1} + 2C_1 + 2\} \tag{23}$$

The first condition guarantees R_ϵ is sufficiently large to apply Proposition 8.30, Proposition 8.31, and Proposition 8.33.

Proof of Theorem 8.1. Let R_ϵ satisfy (23). Let $x, y \in \Sigma^{(0)}$, with $\hat{d}(x, y) \geq R_\epsilon$, fix a $\hat{d}$–geodesic from x to y, and let $m_\pm$ be its trisection vertices. We have two cases:

Case A. $L_1(m) = \emptyset$ for some $m \in \{m_\pm\}$.

Then, by Proposition 8.33, for any choice of a d_1–geodesic from x to y and a vertex p on that geodesic $\hat{d}$–closest to m, we have $d_1(p, gp) \leq D_\epsilon$ for all $g \in A_\epsilon(x, y)$. Since W acts freely on $\Sigma^{(0)}$, which is locally finite, this implies:

$$|A_\epsilon(x, y)| \leq |\{w \in W \mid d_1(p, wp) \leq D_\epsilon\}| = |\bar{\mathcal{N}}_{D_\epsilon}(1)| < \infty$$

Case B. $L_1(m) \neq \emptyset$ for both $m \in \{m_\pm\}$.

Fix a pair $P_- \in L_1(m_-)$ and $P_+ \in L_1(m_+)$. By Proposition 8.31, for any $g \in A_\epsilon(x, y)$:

$$gP_-g^{-1} \in L_2(m_-) \quad \text{and} \quad gP_+g^{-1} \in L_2(m_+)$$

Define the map:

$$\Phi\colon\ A_\epsilon(x, y) \to L_2(m_-) \times L_2(m_+) : g \mapsto (gP_-g^{-1}, gP_+g^{-1})$$

By Proposition 8.30, $|L_1(m)| \leq |L_2(m)| \leq F(C_3)$, so the cardinality of the image of Φ is bounded by $F(C_3)^2$. We now bound the cardinality of each fiber. Suppose $g, h \in A_\epsilon(x, y)$ are in the same fiber, so:

$$gP_-g^{-1} = hP_-h^{-1} \quad \text{and} \quad gP_+g^{-1} = hP_+h^{-1}$$

Equivalently, $h^{-1}g \in N_W(P_-) \cap N_W(P_+)$. By Definition 8.29, $P_- \in L_1(m_-)$ implies $\hat{d}(m_-, \rho(P_-)) \leq C_1$ and, similarly, $\hat{d}(m_+, \rho(P_+)) \leq C_1$. Now:

$$\hat{d}(\rho(P_-), \rho(P_+)) \geq \hat{d}(m_-, m_+) - 2C_1 \geq \hat{d}(x, y)/3 - 2 - 2C_1 \geq R_\epsilon/3 - 2 - 2C_1 \overset{(23)}{>} B_{-1} + 1$$

Then the contrapositive of Corollary 8.15 implies $N_W(P_-) \cap N_W(P_+)$ is a spherical parabolic. It follows that $|A_\epsilon(x, y)|$ is bounded by $F(C_3)^2$ times the maximum cardinality of a spherical parabolic subgroup of (W, S).

Acylindricity for $x, y \in \Sigma^{(0)}$ follows by taking:

$$N_\epsilon := \max\{|\bar{\mathcal{N}}_{D_\epsilon}(1)|, F(C_3)^2 \cdot \max\{|W_T| \mid \text{spherical } T \subset S\}\}$$

For the general case, use the above R_ϵ and N_ϵ to define $R'_\epsilon := R_{\epsilon+1} + 1$ and $N'_\epsilon := N_{\epsilon+1}$. Suppose $x', y' \in \widehat{\Sigma}$ with $\hat{d}(x', y') \geq R'_\epsilon$. Choose closest $\Sigma^{(0)}$ vertices x to x' and y to y'. Then $\hat{d}(x, x'), \hat{d}(y, y') \leq 1/2$, so $\hat{d}(x, y) \geq R_{\epsilon+1}$ and $A_\epsilon(x', y') \subset A_{\epsilon+1}(x, y)$. Thus:

$$\hat{d}(x', y') \geq R'_\epsilon \implies |A_\epsilon(x', y')| \leq |A_{\epsilon+1}(x, y)| \leq N_{\epsilon+1} = N'_\epsilon$$

□

AI use

The authors used several versions of the AI tool ChatGPT as interactive sounding boards to explore possible proof strategies and test their arguments. The conception of the arguments and the writing of this paper were done by the authors.

References

[1] C. Abbott, S. H. Balasubramanya, and D. Osin, *Hyperbolic structures on groups*, Algebr. Geom. Topol. **19** (2019), no. 4, 1747–1835.

[2] C. Abbott, J. Behrstock, and M. G. Durham, *Largest acylindrical actions and stability in hierarchically hyperbolic groups*, Trans. Amer. Math. Soc. Ser. B **8** (2021), 66–104, With an appendix by Daniel Berlyne and Jacob Russell.

[3] C. Abbott, J. Behrstock, and J. Russell, *Advances in hierarchical hyperbolicity problem list*, 2024, Compiled by the participants of the BIRS workshop *Advances in Hierarchical Hyperbolicity* (24w5254), Banff International Research Station; edited by Carolyn Abbott, Jason Behrstock, and Jacob Russell; version dated November 26, 2024, `https://www.wescac.net/HHG_Question_Session-dec-2024.pdf`.

[4] C. Abbott and M. Incerti-Medici, *Hyperbolic projections and topological invariance of sublinearly Morse boundaries*, preprint, `arXiv:2212.09539`.

[5] C. Abbott and S. Zbinden, *A non-loxodromic Morse element in a Morse local-to-global group*, preprint, `arXiv:2502.21176`.

[6] C. R. Abbott and D. Hume, *The geometry of generalized loxodromic elements*, Ann. Inst. Fourier (Grenoble) **70** (2020), no. 4, 1689–1713.

[7] D. V. Alekseevski, P. W. Michor, and Y. A. Neretin, *Rolling of Coxeter polyhedra along mirrors*, Geometric methods in physics, Trends Math., Birkhäuser/Springer, Basel, 2013, pp. 67–86.
[8] Y. Algom-Kfir, *Strongly contracting geodesics in outer space*, Geom. Topol. **15** (2011), no. 4, 2181–2233.
[9] D. Allcock, *Normalizers of parabolic subgroups of Coxeter groups*, Algebr. Geom. Topol. **12** (2012), no. 2, 1137–1143.
[10] D. Allcock, *Reflection centralizers in Coxeter groups*, Transform. Groups **18** (2013), no. 3, 599–613.
[11] T. Aougab, M. G. Durham, and S. J. Taylor, *Pulling back stability with applications to Out*(F_n) *and relatively hyperbolic groups*, J. Lond. Math. Soc. **96** (2017), no. 3, 565–583.
[12] G. N. Arzhantseva, C. H. Cashen, and J. Tao, *Growth tight actions*, Pacific J. Math. **278** (2015), no. 1, 1–49.
[13] P. Azuelos and M. Hagen, *Periodic quasiflats in hierarchically hyperbolic spaces*, preprint, `arXiv:2608.01513`.
[14] S. Balasubramanya, M. Chesser, A. Kerr, J. Mangahas, and M. Trin, *(Non-)recognizing spaces for stable subgroups*, Proc. Amer. Math. Soc. **153** (2025), no. 8, 3197–3207.
[15] W. Ballmann and M. Brin, *Orbihedra of nonpositive curvature*, Publ. Math. Inst. Hautes Études Sci. **82** (1995), no. 1, 169–209.
[16] J. Behrstock and R. Charney, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. **352** (2012), no. 2, 339–356.
[17] J. Behrstock, M. F. Hagen, and A. Sisto, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. **21** (2017), no. 3, 1731–1804.
[18] J. Behrstock, M. F. Hagen, and A. Sisto, *Thickness, relative hyperbolicity, and randomness in Coxeter groups*, Algebr. Geom. Topol. **17** (2017), no. 2, 705–740, With an appendix written jointly with Pierre-Emmanuel Caprace.
[19] M. Bestvina and K. Fujiwara, *A characterization of higher rank symmetric spaces via bounded cohomology*, Geom. Funct. Anal. **19** (2009), no. 1, 11–40.
[20] M. Bestvina, B. Kleiner, and M. Sageev, *The asymptotic geometry of right-angled Artin groups. I*, Geom. Topol. **12** (2008), no. 3, 1653–1699.
[21] S. Blachère, P. Haïssinsky, and P. Mathieu, *Harmonic measures versus quasiconformal measures for hyperbolic groups*, Ann. Sci. Éc. Norm. Supér. (4) **44** (2011), no. 4, 683–721.
[22] R. E. Borcherds, *Coxeter groups, Lorentzian lattices, and* $K3$ *surfaces*, Internat. Math. Res. Notices (1998), no. 19, 1011–1031.
[23] B. H. Bowditch, *Tight geodesics in the curve complex*, Invent. Math. **171** (2008), no. 2, 281–300.
[24] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der mathematischen Wissenschaften, vol. 319, Springer, Berlin, 1999.
[25] B. Brink, *On centralizers of reflections in Coxeter groups*, Bull. London Math. Soc. **28** (1996), no. 5, 465–470.
[26] B. Brink and R. B. Howlett, *Normalizers of parabolic subgroups in Coxeter groups*, Invent. Math. **136** (1999), no. 2, 323–351.
[27] P.-E. Caprace, *Conjugacy of 2-spherical subgroups of Coxeter groups and parallel walls*, Algebr. Geom. Topol. **6** (2006), 1987–2029.
[28] P.-E. Caprace, *Buildings with isolated subspaces and relatively hyperbolic Coxeter groups*, Innov. Incidence Geom. **10** (2009), 15–31.
[29] P.-E. Caprace and K. Fujiwara, *Rank-one isometries of buildings and quasi-morphisms of Kac-Moody groups*, Geom. Funct. Anal. **19** (2010), no. 5, 1296–1319.
[30] P.-E. Caprace and F. Haglund, *On geometric flats in the CAT(0) realization of Coxeter groups and Tits buildings*, Canad. J. Math. **61** (2009), no. 4, 740–761.
[31] P.-E. Caprace and T. Marquis, *Open subgroups of locally compact Kac-Moody groups*, Math. Z. **274** (2013), no. 1-2, 291–313.
[32] P.-E. Caprace and A. Minasyan, *On conjugacy separability of some Coxeter groups and parabolic-preserving automorphisms*, Illinois J. Math. **57** (2013), no. 2, 499–523.
[33] P.-E. Caprace and B. Mühlherr, *Reflection triangles in Coxeter groups and biautomaticity*, J. Group Theory **8** (2005), no. 4, 467–489.
[34] P.-E. Caprace and P. Przytycki, *Bipolar Coxeter groups*, J. Algebra **338** (2011), 35–55.
[35] C. H. Cashen, *Morse, contracting, and strongly contracting sets with applications to boundaries and growth of groups*, Habilitationsschrift, University of Vienna, 2019.
[36] C. H. Cashen, *Morse subsets of CAT(0) spaces are strongly contracting*, Geom. Dedicata **204** (2020), no. 1, 311–314.
[37] C. H. Cashen, P. Dani, A. Edletzberger, and A. Karrer, *RAAGedy right-angled Coxeter groups*, preprint, `arXiv:2506.16789`.
[38] R. Charney and H. Sultan, *Contracting boundaries of* CAT(0) *spaces*, J. Topol. **8** (2015), no. 1, 93–117.
[39] M. Cordes and D. Hume, *Stability and the Morse boundary*, J. Lond. Math. Soc. **95** (2017), no. 3, 963–988.
[40] M. Cordes and I. Levcovitz, *Connectivity of Coxeter group Morse boundaries*, preprint, `arXiv:2503.14085`.
[41] M. W. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008.
[42] V. V. Deodhar, *On the root system of a Coxeter group*, Comm. Algebra **10** (1982), no. 6, 611–630.
[43] C. Druţu, S. Mozes, and M. Sapir, *Divergence in lattices in semisimple Lie groups and graphs of groups*, Trans. Amer. Math. Soc. **362** (2010), no. 5, 2451–2505.
[44] C. Druţu and M. Sapir, *Tree-graded spaces and asymptotic cones of groups*, Topology **44** (2005), no. 5, 959–1058, With an appendix by Denis Osin and Sapir.
[45] M. G. Durham and S. J. Taylor, *Convex cocompactness and stability in mapping class groups*, Algebr. Geom. Topol. **15** (2015), no. 5, 2839–2859.
[46] B. Farb, *Relatively hyperbolic groups*, Geom. Funct. Anal. **8** (1998), no. 5, 810–840.

[47] B. Farb and L. Mosher, *Convex cocompact subgroups of mapping class groups*, Geom. Topol. **6** (2002), 91–152.
[48] A. Genevois, *Hyperbolicities in* CAT(0) *cube complexes*, Enseign. Math. **65** (2019), no. 1-2, 33–100.
[49] A. Genevois, *Coning-off* CAT(0) *cube complexes*, Ann. Inst. Fourier (Grenoble) **71** (2021), no. 4, 1535–1599.
[50] M. F. Hagen, *Cocompactly cubulated crystallographic groups*, J. Lond. Math. Soc. (2) **90** (2014), no. 1, 140–166.
[51] M. F. Hagen, *Weak hyperbolicity of cube complexes and quasi-arboreal groups*, J. Topol. **7** (2014), no. 2, 385–418.
[52] U. Hamenstaedt, *Word hyperbolic extensions of surface groups*, unpublished, `arXiv:math/0505244`.
[53] R. B. Howlett, *Normalizers of parabolic subgroups of reflection groups*, J. London Math. Soc. (2) **21** (1980), no. 1, 62–80.
[54] A. E. Kent and C. J. Leininger, *Shadows of mapping class groups: capturing convex cocompactness*, Geom. Funct. Anal. **18** (2008), no. 4, 1270–1325.
[55] T. Koberda, J. Mangahas, and S. J. Taylor, *The geometry of purely loxodromic subgroups of right-angled Artin groups*, Trans. Amer. Math. Soc. **369** (2017), no. 11, 8179–8208.
[56] D. Krammer, *The conjugacy problem for Coxeter groups*, Groups Geom. Dyn. **3** (2009), no. 1, 71–171.
[57] M. Kudlinska and H. Petyt, *Largest acylindrical actions of free-by-cyclic groups*, preprint, `arXiv:2512.05293`.
[58] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149.
[59] Y. N. Minsky, *Extremal length estimates and product regions in Teichmüller space*, Duke Math. J. **83** (1996), no. 2, 249–286.
[60] Y. N. Minsky, *Quasi-projections in Teichmüller space*, J. Reine Angew. Math. **473** (1996), 121–136.
[61] P. Möller, *A note on almost negative matrices and Gromov-hyperbolic Coxeter groups*, Fund. Math. **268** (2025), no. 1, 1–21.
[62] G. Moussong, *Hyperbolic Coxeter groups*, Ph.D. thesis, Ohio State University, 1988, p. 55.
[63] G. A. Niblo and L. D. Reeves, *Coxeter groups act on* CAT(0) *cube complexes*, J. Group Theory **6** (2003), no. 3, 399–413.
[64] K. Nuida, *On centralizers of parabolic subgroups in Coxeter groups*, J. Group Theory **14** (2011), no. 6, 891–930.
[65] S. Oh, *Quasi-isometry invariants of weakly special square complexes*, Topology and its Applications **307** (2022), 107945.
[66] D. Osin, *Acylindrically hyperbolic groups*, Trans. Amer. Math. Soc. **368** (2016), no. 2, 851–888.
[67] J. Perlmutter, *The Morse local-to-global property for graph products*, `arXiv:2601.09901`.
[68] H. Petyt and D. Spriano, *Unbounded domains in hierarchically hyperbolic groups*, Groups Geom. Dyn. **17** (2023), no. 2, 479–500.
[69] H. Petyt, D. Spriano, and A. Zalloum, *Hyperbolic models for CAT(0) spaces*, Adv. Math. **450** (2024), 1–66.
[70] H. Petyt and A. Zalloum, *Constructing metric spaces from systems of walls*, With an appendix with Davide Spriano, preprint, `arXiv:2404.12057`.
[71] J. Russell, D. Spriano, and H. C. Tran, *The local-to-global property for Morse quasi-geodesics*, Math. Z. **300** (2022), no. 2, 1557–1602.
[72] J. Russell, D. Spriano, and H. C. Tran, *Convexity in hierarchically hyperbolic spaces*, Algebr. Geom. Topol. **23** (2023), no. 3, 1167–1248.
[73] A. Sisto, *Quasi-convexity of hyperbolically embedded subgroups*, Math. Z. **283** (2016), no. 3-4, 649–658.
[74] A. Sisto, *Contracting elements and random walks*, J. Reine Angew. Math. **742** (2018), 79–114.
[75] H. Sultan, *Hyperbolic quasi-geodesics in CAT(0) spaces*, Geom. Dedicata **169** (2014), no. 1, 209–224.
[76] H. Tran, *On strongly quasiconvex subgroups*, Geom. Topol. **23** (2019), no. 3, 1173–1235.
[77] S. Zbinden, *Hyperbolic spaces that detect all strongly-contracting directions*, preprint, `arXiv:2404.12162`.

TU Wien, Institute of Discrete Mathematics and Geometry, Wiedener Hauptstrasse 8-10, 1040 Vienna, Austria, 0000-0002-6340-469X
Email address: `christopher.cashen@tuwien.ac.at`

University of Minnesota, School of Mathematics, 206 Church St. SE, Minneapolis, MN 55455
Email address: `mchu@umn.edu`

University of Oklahoma, Department of Mathematics, 601 Elm Avenue Room 423, Norman, OK 73019
Email address: `jing@ou.edu`