



Geometry & Topology

Volume 27 (2023)

Stable cubulations, bicombings, and barycenters

MATTHEW G DURHAM

YAIR N MINSKY

ALESSANDRO SISTO



Stable cubulations, bicombings, and barycenters

MATTHEW G DURHAM

YAIR N MINSKY

ALESSANDRO SISTO

We prove that the hierarchical hulls of finite sets of points in mapping class groups and Teichmüller spaces are stably approximated by $CAT(0)$ cube complexes, strengthening a result of Behrstock, Hagen and Sisto. As applications, we prove that mapping class groups are semihyperbolic and Teichmüller spaces are coarsely equivariantly bicombable, and both admit stable coarse barycenters. Our results apply to the broader class of “colorable” hierarchically hyperbolic spaces and groups.

20F65, 57K20

1 Introduction

Much of the coarse structure of mapping class groups has the flavor of $CAT(0)$ geometry, in spite of the fact that mapping class groups have no geometric actions on $CAT(0)$ spaces; see Bridson [17]. Manifestations of this include the weakly relatively hyperbolic structure associated to curve complexes—see Masur and Minsky [42]—and the equivariant embedding into finite products of quasitrees found by Bestvina, Bromberg and Fujiwara [13].

A notion of “hulls” of finite sets in mapping class groups was introduced by Behrstock, Kleiner, Minsky and Mosher in [9], and these were more recently shown by Behrstock, Hagen and Sisto [8] to be approximated in a uniform way by finite $CAT(0)$ cube complexes—see also the alternative proof given by Bowditch in [16]. Our goal in this paper is to refine this construction to make it *stable*, in the sense that perturbation of the input data gives rise to bounded change in the cubical structure. As initial applications, we give a construction for equivariant barycenters and a proof that mapping class groups are bicombable.

As in [8], the proof works in a more general context of *hierarchically hyperbolic groups*, a class of groups (and spaces) introduced by Behrstock, Hagen and Sisto [6; 7] which are endowed with a structure similar to the hierarchical family of curve complexes associated to a surface; see Masur and Minsky [43]. See Section 2.2 for the definition of a hierarchically hyperbolic space (HHS).

Our main result, stated informally, is the following:

Theorem A *In a colorable HHS $(\mathcal{X}, \mathfrak{S})$, the coarse hull $H_\theta(F)$ of any finite set F can be approximated in a coarsely equivariant way by a finite $CAT(0)$ cube complex whose dimension is bounded by the complexity of $(\mathcal{X}, \mathfrak{S})$, in such a way that a bounded change in F corresponds to a change of the cubical structure by a bounded number of hyperplane deletions and insertions.*

The colorability assumption — see Definition 2.8 — in Theorem A is apparently quite weak and excludes none of the key examples of HHSs, though there are noncolorable HHGs; see Hagen [36].

For the general context of this result, see the discussion in Section 1.2, where we also give a more precise statement in Theorem 1.4. See Theorem 4.1 for the strongest version. Besides mapping class groups, there are several other classes of spaces and groups that are colorably hierarchically hyperbolic, including

- many cubical groups, including all right-angled Artin and Coxeter groups; see [6] and Hagen and Susse [38];
- Teichmüller spaces with either the Teichmüller or the Weil–Petersson metric; see Durham [27], Eskin, Masur and Rafi [32], and Rafi [48];
- fundamental groups of closed 3–manifolds without Nil or Sol summands [7];
- groups resulting from various combination and small-cancellation-type theorems; see Behrstock, Hagen and Sisto [5; 7], Berlai and Robbio [11], and Robbio and Spriano [51; 56];
- quotients of mapping class groups by suitable large powers of Dehn twists, and other related quotients; see Behrstock, Hagen, Martin and Sisto [4];
- extensions of lattice Veech subgroups of mapping class groups; see Dowdall, Durham, Leininger and Sisto [25; 26];
- extensions of multicurve stabilizer subgroups of mapping class groups; see Russell [52];

- the genus-2 handlebody group; see Chesser [23];
- Artin groups of extra large type; see Hagen, Martin and Sisto [37, Remark 6.16].

With the exception of any hyperbolic and cubical examples from above, our main results and its applications are novel for this wide class of objects.

1.1 Applications

We now discuss our two main applications of Theorem A, namely that mapping class groups and Teichmüller spaces are bicomvable (Corollary D) and admit stable barycenters (Corollary F).

Bicomblings and semihyperbolicity In CAT(0) spaces, geodesics are unique. In geodesic Gromov hyperbolic spaces, all geodesics between any pair of points fellow-travel. In fact, in both of these classes of spaces geodesics are stable under perturbation of their endpoints in the following sense: given points x, x', y and y' with $d(x, y), d(x', y') \leq 1$, all geodesics from x to y fellow-travel those from x' to y' .

The notion of a *bicombling* of a metric space X , introduced by Thurston, generalizes this stability property. Roughly speaking, a bicombling is a transitive family of uniform quasigeodesics with the above parametrized fellow-traveling property under perturbation of endpoints. See Section 6.2 for a precise definition.

Bicomvability is a quasi-isometry invariant which imposes strong constraints on groups, such as property FP_∞ , a quadratic isoperimetric inequality, and the Novikov conjecture; see Alonso and Bridson [1], Baumslag, Gersten, Shapiro and Short [2], Epstein, Cannon, Holt, Levy, Paterson and Thurston [31], Gersten and Short [33], and Storm [57]. Moreover, bicomblings are the key geometric feature of biautomatic structures on groups (where one requires that the bicombling is constructible by a finite state automaton), thereby playing an important role in computational group theory. It is worth noting that bicomvability is decidedly a feature of nonpositive curvature, with the 3-dimensional Heisenberg group not being bicomvable because it does not satisfy a quadratic isoperimetric inequality [31].

The power of our stable cubical models is that they allow us to stably and hierarchically import geometric features of cube complexes into HHSs. In particular, ℓ^1 -geodesics in the cubical models map to *hierarchy paths* (Definition 6.5), which are quasigeodesics that are finely attuned to the HHS structure, in that they project to

uniform, unparametrized quasigeodesics in every hyperbolic space in the hierarchical structure. The stability property of the cubulation then implies that carefully chosen ℓ^1 -geodesics give a bicombing:

Theorem B *Any colorable HHS $(\mathcal{X}, \mathfrak{S})$ admits a coarsely $\text{Aut}(\mathcal{X}, \mathfrak{S})$ -equivariant, discrete, bounded, quasigeodesic bicombing by hierarchy paths with uniform constants.*

If the action by automorphisms is free, then coarse equivariance can be upgraded to equivariance. By the definition of semihyperbolicity [1], we obtain:

Corollary C *Colorable hierarchically hyperbolic groups are semihyperbolic.*

Note that semihyperbolicity has several novel consequences for HHGs. Besides novel consequences of bicombability, these include solvability of the conjugacy problem and the fact that abelian subgroups are undistorted [1].

While many HHSs were known to be bicombable for other reasons, eg many are $\text{CAT}(0)$, this produces bicombings for many new examples, such as extensions of Veech subgroups of mapping class groups.

Our main application is:

Corollary D *For any finite type surface Σ , its mapping class group $\text{MCG}(\Sigma)$ is semihyperbolic and its Teichmüller space $\text{Teich}(\Sigma)$ with either the Teichmüller metric or the Weil–Petersson metric is coarsely $\text{MCG}(\Sigma)$ -equivariantly bicombable by hierarchy paths with uniform constants.*

Note that the HHS notion of hierarchy path that we are using here is more general than the hierarchy paths produced in [27; 43], which are explicitly constructed from hierarchies of tight geodesics in curve graphs.

We remark that semihyperbolicity of $\text{MCG}(\Sigma)$ follows from work in a preprint of Hamenstädt [40]. The result for $\mathcal{T}(\Sigma)$ is new, though we were informed by M Kapovich and K Rafi that they know of a different construction for bicombing $\mathcal{T}(\Sigma)$. Note that $\mathcal{T}(\Sigma)$ with the Weil–Petersson metric is bicombable since its completion is $\text{CAT}(0)$ —see Bridson and Haefliger [18], Tromba [59], and Wolpert [60]—though we note that it is unknown whether Weil–Petersson geodesics are hierarchy paths. Combability of $\text{MCG}(\Sigma)$ follows from work of Mosher [45].

Notably, our bicombing construction applies to both mapping class groups and Teichmüller spaces simultaneously. Moreover, our bicombings are relatively straightforward applications of our more powerful stable cubulation construction. See Section 1.4 for a discussion.

Stable barycenters Another key feature of nonpositively curved spaces is that bounded sets admit (coarse) *barycenters*. Here, we think of barycenters simply as maps assigning a point to any finite subset. Some more properties are required to make this notion meaningful, such as *stability*, which requires the barycenter to vary a bounded amount when the finite set varies a bounded amount, and *coarse equivariance* when a group action is present; see Section 6.1.

In CAT(0) spaces there are a number of useful notions of barycenter which are equivariant and stable, for example center-of-mass constructions and circumcenters. Coarse barycenters are useful in the context of groups for understanding centralizers and solving the conjugacy problem for torsion elements and subgroups. Notably, Gromov hyperbolic spaces admit (coarse) barycenters: a coarse barycenter of a finite set F in a hyperbolic space X can be taken to be one of the standard CAT(0) barycenters in the CAT(0) space which models the hull of F in X , ie a simplicial tree. See Section 1.2 for a discussion of these ideas in the context of the this paper.

We should mention that coarse barycenters for triples of points are used to define coarse medians in the sense of Bowditch [15]; thus playing a central role in the theory of coarse median spaces and its many applications. However it is unclear how to construct barycenters even for pairs of points in a coarse median space, and stability properties appear just as difficult to obtain.

Barycenters in CAT(0) spaces are not in general well behaved under quasi-isometries. Using Theorem A and a construction reminiscent of Niblo and Reeves' normal paths [46], we are able to prove that most HHSs admit equivariant coarse barycenters, which are coarsely invariant under HHS automorphisms:

Theorem E *Let $(\mathcal{X}, \mathfrak{S})$ be a colorable HHS. Then \mathcal{X} admits coarsely $\text{Aut}(\mathcal{X}, \mathfrak{S})$ -equivariant stable barycenters for k points, for any $k \geq 1$.*

We remark that the coarse barycenter we produce for a set F is contained in the hull of F .

As with Theorem B, Theorem E can be applied to mapping class groups and Teichmüller spaces:

Corollary F For any finite type surface Σ , its mapping class group $\text{MCG}(\Sigma)$ and Teichmüller space $\mathcal{T}(\Sigma)$ admit coarsely $\text{MCG}(\Sigma)$ -equivariant stable barycenters for k points, for any $k \geq 1$.

Corollary F is new for arbitrary finite sets of points in $\text{MCG}(\Sigma)$ and $\mathcal{T}(\Sigma)$ with the Teichmüller metric, even without the stability property. The corresponding statement for $\mathcal{T}(\Sigma)$ with the Weil–Petersson metric is an easy consequence of the fact that its completion is $\text{CAT}(0)$. Corollary F, without the stability property, was proven for triples of points in $\text{MCG}(\Sigma)$ by Behrstock and Minsky [10], for orbits of finite order elements of $\text{MCG}(\Sigma)$ in $\text{MCG}(\Sigma)$ by Tao [58], and more generally for orbits of finite subgroups of $\text{MCG}(\Sigma)$ in both $\text{MCG}(\Sigma)$ and $\mathcal{T}(\Sigma)$ with the Teichmüller metric by Durham [28].

In work that appeared simultaneously to ours, Haettel, Hoda and Petyt [35] proved that HHSs are coarse Helly spaces, in the sense of Chalopin, Chepoi, Genevois, Hirai and Osajda [20]. This property has a number of strong consequences, many of which overlap with the results in this paper. In particular, they obtain versions of Theorems B and E along with their corollaries, without the colorability assumption and the hierarchy path conclusion.

Their approach and constructions are very different from ours, using results from the theory of coarse Helly and injective metric spaces, whereas our work relies mostly on hyperbolic and cubical geometry.

1.2 Coarse hulls and their models

Given the technical nature of many of the proofs in this paper, we include here an extended but simplified discussion of the ideas that go into our constructions. The propositions stated in this section will not, however, be used elsewhere in the paper.

Consider first the notion of a *convex hull* in a $\text{CAT}(0)$ space. The convex hull of a finite set F has the following nice property: the map $F \mapsto \text{hull}(F)$ is 1-Lipschitz with respect to the Hausdorff metric on sets. We are interested in generalizing this notion to more coarse hulls (which we will just denote by $\text{hull}(F)$ in each case) in more general spaces.

As a first motivating example, consider the Euclidean plane, $X = \mathbb{R}^2$ with the ℓ^2 metric. The convex hull of two points, $\text{hull}(\{x, y\})$, is just the unique geodesic between them. If, on the other hand, we endow \mathbb{R}^2 with the ℓ^1 metric then the convex hull is the axis-parallel rectangle spanned by x and y . Note that (\mathbb{R}^2, ℓ^1) is not $\text{CAT}(0)$ but is a

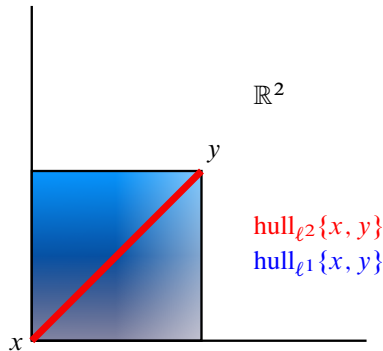


Figure 1: A cartoon of the ℓ^2 -hull (red) and ℓ^1 -hull (blue) of two points in \mathbb{R}^2 . The ℓ^1 -hull reflects the intrinsic product structure of the space.

product of CAT(0) spaces, and this hull is a product of hulls in the CAT(0) factors. See Figure 1. This simple idea is a model for a useful construction in the HHS context.

We can think of an HHS as (coarsely) embedded in a product of hyperbolic spaces, in such a way that it is composed of products of certain factors, intersecting and nesting in a complicated fashion. The reader familiar with the foundational example, namely Masur and Minsky's hierarchy of curve graphs for mapping class groups [42; 43], will lose nothing by keeping it in mind during the ensuing discussion. In that setting, Behrstock, Kleiner, Minsky and Mosher [9] introduced a notion of hull which is essentially a coarse pullback of convex hulls in each hyperbolic factor; see Section 2.2. Behrstock, Hagen and Sisto [8] proved, in the general HHS setting, that these hulls are quasi-isometrically modeled by finite CAT(0) cubical complexes.

Their result is a partial generalization of the situation in Gromov hyperbolic spaces, where Gromov proved that hulls of finite sets of points are quasi-isometrically modeled by finite simplicial trees [34]. However, in the setting of hyperbolic spaces, the modeling trees satisfy additional strong stability properties under perturbation of the set of input points; see Proposition 1.3 below.

Our main theorem — in increasing specificity, Theorems A, 1.4 and 4.1 — endows the modeling cube complexes from [8] with a generalization of the stability properties that Gromov's modeling trees enjoy.

Before giving a full account of our results and an overview of their proofs, it will be beneficial to discuss the situation in hyperbolic spaces and cubical complexes. We will see that our results are a common generalization of the situations from these motivating examples.




Figure 2: Stability of hulls in a tree: the intersection of $\text{hull}(F)$ (red) and $\text{hull}(F')$ (blue) subtrees is a (green) subtree which can be obtained by deleting the boundedly many complementary red and blue subtrees.

Hulls in trees and cube complexes Let X be a simplicial tree. Then the convex hull of any finite set of vertices $F \subset X^0$ is the subtree T_F of X spanned by F . Moreover, the subtree T_F is stable under small perturbations of F , in the following sense; see Figure 2:

Proposition 1.1 *Let X be a simplicial tree. If $F, F' \subset X^0$ satisfy $\#F' = \#F = k$ and $d_{\text{Haus}}(F, F') \leq 1$, then the intersection of their hulls, $T_0 = T_F \cap T_{F'}$, is itself a subtree with both $T_F \setminus T_0$ and $T_{F'} \setminus T_0$ a union of at most k subtrees each of diameter at most 1.*

We will not use this fact, so we leave its proof to the interested reader.

This situation generalizes to when X is a CAT(0) cube complex endowed with the ℓ^1 metric; see Section 2.1 for the relevant definitions. Recall that the ℓ^1 metric on X is completely determined by a special collection \mathcal{H}_X of codimension-1 subspaces called *hyperplanes* (Section 2.1), in the sense that X is precisely the dual cube complex arising from Sageev’s cubulation construction [54] applied to \mathcal{H}_X as a wallspace; see Section 2.1.1.

In the cubical context, the ℓ^1 convex hull of any finite set of vertices $F \subset X^0$ is the cubical subcomplex $\mathcal{Q}_F \subset X$ realized as the dual to the hyperplanes \mathcal{H}_F separating the points in F . In addition, these cubical hulls satisfy the following strong stability property; see Figure 3:

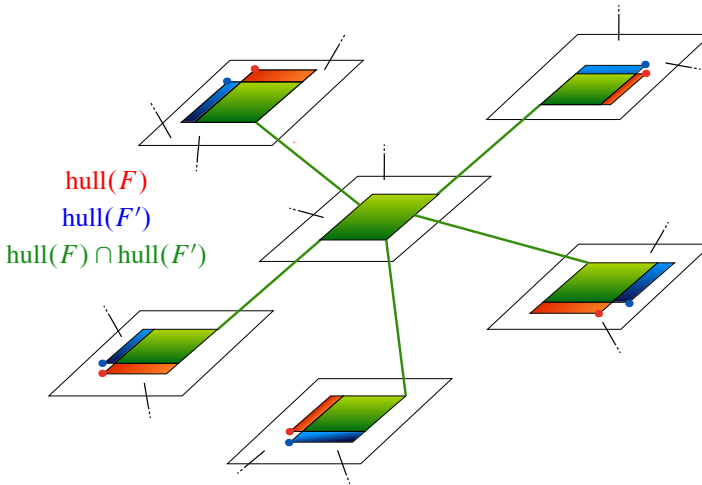


Figure 3: Stability of hulls in the universal cover of $S^1 \vee T^2$: the subcomplex dual to all hyperplanes common to both $\text{hull}(F)$ (red) and $\text{hull}(F')$ (blue) is here realized as the intersection (green) of the ℓ^1 -hulls.

Proposition 1.2 *Let X be a CAT(0) cube complex endowed with the ℓ^1 metric. If $F, F' \subset X^0$ satisfy $\#F, \#F' \leq k$ and $d_{\text{Haus}}(F, F') \leq 1$, then there are convex subcomplexes $X_F \subseteq \mathcal{Q}_F$ and $X_{F'} \subseteq \mathcal{Q}_{F'}$, both dual to the hyperplanes in $\mathcal{H}_0 = \mathcal{H}_F \cap \mathcal{H}_{F'}$, such that $d_{\text{Haus}}(X_F, X_{F'}) \leq 1$. Moreover, both $\mathcal{H}_F \setminus \mathcal{H}_0$ and $\mathcal{H}_{F'} \setminus \mathcal{H}_0$ contain at most k hyperplanes.*

Again, we will not use this proposition, so we omit its proof.

In the cubical structure on a simplicial tree, the hyperplanes correspond to midpoints of edges. Hence Proposition 1.2 generalizes Proposition 1.1. Note that now the diameters of $\mathcal{Q}_F \setminus X_F$ and $\mathcal{Q}_{F'} \setminus X_{F'}$ can be arbitrarily large. However, since the ℓ^1 metric on X is completely determined by its defining hyperplanes, Proposition 1.2 says that \mathcal{Q}_F and $\mathcal{Q}_{F'}$ are metrically and combinatorially related, depending only on k and X — and not on $\text{diam}(F)$. In particular, one can delete boundedly many hyperplanes from the collections \mathcal{H}_F and $\mathcal{H}_{F'}$ to generate a common model; see Section 2.1.2 for a discussion on hyperplane deletions.

Modeling hulls in hyperbolic spaces In coarse geometry, eg when X is the Cayley graph of a finitely generated group, the notion of geodesic is often wobbly, and so our notion of hull needs to be more flexible. Moreover, it will often be more fruitful to construct quasi-isometric *models* of hulls, which we should think of as nice combi-




Figure 4: Stability of hulls in a hyperbolic space: top, the modeling tree (red) for the coarse hull (green) of a finite set F ; bottom, the modeling trees T_F and $T_{F'}$ for $\text{hull}(F)$ and $\text{hull}(F')$ are $(1, K)$ -quasi-isometric after deleting small subtrees (purple).

natural objects which coarsely encode the key geometric features of hulls into their combinatorial structure. The main motivating examples here are hyperbolic spaces, where hulls are modeled by finite simplicial trees.

When X is δ -hyperbolic and $F \subset X$ with $\#F = k$, the right notion of $\text{hull}(F)$ is the *weak hull*, namely the set of all geodesics between points in F . Notice then that the tripod-like δ -slim-triangles condition generalizes to a tree-like slimness for $\text{hull}(F)$. The following proposition is an easy consequence of Gromov’s original arguments [34]; see Figure 4:

Proposition 1.3 *For any $k \in \mathbb{N}$ and $\delta > 0$, there exists $L = L(k, \delta) > 0$ such that the following holds.*

Let X be δ -hyperbolic and $F \subset X$ with $\#F = k$. Then there is a simplicial tree T_F and a $(1, L)$ -quasi-isometric embedding $\phi_F : T_F \rightarrow X$ with $d_{\text{Haus}}(\phi_F(T_F), \text{hull}(F)) < L$. Moreover, if $F' \subset X$ with $\#F' \leq k$ and $d_{\text{Haus}}(F, F') \leq 1$, then there exists a simplicial tree T_0 and a $(1, L)$ -quasi-isometric embedding $\phi_0 : T_0 \rightarrow X$ such that the diagram

$$(1-1) \quad \begin{array}{ccc} T_F & & \\ h_F \downarrow & \searrow \phi_F & \\ T_0 & \xrightarrow{\phi_0} & X \\ h_{F'} \uparrow & \nearrow \phi_{F'} & \\ T_{F'} & & \end{array}$$

commutes up to error at most L , where h_F and $h_{F'}$ are quotient maps which collapse at most L subtrees each of diameter at most L .

Observe that Proposition 1.3 is a generalization of Proposition 1.1, where X is a tree and we can take the trees T_F , $T_{F'}$, and T_0 as before and the maps ϕ_F , $\phi_{F'}$, and ϕ_0 to be inclusions. The main difference here is that a general hyperbolic space is stably *locally* tree-like, and not a tree itself. Hence the need for a model for the hulls.

Our main theorem is a common generalization of the stability properties in Propositions 1.2 and 1.3.

1.3 Stable cubical models for hulls in HHSs

We will deal with *colorable* hierarchically hyperbolic spaces $(\mathcal{X}, \mathfrak{S})$, which means, for the reader familiar with HHSs, that there exists a decomposition of \mathfrak{S} into finitely many families \mathfrak{S}_i such that each \mathfrak{S}_i is pairwise transverse. Colorable HHSs include mapping class groups and Teichmüller spaces of finite-type surfaces.

In fact, colorability is a rather mild condition which is satisfied by all of the main motivating examples. Its definition is inspired by Bestvina, Bromberg and Fujiwara's proof that curve graphs are finitely colorable [12]; see Section 2.2 for a discussion.

Given a finite set of points $F \subset \mathcal{X}$ in an HHS, the standard notions of a hull for F are very difficult to analyze. For example, while little is known about geodesics in the mapping class group, Rafi and Verberne [50] proved that geodesics do not always interact well with the curve graph machinery. In Teichmüller space with the Teichmüller metric, geodesics are unique, but it is an open question of Masur whether the classical convex hull of a set of three points can be the whole space. Moreover, it is a result of Rafi that hulls of two points, ie geodesics, do not behave stably under perturbation [49, Theorem D]. These complications motivate a more flexible definition of hull in this setting.

The *hierarchical hull* of a finite set $F \subset \mathcal{X}$, which we also denote by $\text{hull}(F)$, was introduced in [9] to study subspaces of the asymptotic cones of the mapping class group, on the way to proving that these groups are quasi-isometrically rigid. In hyperbolic spaces and cube complexes, the hierarchical hull coincides with the notions of hull discussed above. In the hierarchical setting, one instead has a notion of projecting F to a family of hyperbolic spaces (eg curve graphs of subsurfaces). In each of these hyperbolic spaces, one then takes the weak hull of the projection — which is coarsely

a tree, as above — and the uses certain hierarchical consistency conditions [7; 9] to fashion these weak hulls in the various spaces into a hull in the ambient HHS which satisfies certain convexity properties [7; 9]. In particular, the hierarchical hull of F is *hierarchically quasiconvex* [7] and contains all of the hierarchy paths between points in F [9].

In [8], Behrstock, Hagen and Sisto proved that the hierarchical hull of a finite set of points is quasi-isometric to a finite CAT(0) cube complex. Their main observation was that the hierarchical consistency conditions are closely related to the consistency conditions on a wallspace from Sageev’s construction of cubical complexes (Section 2.1.1). Their idea was to look at points on the modeling trees in the hyperbolic spaces which are unseen by the other projection data. The preimages of these points under the projection maps turn out to behave like walls in the hull. See Section 1.4 for a sketch of these ideas, and Section 4.2 for a full discussion.

Our main theorem stabilizes their construction, simultaneously generalizing the stability properties from Proposition 1.3 for any hyperbolic space and Proposition 1.2 for cube complexes admitting an HHS structure. The following is a more detailed version of Theorem A.

Theorem 1.4 *Let $(\mathcal{X}, \mathfrak{S})$ be a colorable HHS. Then for each k there exist K and N with the following properties. For any $F \subset \mathcal{X}$ with $\#F \leq k$, there exists a finite CAT(0) cube complex \mathcal{Q}_F and a K -quasimediant, (K, K) -quasi-isometric embedding $\Phi_F : \mathcal{Q}_F \rightarrow \mathcal{X}$ with $d_{\text{Haus}}(\Phi_F(\mathcal{Q}_F), \text{hull}(F)) \leq K$.*

Moreover, if $F' \subseteq \mathcal{X}$ is another subset with $\#F' \leq k$ and $d_{\text{Haus}}(F, F') \leq 1$, there is a finite CAT(0) cube complex \mathcal{Q}_0 and a K -quasimediant, (K, K) -quasi-isometric embedding $\Phi_0 : \mathcal{Q}_0 \rightarrow \mathcal{X}$ such that the diagram

$$(1-2) \quad \begin{array}{ccc} & \mathcal{Q}_F & \\ \eta_F \downarrow & \searrow \Phi_F & \\ & \mathcal{Q}_0 & \xrightarrow{\Phi_0} \mathcal{X} \\ \eta_{F'} \uparrow & \nearrow \Phi_{F'} & \\ & \mathcal{Q}_{F'} & \end{array}$$

commutes up to error at most K , where η_F and $\eta_{F'}$ are hyperplane deletion maps which delete at most N hyperplanes.

See Theorem 4.1 for the full version of the theorem, the details of which are necessary for our applications.

We note that CAT(0) cubical complexes are *median* spaces and HHSs are *coarse median* [7] in the sense of Bowditch [15]. As with the cubical models in [8], our stable cubical models also coarsely preserve the medians, meaning that the maps Φ_F do (as stated in Theorem 4.1). It is worth noting that in view of Russell, Spriano and Tran [53, Corollary 5.12], our cube complexes also approximate coarse median hulls.

1.4 Sketch of proofs

The proof of Theorem 4.1, of which Theorem A is an informal version, is contained in Section 4 and depends crucially on our work in Section 3. Theorems B and E are a consequence of Theorem A and our work in Section 5. We now explain the various parts and how they fit together.

In what follows, we will keep our discussion within the context of mapping class groups and hierarchies of curve graphs [42; 43], though we work in the more general context of HHSs.

Let $F \subset \text{MCG}(\Sigma)$ be a finite subset and consider essential subsurfaces $V \subset \Sigma$ which are not 3–holed spheres. Roughly, the *hierarchical hull* of F , $\text{hull}(F)$, is the set of points of $\text{MCG}(\Sigma)$ whose subsurface projections in each curve graph $\mathcal{C}(V)$ lie close to the weak hull of the subsurface projection $\pi_V(F)$ of F .

In the cubulation construction of [8], the authors build a wallspace for $\text{hull}(F)$.

To do this, they first consider the collection \mathcal{U}_F of *relevant* subsurfaces $V \subset \Sigma$ for which $\text{diam}_V \pi_V(F) > K$ for some fixed threshold $K > 0$. In each of these subsurfaces, they take a tree T_F^V which coarsely models the hull of $\pi_V(F)$ in $\mathcal{C}(V)$, as discussed in Section 1.2. For each such $V \in \mathcal{U}_F$, they then consider the collection of relative projections ρ_V^W of $W \in \mathcal{U}_F$ to $\mathcal{C}(V)$, which correspond to the projection of ∂W to $\mathcal{C}(V)$ and thus are nonempty if V is neither disjoint from nor contained in W . The bounded geodesic image theorem [43] and certain consistency properties of projections — see Behrstock [3] and [9] — imply that each ρ_V^W for such W lies uniformly close to the tree T_F^V .

They then consider, roughly, the complement P_F^V in T_F^V of a regular neighborhood of these projections, which consists of a number of subtrees of T_F^V which are “unseen” by the other subsurfaces in \mathcal{U}_F which interact with V . Any point in $T_F^V \setminus P_F^V$ cuts T_F^V into two subtrees. The partitions of $\text{hull}(F)$ that define the wallspace on $\text{hull}(F)$ come from these subdivision points in the T_F^V , namely one considers the subspaces of $\text{hull}(F)$ whose subsurface projections to $\mathcal{C}(V)$ lie close to either of the subtrees.

While this construction is useful for studying top-dimensional quasiflats, it is unstable under perturbation of F , in that given some other F' with $d_{\text{Haus}}(F, F') \leq 1$, then the cubical models \mathcal{Q}_F and $\mathcal{Q}_{F'}$ might differ by a number of hyperplanes on the order of $\text{diam}(F)$, which is not bounded.

The proof of Theorem A involves stabilizing this process in a number of places. The first step is to robustly stabilize the collection of relevant subsurfaces \mathcal{U}_F (Proposition 2.14), eg so that $|\mathcal{U}_F \Delta \mathcal{U}_{F'}|$ is bounded in terms of the topology of the surface S . We do this by applying work of Bestvina, Bromberg, Fujiwara and Sisto [14], which allows us to stabilize subsurface projections (Theorem 2.9), and then use standard projection complex type arguments.

In Section 3, we stabilize the modeling trees T_F^V for each $V \in \mathcal{U}_F$. Unlike before, it will not do to simply take any Gromov modeling tree, since unboundedly many pieces of it might change in the transition from F to F' when we cut it up using the relative projection data (the P_F^V above). Instead, we use the newly stabilized relative projection data to build a new stable tree. We do this by taking a regular neighborhood of the relative projections, which then group into connected components we call *clusters*. As before, these clusters lie close to any Gromov modeling tree, but we cannot use these trees. Instead, we define a separation graph for these clusters (Definition 3.5), and then prove that the combinatorics of this graph encode how these domain clusters are arranged on *any* Gromov modeling tree. We then construct our stable tree by connecting clusters both internally and externally via minimal spanning networks in $\mathcal{C}(V)$. The stability of the cluster data then is converted into stability of the tree construction in Theorem 3.3, which, in particular, says that the set of long edges of two related trees are in bijection and within bounded Hausdorff distance, with most long edges exactly the same. See Figures 5 and 9 below.

In Section 4, we then plug these stable trees into the cubulation machine from [8]. We must be mindful of how subdivision points change when transitioning from F to F' . In particular, we construct a common refinement of the sets of subdivision points for our two sets F and F' (Proposition 4.12), with the delicate nature of this process necessitating the intricacies in the statement and proof of the stable tree theorem (Theorem 3.3). With this in hand, we prove that this common refinement induces an isomorphism between the resulting cubical models for the hulls of both sets (Proposition 4.13); see Figure 18. This isomorphism depends on a careful hierarchical analysis of when two halfspaces corresponding to two subdivision points intersect (Lemma 4.10). The full

version of the stable cubulation theorem is achieved in Theorem 4.1, which says that the two modeling cube complexes \mathcal{Q}_F and $\mathcal{Q}_{F'}$ become isomorphic when we delete a bounded number of hyperplanes from each, with the bound depending only on $|F|$ and $|F'|$.

In Section 5, we adapt the normal path construction of Niblo and Reeves [46] and analyze how it changes under hyperplane deletion. In particular, for any finite CAT(0) cube complex \mathcal{Q} , we develop a sequence of contractions which take the extremal vertices of \mathcal{Q} (ie its corners) into a “barycentric” cube at the “center” of \mathcal{Q} , and we prove that this contraction sequence is only boundedly perturbed by hyperplane deletions (Theorem 5.1).

Stability of the cubical model and the contraction sequence easily give the barycenter theorem (Theorem E). In the context of a bicombling (Theorem B) when $F = \{x, y\}$, we take the bicombling path from x to y to be the image in $\text{MCG}(\Sigma)$ of the path obtained by following the contraction sequence of x to the barycentric cube, and then traversing the contraction sequence from the barycentric cube to y in reverse order. Once again, stability of the contraction sequence and the cubical models implies that these are uniform quasigeodesics which fellow-travel in a parametrized fashion; see Figure 25. Theorems B and E are proved in Section 6.

1.5 Outline

In Section 2 we collect some background material.

Section 3 takes place entirely in a fixed hyperbolic space, using methods from coarse hyperbolic geometry but with HHS ends in mind. The main result there is Theorem 3.3, and no other result from that section will be used elsewhere.

In Section 4, we prove the precise version of Theorem A, which is Theorem 4.1. Again, no other statement from this section will be used elsewhere. In this section, we use the combinatorial geometry of HHSs.

Section 5 uses the tools of cubical geometry, and it is independent from the previous sections. Its main result is Theorem 5.1, which once again is the only result from here needed in the rest of the paper.

Finally, in Section 6 we put all the pieces together, and we prove Theorems B and E.

Acknowledgements

Durham would like to thank Daniel Groves for many fruitful and insightful discussions about bicombed the mapping class group. We would like to thank the referee for useful comments and for suggesting Lemma 2.15. Durham was partially supported by an AMS Simons travel grant and NSF grant DMS-1906487. Minsky was partially supported by NSF grants DMS-1610827 and DMS-2005328. Sisto was partially supported by the Swiss National Science Foundation (grant 182186).

2 Background

In this section, we will collect and record the various facts about cube complexes and hierarchically hyperbolic spaces that we need.

2.1 CAT(0) cube complexes

We will briefly discuss some basic aspects of CAT(0) cubical geometry. We direct the reader to Sageev's lecture notes [55] for details.

A *cube complex* is a simplicial complex X obtained from a disjoint collection of Euclidean cubes which are glued along their faces by a collection of Euclidean isometries. A cube complex is *nonpositively curved* (NPC) if its vertex links are simplicial flag complexes. An NPC cube complex is *CAT(0)* if it is a 1-connected NPC complex.

A *midcube* of an n -cube $C \subset X$ is an $(n-1)$ -dimensional cube $H' \subset C$ running through the barycenter of C and parallel to one of the faces of C . A *hyperplane* $H \subset X$ is a connected subspace of X such that for all closed cubes C , the intersection $H \cap C$ is either empty or a midcube of C . The *carrier* of H is the union of all of the cubes in X whose intersection with H is a midcube, and it is naturally isomorphic to $H \times [0, 1]$.

Equivalently, there is a natural equivalence relation on the set of edges in the 1-skeleton of X generated by relating two edges if they are opposite edges of some square in X . Any hyperplane can be obtained as the collection of midcubes which intersect the edges in a given equivalence class.

In this paper, we will be considering *finite* cube complexes, namely those with finitely many cubes.

Metrics on cube complexes There are many interesting metrics one can put on a CAT(0) cube complex X . We will be interested in both

- the ℓ^1 or *combinatorial* metric, d_1 , which is generated by the ℓ^1 norm in each cube of X , and can be equivalently defined on the 1-skeleton $X^{(1)}$ as the path metric thereon;
- the *cubical sup metric*, d_∞ , which is the metric generated by the ℓ^∞ or sup norm in each cube in X .

The following is an easy consequence of the observation that, given n , the ℓ^1 and ℓ^∞ norm on an n -cube are bi-Lipschitz equivalent.

Lemma 2.1 *For any $n > 0$, there exists $K = K(n) > 0$ such that if X is an n -dimensional cube complex, then the identity $\text{id}: (X, d_1) \rightarrow (X, d_\infty)$ is a (K, K) -quasi-isometry.*

The differences between these metrics will come up in Sections 5 and 6. See [44] for a detailed discussion of these metrics.

2.1.1 Wallspaces and Sageev’s construction In Section 4, we will adopt the perspective of obtaining cube complexes as duals to wallspaces. Wallspaces were first defined by Haglund and Paulin [39]; see Hruska and Wise [41] for a broader discussion.

Let Y be a nonempty set. A *wall* in Y is a pair of subsets $W = \{\overleftarrow{W}, \overrightarrow{W}\}$ where $Y = \overleftarrow{W} \sqcup \overrightarrow{W}$. In this case, \overleftarrow{W} and \overrightarrow{W} are called *halfspaces*.

Two points $x, y \in Y$ are *separated* by a wall W if x is contained in a different halfspace from y .

A *wallspace* is a set Y with a collection of walls \mathcal{W} on Y such that the number of walls separating any pair of points is finite.

An *orientation* on a wallspace (Y, \mathcal{W}) is an assignment σ such that, for each $W \in \mathcal{W}$, we have $\sigma(W) \in \{\overleftarrow{W}, \overrightarrow{W}\}$. The orientation σ is called *coherent* if, for all $W, W' \in \mathcal{W}$, we have $\sigma(W) \cap \sigma(W') \neq \emptyset$. We call σ *canonical* if there exists $x \in X$ such that $x \in \sigma(W)$ for all but finitely many $W \in \mathcal{W}$.

Given a wallspace (Y, \mathcal{W}) , we can consider the cube complex $X(Y, \mathcal{W})$ constructed as follows. The 0-cubes of $X(Y, \mathcal{W})$ are coherent, canonical orientations of (Y, \mathcal{W}) . Two 0-simplices are connected by a 1-cube if, seen as orientations, they differ on only one

wall. Finally, all subcomplexes of the 1–skeleton isomorphic to the 1–skeleton of an n –cube get filled by an n –cube.

Work of Chatterji and Niblo [21], Chepoi [22], and Nica [47] — building off of work of Sageev [54] — gives that $X(Y, \mathcal{W})$ is a CAT(0) cube complex. We call $X(Y, \mathcal{W})$ the *dual cube complex* associated to the wallspace (Y, \mathcal{W}) .

2.1.2 Hyperplane deletions In Section 5, we will be interested in understanding how cube complexes change under deletions of hyperplanes, so we will use the alternative perspective of obtaining cube complexes from sets of hyperplanes. We briefly explain how this works.

Let X be a CAT(0) cube complex and \mathcal{H}_X its (finite) set of hyperplanes. Then we can identify each hyperplane $H \in \mathcal{H}_X$ with the two halfspaces into which it partitions $X^{(0)}$. As such, we can and will think of $(X^{(0)}, \mathcal{H}_X)$ as a wallspace, and one can show that X is the dual cube complex associated to (X, \mathcal{H}_X) .

Given any subset $\mathcal{H} \subset \mathcal{H}_X$ of hyperplanes in a cube complex X , there is a natural cube complex $X(\mathcal{H})$ defined as the dual cube complex associated to the wallspace defined by \mathcal{H} in X . In particular, each point in $X(\mathcal{H})$ is a choice of coherent, canonical orientations of the half-spaces defined by \mathcal{H} .

With this notation, we can now define hyperplane deletions:

Definition 2.2 Let X be a CAT(0) cube complex obtained with hyperplanes \mathcal{H}_X . For a finite collection of hyperplanes $\mathcal{G} \subset \mathcal{H}_X$, the *hyperplane deletion map* for H is the map

$$\text{Res}_{\mathcal{H}_X \setminus \mathcal{G}}: X \rightarrow X(\mathcal{H}_X \setminus \mathcal{G})$$

obtained by restriction of orientations, where $X(\mathcal{H}_X \setminus \mathcal{G})$ is the dual cube complex associated to the wallspace $(X, \mathcal{H}_X \setminus \mathcal{G})$.

Equivalently, the map $\text{Res}_{\mathcal{H}_X \setminus \mathcal{G}}$ is the quotient map which collapses the $[0, 1]$ factor of each of the carriers of the hyperplanes in \mathcal{G} (recall that the carrier of the hyperplane H is naturally isomorphic to $H \times [0, 1]$). We note that these maps appear elsewhere in the literature, eg [19].

We also record the following fact, which indicates that the isomorphism type of the cube complex coming from a wallspace is determined by the intersection pattern of halfspaces. The proof is elementary.

Lemma 2.3 *Let \mathcal{W} and \mathcal{W}' be wallspaces, and let $\iota: \mathcal{H}_{\mathcal{W}} \rightarrow \mathcal{H}_{\mathcal{W}'}$ be a bijection of their halfspaces which preserves complements and disjointness.*

Denote by j the induced map on walls

$$\{H, H^c\} \mapsto \{\iota(H), \iota(H)^c\}$$

and by $h(x) = \iota \circ x \circ j^{-1}$ the induced map on orientations. Then h , viewed as a map on 0-cubes, induces an isomorphism $h: \mathcal{Y}_{\mathcal{W}} \rightarrow \mathcal{Y}_{\mathcal{W}'}$ between the corresponding CAT(0) cube complexes.

2.2 HHS axioms

We recall from [7] the definition of a hierarchically hyperbolic space.

Definition 2.4 (HHS) The q -quasigeodesic space $(\mathcal{X}, \text{dist}_{\mathcal{X}})$ is a *hierarchically hyperbolic space* if there exists $\delta \geq 0$, an index set \mathfrak{S} , and a set $\{\mathcal{C}W \mid W \in \mathfrak{S}\}$ of δ -hyperbolic spaces $(\mathcal{C}U, \text{dist}_U)$, such that the following conditions are satisfied.

(1) **Projections** There is a set $\{\pi_W: \mathcal{X} \rightarrow 2^{\mathcal{C}W} \mid W \in \mathfrak{S}\}$ of *projections* sending points in \mathcal{X} to sets of diameter bounded by some $\xi \geq 0$ in the various $\mathcal{C}W \in \mathfrak{S}$. Moreover, there exists K such that, for all $W \in \mathfrak{S}$, the coarse map π_W is (K, K) -coarsely Lipschitz and $\pi_W(\mathcal{X})$ is K -quasiconvex in $\mathcal{C}W$.

(2) **Nesting** \mathfrak{S} is equipped with a partial order \sqsubset , and either $\mathfrak{S} = \emptyset$ or \mathfrak{S} contains a unique \sqsubset -maximal element; when $V \sqsubset W$, we say V is *nested* in W . (We emphasize that $W \sqsubset W$ for all $W \in \mathfrak{S}$.) For each $W \in \mathfrak{S}$, we denote by \mathfrak{S}_W the set of $V \in \mathfrak{S}$ such that $V \sqsubset W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \sqsubset W$, there is a specified subset $\rho_W^V \subset \mathcal{C}W$ with $\text{diam}_{\mathcal{C}W}(\rho_W^V) \leq \xi$. There is also a *projection* $\rho_V^W: \mathcal{C}W \rightarrow 2^{\mathcal{C}V}$. (The similarity in notation is justified by viewing ρ_W^V as a coarsely constant map $\mathcal{C}V \rightarrow 2^{\mathcal{C}W}$.)

(3) **Orthogonality** \mathfrak{S} has a symmetric and antireflexive relation called *orthogonality*, and we write $V \perp W$ when V and W are orthogonal. Also, whenever $V \sqsubset W$ and $W \perp U$, we require that $V \perp U$. We require that for each $T \in \mathfrak{S}$ and each $U \in \mathfrak{S}_T$ for which $\{V \in \mathfrak{S}_T \mid V \perp U\} \neq \emptyset$, there exists $W \in \mathfrak{S}_T \setminus \{T\}$ such that, whenever $V \perp U$ and $V \sqsubset T$, we have $V \sqsubset W$. Finally, if $V \perp W$, then V and W are not \sqsubset -comparable.

(4) **Transversality and consistency** If $V, W \in \mathfrak{S}$ are not orthogonal and neither is nested in the other, then we say V and W are *transverse*, denoted by $V \pitchfork W$. There

exists $\kappa_0 \geq 0$ such that if $V \pitchfork W$, then there are sets $\rho_W^V \subseteq CW$ and $\rho_V^W \subseteq CV$, each of diameter at most ξ and satisfying

$$(2-1) \quad \min\{\text{dist}_W(\pi_W(x), \rho_W^V), \text{dist}_V(\pi_V(x), \rho_V^W)\} \leq \kappa_0$$

for all $x \in \mathcal{X}$.

For $V, W \in \mathfrak{S}$ satisfying $V \sqsubset W$ and for all $x \in \mathcal{X}$,

$$(2-2) \quad \min\{\text{dist}_W(\pi_W(x), \rho_W^V), \text{diam}_{CV}(\pi_V(x) \cup \rho_V^W(\pi_W(x)))\} \leq \kappa_0.$$

The preceding two inequalities are the *consistency inequalities* for points in \mathcal{X} .

Finally, if $U \sqsubset V$, then $\text{dist}_W(\rho_W^U, \rho_W^V) \leq \kappa_0$ whenever $W \in \mathfrak{S}$ satisfies either $V \sqsubseteq W$ or $V \pitchfork W$ and $W \not\sqsubset U$.

(5) **Finite complexity** There exists $n \geq 0$, the *complexity* of \mathcal{X} (with respect to \mathfrak{S}), such that any set of pairwise \sqsubset -comparable elements has cardinality at most n .

(6) **Large links** There exist $\lambda \geq 1$ and $E \geq \max\{\xi, \kappa_0\}$ such that the following holds. Let $W \in \mathfrak{S}$ and let $x, x' \in \mathcal{X}$. Let $N = \lambda \text{dist}_W(\pi_W(x), \pi_W(x')) + \lambda$. Then there exists $\{T_i\}_{i=1, \dots, \lfloor N \rfloor} \subseteq \mathfrak{S}_W \setminus \{W\}$ such that for all $T \in \mathfrak{S}_W \setminus \{W\}$, either $T \in \mathfrak{S}_{T_i}$ for some i , or $\text{dist}_T(\pi_T(x), \pi_T(x')) < E$. Also, $\text{dist}_W(\pi_W(x), \rho_W^{T_i}) \leq N$ for each i .

(7) **Bounded geodesic image** There exists $\kappa_0 > 0$ such that for all $W \in \mathfrak{S}$, all $V \in \mathfrak{S}_W \setminus \{W\}$, and all geodesics γ of CW , either

$$\text{diam}_{CV}(\rho_V^W(\gamma)) \leq \kappa_0 \quad \text{or} \quad \gamma \cap \mathcal{N}_{\kappa_0}(\rho_V^W) \neq \emptyset.$$

(8) **Partial realization** There exists a constant α with the following property. Let $\{V_j\}$ be a family of pairwise orthogonal elements of \mathfrak{S} , and let $p_j \in \pi_{V_j}(\mathcal{X}) \subseteq CV_j$. Then there exists $x \in \mathcal{X}$ such that

- $\text{dist}_{V_j}(x, p_j) \leq \alpha$ for all j ;
- for each j and each $V \in \mathfrak{S}$ with $V_j \sqsubset V$, we have $\text{dist}_V(x, \rho_V^{V_j}) \leq \alpha$; and
- if $W \pitchfork V_j$ for some j , then $\text{dist}_W(x, \rho_W^{V_j}) \leq \alpha$.

(9) **Uniqueness** For each $\kappa \geq 0$, there exists $\theta_u = \theta_u(\kappa)$ such that if $x, y \in \mathcal{X}$ and $\text{dist}_{\mathcal{X}}(x, y) \geq \theta_u$, then there exists $V \in \mathfrak{S}$ such that $\text{dist}_V(x, y) \geq \kappa$.

We often refer to \mathfrak{S} , together with the nesting and orthogonality relations, and the projections as a *hierarchically hyperbolic structure* for the space \mathcal{X} .

Where it will not cause confusion, given $U \in \mathfrak{S}$, we will often suppress the projection map π_U when writing distances in $\mathcal{C}U$; ie given $x, y \in \mathcal{X}$ and $p \in \mathcal{C}U$, we write $\text{dist}_U(x, y)$ for $\text{diam}_{\mathcal{C}U}(\pi_U(x) \cup \pi_U(y))$ and $\text{dist}_U(x, p)$ for $\text{diam}_{\mathcal{C}U}(\pi_U(x) \cup \{p\})$. Given $A \subset \mathcal{X}$ and $U \in \mathfrak{S}$ we let $\pi_U(A)$ denote $\bigcup_{a \in A} \pi_U(a)$.

There is a natural notion of automorphism of an HHS, which we now briefly explain. These were originally defined in [7], but we give a more restrictive definition which is essentially equivalent, as explained in [30, Section 2.1]. An automorphism g of an HHS $(\mathcal{X}, \mathfrak{S})$ is an isometry of \mathcal{X} together with a bijection $\mathfrak{S} \rightarrow \mathfrak{S}$, also denoted by $U \mapsto gU$, which preserves nesting and orthogonality, and isometries between corresponding hyperbolic spaces, again still denoted by $g: \mathcal{C}(U) \rightarrow \mathcal{C}(gU)$. We require that $g\pi_U(x) = \pi_{gU}(gx)$ for all $x \in \mathcal{X}$ and $U \in \mathfrak{S}$, and $g\rho_V^U = \rho_{gV}^{gU}$ for all $U, V \in \mathfrak{S}$ where this is defined.

We let $\text{Aut}(\mathcal{X}, \mathfrak{S})$ denote the group of HHS automorphisms of $(\mathcal{X}, \mathfrak{S})$.

We say that a group G is a *hierarchically hyperbolic group* if it acts properly and coboundedly by HHS automorphisms on some HHS $(\mathcal{X}, \mathfrak{S})$.

2.2.1 Some useful facts We now recall results from [7] that will be useful later on.

Definition 2.5 Let $\kappa \geq 0$ and let $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{\mathcal{C}U}$ be a tuple such that for each $U \in \mathfrak{S}$, the U -coordinate b_U has diameter $\leq \kappa$. Then \vec{b} is κ -consistent if for all $V, W \in \mathfrak{S}$,

$$\min\{\text{dist}_V(b_V, \rho_V^W), \text{dist}_W(b_W, \rho_W^V)\} \leq \kappa$$

whenever $V \pitchfork W$ and

$$\min\{\text{dist}_W(x, \rho_W^V), \text{diam}_V(b_V \cup \rho_V^W)\} \leq \kappa$$

whenever $V \sqsubset W$.

The following is [7, Theorem 4.5].

Theorem 2.6 (distance formula) *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then there exists s_0 such that for all $s \geq s_0$, there exist C and K such that for all $x, y \in \mathcal{X}$,*

$$\text{dist}(x, y) \asymp_{K,C} \sum_{U \in \mathfrak{S}} \{\{\text{dist}_U(x, y)\}\}_s,$$

where $\{\{A\}\}_B$ denotes the quantity which is A if $A \geq B$ and 0 otherwise.

We recall the notion of a *hierarchical hull*, which originates in [9] for the setting of mapping class groups, and extends to the HHS setting in [7]. Given a constant θ , for any $F \subset \mathcal{X}$ we define

$$(2-3) \quad H_\theta(F) = \{x \in \mathcal{X} \mid \forall V \in \mathfrak{S}, \pi_V(x) \in \mathcal{N}_\theta(\text{hull}(\pi_V(F)))\},$$

where $\text{hull}(A)$ denotes the union of all geodesics connecting points of A . In words, H_θ is the set of points whose projections in every hyperbolic factor space land in a specified neighborhood of the hull of the image of F . That this is sufficiently nonvacuous is indicated by the following result which, as we will see, is an easy consequence of [7, Theorem 4.4].

Theorem 2.7 *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Given k , there exist θ and κ such that, if $F \subset \mathcal{X}$ is a set of cardinality k then for every $V \in \mathfrak{S}$ the image $\pi_V(H_\theta(F))$ and the hull of $\pi_V(F)$ lie within Hausdorff distance κ of each other.*

Proof By definition, $\pi_V(H_\theta(F))$ lies in a controlled neighborhood of $\text{hull}(\pi_V(F))$, so we are left to show that any point on a geodesic connecting points of $\pi_V(F)$ lies close to $\pi_V(H_\theta(F))$. This follows from [7, Theorem 4.4], which says that any two points of F are connected by a hierarchy path (with uniform constant). These are defined in Definition 6.5 below, but here we only need that said path has projection to any given $\mathcal{C}(W)$ which, as a set, uniformly coarsely coincides with a geodesic between the projections of the endpoints. In particular, for any θ large enough, the path will be contained in $H_\theta(F)$. \square

2.3 Refined projections and stable subsurface collections

We will be working in a broad but restricted class of HHSs:

Definition 2.8 Let $(\mathcal{X}, \mathfrak{S})$ be an HHS and let $G < \text{Aut}(\mathfrak{S})$. We say that $(\mathcal{X}, \mathfrak{S})$ is *G-colorable* if there exists a decomposition of \mathfrak{S} into finitely many families \mathfrak{S}_i such that each \mathfrak{S}_i is pairwise transverse and G acts on $\{\mathfrak{S}_i\}_i$ by permutations. We say that $(\mathcal{X}, \mathfrak{S})$ is *colorable* if it is $\text{Aut}(\mathfrak{S})$ -colorable.

The notion of colorability is inspired by Bestvina, Bromberg and Fujiwara [12], who, essentially, proved that $\text{MCG}(\Sigma)$ and $\mathcal{T}(\Sigma)$ are finitely $\text{MCG}(\Sigma)$ -colorable HHSs; we now explain how their work proves this fact. First of all, the standard HHS structures of $\text{MCG}(\Sigma)$ and $\mathcal{T}(\Sigma)$ both have as index set \mathfrak{S} the set of all essential subsurfaces (including the disconnected ones) that do not have pairs of pants as connected com-

ponents. Except for allowing disconnected subsurfaces, this is the same as the set Y considered in [12, Proposition 5.8]. The proposition yields a certain decomposition $Y = Y^1 \sqcup \dots \sqcup Y^k$, and from an inspection of the first paragraph of the proof one sees that the Y^i are the orbits of a certain finite-index subgroup G of $\text{MCG}(\Sigma)$. We can define our \mathfrak{S}_i to be the orbits in \mathfrak{S} (rather than Y) of this same finite-index subgroup. The fact that distinct elements of each Y^i have intersecting boundaries — as given by [12, Proposition 5.8] — implies that for any two distinct subsurfaces in the same \mathfrak{S}_i have connected components with intersecting boundaries. Since transversality in $\text{MCG}(\Sigma)$ and $\mathcal{T}(\Sigma)$ is defined via intersecting boundaries, we are done.

We sometimes refer to the \mathfrak{S}_i as *BBF families*.

For $A, B \subseteq \mathcal{C}(Y)$, we define $d_Y(A, B) := \text{diam}_{\mathcal{C}(Y)}(A \cup B)$.

Theorem 2.9 [14] *Let $(\mathcal{X}, \mathfrak{S})$ be a G -colorable HHS for $G < \text{Aut}(\mathfrak{S})$ with standard projections $\hat{\pi}_-$ and $\hat{\rho}_-$. There exist $\theta > 0$ and refined projections π_- and ρ_- with the same domains and ranges, respectively, such that:*

- (1) *If X and Y lie in different \mathfrak{S}_j , and $\hat{\rho}_Y^X$ is defined, then $\rho_Y^X = \hat{\rho}_Y^X$.*
- (2) *If $X, Y \in \mathfrak{S}_j$ are distinct, then the Hausdorff distance between ρ_Y^X and $\hat{\rho}_Y^X$ is at most θ .*
- (3) *If $x \in \mathcal{X}$ and $Y \in \mathfrak{S}$, then the Hausdorff distance between $\pi_Y(x)$ and $\hat{\pi}_Y(x)$ is at most θ .*
- (4) *If $X, Y, Z \in \mathfrak{S}_j$ for some j are pairwise distinct and $d_Y(\rho_Y^X, \rho_Y^Z) > \theta$, then $\rho_Z^X = \rho_Z^Y$.*
- (5) *Let $x \in \mathcal{X}$, and $Y, Z \in \mathfrak{S}_j$ for some j be pairwise distinct. If $d_Y(\pi_Y(x), \rho_Y^Z) > \theta$ then $\pi_Z(x) = \rho_Z^Y$.*

Moreover, $(\mathcal{X}, \mathfrak{S})$ equipped with π_- and ρ_- is an HHS, $G < \text{Aut}(\mathfrak{S})$, and it is G -colorable.

Proof The idea is to apply the construction from [14] to the standard projections $\hat{\pi}_-$ and $\hat{\rho}_-$ and distances \hat{d}_- for the sets $\mathfrak{S}_i \cup \mathcal{X}$ for each i , where we think of \mathcal{X} as a collection of single point spaces $x = \{x\}$ for each $x \in \mathcal{X}$.

Given a point $x \in \mathcal{X}$, we define projections ρ_x^- from domains in \mathfrak{S}_i and \mathcal{X} to x as the constant map $\rho_x^- \equiv x$. It is easily checked that $\mathfrak{S}_i \cup \mathcal{X}$, once equipped with the original projections $\hat{\pi}_-$ and $\hat{\rho}_-$ and these additional projections, satisfies the projection

axioms from [12]. The existence of projections and distances π_-, ρ_- and d_- , and that all properties hold for them, is then an immediate consequence of [14, Theorem 4.1]. Finally, the fact that $(\mathcal{X}, \mathfrak{S})$ equipped with these projections is an HHS follows from the fact that the new projections are bounded distance away from the old ones, by items (1), (2) and (3).

The fact that G still acts by automorphisms on the new structure follows from equivariance of the construction of the new projections, meaning [14, Theorem 4.1(3)]. \square

Definition 2.10 We say that a G -colorable HHS $(\mathcal{X}, \mathfrak{S})$ with $G < \text{Aut}(\mathfrak{S})$ has *stable projections* if it is equipped with the projections provided by Theorem 2.9.

For the rest of this section, fix a G -colorable HHS $(\mathcal{X}, \mathfrak{S})$ with $G < \text{Aut}(\mathfrak{S})$ and with stable projections. In particular, we assume that the standard projections for $(\mathcal{X}, \mathfrak{S})$ satisfy the stability properties in Theorem 2.9.

As usual, $d_Y(x_1, x_2)$ denotes $\text{diam}_{\mathcal{C}(Y)}(A_1 \cup A_2)$, where

- $A_i = \pi_Y(x_i)$ if $x_i \in \mathcal{X}$,
- $A_i = \rho_Y^{x_i}$ if $x_i \in \mathfrak{S}$ and either $x_i \sqsubseteq Y$ or $x_i \pitchfork Y$.

For any pair of points $x, y \in \mathcal{X}$ and constant $K > 0$, we let $\text{Rel}_K(x, y) \subset \mathfrak{S}$ denote the collection of $Y \in \mathfrak{S}$ such that $d_Y(x, y) > K$; we also set $\text{Rel}_K^i(x, y) = \text{Rel}_K(x, y) \cap \mathfrak{S}_i$.

Let θ satisfy Theorem 2.9(5). Following eg [9; 12; 24], we now consider a relation on $\text{Rel}_K(x, y)$ — the properties claimed below follow from [7, Proposition 2.8] and Definition 2.4(4). For any $K > 10\theta$, $\text{Rel}_K(x, y)$ is a partially ordered set with order $<$ such that $X < Y$ whenever $X \pitchfork Y$ and one of the following equivalent conditions hold:

$$d_Y(x, \rho_Y^X) \leq \theta, \quad d_X(\rho_X^Y, y) \leq \theta, \quad d_Y(\rho_Y^X, y) \geq K - \theta, \quad d_X(x, \rho_X^Y) \geq K - \theta.$$

When restricted to $\text{Rel}_K^i(x, y)$, the relation $<$ becomes a total order.

For a finite set $F \subset \mathcal{X}$, we define

$$\text{Rel}_K(F) = \bigcup_{x, y \in F} \text{Rel}_K(x, y) \quad \text{and} \quad \text{Rel}_K^i(F) = \text{Rel}_K(F) \cap \mathfrak{S}_i.$$

The following stability lemma follows directly from the construction in [14].

Lemma 2.11 *There exists $K \gg 2\theta$ such that if $x, y, y' \in \mathcal{X}$ satisfy $d_{\mathcal{X}}(y, y') \leq 1$, then for each i ,*

$$|\text{Rel}_K^i(x, y) \Delta \text{Rel}_K^i(x, y')| \leq 2.$$

Proof By contradiction, suppose we have distinct elements

$$Y_0, Y_1, Y_2 \in \text{Rel}_K^i(x, y) \setminus \text{Rel}_K^i(x, y')$$

with $Y_0 < Y_1 < Y_2$. If $K > 10\theta$, applying the definition of $<$ and Theorem 2.9(5), we see that $\pi_{Y_1}(x) = \rho_{Y_1}^{Y_0}$ and $\pi_{Y_1}(y) = \rho_{Y_1}^{Y_2}$. Also, since $\pi_{Y_2}(y)$ and $\pi_{Y_2}(y')$ are uniformly close to each other, if K is sufficiently large then we have $d_{Y_2}(Y_1, y) > \theta$ and hence, one again, $\pi_{Y_1}(y') = \rho_{Y_1}^{Y_2}$. But then $d_Y(x, y') = d_Y(x, y) \geq K$, which contradicts $Y_1 \notin \text{Rel}_K^i(x, y')$. \square

Proposition 2.12 *Let $K \gg 2\theta$ and $F \subset \mathcal{X}$ be any finite set. There exists*

$$M = M(K, \mathfrak{S}, |F|) > 0$$

such that, for any $F' \subset \mathcal{X}$ with $d_{\text{Haus}}(F, F') \leq 1$ and $|F'| \leq |F|$,

$$|\text{Rel}_K(F) \Delta \text{Rel}_K(F')| < M.$$

Proof Assume throughout the proof that K is sufficiently large.

Since there are finitely many colors, it suffices to prove the analogous statement for $\text{Rel}_K^i(F) \Delta \text{Rel}_K^i(F')$ for any given i .

Applying Lemma 2.11 twice, we see that if $d_{\mathcal{X}}(x, x'), d_{\mathcal{X}}(y, y') \leq 1$, then

$$|\text{Rel}_K^i(x, y) \setminus \text{Rel}_K^i(x', y')| \leq 4.$$

For each of the $|F|^2$ pairs $x, y \in F$, we can pick any x', y' with $d_{\mathcal{X}}(x, x'), d_{\mathcal{X}}(y, y') \leq 1$ such that there are at most $4|F|^2$ elements of $\text{Rel}_K^i(F) = \bigcup_{x,y \in F} \text{Rel}_K^i(x, y)$ that are not in $\text{Rel}_K^i(F')$; ie $|\text{Rel}_K^i(F) \setminus \text{Rel}_K^i(F')| \leq 4|F|^2$. Symmetrically, we have $|\text{Rel}_K^i(F') \setminus \text{Rel}_K^i(F)| \leq 4|F'|^2$, and since $|F'| \leq |F|$ by assumption, we finally get $|\text{Rel}_K^i(F) \Delta \text{Rel}_K^i(F')| < 8|F|^2$, as required. \square

2.4 Bounding involved domains

Let $(\mathcal{X}, \mathfrak{S})$ be a G -colorable HHS with stable projections for $G < \text{Aut}(\mathfrak{S})$, as provided by Theorem 2.9.

Let $F, F' \subset \mathcal{X}$ with $|F|, |F'| \leq k$ and $d_{\text{Haus}}(F, F') \leq 1$. We will now prove some stronger stability results about how the set of relevant domains (and their subdomains) changes between F and F' .

For any $K \gg 2\theta$ as above, let $\mathcal{U}(F) = \text{Rel}_K(F)$ and $\mathcal{U}(F') = \text{Rel}_K(F')$. Given $V \in \mathfrak{S}$, let $\mathcal{U}^V(F) = \{W \in \mathcal{U}(F) \mid W \sqsubset V\}$ and define $\mathcal{U}^V(F')$ similarly.

In many of our stability results, we will need to know how domains in $\mathcal{U}(F)$ may differ from those in $\mathcal{U}(F')$. We call such domains *involved*, and they come in two flavors:

Definition 2.13 We say that $V \in \mathcal{U}(F) \cup \mathcal{U}(F')$ is *involved* in the transition between F and F' if

- (1) $\pi_V(F) \neq \pi_V(F')$, or
- (2) $\mathcal{U}^V(F) \neq \mathcal{U}^V(F')$.

Proposition 2.14 *If K is sufficiently large then the following holds. Given $k > 0$ there exists $N_1 = N_1(k, \mathfrak{S}) > 0$ such that, if $|F|, |F'| \leq k$ and $d_{\text{Haus}}(F, F') \leq 1$, then there are at most N_1 domains $V \in \mathcal{U}(F) \cup \mathcal{U}(F')$ involved in the transition between F and F' .*

Proof By Proposition 2.12, it suffices to bound the number of involved domains in $\mathcal{U}(F) \cap \mathcal{U}(F')$. However, we will still have to bound the number of involved domains of type (1) in $\mathcal{U}(F) \cup \mathcal{U}(F')$. We note that, since F and F' lie within Hausdorff distance 1, up to increasing K we can assume that for each $V \in \mathcal{U}(F)$ we have $\text{diam}_{\mathcal{C}V}(\pi_V(F')) \geq K/2$, and similarly for $V \in \mathcal{U}(F')$.

Involved of type (1) Let $x \in F$. We say that $V \in \mathcal{U}(F) \cup \mathcal{U}(F')$ is *exposed to x* if $\pi_V(x)$ is not contained in $\pi_V(F')$. We define exposure for $x \in F'$ similarly (with an abuse, here we are considering F and F' as disjoint, so we should actually define exposure for $x \in F \sqcup F'$).

Observe that $V \in \mathcal{U}(F) \cup \mathcal{U}(F')$ satisfies $\pi_V(F) \neq \pi_V(F')$ if and only if V is exposed to some x in either F or F' . Hence it suffices to bound the number of exposed domains.

Since $|F|, |F'| \leq k$, we may fix a point $x \in F$ and consider domains V which are exposed to x . The case of domains exposed to points in F' follows from a symmetric argument.

Given x and V as above, there is a $y \in F$ such that $d_V(x, y) \geq K/4$ (this is because $\text{diam}_{\mathcal{C}(V)}(\pi_V(F)) \geq K/2$). Since F has at most k elements, we can further fix y with said property.

Suppose for a contradiction that there exist domains $V_1, V_2, V_3 \in (\mathcal{U}(F) \cup \mathcal{U}(F')) \cap \mathfrak{S}_i$ which are exposed to x , where \mathfrak{S}_i is the i^{th} BBF family, making the V_i necessarily

pairwise transverse (this suffices since there are finitely many BBF families). Up to reordering, we have $V_1 \prec V_2 \prec V_3$ in $\text{Rel}_{K/2}(x, y)$.

Since F and F' lie at Hausdorff distance at most 1, there is a pair $x_2, y_2 \in F'$ such that $d(x, x_2), d(y, y_2) \leq 1$, and necessarily we have $\pi_{V_2}(x) \neq \pi_{V_2}(x_2)$ (as we cannot have the containment “ \subseteq ”).

Since $d(x, x_2), d(y, y_2) \leq 1$, by taking $K \gg 2\theta$ sufficiently large, we can ensure that $d_{V_i}(x_2, y_2) > 2\theta$ for $i = 1, 2, 3$. Since $V_1 \prec V_2 \prec V_3 \in \text{Rel}_{K/2}(x, y)$, we also must have the same order $V_1 \prec V_2 \prec V_3$ in $\text{Rel}_{2\theta}(x_2, y_2)$.

Thus by Theorem 2.9, it follows that $\pi_{V_2}(x_2) = \rho_{V_2}^{V_1}$. However, Theorem 2.9 also implies that $\pi_{V_2}(x) = \rho_{V_2}^{V_1}$. This contradicts the fact that $\pi_{V_2}(x) \neq \pi_{V_2}(x_2)$, and completes the proof that there is a bound on domains of type (1).

Involved of type (2) Notice that if $W \in \mathcal{U}(F) \cap \mathcal{U}(F')$ is of type (2), then there necessarily exists an exposed domain $V \in \mathcal{U}(F) \cup \mathcal{U}(F')$ of type (1) with $V \sqsubset W$ (by a \sqsubset -minimality argument: there must exist some $V \sqsubsetneq W$ with $\pi_V(F) \neq \pi_V(F')$, by virtue of V lying in the symmetric difference of $\mathcal{U}(F)$ and $\mathcal{U}(F')$, and hence a \sqsubset -minimal such V exists by Definition 2.4(5)). We therefore bound the number of such containers W for a fixed exposed domain V , of which there is a bounded number by the first part of the proof.

Since $|F|, |F'| \leq k$, it suffices to fix $x, y \in F$ and provide a bound on the number of elements $W \in \text{Rel}_K(x, y)$ which contain a fixed domain V .

In fact, we can conclude with an argument that does not rely on colorability, which we record here as a separate lemma since it might be of independent interest.

Lemma 2.15 *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. Then there exists N such that for any sufficiently large K the following holds. Let $x, y \in \mathcal{X}$ and $V \in \mathfrak{S}$. Then there are at most N elements $W \in \mathfrak{S}$ with $V \sqsubsetneq W$ such that $W \in \text{Rel}_K(x, y)$.*

Proof We fix K large enough that $\text{Rel}_K(x, y)$ is partially ordered for all $x, y \in \mathcal{X}$, and larger than $10\kappa_0$ for κ_0 as in Definition 2.4(4).

First of all, there is a bound on the maximal number of pairwise \sqsubset -comparable domains by Definition 2.4(5), and similarly there is a bound on the maximal number of pairwise orthogonal domains by [29, Lemma 1.4]. Hence, in view of Ramsey’s theorem, there exists N such that any collection of more than N domains contains three pairwise transverse elements.

Suppose, by way of contradiction, that there exist more than N domains W as in the statement, for given x, y and V , and consider $W_1 < W_2 < W_3$ in $\text{Rel}_K(x, y)$. By Definition 2.4(4), we have that $d_{W_2}(\pi_{W_2}(x), \rho_{W_2}^{W_1}) \leq \kappa_0$ and $d_{W_2}(\pi_{W_2}(y), \rho_{W_2}^{W_3}) \leq \kappa_0$. However, since $V \sqsubset W_i$ for each $i = 1, 2, 3$, we have $d_{W_2}(\rho_{W_2}^{W_1}, \rho_{W_2}^V) < \kappa_0$ and $d_{W_2}(\rho_{W_2}^{W_3}, \rho_{W_2}^V) < \kappa_0$, and so $d_{W_2}(\rho_{W_2}^{W_1}, \rho_{W_2}^{W_3}) < 2\kappa_0$ by the triangle inequality. But since $d_{W_2}(x, y) \geq d_{W_2}(\rho_{W_2}^{W_1}, \rho_{W_2}^{W_3}) - 2\kappa_0 > K - 2\kappa_0 > 2\kappa_0$ by assumption, this is a contradiction. \square

As explained above, the lemma concludes the proof of the proposition. \square

Remark 2.16 While not strictly necessary, we can simplify the setup that we deal with in Section 3 thanks to the following: Given an HHS, we can $\text{Aut}(\mathcal{X}, \mathfrak{S})$ -equivariantly change the structure in a way that all $\pi_V(x)$ and ρ_V^U for $U \sqsubset V$ are points, rather than bounded sets, and that moreover the new structure has stable projections if the old one did. This can be achieved, for example, by replacing each $\mathcal{C}(V)$ by the nerve of the covering given by subsets of sufficiently large diameter — which is quasi-isometric to $\mathcal{C}(V)$. In particular, the vertices of the new $\mathcal{C}(V)$ are labeled by bounded sets, and we can redefine $\pi_V(x)$ to be the vertex labeled by $\pi_V(x)$, and similarly for ρ_V^U ; all properties required are straightforward to check.

In Section 3, we will deal with finite subsets of a hyperbolic space. If in Section 4 we did not modify the HHS structure as outlined above, we would instead have to deal with finite collections of bounded subsets. This is possible, but would make the arguments more opaque.

3 Stable trees

In this section we will consider the geometry of trees in a δ -hyperbolic space, in preparation for arguments that will take place in the individual hyperbolic spaces of our hierarchical structure. Our main result will be Theorem 3.3, stated below after some preliminary definitions. This is the only result from this section that will get used later (namely, in Section 4).

Fix a geodesic δ -hyperbolic space \mathcal{Z} . For a finite subset $F \subset \mathcal{Z}$ let $\text{hull}(F) \subset \mathcal{Z}$ be the set of geodesics connecting points of F . Hyperbolicity tells us that $\text{hull}(F)$ can be approximated by a finite tree with accuracy depending only on δ and the cardinality $\#F$. To systematize this for the purposes of this section, we make the following definitions.

Let us fix a function λ which assigns, to any finite subset F of \mathcal{Z} , a minimal network spanning F . That is, $\lambda(F)$ is a 1–complex embedded in \mathcal{Z} with the property that $\lambda(F) \cup F$ is connected, and has minimal length among all such 1–complexes (where the length of a 1–complex embedded in \mathcal{Z} is the sum of the lengths of all edges). Minimality implies $\lambda(F)$ is a tree. Let us similarly define λ' which assigns, to any finite collection A_1, \dots, A_k of subsets of \mathcal{Z} , a minimal network that spans them. That is, $\lambda'(A_1, \dots, A_k)$ is a 1–complex in \mathcal{Z} of minimal length with the property that the quotient of $\lambda' \cup A_1 \cup \dots \cup A_k$ obtained by collapsing each A_i to a point is connected. Minimality again implies that this collapsed graph is a tree. For convenience we assume that $\lambda(\{x_1, \dots, x_k\}) = \lambda'(\{x_1\}, \dots, \{x_k\})$.

The following lemma illustrates a basic property of hyperbolic spaces, and we omit its proof.

Lemma 3.1 *Let \mathcal{Z} be a geodesic δ –hyperbolic space and λ a minimal network function as above. Then there exists $\epsilon_0 = \epsilon_0(k, \delta)$ such that, for all $\epsilon \geq \epsilon_0$, there exists $\epsilon' > \epsilon$ such that if $F \subset \mathcal{Z}$ has cardinality k then*

- *there is a $(1, \epsilon/2)$ –quasi-isometry $\lambda(F) \rightarrow \text{hull}(F)$ which is $\epsilon/2$ –far from the identity;*
- *for any two points $x, y \in \mathcal{N}_\epsilon(\lambda(F))$, any geodesic joining them is in $\mathcal{N}_{\epsilon'}(\lambda(F))$.*

In the rest of this section we consider the following situation. Let a (large but) finite set $\mathcal{Y} \subset \mathcal{N}_{\epsilon/2}(\text{hull}(F))$ be given (see Section 4 for what \mathcal{Y} will be in our setting). It is possible to divide $\lambda(F)$ up into a union of subtrees some of which are close approximations to “clusters” in \mathcal{Y} and the rest interconnect the clusters, but such a construction is not unique, depending on many choices (including the choice of $\lambda(F)$ itself). Our goal in this section is to describe a version of this which is stable, in the sense that small changes in the sets F and \mathcal{Y} only alter the tree and its subtrees in a controlled way — independently of the diameter of F or the cardinality of \mathcal{Y} .

Remark 3.2 For convenience in our discussion we allow ourselves to assume that the points of F are all leaves of $\lambda(F)$. This can be arranged by a slight perturbation, or by considering each point of F as the endpoint of an additional edge of length 0.

Given $E \gg \epsilon$, let $\mathcal{C}_E(\mathcal{Y} \cup F)$ be the graph whose edges connect points of $\mathcal{Y} \cup F$ that are at most E apart. Vertex sets of connected components of \mathcal{C}_E are called *clusters*.

We will choose E to be a suitably large multiple of ϵ' . We note that the relation of E to ϵ' and ϵ is the one sensitive part of the argument, and elsewhere we can be content with order-of-magnitude arguments.

For a simplicial tree T , let $d(v)$ be the valence of each vertex v and let $k(T)$ be the number of leaves, ie vertices of valence 1. We have $\sum_{d(v)>2}(d(v)-2) = k(T) - 2$, for example by an Euler characteristic argument. We call this quantity the *total branching* of T .

The following theorem is the main result of this section.

Theorem 3.3 *Given k, N, δ , and $\epsilon \geq \epsilon_0(k, \delta)$ as in Lemma 3.1 there exists $K > 0$ such that the following holds. Let \mathcal{Z} be a geodesic δ -hyperbolic space and let $F, \mathcal{Y} \subset \mathcal{Z}$ be finite subsets, where $|F| \leq k$ and $\mathcal{Y} \subset \mathcal{N}_{\epsilon/2}(\text{hull}(F))$.*

There exists a metric tree $T = T(F, \mathcal{Y})$ with a decomposition into two forests $T = T_c \cup T_e$ intersecting along a finite set of points and a map $\Xi = \Xi_{F, \mathcal{Y}}: T(F, \mathcal{Y}) \rightarrow \mathcal{Z}$ such that:

- (a) *The total branching of T is bounded by $2k - 4$.*
- (b) *Ξ is a (K, K) -quasi-isometric embedding with image K -Hausdorff close to $\text{hull}(F)$.*
- (c) *For each component τ of T_e we have that $\Xi|_{\tau}$ is a $(1, K)$ -quasi-isometric embedding, and an isometry onto $\Xi(\tau)$ endowed with its path metric.*
- (d) *There is a bijection b between components of T_c and clusters in $\mathcal{C}_E(\mathcal{Y} \cup F)$ such that $\Xi(\tau)$ is K -Hausdorff close to $b(\tau)$ for each component τ of T_c .*

Furthermore, if $F', \mathcal{Y}' \subset \mathcal{Z}$ and $g \in \text{Isom}(\mathcal{Z})$ are such that $|F'| \leq k$, \mathcal{Y}' is finite, $d_{\text{Haus}}(gF, F') \leq 1$, and $|g\mathcal{Y} \Delta \mathcal{Y}'| < N$, then there exists a constant $L = L(N, k, \delta) > 0$ and subsets $T_s \subset T_e(F, \mathcal{Y})$ and $T'_s \subset T_e(F', \mathcal{Y}')$ such that, identifying components of $T_e(F, \mathcal{Y})$ and $T_e(F', \mathcal{Y}')$ with their images in \mathcal{Z} , we have:

- (1) *The components of T_s and T'_s are contained in the edges of $T_e(F, \mathcal{Y})$ and $T_e(F', \mathcal{Y}')$, respectively.*
- (2) *The complements $T_e(F, \mathcal{Y}) \setminus T_s$ and $T_e(F', \mathcal{Y}') \setminus T'_s$ have at most L components, each of diameter at most L .*
- (3) *There is a bijective correspondence between the sets of the components of gT_s and T'_s .*




Figure 5: An example of the stable tree $T = T_c \cup T_e$ provided by Theorem 3.3.

- (4) Under this correspondence, all but L components are exactly the same, and the identical components of T_s and T'_s come from the identical components of $T_e(F, \mathcal{Y})$ and $T_e(F', \mathcal{Y}')$.
- (5) The remaining L components of gT_s are each at Hausdorff distance L to the corresponding component in T'_s .

We call the trees $T(F, \mathcal{Y})$ *stable trees*.

Remark 3.4 (coarse equivariance and its proof) The “furthermore” part of Theorem 3.3 can be interpreted as simultaneously stating two facts. For g the identity, it says that the trees are stable under perturbations of F and \mathcal{Y} . Alternatively, for $F' = gF$ and $\mathcal{Y}' = g\mathcal{Y}$, it says that the construction is coarsely equivariant. In either case, what we have to prove is essentially the following. The construction relies on certain choices, namely the choices of functions λ and λ' as above, and we have to show that these only cause the kinds of perturbations described in the statement of the theorem. From this perspective, it is clear that the proof for a general g is the same as that for $g = 1$, as $gT(F, \mathcal{Y})$ coincides with the tree $T(gF, g\mathcal{Y})$ constructed based on different choices. To save notation and make the proof more readable, we will prove only the case where g is the identity.

3.1 Cluster separation graph

Let $\mathcal{C}_E = \mathcal{C}_E(F \cup \mathcal{Y})$ be as above and let $C_1, C_2, C_3 \subset \mathcal{C}_E^0$ be clusters (ie vertex sets of connected components). We say that C_2 *separates* C_1 from C_3 in \mathcal{Z} if there

exists a minimal \mathcal{Z} -geodesic segment σ with endpoints on C_1 and C_3 which meets the $2\epsilon'$ -neighborhood of C_2 .

Definition 3.5 Let $\mathcal{G}_E = \mathcal{G}_E(F \cup \mathcal{Y})$ be a graph whose vertex set \mathcal{G}_E^0 is the set of clusters of \mathcal{C}_E , and where $[C_1, C_2]$ is an edge whenever there is no cluster separating C_1 from C_2 in \mathcal{Z} . We call \mathcal{G}_E the *separation graph* for \mathcal{C}_E .

Lemma 3.6 *If $E > 4\epsilon'$, then \mathcal{G}_E is connected.*

Proof Let $C, C' \in \mathcal{G}_E^0$. If $C \neq C'$, then $d_{\mathcal{Z}}(C, C') > E$ (here $d_{\mathcal{Z}}$ means the minimal distance, not the diameter of the union). If they are not adjacent in \mathcal{G} , then there is a third cluster B separating them in \mathcal{Z} . Let σ be a minimal geodesic connecting C to C' with $p \in \sigma$ within $2\epsilon'$ of B . Then p is distance at least $E - 2\epsilon'$ from each end of σ since B is at least E from both C and C' . It follows that $d_{\mathcal{Z}}(B, C) \leq d_{\mathcal{Z}}(C, C') - E + 4\epsilon'$, and similarly for $d_{\mathcal{Z}}(B, C')$.

If $d_{\mathcal{Z}}(C, C') \leq 2E - 4\epsilon'$ this gives $d_{\mathcal{Z}}(B, C) \leq E$ which is a contradiction, so C and C' must be connected by an edge in \mathcal{G}_E . For $d_{\mathcal{Z}}(C, C') > 2E - 4\epsilon'$, we have that $d_{\mathcal{Z}}(B, C)$ and $d_{\mathcal{Z}}(B, C')$ are smaller than $d_{\mathcal{Z}}(C, C')$ by at least $E - 4\epsilon'$, so we can proceed inductively. □

For ease of notation, set $\mathcal{C} = \mathcal{C}_E$ and $\mathcal{G} = \mathcal{G}_E$.

Definition 3.7 For any subset A of \mathcal{Z} , let its *shadow* $s(A)$ be the subtree of $\lambda(F)$ obtained by taking the convex hull (in $\lambda(F)$) of all the points in $\lambda(F)$ within distance ϵ from points of A . For a singleton $\{x\}$ we also write $s(x) := s(\{x\})$.

Note that, since $\mathcal{Y} \cup F$ is in $\mathcal{N}_{\epsilon}(\lambda(F))$ by hypothesis, $s(C) \neq \emptyset$ for any nonempty subset $C \subset \mathcal{Y} \cup F$.

The rest of this subsection is devoted to establishing several properties of shadows which will connect the separation properties of clusters in \mathcal{G} to separation properties of their shadows in $\lambda(F)$, thereby allowing us to work with \mathcal{G} and independently of $\lambda(F)$.

The next lemma controls how and when shadows of clusters can intersect.

Lemma 3.8 *Let $E > 7\epsilon$ and let $C, C' \in \mathcal{G}^0$ be distinct clusters. Then*

- (1) $s(C) \cap s(C')$ can contain no leaf of $s(C)$ or $s(C')$;




Figure 6: When there is branching, shadows can overlap in their interiors, but never at their leaves.

- (2) the diameter of $s(C) \cap s(C')$ is bounded by a constant depending on $\#F$, E , and ϵ ;
- (3) if at least one of $s(C)$ and $s(C')$ is an interval along an edge of $\lambda(F)$, then $s(C) \cap s(C') = \emptyset$.

Proof Note first that for any $x \in C$, $s(x)$ is a subtree of diameter at most 3ϵ . This is because any two extreme points of $s(x)$ are within ϵ of x , and $\lambda(F)$ is $(1, \epsilon)$ -quasi-isometrically embedded. Similarly, for any $x, y \in C$,

$$\text{diam}_{\lambda(F)}(s(\{x, y\})) \leq d(x, y) + 3\epsilon.$$

Claim 1 For every $p \in s(C)$, there exists $q \in s(C)$ at distance (in $\lambda(F)$) at most $(E + 3\epsilon)/2$ such that $d(q, C) \leq \epsilon$.

Proof Either $p \in s(x)$ for some $s(x)$ containing an extreme point of $s(C)$, or p separates some $s(x)$ from $s(y)$, for $x, y \in C$. In the first case p is within $3\epsilon/2$ of a point q for which $d(q, x) \leq \epsilon$ and we are done. In the second case, a path in C_E from x to y then yields a sequence of points $x_i \in C$ such that $d(x_i, x_{i+1}) \leq E$ and p is contained in one of the shadows $s(\{x_i, x_{i+1}\})$. Since $\text{diam}_{\lambda(F)}(s(\{x_i, x_{i+1}\})) \leq E + 3\epsilon$, we find that p is within $(E + 3\epsilon)/2$ of an extreme point q of $s(\{x_i, x_{i+1}\})$, so $d(q, x_i) \leq \epsilon$ or $d(q, x_{i+1}) \leq \epsilon$. The claim follows. \square

For (1), suppose that a leaf p of $s(C)$ is in $s(C')$. Note that the leaves of $s(C)$ and $s(C')$ are within ϵ of C and C' , respectively. By the previous paragraph, there is a point q of $s(C')$ within $(E + 3\epsilon)/2$ of p which is ϵ close to C' . Thus we obtain $d(C, C') \leq 2\epsilon + (E + 3\epsilon)/2 < E$, so $C = C'$.

For (2), suppose that $s(C) \cap s(C')$ contains an edge e of length greater than $2(E + 3\epsilon)$. Claim 1 implies that there is a set R in $s(C)$ consisting of points at distance ϵ from




Figure 7: The cluster C is guaranteed valence at least 3 (via S_1, S_2 and S_3) by Lemma 3.10. In this case it has valence 4.

C and whose $(E + 3\epsilon)/2$ -neighborhood covers $s(C)$; there is also a similar set R' in $s(C')$. Since e is in both shadows, it must be that $e \cap R$ and $e \cap R'$ both cut e into intervals of length at most $E + 3\epsilon$. Thus it must be that there is a point $r \in R \cap e$ and $r' \in R' \cap e$ that are distance $(E + 3\epsilon)/2$ apart. Then just as before we obtain $d(C, C') < E$ so $C = C'$. Now the number of edges in $s(C) \cap s(C')$ is bounded by the total branching of the tree, which depends on $\#F$. This gives (2).

Finally, for (3), if one of $s(C)$ and $s(C')$ is an interval contained in an edge of $\lambda(F)$ then it is easy to see that, if they overlap, then one must contain a leaf of the other, thereby violating (1). □

Definition 3.9 From now on we set $E = 8\epsilon'$, so that the conclusions of both Lemmas 3.6 and 3.8 hold.

The following lemma connects the separation properties in \mathcal{G} of a cluster to the separation properties in $\lambda(F)$ of its shadow.

Lemma 3.10 *Let C be a cluster and S_1, \dots, S_m be the components of $\lambda(F) \setminus \text{int}(s(C))$ which meet $s(C)$ at a leaf of $s(C)$. Let G_i be the set of clusters $B \in \mathcal{G}^0 \setminus \{C\}$ such that $s(B) \cap S_i \neq \emptyset$. Then each G_i is in a distinct component of $\mathcal{G} \setminus C$, and moreover the valence of C in \mathcal{G} at least m .*




Figure 8: Clusters with shadows on an edge in $\lambda(F)$ give rise to a path of bivalent vertices in \mathcal{G} .

Proof Note that if B is a cluster in G_i then $s(B)$ is actually disjoint from $s(C)$, since the leaves of $s(C)$ cannot meet $s(B)$ by Lemma 3.8. Moreover, there may be clusters $C' \neq C$ that are not in any G_i ; their shadows meet components of $\lambda(F) \setminus \text{int}(s(C))$ that do not meet leaves of $s(C)$.

Let $A \in G_i$ and $B \notin G_i \cup \{C\}$. A minimal geodesic σ in \mathcal{Z} connecting A to B must be ϵ' -close to the path in $\lambda(F)$ connecting $s(A)$ to $s(B)$, and this path passes through a leaf p of $s(C)$ (namely $s(C) \cap S_i$). Thus there is a point of C within $\epsilon + \epsilon' < 2\epsilon'$ of σ , so C separates A from B in \mathcal{Z} . In particular A and B cannot be adjacent in \mathcal{G} .

Thus G_i cannot be connected to any vertex in $\mathcal{G}^0 \setminus (G_i \cup \{C\})$, which implies distinct G_i are in distinct components of $\mathcal{G} \setminus C$.

To see that the valence is at least m , we must check that each G_i is nonempty. But each S_i must contain a leaf of $\lambda(F)$, which is a point of F , so there must be a cluster whose shadow is in S_i . □

Lemma 3.11 *If e is an edge of $\lambda(F)$, the clusters C whose shadows $s(C)$ are subintervals of e form a path in \mathcal{G} whose interior vertices are bivalent. The ordering of this path matches the ordering of the shadows in e .*

Proof Let $\{C_1, \dots, C_l\}$ be the set of clusters whose shadows are subintervals of e . By Lemma 3.8, $s(C_i) \cap s(C_j) = \emptyset$ for all i and j . We may therefore assume that their indices correspond to the order they appear along e in $\lambda(F)$.

The complement $\lambda(F) \setminus \text{int}(s(C_i))$ has two components for each i , labeled A_i^- and A_i^+ , such that A_i^- contains $s(C_{i-1})$ when $i > 0$ and A_i^+ contains $s(C_{i+1})$ when $i < l$.

By our ordering, no shadows lie between $s(C_i)$ and $s(C_{i+1})$. Lemma 3.10 implies that C_i separates (in \mathcal{G}) the clusters whose shadows lie in A_i^- from those in A_i^+ . In particular, no B can separate C_i from C_{i+1} in \mathcal{G} , so they are adjacent and we obtain a path C_1, \dots, C_l in \mathcal{G} . Moreover for $1 < i < l$ we can see that C_i is bivalent as follows: if $D \in \mathcal{C} \setminus \{C_{i-1}, C_i, C_{i+1}\}$, then one of C_{i-1} or C_{i+1} separates D from C_i in \mathcal{G} , again by Lemma 3.10, and so there can be no edge $[C_i, D]$ and the valence of C_i is exactly 2. \square

Lemma 3.12 *If C has valence 2 in \mathcal{G} but $s(C)$ is not an interval inside an edge of $\lambda(F)$, then C contains a point of F .*

Proof If $s(C)$ is not an interval in an edge of $\lambda(F)$, it has a branch point and hence at least three leaves. At most two of these can be interior to $\lambda(F)$, because otherwise C would have valence at least 3 in \mathcal{G} by Lemma 3.10.

Thus $s(C)$ contains a leaf q of $\lambda(F)$, which is a point of F . This means $d(q, C) \leq \epsilon < E$ (notice that, since q is a leaf, it lies in the convex hull of a subset of $\lambda(F)$ only if it lies in the subset). Hence, we have $q \in C$. \square

Structure of bivalent clusters Let \mathcal{E}^0 denote the set of clusters $C \in \mathcal{G}^0$ which have valence 2 in \mathcal{G} and do not contain a point of F . Lemma 3.12 implies that each $C \in \mathcal{E}^0$ has shadow inside an edge of $\lambda(F)$.

The next lemma gives that almost all clusters are bivalent.

Lemma 3.13 $\#(\mathcal{G}^0 \setminus \mathcal{E}^0) \leq 2k - 2.$

Proof For a cluster $C \in \mathcal{G}^0 \setminus \mathcal{E}^0$, either C contains a point of F , or $s(C)$ contains a branch point of $\lambda(F)$. There are at most k clusters of the former type. The number of clusters of the latter type is bounded by the total branching $b(\lambda(F))$, but to show this we must contend with the fact that shadows can overlap.

Let $W \subset \lambda(F)$ be a connected union of shadows $s(C_1), \dots, s(C_m)$, each of which contains a branch point. By Lemma 3.8, no leaf of $s(C_i)$ can be in $s(C_j)$ for $i \neq j$. Hence all leaves of $s(C_i)$ must be leaves of W and disjoint from each other. Since each $s(C_i)$ has at least two leaves,

$$m \leq \left\lfloor \frac{k(W)}{2} \right\rfloor = \left\lfloor \frac{b(W) + 2}{2} \right\rfloor,$$

where $k(W)$ is the number of leaves and $b(W)$ is the total branching of W . Since $b(W) \geq 1$, this implies $m \leq b(W)$. Summing over all such W we find that the number of clusters with branch points in their shadows is bounded by $b(\lambda(F))$, or $k - 2$. The desired inequality follows. \square

Let \mathcal{E} be the subgraph of \mathcal{G} induced on the vertices \mathcal{E}^0 .

Lemma 3.14 *Let $\mathcal{E}_1, \dots, \mathcal{E}_m$ be the components of \mathcal{E} . For each \mathcal{E}_i there is an edge e_i of $\lambda(F)$ such that \mathcal{E}_i is a path C_1, \dots, C_{r_i} in \mathcal{G} consisting of all elements of \mathcal{E} whose shadows lie in the interior of e_i ; the edges e_i are distinct.*

Proof Since each cluster $D \in \mathcal{E}_i$ is a bivalent vertex of \mathcal{G} with shadow in an edge of $\lambda(F)$ by Lemma 3.12, and Lemma 3.11 implies that all such clusters with shadows on a given edge $e \in \lambda(F)$ form a path in \mathcal{G} , it suffices to prove that no two such edge paths of bivalent clusters in \mathcal{G} are directly connected by an edge.

Suppose $C, D \in \mathcal{E}_i$ are connected by an edge in \mathcal{G} but $s(C)$ and $s(D)$ are not contained in a single edge of $\lambda(F)$. Since $s(C) \cap s(D) = \emptyset$, we may label the components of $\lambda(F) \setminus s(C)$ and $\lambda(F) \setminus s(D)$ by γ_{\pm} and δ_{\pm} , respectively, so that $s(C) \subset \delta_{-}$ and $s(D) \subset \gamma_{-}$. Then the intersection $\gamma_{-} \cap \delta_{-}$ contains a vertex v of $\lambda(F)$ of valence at least 3.

By Lemma 3.10, $\mathcal{G} \setminus C$ is divided into subgraphs $\mathcal{G}(\gamma_{\pm})$ spanned by clusters whose shadows are in γ_{\pm} , respectively, and are separated by C , and similarly $\mathcal{G}(\delta_{\pm})$ are separated by D , respectively. In particular note $C \in \mathcal{G}(\delta_{-})$ and $D \in \mathcal{G}(\gamma_{-})$.

Since v has valence at least 3, there is a component of $\lambda(F) \setminus \{v\}$ that meets neither $s(C)$ or $s(D)$. A leaf of this component is in the shadow of a cluster B which is therefore in $\mathcal{G}(\gamma_{-}) \cap \mathcal{G}(\delta_{-})$.

Since \mathcal{G} is connected, B is connected to C within $\mathcal{G}(\gamma_{-})$ and to D within $\mathcal{G}(\delta_{-})$. Since C and D are bivalent and by hypothesis adjacent in \mathcal{G} , the edge between them is the only edge connecting C to $\mathcal{G}(\gamma_{-})$, and the only edge connecting D to $\mathcal{G}(\delta_{-})$. Hence any path from B to C must pass through this edge and must therefore meet D first. Reversing the roles of C and D we obtain a contradiction. \square

3.2 Constructing the stable tree

In this section, we construct our stable tree $T(F, \mathcal{Y})$ from the structure of \mathcal{G} without referring to $\lambda(F)$ directly. In Proposition 3.17 below, we prove it is quasi-isometric to $\lambda(F)$.




Figure 9: The construction of a stable tree. Note that each complementary component of $\mathcal{G} \setminus \mathcal{E}^0$ determines multiple components of T_e , eg the lavender forest $\lambda'(V^0)$ determined by the component V whose boundary is the bivalent clusters $C_1, C_2,$ and C_3 . Each cluster C then determines a single component $\mu(C)$ of T_c by connecting the points $r(C) = C \cap (T_e \cup F)$.

The two forests Now let us proceed to define the forests $T_c(F, \mathcal{Y})$ and $T_e(F, \mathcal{Y})$. We let $\mathcal{G} = \mathcal{G}_E(F \cup \mathcal{Y})$ be as above.

Let \mathcal{V} denote the set of closures of connected components of $\mathcal{G} \setminus \mathcal{E}^0$. Thus each element of \mathcal{V} is a subgraph connected to the rest of \mathcal{G} along vertices in \mathcal{E}^0 . For each $V \in \mathcal{V}$ let V^0 denote its vertex set, which is a collection of clusters. Some elements of \mathcal{V} are

single edges $[C, D]$, where $C, D \in \mathcal{E}^0$, and others are subgraphs containing vertices in $\mathcal{G}^0 \setminus \mathcal{E}^0$, and we note they are not necessarily trees, though Lemma 3.13 bounds their size.

For each $V \in \mathcal{V}$, let $\lambda'(V^0)$ be the minimal network defined in the beginning of this section, where V^0 is interpreted as a collection of subsets (clusters) in \mathcal{Z} .

Now define

$$T_e = T_e(F, \mathcal{Y}) = \bigsqcup_{V \in \mathcal{V}} \lambda'(V^0).$$

Remark 3.15 The forest T_e is a disjoint union of copies of forests each contained in \mathcal{Z} . It is important to note, however, that these trees might in fact intersect in \mathcal{Z} . With a slight abuse, we will conflate the abstract copies of the $\lambda'(V^0)$ that constitute T_e and their “concrete” counterparts in \mathcal{Z} . Similar comments apply to T_c below. Since the map $\Xi: T \rightarrow \mathcal{Z}$ is just going to be the identity on all the components of T_e and T_c , we will allow ourselves to regard T as a subset of \mathcal{Z} for purposes that do not require understanding the metric of T , eg when measuring the Hausdorff distance between (the image in \mathcal{Z} of) a subset of T and a subset of \mathcal{Z} .

Note that T_e is a forest whose leaves are points of clusters.

Collapsing clusters to points, T_e becomes a connected network N , by the definition of \mathcal{V} . This connected network is a union of trees joined at points that correspond to vertices of \mathcal{E}^0 . Since any vertex in \mathcal{E}^0 disconnects \mathcal{G} , each of these join points disconnects N , so we see that N is a tree.

Now for each cluster $C \in \mathcal{G}^0$, we consider the set of points $r(C) = C \cap (T_e \cup F)$. We let $\mu(C)$ denote the tree $\lambda(r(C))$, and define

$$T_c = T_c(F, \mathcal{Y}) = \bigsqcup_{C \in \mathcal{G}^0} \mu(C).$$

The tree We now define $T(F, \mathcal{Y}) = T_c(F, \mathcal{Y}) \cup T_e(F, \mathcal{Y})$, or $T = T_c \cup T_e$ for short. Note that T is a tree because as above collapsing the subtrees of T_c to points yields a tree; see Figure 9.

Lemma 3.16 *Let $T = T(F, \mathcal{Y}) = T_c \cup T_e$.*

- (1) *The total branching $b = b(T)$ is bounded by $2k - 4$, and the leaves of T are contained in $F \cup \mathcal{Y}$.*

- (2) $\mu(C) \subseteq \mathcal{N}_{O(\epsilon)}(C)$, so $T_c \subset \mathcal{N}_{O(\epsilon)}(\mathcal{Y} \cup F)$.
- (3) For all $p \in T_e$, we have $d_{\mathcal{Z}}(p, \mathcal{Y} \cup F) \geq \frac{1}{b} d_T(p, \partial T_e) - O(\epsilon)$.

Proof To bound $b(T)$, we bound the number of leaves T can have. Leaves of T are leaves of the various components of T_e and T_c , and thus can arise in two ways:

- (a) If a cluster C contains points of F , the points in $C \cap F$ can be leaves of $\mu(C)$ which are also leaves of T . There are at most k such points.
- (b) If a cluster C contains no points of F and a single point q of ∂T_e , which is connected to only one subtree of T_e , then q is a leaf of T ; see Figure 10.

All other vertices of $\partial T_e \cup \partial T_c$ have valence at least 2. Notice that we already showed that all leaves of T are contained in $F \cup \mathcal{Y}$.

Clusters of type (b) must be in $\mathcal{G}^0 \setminus \mathcal{E}^0$ since every cluster in \mathcal{E}^0 belongs to two subgraphs in \mathcal{V} , and hence either has two points in ∂T_e or two subtrees of T_e meeting at a single point. The number of clusters in $\mathcal{G}^0 \setminus \mathcal{E}^0$ that don't contain points of F was bounded in Lemma 3.13 by $k - 2$.

This gives us a bound of $2k - 2$ on the total number of leaves in T , which bounds the total branching by $2k - 4$. This proves part (1).

Now for part (2), consider the minimal network $\mu(C)$ for the cluster C . By Lemma 3.1 and the definition of shadows, $\mu(C)$ is within $O(\epsilon)$ of the shadow of C in $\lambda(F)$, and it follows from Claim 1 (in the proof of Lemma 3.8) that every point of $s(C)$ is within $O(E)$ of C . This proves part (2).

For part (3), let $p \in \lambda'(V^0) \subset T_e$, where $V \in \mathcal{V}$, and let the distance $d_{\mathcal{Z}}(p, \mathcal{Y} \cup F)$ be realized on a point in a cluster C_1 . Write $d_{\mathcal{Z}}(p, C_1) = t$.

Suppose first that $C_1 \in V^0$. The quotient of $\lambda'(V^0)$ obtained by collapsing the clusters of V^0 to points is a tree by minimality of the network, so there is some sequence of components of $\lambda'(V^0)$ which connects p to C_1 , possibly through clusters $C_2, \dots, C_l \in V^0$.

Consider the unique path α in $\lambda'(V^0)$ from p to C_2 . The path α branches at no more than $b = b(T)$ points, so let $\alpha' \subset \alpha$ be the longest unbranched subsegment of α . We thus have $|\alpha'| \geq \frac{1}{b} d_T(p, \partial \lambda'(V^0))$. If $|\alpha'| > t$, we may remove α' from $\lambda'(V^0)$, attach a minimal length path in \mathcal{Z} from p to C_1 (of length t), and obtain a network with smaller total length than $\lambda'(V^0)$ that still connects the clusters in V . This would violate the




Figure 10: The stable tree $T = T(F, \mathcal{Y})$ may have leaves which are not points in F , and some points of F may not be leaves of T . In this example, the ambient space \mathcal{Z} is the whole graph on the left, and T is realized geometrically on the right. The orange cluster points atop spikes at the branch points of the underlying tree create leaves in T . New leaves in T always arise from clusters near branch points of $\lambda(F)$. The pink bivalent cluster points determine bivalent vertices in T , with small neighborhoods thereof folding into the spikes upon inclusion of $T \rightarrow \mathcal{Z}$. By contrast, $\lambda(F)$ contains none of the spikes. Finally, the pairs of nearby points of F in $\lambda(F)$ on the left side of \mathcal{Z} form clusters. The components of T_e connect one point from each pair to a pink cluster, while a component of T_c connects the pair. As a result, some points of F are not leaves of T .

minimality of $\lambda'(V^0)$, so we must have $t \geq |\alpha'|$, and therefore $t \geq \frac{1}{b}d_T(p, \partial\lambda'(V^0))$, as required.

Now consider the possibility that C_1 is a cluster outside of V . Let $s(V)$ denote the shadow of the union of clusters $s(\bigcup_{A \in V^0} A)$, which is the same as the hull in $\lambda(F)$ of the shadows $\{s(A) \mid A \in V^0\}$. We claim that $s(V) \cap s(C) = \emptyset$ for every $C \in \mathcal{G}^0 \setminus V^0$.

Recall from Lemmas 3.8 and 3.10 that the shadow $s(C)$ for each $C \in \mathcal{E}^0$ is disjoint from all other cluster shadows, and that the separation of shadows by $s(C)$ in $\lambda(F)$ is the same as the separation of the corresponding vertices in \mathcal{G} by $C \in \mathcal{G}^0$. In particular, if $V \in \mathcal{V}$ and $C \in \mathcal{E}^0$ then all vertices $D \in V^0$ (other than C itself if C happens to lie in V^0) have shadows $s(D)$ on one side of $s(C)$. Any $V_1 \neq V_2$ in \mathcal{V} are separated in \mathcal{G} by some $C \in \mathcal{E}^0$, including the case when C is the common vertex of V_1 and V_2 .




Figure 11: A basic, complicating example. The ambient space \mathcal{Z} is a bi-infinite line with spikes of height 1 at distance 100 from each other, which we think of as 10-hyperbolic. The cluster points \mathcal{Y} (in orange) sit on the ends of the spikes. The two points in F are very far apart. The spanning tree $\lambda(F)$ for F is a long segment. The stable tree $T(F, \mathcal{Y})$ is also abstractly a long concatenation of segments of length 102. The natural map $T(F, \mathcal{Y}) \rightarrow \mathcal{Z}$ folds the ends of these segments onto the spikes and is therefore not an embedding. It is a quasi-isometric embedding, but the multiplicative constant is at least $102/100$.

Thus the shadows $s(V_1)$ and $s(V_2)$ are either disjoint or overlap exactly on $s(C)$ for this common vertex C . The claim follows.

Now applying Lemma 3.1 again we find that $\lambda'(V^0)$ is in an $O(\epsilon)$ neighborhood of $s(V)$. Thus if C_1 is a cluster in $\mathcal{G}^0 \setminus V^0$ then any \mathcal{Z} -geodesic from p to C_1 has an $O(\epsilon)$ fellow-traveling path in $\lambda(F)$ which must exit $s(V)$ before it arrives at $s(C_1)$. It follows that $d(p, C_1) \geq d(p, C_2) - O(\epsilon)$, for some $C_2 \in V^0$. This reduces to the previous case. □

We are now ready to prove that our stable tree $T(F, \mathcal{Y})$ coarsely behaves like $\lambda(F)$. Unlike Gromov’s trees, stable trees quasi-isometrically embed with multiplicative constants possibly larger than 1; see Figure 11. This is an inconvenient fact for what follows and later in Section 4.

Proposition 3.17 *The natural map $T(F, \mathcal{Y}) \rightarrow \mathcal{Z}$ is a (K_1, K_1) -quasi-isometric embedding, and $T(F, \mathcal{Y})$ lies within Hausdorff distance K_1 of $\lambda(F)$, where $K_1 = K_1(k, \delta, \epsilon)$.*

Proof It follows from Lemma 3.1 that each component of T_c and T_e is $(1, O(\epsilon))$ -quasi-isometric to its shadow in $\lambda(F)$, and moreover is within Hausdorff distance $O(\epsilon)$ of its shadow in $\lambda(F)$.

Consider two distinct clusters $C_1, C_2 \in \mathcal{G}^0$, and their shadows. By Lemma 3.8(2), the shadow intersection $s(C_1) \cap s(C_2)$ has uniformly bounded diameter. If C_1 and C_2 belong to different pieces $V_1, V_2 \in \mathcal{V}$, then there is a cluster $D \in \mathcal{E}^0$ separating V_1 from V_2 in \mathcal{G} . If D is not equal to either C_i then its shadow separates $s(C_1)$ from $s(C_2)$ by Lemma 3.10, and hence the shadows are disjoint. If $C_1 = D$, say, then again the shadows are disjoint, by Lemma 3.8(3).

In particular, the clusters in \mathcal{E}^0 have pairwise disjoint shadows, and moreover by Lemma 3.10 their separation properties in the graph \mathcal{G} are preserved in $\lambda(F)$ — that is, if C_2 separates C_1 from C_3 in \mathcal{G} then $s(C_2)$ separates $s(C_1)$ from $s(C_3)$ in $\lambda(F)$. This means that any $V \in \mathcal{V}$ is associated to a complementary component $c(V)$ of the shadows of \mathcal{E}^0 in $\lambda(F)$, in the following way. For every $C \in \mathcal{E}_0$, any cluster in $V^0 \setminus \{C\}$ has shadow contained in one of the two components of $\lambda(F) \setminus s(C)$, by Lemma 3.10. We let $c(V)$ be the intersection of all these components. Notice that if $V \neq V'$ then $c(V) \neq c(V')$, since in that case some $C \in \mathcal{E}_0$ will separate V from V' in \mathcal{G} .

We now study overlaps of the shadows of the various relevant subtrees of T , showing that said overlaps are bounded.

Let $V \in \mathcal{V}$ be the closure of a component of $\mathcal{G} \setminus \mathcal{E}^0$. If λ_1 and λ_2 are distinct components of $\lambda'(V^0)$, we claim that their shadows in $\lambda(F)$ have an intersection of bounded diameter.

Indeed, if the shadows of λ_1 and λ_2 had overlap of size $\gg \epsilon$, then λ_1 would contain points within $O(\epsilon)$ of λ_2 , at distance $\gg \epsilon$ from each other, and with no branch point of either λ_1 or λ_2 within $O(\epsilon)$ of the geodesic in λ_1 connecting the two points (this uses the bound on the branching of T). A simple surgery would then reduce the total length of $\lambda'(V^0)$, contradicting its minimality.

Now consider a component λ_1 of $\lambda'(V^0)$ and one of the clusters C in V . We claim their shadows also have bounded-diameter intersection. Lemma 3.16 tells us that any point of λ_1 within d of $\mu(C)$ is within $O(d)$ of the boundary of λ_1 . This proves the claim.

Note that the number of clusters in V , and therefore the number of components of $\lambda'(V^0)$, is bounded via Lemma 3.13. Thus the subtree T^V comprising $\lambda'(V^0)$ together with all the components of T_c associated to clusters in V has a decomposition into a bounded number of subtrees, and a map to $\lambda(F)$ (using the shadows) which is a $(1, O(\epsilon))$ -quasi-isometric embedding on each subtree and such that the images of

distinct subtrees have bounded overlap. Under these circumstances it follows that the map $T^V \rightarrow \lambda(F)$ is a (k, k) -quasi-isometric embedding, where k depends on these bounds. Moreover, the image of this map must, up to bounded error, lie in the component of $\lambda(F)$ minus the shadows of those clusters in \mathcal{E}^0 that separate V from the rest of \mathcal{G} (by the preservation of separation properties noted above).

It follows that these maps piece together to give a (K_1, K_1) -quasi-isometry, where $K_1 = K_1(k, \delta, \epsilon)$. □

3.3 Proof of Theorem 3.3

Property (a), the branching bound on $T(F, \mathcal{Y})$, was proved in Lemma 3.16.

Property (b), the quasi-isometry, is given by Proposition 3.17.

Property (c) follows from the construction of T_e and Lemma 3.1.

Regarding property (d), we have a natural bijection between components of T_c and clusters by construction, and each component $\mu(C)$ is contained in a controlled neighborhood of the corresponding cluster C by Lemma 3.16, where by controlled we mean that the corresponding constant depends on δ, ϵ , and E . We are left to argue that C lies in a controlled neighborhood of $\mu(C)$. This is equivalent to showing that $s(C)$ lies in a controlled neighborhood of $\mu(C)$. If this was not true then, in view of the bound on the total branching of $\lambda(F)$, we would have that $s(C)$ contains an interval I in an edge of $\lambda(F)$ of length $\gg E$ not contained in a controlled neighborhood of $\mu(C)$. From Proposition 3.17 we know that T lies within controlled Hausdorff distance of $\lambda(F)$, so I is contained in a union of controlled neighborhoods of the $\mu(C')$ for $C' \neq C$, and controlled neighborhoods of the components of T_e . Neighborhoods of the latter type cannot contain points in I far from its endpoints by Lemma 3.16(3), and the same holds true for neighborhoods of the former type in view of Lemma 3.8(1), a contradiction. Therefore, $s(C)$ and C are contained in a controlled neighborhood of $\mu(C)$, as required.

We now prove the “furthermore” part of the statement. Recall that, for the reasons explained in Remark 3.4, we only treat the case that g is the identity. Let (F', \mathcal{Y}') be a second configuration differing from (F, \mathcal{Y}) as in the statement. We name the constructions arising from (F', \mathcal{Y}') by $\mathcal{C}', \mathcal{G}', \mathcal{E}', \mathcal{V}'$, etc. Also, we write $T = T(F, \mathcal{Y})$ and $T' = T(F', \mathcal{Y}')$, and similarly for T_c, T_e, T'_c , and T'_e . Set $N = \#(\mathcal{Y} \Delta \mathcal{Y}')$.

Claim 1 *The cardinality $|\mathcal{G}^0 \Delta (\mathcal{G}')^0|$ is bounded in terms of k, δ , and N .*

Proof A cluster C is in the symmetric difference $\mathcal{G}^0 \Delta (\mathcal{G}')^0$ only if it is within E of a point of $F \cup F' \cup (\mathcal{Y} \Delta \mathcal{Y}')$, of which there are at most $2k + 2 + N$. Now each point of a cluster in \mathcal{C} is within ϵ of some point in $\lambda(F)$, and there is a number R depending only on the total branching of $\lambda(F)$ such that among any R points in $\lambda(F)$ within a ball of radius $E + \epsilon$ (in $\lambda(F)$), there must be two which are less than ϵ apart (and the same is true for \mathcal{C}' and $\lambda(F')$). Thus, if there are more than $(2k + 2 + N)R$ elements in $\mathcal{G}^0 \Delta (\mathcal{G}')^0$, then two are closer than E apart, which is a contradiction. \square

Claim 2 *The symmetric difference of the edge sets of \mathcal{G} and \mathcal{G}' has cardinality bounded in terms of $k, \delta,$ and N .*

Proof By Lemma 3.13, the maximal valence of any vertex of \mathcal{G} is bounded, and so the number of edges incident to elements of $\mathcal{G}^0 \Delta (\mathcal{G}')^0$ is bounded. Therefore it suffices to consider the case where $C, D \in \mathcal{G}^0 \cap (\mathcal{G}')^0$ with $[C, D]$ an edge in \mathcal{G} but not in \mathcal{G}' . This implies there is a $B' \in (\mathcal{G}')^0 \setminus \mathcal{G}^0$ separating C from D when no such cluster in \mathcal{G}^0 did so before.

Since B' separates C from D , there is a point $q \in B' \cap ((\mathcal{Y}' \Delta \mathcal{Y}) \cup (F' \Delta F))$ which lies at distance at most ϵ from a \mathcal{Z} -geodesic γ joining C and D . The shadow $s(q)$ on $\lambda(F)$ must therefore be in a 10ϵ -neighborhood of the interval in $\lambda(F)$ between $s(C)$ and $s(D)$.

Each such q can only affect a bounded number of such edges (C, D) in this way, because the shadows of edges in each component of \mathcal{E} are arranged sequentially and disjointly along edges of $\lambda(F)$ by Lemma 3.14, and $\mathcal{G}^0 \setminus \mathcal{E}^0$ is bounded (again by Lemma 3.13). Since there are only boundedly many such q , this bounds the number of edges in the symmetric difference. \square

Let $\pi_0(\mathcal{T})$ denote the set of components of a forest \mathcal{T} . By Claim 2, there is a bound on the number of collections of clusters in $\mathcal{V} \Delta \mathcal{V}'$, and this gives us a bound, say K , on $|\pi_0(T_e) \Delta \pi_0(T'_e)|$.

Now the components of T_c and T'_c correspond to the elements of \mathcal{G}^0 and $(\mathcal{G}')^0$ respectively, whose symmetric difference is bounded by Claim 1. Moreover for each element C of $\mathcal{G}^0 \cap (\mathcal{G}')^0$, the points $r(C)$ used to determine the component of T_c (resp. T'_c) associated to C are determined by the components of T_e (resp. T'_e) adjacent to C . Thus, together with the bound on $|\pi_0(T_e) \Delta \pi_0(T'_e)|$, we obtain a bound on $|\pi_0(T_c) \Delta \pi_0(T'_c)|$.

We now increase K in a controlled way a few times, with the result of each step depending only on $k, \delta,$ and ϵ .

By item (b), we can increase K to ensure that $d_{\text{Haus}}(T, T') < K$. By Lemma 3.18 (below) we can further assume that $d_{\text{Haus}}(\mathcal{B} \cup T_c \cup F, \mathcal{B}' \cup T'_c \cup F') \leq K$, where \mathcal{B} and \mathcal{B}' are the sets of branch points of T and T' . We have to be careful in using Lemma 3.18 because the sets of leaves of T and T' need not be within bounded Hausdorff distance of each other, since they might contain more than F and F' ; see Figure 10. However, we can apply the lemma after slightly modifying T and T' by adding “spikes” (meaning edges attached at T , or T' , at one side and having a leaf on the other side) of length, say, 1 to ensure that the sets of leaves of the new trees that we obtain do lie within controlled Hausdorff distance. Such spikes only need to be added close to T_c and T'_c by Lemma 3.16(1), yielding the required Hausdorff distance estimate (indeed, $\mathcal{B} \cup T_c \cup F$ is Hausdorff close to the union of T_c and the set of branch points and leaves of the modification of T , with T_c being Hausdorff close to T'_c in view of property (d), and the latter set being close to the corresponding one for T' by Lemma 3.18).

We can then increase K once more to ensure that $T_c \cup (\mathcal{Y}' - \mathcal{Y})$ and $T'_c \cup (\mathcal{Y} - \mathcal{Y}')$ also lie at Hausdorff distance bounded by K . This can be done since T_c, \mathcal{Y} and T'_c, \mathcal{Y}' are at bounded Hausdorff distance by Lemma 3.16(2). Finally, we also require that $K > K_1$ as in Proposition 3.17.

Now let $\sigma = \sigma(K, \delta)$ be the fellow-traveling constant for $(1, K)$ -quasigeodesics with endpoints at distance at most K in a δ -hyperbolic space. This constant will be relevant later because geodesics in our trees T and T' are $(1, K)$ -quasigeodesics in \mathcal{Z} by item (b) and our choice of K .

For ease of notation, we will refer to components in $\mathcal{U} = \pi_0(T_e) \cap \pi_0(T'_e)$ as “unchanged” components, and the remaining components as “changed”. We note that there are at most K changed components in each of T_e and T'_e .

For each component $E \in \pi_0(T_e) \setminus \pi_0(T'_e)$, let $E_{\mathcal{Y}'} = \text{hull}_E(E \cap \mathcal{N}_K(\mathcal{Y}'))$, where the hull of this intersection is taken in the tree E , while the neighborhood is taken in \mathcal{Z} . Now define $\mathcal{R} = \bigcup_{E \in \pi_0(T_e) \setminus \pi_0(T'_e)} E_{\mathcal{Y}'}$, and define \mathcal{R}' similarly. We now collect the “unstable parts” of the trees along with the unchanged parts. Set

$$U = T_c \cup \mathcal{R} \cup \mathcal{U} \cup \mathcal{B} \cup F \quad \text{and} \quad U' = T'_c \cup \mathcal{R}' \cup \mathcal{U} \cup \mathcal{B}' \cup F'.$$

Note that we included F and F' in these sets to ensure that they lie within bounded Hausdorff distance of each other, but this is inconsequential for the purposes of considering the complementary forests, which is what we want to do next.

Let $K' = K'(K, N, k) \geq 10K$ be such that

- (i) $d_{\text{Haus}}(T, T') < K$ and $d_{\text{Haus}}(U, U') < 10K$,
- (ii) $|\#\pi_0(T \setminus U)| < K'$ and $|\#\pi_0(T' \setminus U')| < K'$,
- (iii) $(T \setminus U) \subset T_e$ and $(T' \setminus U') \subset T'_e$,
- (iv) $|\#\pi_0(T_e \cap U) \setminus \mathcal{U}| < K'$ and $|\#\pi_0(T'_e \cap U') \setminus \mathcal{U}'| < K'$, and
- (v) for any $C \in (\pi_0(T_e \cap U) \cup \pi_0(T'_e \cap U')) \setminus \mathcal{U}$, we have $\text{diam } C < K'$.

Regarding property (i), the “10” in “10K” is there to keep into account the fact that we took hulls.

Property (ii) can be shown observing that the intersection of U with each changed component can only have a bounded number of components because of the bound N on $|\mathcal{Y} \triangle \mathcal{Y}'|$ and the bound on $|\mathcal{B}|$ given by Lemma 3.16. This same observation shows item (iv). Property (iii) holds by construction.

Property (v) is nontrivial for \mathcal{R} and \mathcal{R}' , in which case it holds since C is a union of boundedly many components, each of bounded diameter. In particular, any set $E_{\mathcal{Y}}$ is a union of hulls in E of intersections with balls centered at an element of \mathcal{Y}' , which have bounded diameter, and the union consists of at most N elements.

Now let $L_1 = L_1(K', \delta)$, $L_2 = L_2(K', \delta) > 0$ be the constants given by Lemma 3.19 (below) with T, U and T', U' satisfying the conditions of that lemma via (i) and (ii).

Let $\mathcal{L} \subset \pi_0(T \setminus U)$ be the set of all components of $T \setminus U$ of diameter greater than L_1 .

Define $\mathcal{L}' \subset \pi_0(T' \setminus U')$ to be the set of components of $T' \setminus U'$ that lie within Hausdorff distance L_2 of an element of \mathcal{L} . Since there are at most K' components of $T \setminus U$, this bounds the cardinality of $\pi_0(T \setminus U) \setminus \mathcal{L}$, and any component in this set has diameter at most L_1 . Similarly, the numbers and diameters of the components of $(T' \setminus U')$ not in \mathcal{L}' are also bounded, this time the bound on the diameter being L_3 by the “moreover” part of Lemma 3.19.

Lemma 3.19 provides a bijection $\rho: \mathcal{L} \rightarrow \mathcal{L}'$ which sends any component in \mathcal{L} to the unique component in \mathcal{L}' which is within Hausdorff distance L_2 . That is, for any $C \in \pi_0(\mathcal{L})$, we have $d_{\text{Haus}}(C, \rho(C)) < L_2$.

Now let $T_{s,0}$ be the union of all elements of $\mathcal{L} \cup \mathcal{U}$, and define $T'_{s,0}$ similarly. We observe that we have by construction and (iii) above that $T_{s,0} \subset T_e$ and $T'_{s,0} \subset T'_e$. Moreover, the number of components of $T_e \setminus T_{s,0}$ and their diameters are bounded by

$2K'(L_1 + 10L_2 + K' + 1)$, and similarly for $T'_e \setminus T'_{s,0}$. In fact, the number of such components is bounded by $2K'$, since each is a union of

- changed components of $T_e \cap U$, and there are at most K' of those by (iv), and
- components of $T \setminus U$ of diameter at most L_1 , and again there are at most K' of those.

The bound on the diameter also follows from this description.

To obtain the sets T_s and T'_s required by the theorem, it suffices now to remove the branch points from the unchanged components contained in $T_{s,0}$ and $T'_{s,0}$ to ensure (1) (which at this point is not satisfied only because of the unchanged components, since we included the branch points in U and U'), while all other properties have been checked above. \square

Two supporting lemmas The following two lemmas were used in the proof of Theorem 3.3 above. To simplify notation, we will not distinguish between a tree quasi-isometrically embedded in a metric space, and the image of said tree.

Lemma 3.18 *For each K and δ there exists L_0 such that the following holds. Let T and T' be trees (K, K) -quasi-isometrically embedded in the δ -hyperbolic metric space \mathcal{Z} , with $d_{\text{Haus}}(F_0, F'_0) \leq K$, where F_0 and F'_0 are the sets of leaves of T and T' respectively. Then the sets of branch points \mathcal{B} and \mathcal{B}' of T and T' satisfy $d_{\text{Haus}}(\mathcal{B} \cup F_0, \mathcal{B}' \cup F'_0) \leq L_0$.*

Proof The set $\mathcal{B} \cup F_0$ can be coarsely characterized as the set of points x of T such that there are f_1, f_2 and f_3 in F_0 (not necessarily distinct) with the property that the Gromov product at x between any f_i and f_j is small, and similarly for T' . We leave the details to the reader. \square

Lemma 3.19 *For each K there exist L_1, L_2 and L_3 such that the following holds. Let T and T' be trees (K, K) -quasi-isometrically embedded in the metric space \mathcal{Z} , with $d_{\text{Haus}}(T, T') \leq K$. Also, let $U \subseteq T$ and $U' \subseteq T'$ be subforests such that $d_{\text{Haus}}(U, U') \leq K$ and all branch points of T (resp. T') are contained in U (resp. U').*

Then for each component C of $T \setminus U$ of diameter at least L_1 , there exists a unique component C' of $T' \setminus U'$ within Hausdorff distance L_2 of C . Moreover, every component C' of $T' \setminus U'$ of diameter at least L_3 arises in this way.

Proof We will conflate components of $T \setminus U$ with their closures, so we can talk about their leaves, and similarly for $T' \setminus U'$.

The main observation is that there exists $K_2 = K_2(K)$ such that the following holds. Let C be a component of $T \setminus U$ and let x and y be (not necessarily distinct) points in C that, in the metric of C , are at least K_2 from all the leaves of C . Then there exists a unique component C' of $T' \setminus U'$ which is within K of both x and y .

To prove this, suppose by contradiction that $K_2 \gg K$ and that there are distinct components of $T' \setminus U'$ that contain points x' and y' that are within K of x and y , respectively. Then there exists some $p' \in U'$ on the geodesic $[x', y']$ in T' from x' to y' . Let $p \in U$ be such that $d(p, p') < K$. Then p lies within $\sigma = \sigma(K) > 0$ from the geodesic from x to y , since considering points in T that are within K of those along $[x', y']$ yields a quasigeodesic in T . Since x and y are at least K_2 from the leaves of C , they cannot lie close to any point of U , in particular p . We can then deduce that either p lies along the geodesic $[x, y]$ in T from x to y , or there is a branch point of T along $[x, y]$. In either case, x and y do not lie in the same component of $T \setminus U$, a contradiction. (Recall that U contains all branch points of T by hypothesis.)

Consider now a component C of $T \setminus U$ of diameter sufficiently large that it contains a point which is at least K_2 from all the leaves of C . By the observation above, all such points are close to a unique component C' of T' , and since the set of all such points has bounded Hausdorff distance from C , we have that C is contained in a uniform neighborhood of C' . Moreover, if C has sufficiently large diameter, then we can apply the same reasoning to C' and deduce that C' contains points that are within K of a unique component C'' of $T \setminus U$, and that C' is contained in a uniform neighborhood of C'' . But the above observation implies that $C'' = C$, and it follows that C and C' lie within uniformly bounded Hausdorff distance.

Finally, the “moreover” part follows from a similar back-and-forth using the previous part of the statement. Namely C' has sufficiently large diameter and is within bounded Hausdorff distance of a component C of $T \setminus U$, then C also has large diameter and is thus in turn within bounded Hausdorff distance of some component of $T' \setminus U'$, which needs to be C' . □

4 Stable cubulations

Fix a G -colorable HHS $(\mathcal{X}, \mathfrak{S})$ for $G < \text{Aut}(\mathfrak{S})$, and let $F \subset \mathcal{X}$ be a finite set.

In this section, we use the stable trees constructed in Section 3 for the projections of F to the relevant domains to define a wallspace on the hull $H_\theta(F)$. This wallspace

can then be plugged into Sageev’s machine to produce a cube complex which, by an argument from [8], coarsely models the hull of F in \mathcal{X} . Stability of the tree construction then induces stability in the cubulations under perturbations of F . We refer the reader to Section 2.1 for some background and references on cube complexes, wallspaces, and hyperplane deletions, as well as to Section 2.2 for background on HHSs.

The main result of this section, and the only statement from this section that we will use in the rest of the paper, is the following precise version of Theorem A.

Theorem 4.1 *Let $(\mathcal{X}, \mathfrak{S})$ be a G -colorable HHS for $G < \text{Aut}(\mathfrak{S})$. Then for each k there exist K and N with the following properties. To each subset $F \subseteq \mathcal{X}$ of cardinality at most k one can assign a triple $(\mathcal{Q}_F, \Phi_F, \psi_F)$ satisfying*

- (1) \mathcal{Q}_F is a CAT(0) cube complex of dimension at most the maximal number of pairwise orthogonal domains of $(\mathcal{X}, \mathfrak{S})$;
- (2) $\Phi_F : \mathcal{Q}_F \rightarrow H_\theta(F)$ is a K -median (K, K) -quasi-isometry;
- (3) $\psi_F : F \rightarrow (\mathcal{Q}_F)^{(0)}$ satisfies $d_{\mathcal{X}}(\Phi_F \circ \psi_F(f), f) \leq K$ for each $f \in F$.

Moreover, suppose that $F' \subseteq \mathcal{X}$ is another subset of cardinality at most k , $g \in G$, and $d_{\text{Haus}}(gF, F') \leq 1$. Choose any map $\iota_F : F \sqcup F' \rightarrow F$ such that $\iota_F(f) = f$ if $f \in F$ and $d_{\mathcal{X}}(g(\iota_F(f)), f) \leq 1$ if $f \in F'$. Also, choose a map $\iota_{F'} : F \sqcup F' \rightarrow F'$ such that $\iota_{F'}(f) = f$ if $f \in F'$ and $d_{\mathcal{X}}(g(f), \iota_{F'}(f)) \leq 1$ if $f \in F$. Then the following holds:

There is a third CAT(0) cube complex \mathcal{Q}_0 and K -median (K, K) -quasi-isometric embedding Φ_0 such that the diagram

$$(4-1) \quad \begin{array}{ccccc} & & F & \xrightarrow{\psi_F} & \mathcal{Q}_F \\ & \nearrow \iota_F & & & \searrow g \circ \Phi_F \\ & & F \sqcup F' & & \mathcal{Q}_0 \\ & \searrow \iota_{F'} & & & \nearrow \Phi_0 \\ & & F' & \xrightarrow{\psi_{F'}} & \mathcal{Q}_{F'} \end{array}$$

commutes up to error at most K , where η and η' are hyperplane deletion maps that delete at most N hyperplanes. The left side commutes exactly, ie $\eta \circ \psi_F \circ \iota_F = \eta' \circ \psi_{F'} \circ \iota_{F'}$.

The notion of K -median refers to the coarse median structure on \mathcal{X} in the sense of [15], and we only define it later where it is needed, since we obtain it directly from [8].

The “1” in “ $d_{\text{Haus}}(gF, F') \leq 1$ ” could be changed with any other constant up to changing K in the “moreover” part. We decided to keep the statement slightly simpler and not introduce further quantifiers since one can always replace an HHS with a graph, or rescale the metric.

The remainder of this section is devoted to the proof of this theorem.

Standing assumptions For this section, we fix an HHS $(\mathcal{X}, \mathfrak{S})$ as in Theorem 4.1, which, in light of Theorem 2.9, we can assume to have stable projections. Furthermore, we can assume that it has the property that all $\pi_V(x)$ and all ρ_V^U for $U \sqsubset V$ are single points; see Remark 2.16.

4.1 Subdivision sets for stable trees

In this section, we establish a formalism for subdividing trees. In Section 4.2 these subdivision points will give the walls in our cubulation. This mostly follows the strategy of [8], except that we need to take greater care in making choices for the subdivision.

Definition 4.2 Let $M' > M > 0$. An (M, M') -subdivision of a tree T is a collection of points $\mathbf{p}(T) \subset T$ satisfying:

- (1) The points $\mathbf{p}(T)$ are contained in the interiors of edges of T . We set

$$\mathbf{p}(e) = \mathbf{p}(T) \cap \text{int}(e)$$

for each edge e of T .

- (2) The $M'/2$ -neighborhood of $\mathbf{p}(e) \cup \partial e$ covers e .
- (3) All points of $\mathbf{p}(e) \cup \partial e$ are at least M apart in e .

In other words, the spacing between points of $\mathbf{p}(e) \cup \partial e$ along e is at least M and strictly less than M' .

We additionally say that $\mathbf{p}(T)$ is (M, M') -evenly spaced if $M' \geq 8M$ and the spacing between successive points of $\mathbf{p}(e)$ is exactly M for each edge e .

We will specify M and M' later, though for now it suffices to assume they are large relative to the various HHS constants.

We fix, once and for all, a subdivision operator $\wp_{M, M'}$ which to any tree T associates a fixed (M, M') -evenly spaced subdivision $\wp_{M, M'}(T)$ of (the edges of) T . Often the constants M and M' are fixed, and we simply write $\wp(T)$. Similarly we can define $\wp_{M, M'}$ on a forest as the union of subdivisions $\wp_{M, M'}$ on its components.

We now explain how to associate a collection of subdivisions on stable trees to a set of points $F \subset \mathcal{X}$.

Fix some large K (depending on M , to be specified later) and write $\mathcal{U}(F) = \text{Rel}_K(F)$ for any finite set $F \subset \mathcal{X}$.

For any $V \in \mathfrak{S}$, define

$$\mathcal{Y}^V = \{\rho_V^W \mid W \in \mathcal{U}(F), W \sqsupseteq V\} \quad \text{and} \quad F^V = \pi_V(F).$$

Note that whenever K is larger than the bounded geodesic image constant κ_0 , we have that \mathcal{Y}^V is contained in the κ_0 -neighborhood of the hull of F^V in $\mathcal{C}(V)$. This ensures that \mathcal{Y}^V satisfies the requirements of Theorem 3.3 for any $\epsilon \geq 2\kappa_0$. We fix such ϵ as in Theorem 3.3, and we will always apply that theorem with this ϵ .

Note that if $V \notin \mathcal{U}(F)$ then $\mathcal{Y}^V \cup F^V$ has uniformly bounded diameter.

Let $V \in \mathcal{U}(F)$. From V , we get corresponding sets of projections F^V and \mathcal{Y}^V in $\mathcal{C}(V)$, which we may consider independently of the set F . Doing so, we obtain from Theorem 3.3 a fixed stable tree

$$T_F^V := T(F^V, \mathcal{Y}^V)$$

and we denote its decomposition by $T_c^V \cup T_e^V$.

Let $M' > 8M > 0$ be subdivision constants (to be specified later). Applying $\wp_{M, M'}$ to each forest T_e^V in T_F^V , we call the resulting subdivision $\wp^V(F)$.

Remark 4.3 We emphasize that the distance between the subdivision points is measured in the various trees themselves, not in the corresponding $\mathcal{C}(V)$ where the trees quasi-isometrically embed. This is something that will require care throughout this section, but the fact that the trees are quasi-isometrically embedded in the corresponding $\mathcal{C}(V)$ will ensure that no real issues arise.

The disjoint union over all $V \in \mathcal{U}(F)$ gives us

$$\wp \equiv \wp(F) = \bigsqcup_{V \in \mathcal{U}(F)} \wp^V(F).$$

More generally, we will consider (M, M') subdivisions of the forests T_e^V which are not necessarily obtained from our subdivision operator $\wp_{M, M'}$, and in particular may not be evenly spaced (they will arise by taking subsets). In this case if we name the full configuration \mathbf{p} , we will again denote the restriction to T_e^V as $\mathbf{p}_V(F)$.




Figure 12: An evenly spaced subdivision of the stable tree T_F^V . Note that subdivision points are far away from leaves, components of T_c , and branch points.

Remark 4.4 Our choice of constants $M' > 8M$ gives that any point $p \in \wp^V(F)$ is at least distance $4M$ from any point $\rho_V^W \in \mathcal{Y}^V$ for $W \sqsubset V$ or leaf $f \in F^V$.

4.2 The cube complex

Following the scheme of [8], given a union \mathbf{p} of (M', M) subdivisions as above, we now describe a cube complex $\mathcal{Q}_{F,\mathbf{p}}$ and a map

$$\Phi_{F,\mathbf{p}}: \mathcal{Q}_{F,\mathbf{p}} \rightarrow \mathcal{X}$$

which turns out to be a quasi-isometric embedding whose image is within bounded Hausdorff distance of $H_\theta(F)$, when the constants M, M' and K are chosen suitably (see below).

To build the cube complex $\mathcal{Q}_{F,\mathbf{p}}$, we build a wallspace structure on the hull $H_\theta(F)$ in which each $p \in \mathbf{p}$ corresponds to a wall, ie a partition of $H_\theta(F)$ into two sets.

For each $V \in \mathfrak{S}$, we let $\beta_F^V: \mathcal{C}(V) \rightarrow T_F^V$ be, roughly, a closest-point projection map to T_F^V , and more precisely any fixed map such that $d_V(x, \Xi(\beta_F^V(x)))$ is minimal for

all $x \in \mathcal{C}(V)$, where we recall that Ξ is the quasi-isometric embedding of T_F^V in $\mathcal{C}(V)$. With slight abuse of notation, for $x \in \mathcal{X}$ we write $\beta_F^V(x)$ for $\beta_F^V(\pi_V(x))$.

Remark 4.5 To save notation, we will often omit the map Ξ , thereby identifying points and subsets of some T_F^V with their image in $\mathcal{C}(V)$ (as one often does with quasi-geodesics). For example, for $p \in \beta^V(F)$ and $x \in \mathcal{C}(V)$ we will write $d_V(\beta_F^V(p), x)$ rather than $d_V(\Xi(\beta_F^V(p)), x)$.

Given $p \in \mathbf{p}_V(F)$, let $T_{p,+}^V$ denote one of the components of $T_F^V \setminus \{p\}$, and let $T_{p,-}^V$ be the union of the other component and $\{p\}$ (we arbitrarily choose the first component). Let

$$W_{p,\pm}^V = (\beta_F^V)^{-1}(T_{p,\pm}^V) \cap H_\theta(F).$$

Note that $W_{p,+}^V$ and $W_{p,-}^V$ form a partition of $H_\theta(F)$. Let $\mathcal{L}_p^V = \{W_{p,+}^V, W_{p,-}^V\}$ be the wall associated to p . We call $T_{p,\pm}^V$ the *half-trees* associated to the wall \mathcal{L}_p^V .

Let $\mathcal{Q} = \mathcal{Q}_{F,\mathbf{p}}$ be the CAT(0) cube complex dual to the wallspace $\{\mathcal{L}_p^V\}$.

To define $\Phi_F = \Phi_{F,\mathbf{p}}: \mathcal{Y} \rightarrow \mathcal{X}$, note that it suffices to define Φ_F on the 0-skeleton of \mathcal{Y} . Let $x \in \mathcal{Y}^{(0)}$; we view x as a coherent orientation of the walls \mathcal{L}_p^V ; see Section 2.1.1. That is, for each $p \in \mathbf{p}$, we have $x(p)$ equal to one of $W_{p,+}^V$ or $W_{p,-}^V$.

Coarsely, we would like to define $\Phi(x)$ by

$$\Phi(x) \approx \bigcap_{p \in \mathbf{p}} x(p).$$

This is done in [8] by considering the projections of $x(p)$ to the factor trees T_F^Y . That is, we set

$$S_{p,Y}(x) = \text{hull}(\beta_F^Y(x(p))),$$

where $Y \in \mathfrak{S}$ and hull denotes convex hull in the tree T_F^Y . Note that $p \in \mathbf{p}_V(F)$ for some V , which is typically different from Y .

We now define the intersection

$$(4-2) \quad b_Y = b_Y(x) = \bigcap_{p \in \mathbf{p}} S_{p,Y}(x).$$

These b_Y will serve as (coarse) coordinates for the map Φ , as given by the theorem that we state below, after we explain how it can be extracted from [8].

Remarks on the construction in [8] We now summarize various results proven in [8] regarding cubulations of hulls in HHSs. We note that the construction in [8] of the CAT(0) cube complexes approximating hulls is the same as the one we just described,

with one difference. The difference is that, rather than using T_F^V to approximate the hull of $\pi_V(F)$ in $\mathcal{C}(V)$, in [8] the authors use any choice of tree contained in $\mathcal{C}(V)$ that uniformly approximates the hull and is quasi-isometrically embedded with multiplicative constant 1. This choice is not due to the fact that the multiplicative constant being 1 is needed for the arguments, but more simply due to the fact that one such tree exists, and so it is more convenient to use it. For our purposes, we have to be more careful in the choice of the tree, and as a result we cannot guarantee that the multiplicative constant is 1 with our construction (recall Figure 11). However, this does not affect the arguments of [8], except that the subdivision constant M has to be chosen large compared to the quasi-isometric embedding constants, so that the subdivision points are sufficiently far apart in the various hyperbolic spaces $\mathcal{C}(V)$.

Another remark about the statement below is that the constant dependencies that we give below are not explicitly stated in [8], but can be recovered as follows. In [8, Section 2], the constant M is chosen large compared to various HHS constants and $k = |F|$, so that the construction has all the stated properties for any sufficiently large M . Regarding M' , in [8, Section 2.1] it is taken to be $10Mk$, as can be seen from point (4) of the construction of the walls. The reason for the constant $10Mk$ is that one can choose subdivision points that make the diameter of the complementary components at most that quantity, but with any other bound one would obtain the same properties (eg that the CAT(0) cube complex quasi-isometrically embeds in the HHS), with different constants. Regarding K , in [8, Section 2.1] it is chosen to be $100Mk$, and similar remarks to those regarding M' apply.

Properties of the cubulation With this in mind, we now state various results about the construction we explained above, and point out where the arguments for those can be found in [8].

Theorem 4.6 *Given an HHS $(\mathcal{X}, \mathfrak{S})$ and an integer k , there exist $M_0 \geq 1$ and functions $M'_0: \mathbb{R} \rightarrow \mathbb{R}$ and $K_0: \mathbb{R} \rightarrow \mathbb{R}$ with the following property. Whenever $M \geq M_0$, $M' \geq M'_0(M)$, and $K \geq K_0(M)$, there exists ξ such that for every $F \subseteq \mathcal{X}$ with $|F| \leq k$, the following hold:*

- (1) [8, Lemma 2.6, paragraph “Definition of p_A ” in proof of Theorem 2.1] For every $x \in \mathcal{Q}_{F, \mathbf{p}}^{(0)}$, $b_Y(x)$ is nonempty and $\text{diam}_{\mathcal{C}(Y)}(b_Y(x)) \leq \xi$.
- (2) [8, Lemma 2.7, paragraph “Definition of p_A ” in proof of Theorem 2.1] For every $x \in \mathcal{Q}_{F, \mathbf{p}}^{(0)}$, there exists a point in $H_\theta(F)$, denoted by $\Phi(x)$, whose projections to all $\mathcal{C}(Y)$ are within distance ξ of $b_Y(x)$.

- (3) [8, Theorem 2.1] Φ is a ξ -median (ξ, ξ) -quasi-isometry to $H_\theta(F)$.
- (4) [8, Theorem 2.1] The dimension of $\mathcal{Q}_{F, \mathbf{p}}$ is bounded by the maximal number of pairwise orthogonal domains in \mathfrak{S} .

We will also need some more technical properties of the trees T_F^V and projections ρ_V^U related to the HHS consistency axioms.

Proposition 4.7 *Given an HHS $(\mathcal{X}, \mathfrak{S})$ and an integer k , there exists κ such that given K sufficiently large (depending only on $(\mathcal{X}, \mathfrak{S})$):*

- (1) [8, Lemma 2.3] *If $U, V \in \mathcal{U}(F)$ and $U \pitchfork V$, then ρ_V^U lies κ -close in $\mathcal{C}(V)$ to a point of $\pi_V(F)$.*
- (2) [8, Lemma 2.5] *If $U, V \in \mathcal{U}(F)$, $V \sqsubseteq U$, and $q \in \mathbf{p}(U)$, then $\rho_V^U(q)$ lies κ -close in $\mathcal{C}(V)$ to a point of $\pi_V(F)$.*

Remark 4.8 In the rest of this section, whenever we use constants M, M' and K we will assume that they are chosen as in Theorem 4.6, and so that all supporting lemmas in [8] apply. Moreover, we will impose further requirements as needed. We note that the role of K is often hidden in the statements, since it affects the set $\mathcal{U}(F)$, which in the various statements often only plays a role implicitly.

4.3 Deleting subdivision points

Now we consider how the construction of $\mathcal{Q}_{F, \mathbf{p}}$ is affected by the deletion of points in \mathbf{p} . If $\mathbf{p}_0 \subset \mathbf{p}$, there is a hyperplane-deletion map $h: \mathcal{Q}_{F, \mathbf{p}} \rightarrow \mathcal{Q}_{F, \mathbf{p}_0}$. That is, for $x \in \mathcal{Q}_{F, \mathbf{p}}$, the image $h(x) \in \mathcal{Q}_{F, \mathbf{p}_0}$ is just the orientation on the remaining walls: $h(x)(p) = x(p)$ for $p \in \mathbf{p}_0$. We note that the subdivisions in the following proposition need not be evenly spaced.

Proposition 4.9 *For every k and n , and M, M' and K as in Theorem 4.6, there exist K' and M'' such that if F has cardinality at most k , \mathbf{p} is an (M, M') -subdivision, and $\mathbf{p}_0 \subset \mathbf{p}$ satisfies $|\mathbf{p} \setminus \mathbf{p}_0| \leq n$, then \mathbf{p}_0 is an (M, M'') -subdivision satisfying the conclusions of Theorem 4.6, and the diagram*

$$(4-3) \quad \begin{array}{ccc} \mathcal{Q}_{F, \mathbf{p}} & & \\ h \downarrow & \searrow \Phi_{F, \mathbf{p}} & \\ \mathcal{Q}_{F, \mathbf{p}_0} & \xrightarrow{\Phi_{F, \mathbf{p}_0}} & \mathcal{X} \end{array}$$

commutes up to error K' .

Proof As above, the map $\Phi_{F,\mathbf{p}}(x)$ is determined by the coordinates

$$b_U = \bigcap_{p \in \mathbf{p}} S_{p,U}(x),$$

whereas $\Phi_{F,\mathbf{p}_0}(h(x))$ is determined by

$$b_{0,U} = \bigcap_{p \in \mathbf{p}_0} S_{p,U}(x).$$

Note that $b_U \subset b_{0,U}$, and in the other direction the diameter of $b_{0,U}$ is bounded by Theorem 4.6(1) because the new set \mathbf{p}_0 is an (M, M'') -subdivision for some $M'' \geq M'$ (and in particular $M'' \geq M'_0(M)$, so that the conclusion of Theorem 4.6 holds for \mathbf{p}_0) depending only on M, M' and the number n of deletions.

The bound on $d(\Phi_{F,\mathbf{p}}, \Phi_{F,\mathbf{p}_0} \circ h)$ then follows from Theorem 2.6 (the distance formula) since the coordinates of $\Phi_{F,\mathbf{p}_0}(h(x))$ and $\Phi_{F,\mathbf{p}}(x)$ coarsely coincide with b_U and $b_{0,U}$. □

4.4 Intersection conditions

Recall that Lemma 2.3 explains how a bijection between halfspaces that preserves intersection properties induces an isomorphism of the corresponding cube complexes. In view of this, we are interested in knowing when two of our halfspaces intersect.

We fix the setup of Section 4.2. The next lemma is the main technical support for Proposition 4.13.

Lemma 4.10 *There exists M_1 , depending on $(\mathcal{X}, \mathfrak{S})$ and $|F|$, such that the following holds. Let \mathbf{p} be an (M, M') -subdivision with $M \geq M_1$. Consider two halfspaces, $W_{p,\sigma}^V$ and $W_{q,\tau}^Z$, with associated half-trees $T_{p,\sigma}^V$ and $T_{q,\tau}^Z$, where $p \in \mathbf{p}_V(F)$, $q \in \mathbf{p}_Z(F)$ and $\tau, \sigma \in \{\pm\}$.*

Then $W_{p,\sigma}^V$ and $W_{q,\tau}^Z$ intersect if and only if one of the following holds, up to switching the roles of the half-spaces:

- (1) $V \perp Z$.
- (2) $V = Z$, and $T_{p,\sigma}^V \cap T_{q,\tau}^Z \neq \emptyset$.
- (3) $V \pitchfork Z$, and $T_{q,\tau}^Z$ contains $\beta_F^Z(\rho_V^V)$.
- (4) $V \sqsubset Z$, and $T_{q,\tau}^Z$ contains $\beta_F^Z(\rho_V^V)$.
- (5) $V \sqsubset Z$, and $T_{p,\sigma}^V$ contains $\beta_F^V(\rho_V^Z(q))$.




Figure 13: The first possibility in case (2), when $V = Z$, where p and q choose different ends of T_F^V , and $\text{diam}(T_{p,\sigma}^V \cap T_{q,\tau}^V) \geq d_{T_F^V}(p, q) \geq M$.

We note that the last three cases of the lemma boil down to the consistency inequalities for HHSs (and could even be seen as interpretations thereof).

Proof In an effort to enhance readability of the proof, we will make coarse, comparative arguments which keep track of dependencies of constants and their relative size, instead of precise quantities.

It follows from Theorem 2.7 and the fact that the various T_F^V are quasi-isometrically embedded (Theorem 3.3) that the image $\beta_F^V(H_\theta(F))$ is D -dense in the tree T_F^V (with respect to its path metric), where D only depends on the HHS structure and $|F|$. We may assume that M has been chosen greater than $10D$. Moreover, throughout the proof we will further specify conditions on M , requiring it to be suitably larger than other constants appearing in the argument. It is important to notice that, in each case, these constants depend only on $(\mathcal{X}, \mathfrak{S})$ and $|F|$. We also remark that we need to be careful when comparing distances in the trees T_F^V and the ambient spaces $\mathcal{C}(V)$.

Case (1) Suppose $V \perp Z$. The fact that in this case all pairs of halfspaces intersect is [8, Lemma 2.13].

Case (2) Suppose $V = Z$. If the halfspaces intersect, then any point x in their intersection has $\beta_F^V(x) \in T_{p,\sigma}^V \cap T_{q,\tau}^V$. Conversely, suppose $T_{p,\sigma}^V \cap T_{q,\tau}^V \neq \emptyset$. If the intersection contains both p and q , then

$$\text{diam}_{T_F^V}(T_{p,\sigma}^V \cap T_{q,\tau}^V) \geq d_{T_F^V}(p, q) \geq M.$$




Figure 14: The other possibility in case (2), when $V = Z$, where p and q now choose the same ends of T_F^V .

If not, then the intersection contains all of $T_{p,\sigma}^V$ or all of $T_{q,\tau}^V$, and again

$$\text{diam}_{T_F^V}(T_{p,\sigma}^V \cap T_{q,\tau}^V) \geq M.$$

Since $\beta_F^V(H_\theta(F))$ is D -dense in T_F^V and $M > D$, we can then find $x \in H_\theta(F)$ with $\beta_F^V(x) \in T_{p,\sigma}^V \cap T_{q,\tau}^V$, so that $x \in W_p^V \cap W_q^Z$, as required.

Case (3) Suppose $V \pitchfork Z$. If the halfspaces $W_{p,\sigma}^V$ and $W_{q,\tau}^Z$ intersect, then any point x in their intersection has $\beta_F^V(x) \in T_{p,\sigma}^V$ and $\beta_F^Z(x) \in T_{q,\tau}^Z$. We claim that

$$\beta_F^Z(\rho_Z^V) \notin T_{q,\tau}^Z \implies d_Z(\rho_Z^V, \pi_Z(x)) > \kappa_0,$$

where κ_0 is the constant in the transverse consistency inequality (2-1), and similarly with the roles of Z and V interchanged. Indeed, recall that Proposition 4.7(1) says that (assuming M is sufficiently large) ρ_Z^V lies within κ of a leaf of T_F^Z , where κ only depends on the HHS structure and $|F|$. Since q is at least M away from the leaves, and we can assume that M is sufficiently large compared to κ and the quasi-isometric embedding constants of T_F^Z , we see that if $\beta_F^Z(\rho_Z^V)$ is not in $T_{q,\tau}^Z$ then it must be at least $M/2$ from it (as measured in the metric of T_F^Z). Thus $\beta_F^Z(\rho_Z^V)$ is at least $M/2$ from $\beta_F^Z(x)$ in T_F^Z , and since we can assume that M is sufficiently large compared to κ , the quasi-isometric embedding constants of T_F^Z , and the distance between $\pi_Z(x)$ and (the image in $C(Z)$ of) T_F^Z , we have the desired inequality. The same holds with V and Z interchanged.

However, the transverse consistency inequality (2-1) says that we cannot have both $d_Z(\rho_Z^V, \pi_Z(x)) > \kappa_0$ and $d_V(\rho_V^Z, \pi_V(x)) > \kappa_0$. Hence one of $\beta_F^Z(\rho_Z^V) \in T_{q,\tau}^Z$ or $\beta_F^V(\rho_V^Z) \in T_{p,\sigma}^V$ must hold, which is what we wanted to prove.




Figure 15: Case (3), when $V \pitchfork Z$.

Conversely, suppose that $T_{q,\tau}^Z$ contains $\beta_F^Z(\rho_Z^V)$. Because $\beta_F^V(H_\theta(F))$ is D -dense in T_F^V , p is far from any leaf of T_F^Z , and $\beta_F^Z(\rho_Z^V)$ is close to a leaf (Proposition 4.7(1)), we may choose $x \in H_\theta(F)$ with $\beta_F^V(x) \in T_{p,\sigma}^V$ and $d_{T_F^V}(\beta_F^V(x), \beta_F^V(\rho_V^Z)) > M/2$ (again, with distance in T_F^V).

Then $x \in W_{p,\sigma}^V$ by construction, and we claim also $x \in W_{q,\tau}^Z$. Indeed, for M sufficiently large, $d_V(x, \rho_V^Z) > \kappa_0$, so transverse consistency (2-1) implies that $d_Z(\pi_Z(x), \rho_Z^V) < \kappa_0$. Again for M sufficiently large, this gives $d_{T_F^Z}(\beta_F^Z(x), \beta_F^Z(\rho_Z^V)) \leq M/2$, and since the $M/2$ -neighborhood in T_F^Z of $\beta_F^Z(\rho_Z^V)$ is contained in $T_{q,\tau}^Z$, we conclude $\beta_F^Z(x) \in T_{q,\tau}^Z$.

Cases (4) and (5) Suppose $V \sqsubset Z$. Cases (4) and (5) are similar to the previous case, instead using the nested consistency inequality (2-2).

Let $x \in W_{p,\sigma}^V \cap W_{q,\tau}^Z$. Because the partition points in T_F^Z are M away from ρ_Z^V by assumption, we have again that either

- $T_{q,\tau}^Z$ contains the $M/2$ -neighborhood of $\beta_F^Z(\rho_Z^V)$, or
- $d_{T_F^Z}(T_{q,\tau}^Z, \beta_F^Z(\rho_Z^V)) > M/2$.

In the first case, we are done (obtaining case (4)). In the second case, for M sufficiently large, $d_V(\pi_V(x), \rho_V^Z(\pi_Z(x))) < \kappa_0$ by the nested consistency inequality (2-2). Moreover, for M sufficiently large, there is a geodesic (in $C(Z)$) from q to $\pi_Z(x)$ which is farther than κ_0 from ρ_Z^V , where κ_0 is the constant in the bounded geodesic image property (axiom (7) in Definition 2.4). In fact, the distance from $\pi_Z(x)$ to $T_{q,\tau}^Z$




Figure 16: Case (4), when $V \sqsubset Z$ and $T_{q,\tau}^Z$ contains $\beta_F^Z(\rho_Z^V)$.

is bounded in terms of $(\mathcal{X}, \mathfrak{S})$ and $|F|$, as is the quasiconvexity constant (of the image in $\mathcal{C}(Z)$) of $T_{q,\tau}^Z$, and thus so is the distance of any point along this geodesic from $T_{q,\tau}^Z$. On the other hand, we can estimate $d_Z(T_{q,\tau}^Z, \rho_Z^V)$ in terms of M and the quasi-isometric embedding constants of T_F^Z (which is independent of M). Hence by choosing M large enough, we may use the bounded geodesic image property to conclude that $d_V(\rho_V^Z(\pi_Z(x)), \rho_V^Z(q)) < \kappa_0$. Since $\pi_V(x)$ is within some distance of $T_{p,\sigma}^V$ which is bounded in terms of $(\mathcal{X}, \mathfrak{S})$ and $|F|$, we see that $d_V(\beta_F^V(\rho_V^Z(q)), T_{p,\sigma}^V) < M/2$ for M sufficiently large. However, this implies that $\beta_F^V(\rho_V^Z(q))$ is in fact contained in $T_{p,\sigma}^V$, since $\rho_V^Z(q)$ is within κ of a leaf of T_F^V by Proposition 4.7(2). Hence we are done in either case.

Now in the converse direction, suppose first that (4) holds; namely that $T_{q,\tau}^Z$ contains $\beta_F^Z(\rho_Z^V)$. Again since partition points are M away from ρ points, we know that $T_{q,\tau}^Z$ contains the $M/2$ -neighborhood (in T_F^V) of $\beta_F^Z(\rho_Z^V)$. Using D -density of $\beta_F^V(H_\theta(F))$ in T_F^V and the fact that p is far from any leaf of T_F^V , we can choose $x \in H_\theta(F)$ with $\beta_F^V(x) \in T_{p,\sigma}^V$ and $\beta_F^V(x)$ at least $M/2$ from a leaf of $T_{p,\sigma}^V$.

We claim that this implies that $\beta_F^Z(x)$ has to lie within $M/2$ of $\beta_F^Z(\rho_Z^V)$, as usual for M large. If not, another application of the bounded geodesic image axiom would imply that $\rho_V^Z(\pi_Z(x))$ is within κ_0 of a leaf, and the nested consistency inequality again gives that $\rho_V^Z(\pi_Z(x))$ and $\pi_V(x)$ are within κ_0 of each other, showing that $\pi_V(x)$ lies within $2\kappa_0$ of a leaf. Since the constants involved in the preceding argument depend only on $(\mathcal{X}, \mathfrak{S})$ and $|F|$, it follows that the distance (in T_F^V) between $\beta_F^V(x)$ and a leaf is bounded independently of M , contradicting the choice of x .




Figure 17: A cartoon for case (5), when $V \sqsubseteq Z$ and $T_{p,\sigma}^V$ contains $\beta_F^V(\rho_V^Z(q))$.

As a result, we find that $\beta_F^Z(x)$ is in $T_{q,\tau}^Z$ and hence $x \in W_{p,\sigma}^V \cap W_{q,\tau}^Z$, as required.

Suppose now as in (5) that $T_{p,\sigma}^V$ contains $\beta_F^V(\rho_V^Z(q))$, and assume also that (4) does not hold, so that $T_{q,\sigma}^Z$ avoids the $M/2$ -neighborhood of $\beta_F^V(\rho_V^Z(q))$ in T_F^V . In this case, we can take $x \in H_\theta(F)$ with $\beta_F^Z(x) \in T_{q,\tau}^Z$ and $d_{T_F^Z}(\beta_F^Z(x), q) \leq D$. Then since $\beta_F^Z(\rho_V^Z(q))$ is at least $M/2$ from the geodesic in T_F^Z between $\beta_F^Z(x)$ and q , and since T_F^Z is quasi-isometrically embedded in $\mathcal{C}(Z)$, we may apply the bounded geodesic image axiom and the nested consistency inequality to obtain that $d_V(\beta_F^V(x), \beta_F^V(\rho_V^Z(q))) < M/2$, where as before we choose M as large as necessary. Since $\rho_V^Z(q)$ lies within κ of a leaf of T_F^V by Proposition 4.7(2) with κ depending only on \mathfrak{S} and $|F|$, the $M/2$ -neighborhood of $\beta_F^V(\rho_V^Z(q))$ in T_F^V must be contained in $T_{p,\sigma}^V$, and hence $\beta_F^V(x)$ is contained in $T_{p,\sigma}^V$, showing $x \in W_{p,\sigma}^V \cap W_{q,\tau}^Z$ as required. \square

4.5 Refining the subdivisions

In this section, we analyze the difference between (M, M') -evenly spaced subdivisions $\mathbf{p} = \wp(F)$ and $\mathbf{p}' = \wp(F')$ (using the fixed subdivision mechanism \wp of Section 4.1). The main result here, Proposition 4.12, is that there are bounded refinements \mathbf{p}_0 and \mathbf{p}'_0 which are related by an “order-preserving” bijection $j : \mathbf{p}_0 \rightarrow \mathbf{p}'_0$ (in the sense specified below) which only moves points a bounded distance. This is where we will use the “evenly spaced” condition, which guarantees, roughly, that along corresponding edges we have the same number of partition points up to an additive, rather than multiplicative constant.

In the next subsection, we prove that these refined subdivisions determine isomorphic cube complexes when put through our cubulation machine.

Coarse separation We will need the following short discussion about coarse separation in quasi-isometrically embedded trees. Suppose $j : T \rightarrow X$ is a (χ, χ) -quasi-isometric embedding of a tree T into a δ -hyperbolic space X , and $p \in T$ is a point inside an edge, distance more than $\mu > 0$ from the endpoints. Let $T_{p,\pm}(\mu)$ be the two components of the complement of a μ -neighborhood of p . For any ϵ there is a $\mu(\epsilon, \delta, \chi)$ such that the images $j(T_{p,\pm}(\mu))$ are at least ϵ apart from each other. We want to compare this separation for two nearby trees:

Lemma 4.11 *Let T_1 and T_2 be trees with (χ, χ) -quasi-isometric embeddings $u_i : T_i \rightarrow X$ into a δ -hyperbolic space X whose images are within Hausdorff distance ϵ . There exists $\mu_0 = \mu_0(\delta, \chi, \epsilon) > 10\epsilon$ such that for all $\mu \geq \mu_0$ the following holds.*

Suppose that $p_i \in T_i$ are points with $d(u_1(p_1), u_2(p_2)) \leq \mu$ and each p_i is in a segment e_i contained in an edge of T_i such that $u_i|_{e_i}$ is a $(1, \chi)$ -quasi-isometric embedding. Moreover, assume that p_i lies at distance more than 2μ from the endpoints of e_i . Then the labels of the components can be chosen so that

- $u_1((T_1)_{p_1,+}(2\mu))$ is in an ϵ -neighborhood of $u_2((T_2)_{p_2,+})$,
- $u_1((T_1)_{p_1,+}(2\mu))$ is μ -far from $u_2((T_2)_{p_2,-})$, and
- the same holds if we swap $+$ and $-$.

In what follows we apply this to trees T_F^V , and so that half-trees that here would be $(T_F^V)_{p,\sigma}$ correspond to what we wrote as $T_{p,\sigma}^V$ in Section 4.1. We will use both notations, and hope not to confuse the reader.

Common refinements from close components We are now ready to prove the refinement statement. We note that this is the main point at which we use the full power of Theorem 3.3, which provides not only that each of the pairs of trees T_F^V and $T_{F'}^V$ have Hausdorff close images in $\mathcal{C}(V)$, but that they are identical on almost every component, and that their different components can be cut into large pieces where they are coarsely identical.

Proposition 4.12 *There exists $M_2 = M_2(|F|, \mathfrak{S}) > 0$ such that if $M > M_2$, there exists $R > 0$ and:*

- (1) Subsets $p_0 \subset p$ and $p'_0 \subset p'$ with
 - $p_0(V)$ or $p'_0(V) \neq \emptyset$ only if $V \in \mathcal{U}(F) \cap \mathcal{U}(F')$, and
 - $|p \setminus p_0|, |p' \setminus p'_0| < R$.
- (2) A bijection $j: p_0 \rightarrow p'_0$ with $d_V(p, j(p)) < 2M/3$ for any $p \in p_0(V)$.
- (3) For every $p \in p_0$, a bijection j_p between the half-trees defined by p and those defined by $j(p)$ with the following property. For any V and $p, q \in p_0(V)$, if the half-tree $(T_F^V)_{p,\sigma}$ contains q , then $j_p((T_F^V)_{p,\sigma})$ contains $j(q)$.
- (4) Moreover, let f and f' be such that either $f \in F, f' \in F'$, and $d_X(f, f') \leq 1$, or $f = f' = \rho_V^U$ for some $U, V \in \mathcal{U}(F) \cap \mathcal{U}(F')$ with $U \sqsubset V$ or $U \pitchfork V$. If $\beta_F^V(f) \in (T_F^V)_{p,\sigma}$, then $\beta_{F'}^V(f') \in j_p((T_F^V)_{p,\sigma})$.

Proof We work in each $V \in \mathfrak{S}$ separately, constructing $p_0(V), p'_0(V)$ and the bijection, and combine the results.

First, if $V \in \mathcal{U}(F) \Delta \mathcal{U}(F')$, then $\text{diam}_V(F)$ and $\text{diam}_V(F')$ are uniformly bounded and we set $p_0(V) = p'_0(V) = \emptyset$. By Proposition 2.12, there are boundedly many such V ; hence this involves deleting at most boundedly many subdivision points.

Next, if $V \in \mathcal{U}(F) \cap \mathcal{U}(F')$ is not involved in the transition from F to F' (see Definition 2.13), then all of the relevant data is constant. Hence, by our formalism for choosing subdivisions (see the definition of \wp in Section 4.1), we have $p(V) = p'(V)$, and we set $p_0(V) = p'_0(V) = p(V)$.

Hence we may restrict our attention to a fixed $V \in \mathcal{U}(F) \cap \mathcal{U}(F')$ for which either $\pi_V(F) \neq \pi_V(F')$ or $\mathcal{U}^V(F) \neq \mathcal{U}^V(F')$ (or both). We note that Proposition 2.14 bounds the number of such V solely in terms of n and \mathfrak{S} .

Fix such a V . Recall that within $\mathcal{C}(V)$ we have $F^V = \pi_V(F)$ and

$$\mathcal{Y}^V = \{\rho_V^Y \mid Y \in \mathcal{U}^V(F)\},$$

and similarly for F'^V and \mathcal{Y}'^V . By Proposition 2.14, $\#(\mathcal{Y}^V \Delta \mathcal{Y}'^V) < N_1$, where $N_1 = N_1(\mathfrak{S}, k) > 0$.

Recall that Theorem 3.3 provides a constant $L = L(k, \mathfrak{S}) > 0$ and the following:

- (1) Stable trees with decompositions

$$T_F^V = T_c(F^V, \mathcal{Y}^V) \cup T_e(F^V, \mathcal{Y}^V), \quad T_{F'}^V = T_c(F'^V, \mathcal{Y}'^V) \cup T_e(F'^V, \mathcal{Y}'^V);$$

we write these as $T_c^V(F), T_e^V(F)$ and $T_c^V(F'), T_e^V(F')$ for short.

- (2) Two stable subsets, $T_s \subset T_e^V(F)$ and $T'_s \subset T_{e'}^V(F')$, such that
 - (a) T_s and T'_s are contained in the interiors of the edges of $T_e^V(F)$ and $T_{e'}^V(F')$, respectively;
 - (b) the complements $T_e^V(F) \setminus T_s$ and $T_{e'}^V(F') \setminus T'_s$ each have at most L components, each of which has diameter at most L ;
 - (c) there is a bijective correspondence between the components of T_s and T'_s ;
 - (d) this bijective correspondence is the identity on all but at most L components, with identical components of T_s, T'_s coming from identical components of $T_e^V(F)$ and $T_{e'}^V(F')$;
 - (e) the remaining components of T_s are within Hausdorff distance L of their corresponding components in T'_s .

We assume $M > \max(4L, 4\mu_0, 4\kappa)$, where μ_0 is given by Lemma 4.11 when $\epsilon = L$ and the quasi-isometry constants of the trees match those for T_F^V and $T_{F'}^V$, and κ is the constant in Proposition 4.7.

Consider the sets of subdivision points $\mathbf{p}_1(V) = \mathbf{p}(V) \cap T_s$ and $\mathbf{p}'_1(V) = \mathbf{p}'(V) \cap T'_s$ which are contained in the stable subsets. Since $\mathbf{p}(V) \subset T_e^V(F)$ and $\mathbf{p}'(V) \subset T_{e'}^V(F')$, items (2a) and (2b), and our choice of subdivision width $M > 4L$, imply $\mathbf{p}(V) \setminus \mathbf{p}_1(V)$ and $\mathbf{p}'(V) \setminus \mathbf{p}'_1(V)$ both have cardinality bounded above by L .

By item (2d), the induced subdivisions $\mathbf{p}_1(V)$ and $\mathbf{p}'_1(V)$ agree on all but at most L components of T_s and T'_s , respectively, as these components are segments in the components of $T_e^V(F)$ and $T_{e'}^V(F')$ which are equal and hence have the same subdivisions by our setup (see Section 4.1).

On the remaining L components we can make bounded adjustments. Let e and e' be edge components of T_s and T'_s related by the correspondence. The closest subdivision point in e to an endpoint is at most $M'/2$ away and at least M , by definition of the subdivisions, and similarly for e' . Since $d_{\text{Haus}}(e, e') < L$ by item (2e), and e and e' are quasi-isometrically embedded with multiplicative constant 1 and additive constant depending on $(\mathcal{X}, \mathfrak{S})$ and n (Theorem 3.3(c)), the difference between the number of subdivision points in e and e' is bounded in terms of $M, M', L, (\mathcal{X}, \mathfrak{S})$, and n .

We write $\mathbf{p}_1(e) = \mathbf{p}_1(V) \cap e$, and similarly for $\mathbf{p}'_1(e')$. Note that these sets are naturally ordered once we choose endpoints of e and e' , and we order e and e' by declaring those endpoints to be minimal. We choose endpoints of e and e' within distance bounded in terms of L and the constants of the quasi-isometric embeddings of e and e' in $\mathcal{C}(V)$. We will assume that M is larger than this bound.

After deleting from $p_1(e) \cup p'_1(e')$ a number of points bounded by this constant (all of which occur near the endpoints of e and e') and using the fact that $p_1(e)$ and $p'_1(e')$ are M -evenly spaced, we can obtain refinements $p_0(e)$ and $p'_0(e')$, which admit an order-preserving (with respect to the aforementioned order) bijection $j_e: p_0(e) \rightarrow p'_0(e')$ which satisfies

$$d_V(p, j_e(p)) < L + \frac{M}{2} + \zeta < \frac{2M}{3},$$

where $e' \rightarrow \mathcal{C}(V)$ is a $(1, \zeta)$ -quasi-isometric embedding by Theorem 3.3(c), and we have chosen M sufficiently large to guarantee the inequality. Indeed, for each p we find a nearest point in e' , which is at most L away, and then move along e' at most $M/2 + \zeta$ to a point of $p'_0(e')$. For later purposes, we can assume that points in $p_0(e)$ do not lie within $2M$ of the endpoints of e , and similarly for $p'_0(e')$.

If we set $p_0(V) = \bigcup_{e \in \pi_0(T_S)} p_0(e)$ and define $p'_0(V)$ similarly, then the j_e maps combine to give a bijection $j_V: p_0(V) \rightarrow p'_0(V)$ which moves points by distance at most $2M/3$, as required for item (2) of the proposition.

To define the map j_p between half-trees we use Lemma 4.11. Namely, we pair the half-tree $(T_F^V)_{p,\sigma}$ with the half-tree $(T_{F'}^V)_{j_V(p),\tau}$ that contains $(T_F^V)_{p,\sigma}(2M)$ in its L -neighborhood.

For item (3) of the proposition, let $p, q \in p_0(V)$. There are two cases.

Suppose first that p and q lie in the same edge-component e of T_S . Recall that we chose the bijection j_e to be order-preserving, with respect to the order along e and e' determined by choosing endpoints $e^- \in e$ and $(e')^- \in e'$ which are a small distance (less than M) apart to be minimal in the orders. Let $(T_F^V)_{p,\sigma}$ denote the half-tree of p containing e^- and let $(T_{F'}^V)_{j_V(p),\tau}$ denote the half-tree of $j_V(p)$ containing $(e')^-$. Since we have $2M$ spacing now between p and the endpoints of e , e^- is in $(T_F^V)_{p,\sigma}(2M)$, and Lemma 4.11 says that there is exactly one half-tree at $j_V(p)$ which comes within M of $(T_F^V)_{p,\sigma}(2M)$. Since e^- is within M of $(e')^-$, it follows that $(T_{F'}^V)_{j_V(p),\tau}$ is in fact the paired half-tree provided by Lemma 4.11, which is our $j_p((T_F^V)_{p,\sigma})$. We conclude that $q < p$ in the order along e if and only if $q \in (T_F^V)_{p,\sigma}$ and $j_V(q) < j_V(p)$ along e' if and only if $j_V(q) \in j_p((T_F^V)_{p,\sigma})$. Since j_e is order-preserving, (3) follows in this case.

Suppose now that p and q do not lie in the same edge-component of T_S . Suppose q is contained in the half-tree $(T_F^V)_{p,\sigma}$, and let $(T_{F'}^V)_{q,\tau} = j_p((T_F^V)_{p,\sigma})$. In this case we have that q lies in $(T_F^V)_{p,\sigma}(2M)$, rather than just in $(T_F^V)_{p,\sigma}$. By Lemma 4.11

there is only one half-tree of $T_{F'}^V$ at $j_V(p)$ that comes within M of $(T_F^V)_{p,\sigma}(2M)$, and said half-tree must be $(T_{F'}^V)_{j_V(p),\tau}$. Since $j_V(q)$ lies within M of q , we have $j_V(q) \in (T_{F'}^V)_{j_V(p),\tau}$, as required.


The argument for part (4) of the proposition is essentially the same as the argument for the second case of part (3), since all we used there is that the point q of T_F^V is not close to p , but it is close to a corresponding point in $T_{F'}^V$, and the analogous properties hold in all the listed cases. □

4.6 Refinements give isomorphic cube complexes

Consider the refined subdivisions p_0 and p'_0 for F and F' , respectively, that are produced by Proposition 4.12. These are (M, M'') spaced subdivisions (though no longer evenly spaced) so each of the sets of data (F, p_0) and (F', p'_0) — and their associated collections of stable trees — can be plugged into our cubulation machine to produce cube complexes \mathcal{Q}_{F,p_0} and \mathcal{Q}_{F',p'_0} , respectively. We also assume that $M > \max\{M_1, M_2\}$, where M_1 and M_2 are the constants from Lemma 4.10 and Proposition 4.12, respectively, along with our other base assumptions about M .

Our next result says that these cube complexes are abstractly isomorphic and admit coarsely compatible quasi-isometric embeddings into \mathcal{X} . Using Proposition 4.9, we will be able to conclude that the right hand side of diagram (4-1) from Theorem 4.1 commutes.

Proposition 4.13 *There exists $M_3 = M_3(|F|, \mathfrak{S}) > 0$, such that if $M > M_3$, there exists $B > 0$ and a cubical isomorphism $\hat{h}: \mathcal{Q}_{F,p_0} \rightarrow \mathcal{Q}_{F',p'_0}$ such that the diagram*



commutes up to error B .

Proof We will define the required cubical isomorphism \hat{h} and then prove that the middle triangle commutes up to bounded error. This suffices for the proposition because Proposition 4.9 says that the top and bottom triangles commute up to bounded error.

In order to use Lemma 2.3 to define an isomorphism between $\mathcal{Q}_{F, \mathbf{p}_0}$ and $\mathcal{Q}_{F', \mathbf{p}'_0}$, we need a bijection between the corresponding collections of half-spaces which preserves complements and disjointness.

Let $\mathcal{W}_{F, \mathbf{p}_0}$ and $\mathcal{W}_{F', \mathbf{p}'_0}$ be the sets of half-spaces and construct a bijection

$$\iota: \mathcal{W}_{F, \mathbf{p}_0} \rightarrow \mathcal{W}_{F', \mathbf{p}'_0}$$

as follows. By Proposition 4.12, we need only consider $V \in \mathcal{U}(F) \cap \mathcal{U}(F')$. Any $p \in \mathbf{p}_0(V)$ is contained in an edge of T_F^V (in fact in T_e^V) and is at least M from its endpoints.

Let $j: \mathbf{p}_0 \rightarrow \mathbf{p}'_0$ be the bijection provided by Proposition 4.12, together with the corresponding maps j_p pairing the half-trees at $p \in \mathbf{p}_0$ with those at $j(p)$.

To define ι , let $p \in \mathbf{p}_0$ and $\sigma \in \{\pm\}$. If $j_p(T_{p, \sigma}^V) = T_{j_V(p), \tau}^V$ we define $\iota(W_{p, \sigma}^V) = W_{j_V(p), \tau}^V$. It is straightforward to confirm that ι respects complementation as in condition (1) of Lemma 2.3.

We now confirm condition (2) of Lemma 2.3 for ι by using the various characterizations of half-space intersections given in Lemma 4.10. We remark that the figures from the proof of Lemma 4.10 are again relevant.

Let $p \in \mathbf{p}_0(V)$ and $q \in \mathbf{p}_0(Z)$ for $Z, V \in \mathcal{U}(F) \cap \mathcal{U}(F')$, and suppose that the half-spaces corresponding to the half-trees $T_{p, \sigma}^V$ and $T_{q, \tau}^Z$ intersect nontrivially, where $\sigma, \tau \in \{\pm\}$. There are five cases, up to switching the roles of V and Z .

Case $Z \perp V$ This case follows immediately from the construction and Lemma 4.10(1), since all relevant pairs of half-spaces intersect in this case.

Case $Z = V$ In this case, Lemma 4.10(2) implies that $T_{p, \sigma}^V \cap T_{q, \tau}^V \neq \emptyset$ (recall Figures 13 and 14).

In particular, up to switching the roles of p and q , we have $q \in T_{p, \sigma}^V$. But then, in view of Proposition 4.12(3), $j(q) \in j_p(T_{p, \sigma}^V)$. In particular, $j_p(T_{p, \sigma}^V) \cap j_q(T_{q, \tau}^V) \neq \emptyset$, so that the corresponding half-spaces intersect, again by Lemma 4.10(2).

Case $Z \pitchfork V$ In this case, Lemma 4.10(3) implies, up to switching the roles of Z and V , that $(T_F^V)_{p, \sigma}$ contains $\beta_F^V(\rho_V^Z)$ (recall Figure 15). It follows immediately from part (4) of Proposition 4.12 that $\beta_{F'}^V(\rho_V^Z) \in j_p(T_{p, \sigma}^V)$, as required.

Case $V \sqsubset Z$ By Lemma 4.10, there are two subcases, up to switching the roles of V and Z : (a) when $T_{q,\tau}^Z$ contains $\beta_F^Z(\rho_Z^V)$, and (b) when $T_{p,\sigma}^V$ contains $\beta_F^V(\rho_V^Z(q))$.

In case (a), since $V \in \mathcal{U}(F) \cap \mathcal{U}(F')$ part (4) of Proposition 4.12 immediately gives that $\beta_{F'}^Z(\rho_Z^V)$ lies in $j_q(T_{q,\tau}^Z)$.

Suppose now that (b) holds. We prove that $\beta_{F'}^V(\rho_V^Z(j(q)))$ is contained in $j_p(T_{p,\sigma}^V)$ (recall Figure 17). Recall from Proposition 4.7 that $\rho_V^Z(q)$ lies close to some $\pi_V(f)$ with bound in terms of \mathfrak{S} and $|F|$. Since $d_V(q, j(q)) < 2M/3$ and both q and $j(q)$ are at least $M - 2K$ from ρ_Z^V , with the $2K$ coming from the facts that $T_c^Z(F)$ is within Hausdorff distance K of the \mathcal{Y}^Z and the edge of $T_e^Z(F)$ containing q is $(1, K)$ -quasi-isometrically embedded in $\mathcal{C}(Z)$ (Theorem 3.3(d)), choosing M sufficiently large, we can guarantee that any geodesic in $\mathcal{C}(V)$ between p and $j_V(p)$ avoids the κ_0 -neighborhood of ρ_Z^V , for κ_0 the constant of the bounded geodesic image property, which then bounds $d_V(\rho_V^Z(q), \rho_V^Z(j_Z(q))) < \kappa_0$. Hence $\rho_V^Z(j_Z(q))$ is close to both $\rho_V^Z(q)$ and a leaf of $T_{F'}^V$, and thus $\beta_{F'}^V(\rho_V^Z(j(q))) \subset j_p(T_{p,\sigma}^V)$, as required.

Since the wallspace map ι satisfies the conditions of Lemma 2.3, we obtain a cubical isomorphism $\hat{h}: \mathcal{Q}_{F,p_0} \rightarrow \mathcal{Q}_{F',p'_0}$.

Commutativity It remains to prove that

$$\Phi_{F,p_0}: \mathcal{Q}_{F,p_0} \rightarrow \mathcal{X} \quad \text{and} \quad \Phi_{F',p'_0} \circ h: \mathcal{Q}_{F,p_0} \rightarrow \mathcal{X}$$

are the same up to a bounded error depending only on k and the ambient HHS structure.

By the distance formula (Theorem 2.6), it suffices to show that for each 0-cube $x \in \mathcal{Q}_{F,p_0}$, its respective images $\Phi_{F,p_0}(x)$ and $\Phi_{F',p'_0} \circ h(x)$ have coarsely the same projections to $\mathcal{C}(V)$ for each $V \in \mathcal{U}(F) \cap \mathcal{U}(F')$.

Recall from the end of Section 4.2 that the maps Φ_{F,p_0} and Φ_{F',p'_0} are defined domainwise by intersecting certain collections of half-trees of T_F^V and $T_{F'}^V$, for each $V \in \mathfrak{S}$, and hence the same is true for $\Phi_{F',p'_0} \circ h$. By chasing the relevant definitions, of h and j especially, we see that the collections of half-trees involved in the definition of $\Phi_{F,p_0}(x)$ and $\Phi_{F',p'_0} \circ h(x)$ are in bijection with each other, with corresponding half-trees lying within bounded Hausdorff distance depending only on k and the ambient HHS structure. Hence their intersections in T_F^V and $T_{F'}^V$ coarsely coincide, as required. This completes the proof of the proposition. □




Figure 18: A simple example of how subdivision bijections determine hyperplane deletions. The relevance of X for F' but not F requires deleting a subdivision point to obtain the bijection $j_U: \mathbf{p}_U \rightarrow \mathbf{p}'_U$ (indicated in pink), and similarly with Y for $j_V: \mathbf{p}_V \rightarrow \mathbf{p}'_V$. Deleting subdivision points results in hyperplane deletions when passing from \mathcal{Q}_F and $\mathcal{Q}_{F'}$ to \mathcal{Q}_0 . Note that since $X, Y \notin \mathcal{U}_F \cap \mathcal{U}_{F'}$, neither domain determines any subdivision points.

4.7 Proof of Theorem 4.1

Let F and F' be as in the statement.

The CAT(0) cube complexes $\mathcal{Q}_F, \mathcal{Q}_{F'}$ are constructed in Section 4.2, using subdivisions $\mathbf{p} = \wp(F), \mathbf{p}' = \wp(F')$ within each of the relevant stable trees produced in Section 3. Items (1) and (2) are proven in [8], as recalled in Theorem 4.6 above.

For (3), we may define a map $\psi_F : F \rightarrow \mathcal{Q}_F$ as in [8, Proof of Theorem 2.1]. For each $f \in F$, let $\psi_F(f)$ be the orientation of the walls on hull $H_\theta(F)$ obtained by choosing, for each wall $(W_{p,+}^V, W_{p,-}^V)$, the halfspace containing f . We define $\psi_{F'}$ similarly. That ψ_F and $\psi_{F'}$ satisfies (3) now follows from Theorem 4.6.

We now prove the stability statements. We will only consider the case that g is the identity for the same reason as in Remark 3.4, namely that what we have to prove is that the choices that we made along the way only affect the output in the way predicted by the statement, and such choices can be “translated” by automorphisms. The choices we are referring to are those of the trees T_F^V , of the evenly spaced subdivisions, and of points projecting coarsely in specified places in the various $\mathcal{C}(Y)$, as in Theorem 4.6.(2).

Let $\mathbf{p}_0 \subset \mathbf{p}$ and $\mathbf{p}'_0 \subset \mathbf{p}'$ be the refinements provided by Proposition 4.12, where at most N subdivision points are deleted.

The cube complex \mathcal{Q}_0 as in the statement of Theorem 4.1 can be taken to be either of the cube complexes $\mathcal{Q}_{F,\mathbf{p}_0}$ or $\mathcal{Q}_{F',\mathbf{p}'_0}$ that are produced by Proposition 4.13. So take $\mathcal{Q}_0 = \mathcal{Q}_{F',\mathbf{p}'_0}$ and let $\Phi_0 : \mathcal{Q}_0 \rightarrow \mathcal{X}$ be the map $\Phi_{F',\mathbf{p}'_0} : \mathcal{Q}_{F',\mathbf{p}'_0} \rightarrow \mathcal{X}$ given by Proposition 4.13. Finally, let $\eta' \equiv h' : \mathcal{Q}_{F'} \rightarrow \mathcal{Q}_0$ be as in Proposition 4.13 and define $\eta : \mathcal{Q}_F \rightarrow \mathcal{Q}_0$ by $\eta = \hat{h} \circ h$.

By Proposition 4.13, these maps each satisfy the required properties and the right part of diagram (4-1) commutes up to bounded error, as required.

To see exact commutativity of the left part of the diagram, it remains to prove that $\eta \circ \psi_F \circ \iota_F = \eta' \circ \psi_{F'} \circ \iota_{F'}$. Recall that $\psi_F : F \rightarrow \mathcal{Q}_F$ is defined by sending f to the coherent orientation on the wallspace defined by \mathbf{p} which, for each $p \in \mathbf{p}(V)$, chooses the half-tree of $T_F^V \setminus p$ containing $\beta_F^V(f)$, for each $V \in \mathcal{U}(F)$. Since h is a hyperplane deletion map, $h \circ \psi_F$ makes, for $p \in \mathbf{p}_0$, the same choice as ψ_F .

Fix $f \in F$ (the argument for $f \in F'$ is similar). Then the two sides of the equation are coherent orientations on the wallspace defining $\mathcal{Q}_{F',\mathbf{p}'_0}$, and we have to check that they coincide on every halfspace. Pick $p \in \mathbf{p}_0$ and let $p' = j_V(p)$, the map defined in Proposition 4.12. As above the orientation of $h \circ \psi_F \circ \iota_F(f) = h \circ \psi_F(f)$ on the wall associated to p is the one that chooses the half-tree of $T_F^V \setminus p$, call it $(T_F^V)_{p,+}$, that contains $\beta_F^V(f)$. The map \hat{h} , by the construction in Proposition 4.13, takes this to the orientation that chooses the half-tree of p' in $T_{F'}^V$ given by $j_p((T_F^V)_{p,+})$ (where j_p is the bijection of half-trees provided by Proposition 4.12). Letting $f' = \iota_{F'}(f)$ we have $d_{\mathcal{X}}(f', f) \leq 1$, so part (4) of Proposition 4.12 tells us that $\beta_{F'}^V(f') \in j_p((T_F^V)_{p,+})$,

which means that $j_P((T_F^V)_{P,+})$ is the half-tree selected by $\eta' \circ \psi_{F'}(f')$. Thus we conclude that $\eta(\psi_F(\iota_F(f))) = \eta'(\psi_{F'}(\iota_{F'}(f)))$.

This completes the proof of Theorem 4.1. □

5 Generalizing normal paths to find barycenters

In this section, we describe a variation on the “normal paths” construction due to Niblo and Reeves [46]. We remark that they used these normal paths to build a biautomatic structure for any cubical group, with the paths playing the central role of the bicombing in that structure.

For the case of two points, our construction gives a “symmetrized” version of the Niblo–Reeves construction. The main difference with their construction, however, is that ours is adapted to allow for multiple points, which we need for our barycenter application (Theorem E).

The reader may want to refer to Section 2.1 for the various notions and notations relating to cube complexes that we will use throughout this section

Let X be a CAT(0) cube complex, \mathcal{H} its set of hyperplanes, and $f : P \rightarrow X^{(0)}$ a (not necessarily injective) map from a finite set P into the vertices of X . Roughly, we will find a barycenter for the set $f(P)$ in X by an iterative sequence of contractions which behaves stably under hyperplane deletions. The main result of this section, and the only statement from this section that we need to prove our main theorems, is the following:

Theorem 5.1 *Let X be a CAT(0) cube complex, and let \mathcal{H} be its set of hyperplanes.*

For each $f : P \rightarrow X^{(0)}$, where P is a finite set, there is a finite sequence

$$\{f_i : P \rightarrow X^{(0)} \mid i = 0, \dots, n = n_f\}$$

with the following properties:

- (1) $f_0 = f$ and $\text{diam}_\infty(f_n(P)) \leq 1$.
- (2) For each $p \in P$ and $0 \leq i \leq n - 1$ we have $d_\infty(f_i(p), f_{i+1}(p)) \leq 1$.
- (3) For each p , there is an ℓ^1 -geodesic going through the vertices

$$f_0(p), f_1(p), \dots, f_n(p)$$

in this order.

- (4) No hyperplane separates every point of $f(P)$ from a point in $f_n(P)$.

- (5) If $g: Q \rightarrow P$ is surjective, then $f_i \circ g = (f \circ g)_i$ for all i .
- (6) If G is any hyperplane of X , the hyperplane deletion map $\text{Res}_{\mathcal{H} \setminus G}: X \rightarrow X(\mathcal{H} \setminus G)$ satisfies

$$|n_f - n_{\text{Res}_{\mathcal{H} \setminus G} \circ f}| \leq 1 \quad \text{and} \quad d_\infty(\text{Res}_{\mathcal{H} \setminus G}(f_i(p)), (\text{Res}_{\mathcal{H} \setminus G} \circ f)_i(p)) \leq 1.$$

Recall that d_∞ is the metric generated by the sup metric on each cube in the ambient complex.

The proof of Theorem 5.1 occurs in parts over the remainder of this section. We tie them together in Section 5.4 below.

We will mostly ignore the ambient cube complexes and focus on the hyperplane set \mathcal{H} and regard maps f as above as maps associating to $p \in P$ an orientation on \mathcal{H} . For each pair (f, \mathcal{H}) we consider a number of operations.

First, let \mathcal{H}_f be the set of hyperplanes of X that separate $f(P)$. That is, \mathcal{H}_f is the set of hyperplanes $H \in \mathcal{H}$ for which there exist $p, p' \in P$ such that $f(p)$ and $f(p')$ are on different sides of H .

Let

$$\text{Trim}(f, \mathcal{H}) = (\text{Res}_{\mathcal{H}_f}(f), \mathcal{H}_f)$$

be the restriction to \mathcal{H}_f . Note that the quotient complex $X(\mathcal{H}_f)$ actually embeds in $X(\mathcal{H})$ (with image the subcomplex spanned by all vertices that lie in the intersection of all the halfspaces of $X(\mathcal{H})$ that contain $f(P)$), and that it is finite even if $X(\mathcal{H})$ is not.

If G is a collection of mutually crossing hyperplanes, we let

$$\text{del}_G(f, \mathcal{H}) = (\text{Res}_{\mathcal{H} \setminus G}(f), \mathcal{H} \setminus G)$$

This “deletion map” corresponds to composing f with the quotient by G ; that is,

$$\text{Res}_{\mathcal{H} \setminus G} \circ f: P \rightarrow X(\mathcal{H} \setminus G).$$

We will also write $\text{Res}_{\mathcal{H} \setminus G}(f)$ as $f|_{\mathcal{H} \setminus G}$, in a slight abuse of notation.

5.1 Extremal and transitional hyperplanes

As stated above, our generalized normal paths give a series of contractions of our set $f(P)$ to a bounded diameter set in X . This is accomplished by iteratively jumping the points of $f(P)$ (and their subsequent contracted images) over a sequence of hyperplanes in \mathcal{H}_f . Thus we are led to understand which hyperplanes are next in line to be jumped.




Figure 19: H is extremal but not transitional, while H' and H'' are both transitional.

Definition 5.2 Let X be a cube complex and \mathcal{H} its set of hyperplanes.

- A point $p \in X$ is *adjacent* to a hyperplane $H \in \mathcal{H}$ if there are no hyperplanes separating p from H .
- A hyperplane $H \in \mathcal{H}_f$ is *extremal* if on one side of H every point of $f(P)$ is adjacent to H in X . We let $E(f, \mathcal{H}) \subset \mathcal{H}_f$ denote the set of extremal hyperplanes.
- A hyperplane $H \in E(X, f)$ is *transitional* if it is extremal and on one side of H not every point of $f(P)$ is adjacent to H ; we let $T(f, \mathcal{H}) \subset E(f, \mathcal{H})$ be the set of transitional hyperplanes.

If H is transitional, we write $P = P_0(H) \sqcup P_1(H) = P_0 \sqcup P_1$ where $f(P_0)$ is adjacent to H on one side, $f(P_1)$ is on the other side, and at least one point of $f(P_1)$ is not adjacent to H .

We note that $E(f, \mathcal{H})$ is always nonempty when P is nonempty. In fact, for any $p \in P$, it is readily shown that H is extremal whenever H is a hyperplane in \mathcal{H}_f such that the number of hyperplanes separating $f(p)$ from H is maximal.

Moreover, for any hyperplane $H \in E(f, \mathcal{H}) \setminus T(f, \mathcal{H})$, every point of $f(P)$ must be adjacent to H .

Lemma 5.3 The set $f(P)$ is contained in a single cube of X if and only if $T(f, \mathcal{H}) = \emptyset$.

Proof First of all, we observe that $f(P)$ being contained in a single cube is equivalent to \mathcal{H}_f being mutually crossing. Moreover, it is clear from the definitions that if \mathcal{H}_f is mutually crossing, then $T(f, \mathcal{H}) = \emptyset$.

In the other direction, we suppose that $f(P)$ is not contained in a single cube and then produce a transitional hyperplane. The fact that $f(P)$ is not contained in a single cube implies that there exists $p \in P$ and a hyperplane $H \in \mathcal{H}_f$ with $f(p)$ not adjacent to H .




Figure 20: An example of a single move. $f(P)$ consists of three points, and the points of $f_1(P)$, in matching colors, are on the other side of the transitional hyperplanes, which are indicated in orange.

Choose now $p \in P$ and $H \in \mathcal{H}_f$ so that the number of hyperplanes separating $f(p)$ from H is maximal; the argument above implies that the number of these hyperplanes is positive, so $f(p)$ is not adjacent to H . As observed before the lemma, it follows that H is extremal, and since $f(p)$ is not adjacent to H , we have that H is in fact transitional, as required. \square

5.2 The move sequence

We now build our sequence of contractions, which we call moves.

Roughly speaking, a *move* is an operation on (f, \mathcal{H}) in which, for every $H \in \mathbf{T}(f, \mathcal{H})$, the points of $P_0(H)$ cross H to the opposite side. The resulting pair

$$(f_1, \mathcal{H}) = \text{Move}(f, \mathcal{H})$$

is a map for which $\mathcal{H}_{f_1} = \mathcal{H}_f \setminus \mathbf{T}(f, \mathcal{H})$, so the image of the new map $f_1(P)$ is strictly contained within the subcomplex spanned by $f(P)$.

To define f_1 , we need the following observation:

Lemma 5.4 For each $p \in P$, the set

$$J(p) = \{H \in \mathbf{T}(f, \mathcal{H}) \mid p \in P_0(H)\}$$

is mutually crossing.

Proof Suppose, by way of contradiction, that $H_1, H_2 \in J(p)$ and that H_1 does not cross H_2 . Let s_1 be the side of H_1 on which $f(P_0(H_1))$ lies (and hence is adjacent)




Figure 21: An example of a move sequence terminating in all points lying in a single cube. Concentric circles around a vertex indicate, in this case, a point p with $f_i(p) = f_{i+1}(p)$.

and let s_2 be the corresponding side for H_2 . Since $p \in P_0(H_1) \cap P_0(H_2)$, we have that s_1 and s_2 intersect. Since H_1 does not cross H_2 , up to swapping indices we have that H_2 lies in s_1 and separates H_1 from $f(P_0(H_2))$, which contains p and thus contradicts the choice of s_1 . □

We can now define $f_1(p)$. By Lemma 5.4, $f(p)$ is the corner of a unique maximal cube, each of whose midcubes is contained in some hyperplane of $J(p)$, and we can choose $f_1(p)$ to be the diagonally opposite corner. Equivalently, $f_1(p)$ can be defined as the point obtained by flipping the orientation that $f(p)$ gives to all the hyperplanes in $J(p)$. Either definition gives the new map f_1 .

Note that there may be p for which $J(p) = \emptyset$, and for these $f_1(p) = f(p)$. Moreover, since the operation moves all points of $P_0(H)$ across H for each $H \in T(f, \mathcal{H})$, we have the following consequence:

Lemma 5.5 $\mathcal{H}_{f_1} = \mathcal{H}_f \setminus T(f, \mathcal{H})$.

Now we are ready to consider the *move sequence* of (f, \mathcal{H}) . For $i \geq 0$ define

$$(f_i, \mathcal{H}) = \text{Move}^i(f, \mathcal{H}).$$

Lemma 5.6 *The move sequence $\{f_i\}$ is eventually constant, with constant image all contained in one cube.*

Proof By Lemma 5.5, the set \mathcal{H}_{f_i} becomes strictly smaller for each i as long as $T(f_i, \mathcal{H}) \neq \emptyset$, so after a finite number of steps we must have $T(f_n, \mathcal{H}) = \emptyset$. Thereafter, the f_i are all the same and their images are all contained in a cube by Lemma 5.3. □




Figure 22: An illustration of Proposition 5.7, where P has two points.

Our main result about this sequence is its stability under the operation del_G . This is part of Theorem 5.1, which we rephrase here in our new language:

Proposition 5.7 *Let G be a mutually crossing set in \mathcal{H}_f . Let*

$$(f_i, \mathcal{H}) = \text{Move}^i(f, \mathcal{H})$$

and

$$(f'_i, \mathcal{H} \setminus G) = \text{Move}^i(\text{del}_G(f, \mathcal{H}))$$

be the move sequences for (f, \mathcal{H}) and $\text{del}_G(f, \mathcal{H})$. Then for each $p \in P$,

$$d_\infty(\text{Res}_{\mathcal{H} \setminus G} \circ f_i(p), f'_i(p)) \leq 1.$$

The proposition is in fact a generalization of what we need, since it deals with a mutually crossing set.

5.3 Proof of stability of move sequences

We begin by studying the structure of the extremal and transitional hyperplanes for a pair (f, \mathcal{H}) , and the way in which they are affected by hyperplane deletions.

Lemma 5.8 *Every $J \in E(f, \mathcal{H}) \setminus T(f, \mathcal{H})$ crosses every hyperplane in $\mathcal{H}_f \setminus \{J\}$.*

Proof Since J is extremal but not transitional, $f(P)$ is adjacent to it on both sides; see Figure 19. This means that no other hyperplane can separate any $f(p)$ from J , and this implies that every $H \in \mathcal{H}_f \setminus \{J\}$ crosses J . \square

The next lemma explains how the extremal and transitional hyperplanes change after a deletion step:

Lemma 5.9 *Let G be a mutually crossing hyperplane set in \mathcal{H}_f . Then*

$$(5-1) \quad E(f, \mathcal{H}) \setminus G \subset E(\text{del}_G(f, \mathcal{H}))$$

and

$$(5-2) \quad E(\text{del}_G(f, \mathcal{H})) \setminus E(f, \mathcal{H}) = T(\text{del}_G(f, \mathcal{H})) \setminus T(f, \mathcal{H}).$$

Moreover, if $G \cap E(f, \mathcal{H}) = \emptyset$, then

$$(5-3) \quad E(f, \mathcal{H}) = E(\text{del}_G(f, \mathcal{H}))$$

and

$$(5-4) \quad T(f, \mathcal{H}) = T(\text{del}_G(f, \mathcal{H})).$$

Proof The inclusion (5-1) is clear from the definitions.

For (5-2), if $J \in E(\text{del}_G(f, \mathcal{H})) \setminus E(f, \mathcal{H})$, then $f(P)$ is not adjacent to J on either side, but on at least one side the only hyperplanes separating J from $f(P)$ are in G . In fact this happens on exactly one side since G is a mutually crossing set and its members cannot be separated by J . In particular this means $J \in T(\text{del}_G(f, \mathcal{H}))$, and therefore $J \in T(\text{del}_G(f, \mathcal{H})) \setminus T(f, \mathcal{H})$ since $T(f, \mathcal{H}) \subset E(f, \mathcal{H})$. This situation is indicated in Figure 23, left.

Conversely if $J \in T(\text{del}_G(f, \mathcal{H})) \setminus T(f, \mathcal{H})$, then either $f(P)$ is not adjacent to J on either side, in which case we are in the same situation as above, or $f(P)$ is adjacent to J on both sides. But in the latter case this adjacency remains true after deletion of G , which contradicts $J \in T(\text{del}_G(f, \mathcal{H}))$.

Thus, $J \in E(\text{del}_G(f, \mathcal{H})) \setminus E(f, \mathcal{H})$ if and only if $J \in T(\text{del}_G(f, \mathcal{H})) \setminus T(f, \mathcal{H})$, which gives (5-2).

In the description of $J \in E(\text{del}_G(f, \mathcal{H})) \setminus E(f, \mathcal{H})$, we note that the hyperplanes of G separating J from $f(P)$ cannot themselves be separated from $f(P)$ (on the side not containing J) by any other hyperplanes, because this would contradict $J \in E(\text{del}_G(f, \mathcal{H}))$.




Figure 23: Changes caused by deleting G . Left, J is not in $E(f, \mathcal{H})$ but is in $T(\text{del}_G(f, \mathcal{H}))$. Middle, J is in $E \setminus T$ both before and after. Right, J is in $T(f, \mathcal{H})$ but not in $T(\text{del}_G(f, \mathcal{H}))$.

Thus those hyperplanes are themselves extremal. We conclude that, if $G \cap E(f, \mathcal{H}) = \emptyset$, then $E(\text{del}_G(f, \mathcal{H})) \setminus E(f, \mathcal{H}) = \emptyset$, giving (5-3).

Finally to show (5-4) when $G \cap E(f, \mathcal{H}) = \emptyset$, note first that (5-2) and (5-3) imply that $T(\text{del}_G(f, \mathcal{H})) \subset T(f, \mathcal{H})$. Now if $J \in T(f, \mathcal{H}) \setminus T(\text{del}_G(f, \mathcal{H}))$, then on the side of J where $f(P)$ is not adjacent there must only be hyperplanes of G separating J from $f(P)$, whose deletion makes $f(P)$ adjacent on that side; see Figure 23, right. But this contradicts the assumption that $G \cap E(f, \mathcal{H}) = \emptyset$. \square

We can now obtain the following lemma, which in the simplest case shows that moves and deletions commute.

Lemma 5.10 *In the notation of Proposition 5.7, if $G \cap E(f, \mathcal{H}) = \emptyset$, then the following diagram commutes:*

$$\begin{array}{ccc}
 (f, \mathcal{H}) & \xrightarrow{\text{del}_G} & (f', \mathcal{H} \setminus G) \\
 \downarrow \text{Move} & & \downarrow \text{Move} \\
 (f_1, \mathcal{H}) & \xrightarrow{\text{del}_G} & (f'_1, \mathcal{H} \setminus G)
 \end{array}$$

Proof By Lemma 5.9, $T(f, \mathcal{H}) = T(\text{del}_G(f, \mathcal{H}))$. This means that the Move operation on both (f, \mathcal{H}) and $(f', \mathcal{H} \setminus G)$ affects exactly the same set of hyperplanes, and in exactly the same way. That is, for $J \in T(f, \mathcal{H})$, the subset $P_0 \subset P$ whose f -image is on the adjacent side of J is also the subset whose f' -image is on the adjacent side, since the deletion of G does not affect this. The lemma follows. \square

Now we consider the general situation, where some hyperplanes of G may be in $E(f, \mathcal{H})$.





Figure 24: The contradictions arising from hyperplanes in G' not being mutually crossing.

Lemma 5.11 *Let G be a mutually crossing set in \mathcal{H} and define*

$$G' = (G \cup T(\text{del}_G(f, \mathcal{H}))) \setminus T(f, \mathcal{H}) \quad \text{and} \quad K = T(f, \mathcal{H}) \setminus (G \cup T(\text{del}_G(f, \mathcal{H}))).$$

Then G' is a mutually crossing set, every hyperplane $H \in K$ crosses all hyperplanes of $\mathcal{H}_f \setminus (G \cup \{H\})$, and there exists a map $g: P \rightarrow X(\mathcal{H}_f \setminus (G \cup G' \cup K))$ such that the following diagram commutes:



Proof First we prove that $G' = \{G'_1, \dots, G'_k\}$ is a mutually crossing set. Suppose that G'_i and G'_j do not cross. Then they cannot both be in G by hypothesis.

Assume first that G'_i and G'_j are in $T(\text{del}_G(f, \mathcal{H})) \setminus T(f, \mathcal{H})$, which is the same as $E(\text{del}_G(f, \mathcal{H})) \setminus E(f, \mathcal{H})$ by Lemma 5.9. Let s_i be the side of G'_i that contains G'_j , and define s_j similarly (Figure 24, left). Then $f'(P)$ cannot be adjacent to G'_i on the side s_i , since part of $f'(P)$ is separated from G'_i by G'_j . Therefore, since $G'_i \in E(\text{del}_G(f, \mathcal{H}))$, $f'(P)$ must be adjacent to G'_i on the opposite side, \bar{s}_i . Similarly $f'(P)$ must be adjacent to G'_j on the opposite side, \bar{s}_j . Note that \bar{s}_i and \bar{s}_j are disjoint. On the other hand, since G'_i and G'_j are not in $E(f, \mathcal{H})$, there must be some $G_k \in G$ in \bar{s}_i separating

G'_i from some point of $f(P)$, and similarly $G_l \in G$ in \bar{s}_j separating G'_j from some point of $f(P)$. But this is not possible since G_k crosses G_l .

Now assume $G'_i \in E(\text{del}_G(f, \mathcal{H})) \setminus E(f, \mathcal{H})$ and $G'_j \in G \setminus T(f, \mathcal{H})$ (Figure 24, right). If $G'_j \in E(f, \mathcal{H}) \setminus T(f, \mathcal{H})$ then it crosses G'_i by Lemma 5.8, so we may assume $G'_j \notin E(f, \mathcal{H})$. Define s_i as before. Now since $G'_j \notin E(f, \mathcal{H})$, on the side of G'_j contained in s_i , there must be another hyperplane H separating G'_j from a point of $f(P)$. This H cannot be in G since G'_j crosses H , so it is not deleted and hence $f'(P)$ is not adjacent to G'_i on the s_i side. As above there must therefore be a $G_k \in G$ on the \bar{s}_i side. This G_k cannot cross G'_j ; again a contradiction.

To see that each hyperplane of K crosses all other hyperplanes of $\mathcal{H}_f \setminus G$, note that K is contained in $E(\text{del}_G(f, \mathcal{H})) \setminus T(\text{del}_G(f, \mathcal{H}))$ by (5-1) of Lemma 5.9, and then use Lemma 5.8.

To finish the argument, we claim that all we have to check is that the set of hyperplanes that are either transitional for a Move operation or deleted along each side of the diagram is the same. This is because of the relations

$$\text{Trim} \circ \text{Move}(f, \mathcal{H}) = \text{del}_{T(f, \mathcal{H})} \circ \text{Trim}(f, \mathcal{H})$$

and

$$\text{Trim} \circ \text{del}_G = \text{del}_G \circ \text{Trim}.$$

which follow directly from the definitions. With these relations, we can simplify each side of the diagram to a single $\text{Trim} \circ \text{del}_V$ where V is the union of hyperplanes from all the deletion and Move steps on that side.

Thus, comparing the left side of the diagram with the top arrow and right side, it remains to check that

$$(5-5) \quad T(f, \mathcal{H}) \cup G' = G \cup T(\text{del}_G(f, \mathcal{H})) \cup K.$$

Using the definitions of G' and K , we see that both sides are equal to

$$G \cup T(\text{del}_G(f, \mathcal{H})) \cup T(f, \mathcal{H}). \quad \square$$

Every hyperplane $H \in \mathcal{H}_f$ becomes extremal at some point along the move sequence. We want to understand how deletions affect when this occurs. Note that when a hyperplane $H \in \mathcal{H}_f$ becomes extremal, it need not become transitional, and what happens can change with the deletion of a nearby hyperplane.

For a hyperplane $H \in \mathcal{H}_f$, define

$$e_H(f, \mathcal{H})$$

to be the first index i such that $H \in \mathbf{E}(f_i, \mathcal{H})$.

Lemma 5.12 *For any mutually crossing set $G \subset \mathcal{H}_f$, and any $H \in \mathcal{H}_f \setminus G$,*

$$(5-6) \quad e_H(f, \mathcal{H}) = e_H(\text{del}_G(f, \mathcal{H})) + \delta$$

for some $\delta \in \{0, 1\}$.

Proof If $e_H(f, \mathcal{H}) = 0$ then H is already in $\mathbf{E}(f, \mathcal{H})$, which implies $H \in \mathbf{E}(\text{del}_G(f, \mathcal{H}))$ by (5-1) of Lemma 5.9. The equality (5-6) follows with $\delta = 0$.

Thus we may assume $e_H(f, \mathcal{H}) > 0$. Suppose that $e_H(\text{del}_G(f, \mathcal{H})) = 0$. This means that $H \in \mathbf{E}(\text{del}_G(f, \mathcal{H})) \setminus \mathbf{E}(f, \mathcal{H})$, which implies (as in the proof of Lemma 5.11) that there are some elements G_i of G which separate H from $f(P)$ on one side s_i , so that $f(P)$ is adjacent to G_i on the s_i side. But $f(P)$ is not adjacent to G_i on the other side because H is there, which means $G_i \in \mathbf{T}(f, \mathcal{H})$. But this implies that, in $\text{Move}(f, \mathcal{H})$, all G_i as above are no longer in the set of separating hyperplanes, and hence $H \in \mathbf{E}(\text{Move}(f, \mathcal{H}))$, so $e_H(f, \mathcal{H}) = 1$. This gives (5-6) with $\delta = 1$.

From now on we can assume $e_H(f, \mathcal{H}) > 0$ and $e_H(\text{del}_G(f, \mathcal{H})) > 0$, and prove the statement by induction on the cardinality of \mathcal{H}_f (the case $|\mathcal{H}_f| = 2$ is easy, and already covered by the previous paragraphs).

Let $(f', \mathcal{H}') = \text{del}_G(f, \mathcal{H})$, $(f_1, \mathcal{H}) = \text{Move}(f, \mathcal{H})$, and $(f'_1, \mathcal{H}') = \text{Move}(f', \mathcal{H}')$. By definition (since $e_H(f, \mathcal{H}) > 0$ and $e_H(f', \mathcal{H}') > 0$),

$$e_H(f, \mathcal{H}) = e_H(f_1, \mathcal{H}) + 1 \quad \text{and} \quad e_H(f', \mathcal{H}') = e_H(f'_1, \mathcal{H}') + 1.$$

Thus it will suffice to prove

$$(5-7) \quad e_H(f_1, \mathcal{H}) = e_H(f'_1, \mathcal{H}') + \delta.$$

Consider the warmup case when $G \cap \mathbf{E}(f, \mathcal{H}) = \emptyset$. By Lemma 5.10,

$$(f'_1, \mathcal{H}') = \text{del}_G(f_1, \mathcal{H}).$$

Since $|\mathcal{H}_{f_1}| < |\mathcal{H}_f|$, the inductive hypothesis gives us

$$e_H(f_1, \mathcal{H}) = e_H(f'_1, \mathcal{H}') + \delta$$

for $\delta = 0$ or 1 , proving (5-7) and hence (5-6).

Now in the general case, we use the diagram of Lemma 5.11. Note first that the value of e_H is not affected by a Trim operation. This is because Trim does not affect the set \mathcal{H}_f , or the membership in \mathbf{E} or \mathbf{T} .

The value of e_H is also unaffected by the del_K arrow on the right side. This is because each hyperplane $H \in K$ crosses every hyperplane in $\mathcal{H}_f \setminus (G \cup \{H\})$, which implies that any hyperplane in K cannot affect the membership in \mathbf{E} or \mathbf{T} of any other hyperplane in $\mathcal{H}_f \setminus G$. Therefore, we see that del_K commutes with the Move sequence on $(f'_1, \mathcal{H}_f \setminus G)$.

The remaining arrow is labeled by $\text{del}_{G'}$, and G' is a mutually crossing set. Thus by induction we know

$$e_H(f_1, \mathcal{H}) = e_H(\text{del}_{G'}(f_1, \mathcal{H})) + \delta$$

for $\delta = 0$ or 1 . Again the equality (5-7) follows. □

We are now ready to prove the stability result for move sequences.

Proof of Proposition 5.7 We need to prove the following statement: for each i and $p \in P$, the set of $H \in \mathcal{H}_f \setminus G$ such that H separates $f_i(p)$ from $f'_i(p)$ is mutually crossing.

Note that when a sequence $(f_i(p))$ crosses a hyperplane H , it can only happen in the transition from f_j to f_{j+1} where $j = e_H(f, \mathcal{H})$. Moreover it must be that $H \in \mathbf{T}(f_j, \mathcal{H})$, and that $f_j(p)$ is on the side of H where $f_j(P)$ is adjacent. If the sequence $(f_i(p))$ does not cross H , then H will either remain in \mathcal{H}_{f_i} for all $i \geq j$, or it will be crossed only by points on the other side and not by $f_i(p)$. The same holds for the sequence (f'_i) .

Now suppose that H_1 and H_2 separate $f_i(p)$ from $f'_i(p)$, for some i , and H_1 does not cross H_2 . Since $f_0(p)$ and $f'_0(p)$ are on the same side of both hyperplanes, we may assume that H_1 separates them from H_2 . Hence there exists a (different) time i at which one of them is still on the original side of H_1 , whereas the other has crossed both hyperplanes; fix such an i .

Suppose first that $f'_i(p)$ is the one which lies on the other side of H_2 . Let $j = e_{H_1}(f', \mathcal{H})$. Then since H_1 does not cross H_2 , $f'_{j+1}(p)$ has not yet crossed H_2 . Thus we must have

$$j = e_{H_1}(f', \mathcal{H}) < e_{H_2}(f', \mathcal{H}) < i$$

and $H_1 \in \mathbf{T}(f'_j, \mathcal{H})$, with $f'_j(p)$ on the side where $f'_j(P)$ is adjacent to H_1 .

Now $k = e_{H_1}(f, \mathcal{H}) \leq j + 1$ by Lemma 5.12. Since $j + 1 < i$ and $f_i(p)$ has not crossed H_1 , it must be that the side of H_1 where $f_k(P)$ is adjacent is the one opposite from $f_k(p)$, the one containing H_2 .

But this means that H_2 can no longer be in \mathcal{H}_{f_k} , which can only be if

$$e_{H_2}(f, \mathcal{H}) \leq k - 1 \leq j.$$

Thus, again by Lemma 5.12, $e_{H_2}(f', \mathcal{H}) \leq j$, which is a contradiction.

We conclude that H_1 does not cross H_2 , which is what we wanted.

The case where $f_i(p)$ crosses the hyperplanes and $f'_i(p)$ does not is handled similarly. The main difference is that in this case, instead of using $e_{H_1}(f, \mathcal{H}) \leq e_{H_1}(f', \mathcal{H}) + 1$, we use $e_{H_1}(f', \mathcal{H}) \leq e_{H_1}(f, \mathcal{H})$, which creates a “+1” there, which is then lost in the other application of Lemma 5.12. □

5.4 Completing the proof of Theorem 5.1

Property (1) follows from Lemma 5.3, while property (2) follows from the construction, where for both properties we use the fact that cubes have diameter 1 in the d_∞ metric. Property (5) is also easily seen to hold by construction, and more specifically it follows from the fact that the sets $E(f, \mathcal{H})$ and $T(f, \mathcal{H})$ only depend on the image of f (and hence that this will remain true throughout the move sequence).

For property (3), observe that for a fixed $p \in P$, no hyperplane H separates $f_i(p)$ from $f_{i+1}(p)$ for two different values of i . It follows that any combinatorial path obtained by concatenating a choice of geodesics from $f_i(p)$ to $f_{i+1}(p)$ for $0 \leq i < n$ is an ℓ^1 -geodesic in X . This proves (3).

For property (4), by definition of the contraction sequence, the hyperplanes that separate any given $f(p)$ from $f_n(q)$ for $p, q \in P$ are contained in \mathcal{H}_f . On the other hand, a hyperplane $H \in \mathcal{H}_f$ cannot separate every point in $f(P)$ from any fixed vertex of X , because there are elements of $f(P)$ on both sides of H , by definition of \mathcal{H}_f . Hence, there can be no hyperplane separating $f(P)$ from a point in $f_n(P)$. This gives property (4).

Property (6) is a direct consequence of Proposition 5.7 and Lemma 5.6. □

6 Proofs of the main theorems

We are now almost ready to prove our main theorems, Theorem E and Theorem B. The main bit of work here is Proposition 6.1, which compiles our preceding stability results into a useful form for our current purposes.

Proposition 6.1 *Let $(\mathcal{X}, \mathfrak{S})$ be a G -colorable HHS for $G < \text{Aut}(\mathfrak{S})$. For any $k \in \mathbb{N}$, there exists $K_3 = K_3(k, \mathfrak{S}) > 0$ such that the following holds.*

Suppose that $F, F' \subset \mathcal{X}$ are finite subsets satisfying $|F|, |F'| \leq k$, let $g \in G$, and suppose that $d_{\text{Haus}}(gF, F') \leq 1$. Choose any map $\iota_F: F \sqcup F' \rightarrow F$ such that $\iota_F(f) = f$ if $f \in F$ and $d_{\mathcal{X}}(g\iota_F(f), f) \leq 1$ if $f \in F'$. Also, choose a map $\iota_{F'}: F \sqcup F' \rightarrow F'$ such that $\iota_{F'}(f) = f$ if $f \in F'$ and $d_{\mathcal{X}}(g\iota_{F'}(f), f) \leq 1$ if $f \in F$. Consider:

- The cube complexes $\mathcal{Q}_F, \mathcal{Q}_{F'}$ produced by Theorem 4.1 with associated maps $\Phi_F, \Phi_{F'}$ to \mathcal{X} , and $\psi_F, \psi_{F'}$ from F, F' to $\mathcal{Q}_F, \mathcal{Q}_{F'}$.
- The sequences of contractions $\{(\psi_F)_i = \psi_i\}_{i \leq n_{\psi_F}}$ and $\{(\psi_{F'})_i = \psi'_i\}_{i \leq n_{\psi_{F'}}}$ produced by Theorem 5.1. Set $n_{\psi_F} = n_F$ and $n_{\psi_{F'}} = n_{F'}$.

Then

- (1) $|n_F - n_{F'}| < K_3$,
- (2) for each $i \in \{1, \dots, \max\{n_F, n_{F'}\}\}$ and any $f \in F \sqcup F'$,

$$d_{\mathcal{X}}(g \circ \Phi_F \circ \psi_i(\iota_F(f)), \Phi_{F'} \circ \psi'_i(\iota_{F'}(f))) \leq K_3,$$
- (3) $\text{diam}_{\mathcal{X}}(\Phi_F(\psi_{n_F}(F))) < K_3$.

More visually, item (2) says that the diagram

$$(6-1) \quad \begin{array}{ccccc} & & F & \xrightarrow{\psi_i = (\psi_F)_i} & \mathcal{Q}_F & & \\ & \nearrow \iota_F & & & & \searrow g \circ \Phi_F & \\ F \sqcup F' & & & & & & \mathcal{X} \\ & \searrow \iota_{F'} & F' & \xrightarrow{\psi'_i = (\psi_{F'})_i} & \mathcal{Q}_{F'} & \nearrow \Phi_{F'} & \end{array}$$

coarsely commutes.

Proof We will use the output and notation of Theorem 4.1, and in particular the CAT(0) cube complex \mathcal{Q}_0 obtained from both \mathcal{Q}_F and $\mathcal{Q}_{F'}$ by collapsing at most $N = N(k, \mathfrak{S}) > 0$ hyperplanes, with hyperplanes collapse maps h and h' .

We will also use the notation of Theorem 5.1, in particular the notation $\{f_i \mid i \leq n_f\}$ for the sequence of maps starting with f and ending with a map with bounded image.

We have $\psi := h \circ \psi_F \circ \iota_F = h' \circ \psi_{F'} \circ \iota_{F'}$, as stated in Theorem 4.1. By Theorem 5.1, composing, say, ψ_F with a hyperplane deletion map affects the length of the corresponding sequence of maps by at most 1. In particular, $|n_F - n_\psi| \leq N$ and, similarly, $|n_{F'} - n_\psi| \leq N$ (notice that $n_{\psi_F} = n_{\psi_F \circ \iota_F}$ by Theorem 5.1(5), and a similar statement holds for F'). Hence, conclusion (1) holds for any K_3 larger than $2N$.

We now prove conclusion (2). By Theorem 4.1, diagram (4-4) commutes with error at most $K = K(k, \mathfrak{S})$. For convenience, we reproduce the diagram here:

$$(6-2) \quad \begin{array}{ccccc} & & F & \xrightarrow{\psi_F} & Q_F \\ & \nearrow \iota_F & & & \downarrow \eta \\ F \sqcup F' & & & & Q_0 \\ & \searrow \iota_{F'} & & & \uparrow \eta' \\ & & F' & \xrightarrow{\psi_{F'}} & Q_{F'} \end{array} \quad \begin{array}{ccc} & & \searrow g \circ \Phi_F \\ & & \chi \\ & \xrightarrow{\Phi_0} & \\ & & \nearrow \Phi_{F'} \end{array}$$

For any $f \in F \sqcup F'$,

$$d_\chi(g \circ \Phi_F \circ (\psi_F)_i(\iota_F(f)), \Phi_0 \circ \eta \circ (\psi_F)_i(\iota_F(f))) \leq K.$$

By Theorem 5.1, $d_\infty(\eta \circ (\psi_F)_i(\iota_F(f)), (\eta \circ \psi_F)_i(\iota_F(f))) \leq N$, and hence

$$d_\chi(\Phi_0 \circ \eta \circ (\psi_F)_i(f), \Phi_0 \circ (\eta \circ \psi_F)_i(f)) \leq K' = K'(k, \mathfrak{S})$$

since Φ_0 is a quasi-isometric embedding with controlled constants (and the dimension of Q_0 is bounded in terms of \mathfrak{S} by Theorem 4.1, so that the ℓ^∞ and ℓ^1 metrics on it are uniformly quasi-isometric). The triangle inequality then gives

$$d_\chi(g \circ \Phi_F \circ (\psi_F)_i(\iota_F(f)), \Phi_0 \circ (\eta \circ \psi_F)_i(\iota_F(f))) \leq K' + K.$$

Similarly, we get

$$d_\chi(\Phi_{F'} \circ (\psi_{F'})_i(\iota_{F'}(f)), \Phi_0 \circ (\eta' \circ \psi_{F'})_i(\iota_{F'}(f))) \leq K' + K.$$

By Theorem 5.1(5), $(\eta' \circ \psi_{F'})_i(\iota_{F'}(f)) = (\eta' \circ \psi_{F'} \circ \iota_{F'})_i(f) = (\eta \circ \psi_F \circ \iota_F)_i(f)$, and hence conclusion (2) holds for any K_3 larger than $2(K' + K)$.

Finally, to prove (3), we now bound the diameter of $\Phi_F(\psi_{n_F}(F))$. Similarly to above, Φ_F is a quasi-isometric embedding with constants controlled in terms of k and \mathfrak{S} even when we endow Q_F with the ℓ^∞ metric. Since $\text{diam}_{Q_F}(\psi_{n_F}(F)) \leq 1$ (in the ℓ^∞ metric) by Theorem 5.1, this gives the required bound on $\text{diam}_\chi(\Phi_F(\psi_{n_F}(F)))$. \square

6.1 Barycenters: proof of Theorem E

Our next goal is to prove Theorem E. To do so, we'll need the precise definition of stable barycenter:

Definition 6.2 For a metric space X a *stable barycenter map* for k points is a map $\tau : X^k \rightarrow X$ which is

- *permutation invariant*, meaning $\tau \circ \pi = \tau$ for any $\pi : X^k \rightarrow X^k$ that is a permutation of the factors;
- *coarsely Lipschitz*, meaning there exists $\kappa_1 > 0$ such that for $x, x' \in X^k$,

$$d_X(\tau(x), \tau(x')) \leq \kappa_1 d_{X^k}(x, x') + \kappa_1.$$

We further say that τ is *coarsely equivariant* with respect to a group Γ acting on X by isometries if there exists $\kappa_1 > 0$ such that for all $g \in \Gamma$

$$d_X(g\tau(x), \tau(gx)) \leq \kappa_1,$$

where Γ acts on X^k diagonally.

We now prove that colorable HHSs admit stable coarsely equivariant barycenters, with the following version slightly more general than Theorem E:

Theorem 6.3 *Let $(\mathcal{X}, \mathfrak{S})$ be a G -colorable HHS for $G < \text{Aut}(\mathfrak{S})$. Then \mathcal{X} admits coarsely G -equivariant stable barycenters for k points, for any $k \geq 1$. Moreover, the coarse barycenter of a set F is contained in the hierarchical hull of F .*

Proof We use the notation from the statement of Proposition 6.1.

To define a barycenter $\tau(f_1, \dots, f_k)$, we consider $F = \{f_i\}$, set $x_F = \Phi_F(\psi_{n_F}(F))$, and let $\tau(f_1, \dots, f_k)$ be an arbitrary point in x_F ; we make the choice depending on the set F only, so that permutation invariance is achieved.

This choice does not matter for our purposes since $\text{diam}_{\mathcal{X}}(\Phi_F(\psi_{n_F}(F))) < K_3$ by Proposition 6.1.

Now suppose that (f'_1, \dots, f'_k) is such that $d_{\text{Haus}}(\{f'_i\}, \{f_i\}) \leq 1$, and set $F' = \{f'_i\}$. Without loss of generality, assume that $n_F \geq n_{F'}$, where we note that $n_F - n_{F'} < K_3$ by part (1) of Proposition 6.1. Part (2) of Proposition 6.1 now implies that for any $f \in F \sqcup F'$,

$$d_{\mathcal{X}}(\Phi_F(\psi_{n_F}(\iota_F(f))), \Phi_{F'}(\psi'_{n_{F'}}(\iota_{F'}(f)))) < K_3.$$

But $\psi'_{n_{F'}} = \psi'_{n_F}$ since $n_F \geq n_{F'}$, so we can conclude that

$$d_{\mathcal{X}}(\Phi_F(\psi_{n_F}(\iota_F(f))), \Phi_{F'}(\psi'_{n_{F'}}(\iota_{F'}(f)))) < K_3.$$

Finally, the fact that $\text{diam}_{\mathcal{X}}(\Phi_F(\psi_{n_F}(F))) < K_3$ and $\text{diam}_{\mathcal{X}}(\Phi_{F'}(\psi'_{n_{F'}}(F'))) < K_3$ gives that

$$\text{diam}_{\mathcal{X}}(x_F \cup x_{F'}) < 3K_3.$$

Setting $\kappa_1 = 3K_3$, we get that τ is κ_1 -coarsely Lipschitz.

Finally, coarse equivariance follows similarly, applying Proposition 6.1 with $F' = gF$, as follows. First, as above we can assume $n_F \geq n_{gF}$, for otherwise we can swap the roles of F and gF , by considering the automorphism g^{-1} . We still have $n_F - n_{gF} < K_3$. Part (2) of Proposition 6.1 implies that for any $f \in F \sqcup gF$,

$$d_{\mathcal{X}}(g \circ \Phi_F(\psi_{n_F}(\iota_F(f))), \Phi_{gF}(\psi'_{n_{gF}}(\iota_{gF}(f)))) < K_3.$$

As above, we conclude that

$$\text{diam}_{\mathcal{X}}(g(x_F) \cup x_{gF}) < 3K_3,$$

which completes the proof. □

6.2 Bicombability: Proof of Theorem B

We begin with the formal definition of bicombing which is appropriate for our context; see [1]. In the following definition, we adopt the convention that if $\phi: [0, a] \rightarrow X$ is a map, then we trivially extend ϕ by $\phi(t) = \phi(a)$ for all $t > a$.

Definition 6.4 A discrete, bounded, quasigeodesic bicombing of a metric space X consists of a family of discrete paths $\{\Omega_{x,y}\}_{x,y \in X}$ which are, for some constant $\kappa_2 > 0$,

- (1) *quasigeodesic*, meaning that for any $x, y \in X$ with $d = d_X(x, y)$, there exists $n_{x,y} \leq \kappa_2 d + \kappa_2$ such that the path $\Omega_{x,y}: \{0, \dots, n_{x,y}\} \rightarrow X$ is a (κ_2, κ_2) -quasi-isometric embedding with $\Omega_{x,y}(0) = x$ and $\Omega_{x,y}(n_{x,y}) = y$; and
- (2) *fellow-traveling*, meaning that if $x', y' \in X$ with $d' = d_X(x', y')$ and

$$d_X(x, x'), d_X(y, y') \leq 1,$$

then for all $t \in \{0, \dots, \max\{n_{x,y}, n_{x',y'}\}\}$,

$$d_X(\Omega_{x,y}(t), \Omega_{x',y'}(t)) \leq \kappa_2.$$




Figure 25: A simple example of bicombing paths, building on the hierarchical setup from Figure 18. Deleting the orange hyperplanes from \mathcal{Q}_F and $\mathcal{Q}_{F'}$ results in perturbing the contraction paths in \mathcal{Q}_0 .

In addition, we say that $\{\Omega_{x,y}\}_{x,y \in X}$ is Γ -coarsely equivariant with respect to a group $\Gamma < \text{Isom}(X)$ if for any $g \in \Gamma$ and $x, y \in X$ and $t \in \{0, \dots, \max\{n_{x,y}, n_{x',y'}\}\}$,

$$d_X(g \cdot \Omega_{x,y}(t), \Omega_{g \cdot x, g \cdot y}(t)) < \kappa_2.$$

Finally, we recall the following definition from [7], which was inspired by the paths constructed in [43]:

Definition 6.5 For $D \geq 1$, a path γ in \mathcal{X} is a D -hierarchy path if

- (1) γ is a (D, D) -quasigeodesic,
- (2) for each $W \in \mathfrak{S}$, $\pi_W \circ \gamma$ is an unparametrized (D, D) -quasigeodesic.

We can now prove that colorable HHSs admit discrete, bounded, quasigeodesic, coarsely equivariant bicombings by hierarchy paths.

Theorem 6.6 Let $(\mathcal{X}, \mathfrak{S})$ be a G -colorable HHS with $G < \text{Aut}(\mathfrak{S})$. Then there exists $D > 0$ such that $(\mathcal{X}, \mathfrak{S})$ admits a coarsely G -equivariant, discrete, bounded, quasigeodesic bicombing by D -hierarchy paths.

Proof Let $(\mathcal{X}, \mathfrak{S})$ be a colorable HHS. We again use the notation from the statement of Proposition 6.1, where now $F = \{x, y\}$ and $F' = \{x', y'\}$ with $d_{\mathcal{X}}(x, x') \leq 1$ and $d_{\mathcal{X}}(y, y') \leq 1$. We make a blanket observation that $k = 2$ and so the constant K_3 in Proposition 6.1 depends only on $(\mathcal{X}, \mathfrak{S})$.

Coarse equivariance can be obtained using the argument below, setting $F' = gF$. We omit the details for readability.

Construction of the bicombing paths Let $\psi = \psi_0$ and define a map

$$\omega_{x,y}: \{0, \dots, 2n_{x,y}\} \rightarrow \mathcal{Q}_F$$

by

$$\omega_{x,y}(i) = \begin{cases} \psi_i(x) & \text{if } i \in \{0, \dots, n_{x,y}\}, \\ \psi_{2n_{x,y}-i}(y) & \text{if } i \in \{n_{x,y} + 1, \dots, 2n_{x,y}\}. \end{cases}$$

We claim that $\omega_{x,y}$ is a (C, C) -quasigeodesic in the ℓ^1 metric on \mathcal{Q}_F , for some uniform C . First, the points $\omega_{x,y}(0), \dots, \omega_{x,y}(n_{x,y})$ appear on an ℓ^1 -geodesic γ_1 from $\omega_{x,y}(0)$ to $\omega_{x,y}(n_{x,y})$ in the given order by Theorem 5.1(3), and the same holds for $\omega_{x,y}(n_{x,y} + 1), \dots, \omega_{x,y}(2n_{x,y})$ for some ℓ^1 -geodesic γ_2 from $\omega_{x,y}(n_{x,y} + 1)$ to $\omega_{x,y}(2n_{x,y})$. Moreover, consecutive $\omega_{x,y}(i)$ are uniformly close to each other, since they are at distance at most 1 in the ℓ^∞ metric, which is uniformly quasi-isometric to the ℓ^1 metric with constant only depending on the dimension of \mathcal{Q}_F , which in turn only depends on \mathfrak{S} .

Let γ be the concatenation of γ_1 , an ℓ^1 -geodesic from $\omega_{x,y}(n_{x,y})$ to $\omega_{x,y}(n_{x,y} + 1)$, and γ_2 . Since no hyperplane can separate $\{x, y\}$ from $\psi_{n_{x,y}}(x)$ or $\psi_{n_{x,y}}(y)$, again by Theorem 5.1, we see that γ crosses each hyperplane at most once, and is therefore an ℓ^1 -geodesic. Since $\omega_{x,y}(n_{x,y})$ and $\omega_{x,y}(n_{x,y} + 1)$ are just opposite corners of a cube, we see that the $\omega_{x,y}(i)$ appear along an ℓ^1 -geodesic in the given order, and with uniformly spaced gaps. This shows that $\omega_{x,y}$ is a (C, C) -quasigeodesic in the ℓ^1 metric, for C depending only on the dimension of \mathcal{Q}_F and hence only on \mathfrak{S} .

It follows then that the composition

$$\Omega_{x,y} = \Phi_F \circ \omega_{x,y}: \{0, \dots, 2n_{x,y}\} \rightarrow \mathcal{X}$$

is a (K_4, K_4) -quasigeodesic in \mathcal{X} with $K_4 = K_4(\mathfrak{S})$. We can perturb it a uniformly bounded amount at the endpoints to make sure that the endpoints are x and y ; with a slight abuse of notation we still denote the perturbation by $\Omega_{x,y}$ and the quasi-isometry constants by K_4 .

This proves that Definition 6.4(1) holds for the family $\{\Omega_{x,y}\}_{x,y \in \mathcal{X}}$.

Fellow-traveling We now prove the fellow-traveling condition in Definition 6.4(2) holds. Once again adopting our previous notation, we want to prove that there exists $\kappa_2 = \kappa_2(\mathcal{X}, \mathfrak{S}) > 0$ such that for any $t \in \{0, \dots, \max\{2n_{x,y}, 2n_{x',y'}\}\}$,

$$(6-3) \quad d_{\mathcal{X}}(\Omega_{x,y}(t), \Omega_{x',y'}(t)) < \kappa_2.$$

Without loss of generality, suppose that $n_{x,y} \geq n_{x',y'}$ and recall that Proposition 6.1(1) gives that $\delta = n_{x,y} - n_{x',y'} < K_3$, where K_3 depends only on $(\mathcal{X}, \mathfrak{S})$. There are four cases to consider:

- (i) When $0 \leq i \leq n_{x',y'}$, where $\Omega_{x,y}$ and $\Omega_{x',y'}$ are defined using x and x' , respectively.
- (ii) When $n_{x',y'} < j \leq n_{x,y}$, where $\Omega_{x,y}$ is defined using x whereas $\Omega_{x',y'}$ is defined using y' .
- (iii) When $n_{x,y} < q \leq 2n_{x',y'}$, where both $\Omega_{x,y}$ and $\Omega_{x',y'}$ are nonconstant and defined using y and y' , respectively.
- (iv) When $2n_{x',y'} < r \leq 2n_{x,y}$, where $\Omega_{x,y}$ is nonconstant but

$$\Omega_{x',y'}(j) = \Omega_{x',y'}(2n_{x',y'})$$

for all such j .

In what follows, we will repeatedly use the fact that $\Omega_{x,y}$ and $\Omega_{x',y'}$ are (K_4, K_4) -quasigeodesics. Also, set $K_5 = K_4 \cdot (2\delta) + K_4$.

In case (i), equation (6-3) follows immediately from Proposition 6.1(2) with $\kappa_2 = K_3$.

In case (ii),

$$d_{\mathcal{X}}(\Omega_{x,y}(n_{x',y'}), \Omega_{x,y}(j)) < K_5 \quad \text{and} \quad d_{\mathcal{X}}(\Omega_{x',y'}(n_{x',y'}), \Omega_{x',y'}(j)) < K_5,$$

while Proposition 6.1(2) gives

$$d_{\mathcal{X}}(\Omega_{x,y}(j + 2\delta), \Omega_{x',y'}(j)) < K_3,$$

so the triangle inequality implies that (6-3) holds in this case with $\kappa_2 = 2K_5 + K_3$.

In case (iii),

$$d_{\mathcal{X}}(\Omega_{x,y}(q), \Omega_{x,y}(q + 2\delta)) < K_5$$

and Proposition 6.1(2) provides

$$d_{\mathcal{X}}(\Omega_{x,y}(q + 2\delta), \Omega_{x',y'}(q)) < K_3,$$

so the triangle inequality implies that (6-3) holds with $\kappa_2 = K_3 + K_5$.

Finally, in case (iv),

$$d_{\mathcal{X}}(\Omega_{x,y}(2n_{x',y'}), \Omega_{x,y}(2n_F)) < K_5$$

and Proposition 6.1(2) provides

$$d_{\mathcal{X}}(\Omega_{x,y}(2n_{x',y'}), \Omega_{x',y'}(2n_{x',y'})) < K_3.$$

Since $\Omega_{x',y'}(2n_F) = \Omega_{x',y'}(2n_{x',y'})$ by convention, the triangle inequality implies that (6-3) holds with $\kappa_2 = K_3 + K_5$.

Hence we may set $\kappa_2 = 2K_5 + K_3$ to complete the proof of the fellow-traveling condition in Definition 6.4(2). This completes the proof that these paths give a bicombing.

Hierarchy paths To finish the proof, we now show that $\Omega_{x,y}$ is a D -hierarchy path for some $D = D(\mathfrak{S}) > 0$ (Definition 6.5). We will use that Φ_F is a K -median map (Theorem 4.1(2)); we now recall what this means.

In a CAT(0) cube complex \mathcal{Q} one can define a map $m_{\mathcal{Q}}: \mathcal{Q}^3 \rightarrow \mathcal{Q}$ (called median), and the only property of this map that we need here is that if x , y , and z appear in this order along an ℓ^1 -geodesic, then $m(x, y, z) = y$. Also, in an HHS \mathcal{X} , there is a map $m_{\mathcal{X}}: \mathcal{X}^3 \rightarrow \mathcal{X}$ called coarse median, whose definition we do not need, and Φ_F being K -median means that for all $x, y, z \in \mathcal{Q}_F$,

$$d_{\mathcal{X}}(\Phi_F(m_{\mathcal{Q}_F}(x, y, z)), m_{\mathcal{X}}(\Phi_F(x), \Phi_F(y), \Phi_F(z))) \leq K.$$

This inequality implies that, for all $i < j < k$, we have that $\Omega_{x,y}(j)$ lies uniformly close to $m_{\mathcal{X}}(\Omega_{x,y}(i), \Omega_{x,y}(j), \Omega_{x,y}(k))$. This is enough to guarantee that the quasi-isometric embedding $\Omega_{x,y}$ is a hierarchy path by [8, Lemma 1.37]. \square

References

- [1] **J M Alonso, M R Bridson**, *Semihyperbolic groups*, Proc. London Math. Soc. 70 (1995) 56–114 MR Zbl
- [2] **G Baumslag, S M Gersten, M Shapiro, H Short**, *Automatic groups and amalgams — a survey*, from “Algorithms and classification in combinatorial group theory” (G Baumslag, C F Miller, III, editors), Math. Sci. Res. Inst. Publ. 23, Springer (1992) 179–194 MR Zbl
- [3] **J A Behrstock**, *Asymptotic geometry of the mapping class group and Teichmüller space*, Geom. Topol. 10 (2006) 1523–1578 MR Zbl

- [4] **J Behrstock, M Hagen, A Martin, A Sisto**, *A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups*, preprint (2020) arXiv 2005.00567
- [5] **J Behrstock, MF Hagen, A Sisto**, *Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups*, Proc. Lond. Math. Soc. 114 (2017) 890–926 MR Zbl
- [6] **J Behrstock, MF Hagen, A Sisto**, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. 21 (2017) 1731–1804 MR Zbl
- [7] **J Behrstock, M Hagen, A Sisto**, *Hierarchically hyperbolic spaces, II: Combination theorems and the distance formula*, Pacific J. Math. 299 (2019) 257–338 MR Zbl
- [8] **J Behrstock, MF Hagen, A Sisto**, *Quasiflats in hierarchically hyperbolic spaces*, Duke Math. J. 170 (2021) 909–996 MR Zbl
- [9] **J Behrstock, B Kleiner, Y Minsky, L Mosher**, *Geometry and rigidity of mapping class groups*, Geom. Topol. 16 (2012) 781–888 MR Zbl
- [10] **J A Behrstock, Y N Minsky**, *Centroids and the rapid decay property in mapping class groups*, J. Lond. Math. Soc. 84 (2011) 765–784 MR Zbl
- [11] **F Berlai, B Robbio**, *A refined combination theorem for hierarchically hyperbolic groups*, Groups Geom. Dyn. 14 (2020) 1127–1203 MR Zbl
- [12] **M Bestvina, K Bromberg, K Fujiwara**, *Constructing group actions on quasi-trees and applications to mapping class groups*, Publ. Math. Inst. Hautes Études Sci. 122 (2015) 1–64 MR Zbl
- [13] **M Bestvina, K Bromberg, K Fujiwara**, *Proper actions on finite products of quasi-trees*, Ann. H. Lebesgue 4 (2021) 685–709 MR Zbl
- [14] **M Bestvina, K Bromberg, K Fujiwara, A Sisto**, *Acyindrical actions on projection complexes*, Enseign. Math. 65 (2019) 1–32 MR Zbl
- [15] **B H Bowditch**, *Coarse median spaces and groups*, Pacific J. Math. 261 (2013) 53–93 MR Zbl
- [16] **B Bowditch**, *Convex hulls in coarse median spaces*, preprint (2018) Available at <http://homepages.warwick.ac.uk/~masgak/papers/hulls-cms.pdf>
- [17] **MR Bridson**, *Semisimple actions of mapping class groups on CAT(0) spaces*, from “Geometry of Riemann surfaces” (FP Gardiner, G González-Diez, C Kourouniotis, editors), London Math. Soc. Lecture Note Ser. 368, Cambridge Univ. Press (2010) 1–14 MR Zbl
- [18] **MR Bridson, A Haefliger**, *Metric spaces of non-positive curvature*, Grundle. Math. Wissen. 319, Springer (1999) MR Zbl
- [19] **P-E Caprace, M Sageev**, *Rank rigidity for CAT(0) cube complexes*, Geom. Funct. Anal. 21 (2011) 851–891 MR Zbl

- [20] **J Chalopin, V Chepoi, A Genevois, H Hirai, D Osajda**, *Helly groups*, preprint (2020) arXiv 2002.06895
- [21] **I Chatterji, G Niblo**, *From wall spaces to CAT(0) cube complexes*, Internat. J. Algebra Comput. 15 (2005) 875–885 MR Zbl
- [22] **V Chepoi**, *Graphs of some CAT(0) complexes*, Adv. in Appl. Math. 24 (2000) 125–179 MR Zbl
- [23] **M Chesser**, *Stable subgroups of the genus 2 handlebody group*, Algebr. Geom. Topol. 22 (2022) 919–971 MR Zbl
- [24] **M T Clay, C J Leininger, J Mangahas**, *The geometry of right-angled Artin subgroups of mapping class groups*, Groups Geom. Dyn. 6 (2012) 249–278 MR Zbl
- [25] **S Dowdall, M G Durham, C J Leininger, A Sisto**, *Extensions of Veech groups, I: A hyperbolic action*, preprint (2020) arXiv 2006.16425
- [26] **S Dowdall, M G Durham, C J Leininger, A Sisto**, *Extensions of Veech groups, II: Hierarchical hyperbolicity and quasi-isometric rigidity*, preprint (2021) arXiv 2111.00685
- [27] **M G Durham**, *The augmented marking complex of a surface*, J. Lond. Math. Soc. 94 (2016) 933–969 MR Zbl
- [28] **M G Durham**, *Elliptic actions on Teichmüller space*, Groups Geom. Dyn. 13 (2019) 415–465 MR Zbl
- [29] **M G Durham, M F Hagen, A Sisto**, *Boundaries and automorphisms of hierarchically hyperbolic spaces*, Geom. Topol. 21 (2017) 3659–3758 MR Zbl
- [30] **M G Durham, M F Hagen, A Sisto**, *Correction to the article “Boundaries and automorphisms of hierarchically hyperbolic spaces”*, Geom. Topol. 24 (2020) 1051–1073 MR Zbl
- [31] **D B A Epstein, J W Cannon, D F Holt, S V F Levy, M S Paterson, W P Thurston**, *Word processing in groups*, Jones and Bartlett, Boston, MA (1992) MR Zbl
- [32] **A Eskin, H Masur, K Rafi**, *Large-scale rank of Teichmüller space*, Duke Math. J. 166 (2017) 1517–1572 MR Zbl
- [33] **S M Gersten, H B Short**, *Rational subgroups of biautomatic groups*, Ann. of Math. 134 (1991) 125–158 MR Zbl
- [34] **M Gromov**, *Hyperbolic groups*, from “Essays in group theory” (S M Gersten, editor), Math. Sci. Res. Inst. Publ. 8, Springer (1987) 75–263 MR Zbl
- [35] **T Haettel, N Hoda, H Petyt**, *Coarse injectivity, hierarchical hyperbolicity, and semi-hyperbolicity*, preprint (2020) arXiv 2009.14053
- [36] **M Hagen**, *Non-colorable hierarchically hyperbolic groups*, Internat. J. Algebra Comput. 33 (2023) 337–350 MR Zbl
- [37] **M Hagen, A Martin, A Sisto**, *Extra-large type Artin groups are hierarchically hyperbolic*, Math. Ann. (2022)

- [38] **M F Hagen, T Susse**, *On hierarchical hyperbolicity of cubical groups*, Israel J. Math. 236 (2020) 45–89 MR Zbl
- [39] **F Haglund, F Paulin**, *Simplicité de groupes d'automorphismes d'espaces à courbure négative*, from “The Epstein birthday schrift” (I Rivin, C Rourke, C Series, editors), Geom. Topol. Monogr. 1, Geom. Topol., Coventry (1998) 181–248 MR Zbl
- [40] **U Hamenstädt**, *Geometry of the mapping class group, II: A biautomatic structure*, preprint (2009) arXiv 0912.0137
- [41] **G C Hruska, D T Wise**, *Finiteness properties of cubulated groups*, Compos. Math. 150 (2014) 453–506 MR Zbl
- [42] **H A Masur, Y N Minsky**, *Geometry of the complex of curves, I: Hyperbolicity*, Invent. Math. 138 (1999) 103–149 MR Zbl
- [43] **H A Masur, Y N Minsky**, *Geometry of the complex of curves, II: Hierarchical structure*, Geom. Funct. Anal. 10 (2000) 902–974 MR Zbl
- [44] **B Miesch**, *Injective metrics on cube complexes*, preprint (2014) arXiv 1411.7234
- [45] **L Mosher**, *Mapping class groups are automatic*, Ann. of Math. 142 (1995) 303–384 MR Zbl
- [46] **G A Niblo, L D Reeves**, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology 37 (1998) 621–633 MR Zbl
- [47] **B Nica**, *Cubulating spaces with walls*, Algebr. Geom. Topol. 4 (2004) 297–309 MR Zbl
- [48] **K Rafi**, *A combinatorial model for the Teichmüller metric*, Geom. Funct. Anal. 17 (2007) 936–959 MR Zbl
- [49] **K Rafi**, *Hyperbolicity in Teichmüller space*, Geom. Topol. 18 (2014) 3025–3053 MR Zbl
- [50] **K Rafi, Y Verberne**, *Geodesics in the mapping class group*, Algebr. Geom. Topol. 21 (2021) 2995–3017 MR Zbl
- [51] **B Robbio, D Spriano**, *Hierarchical hyperbolicity of hyperbolic-2-decomposable groups*, preprint (2020) arXiv 2007.13383
- [52] **J Russell**, *Extensions of multicurve stabilizers are hierarchically hyperbolic*, preprint (2021) arXiv 2107.14116
- [53] **J Russell, D Spriano, H C Tran**, *Convexity in hierarchically hyperbolic spaces*, Algebr. Geom. Topol. 23 (2023) 1167–1248 MR
- [54] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. 71 (1995) 585–617 MR Zbl
- [55] **M Sageev**, *CAT(0) cube complexes and groups*, from “Geometric group theory” (M Bestvina, M Sageev, K Vogtmann, editors), IAS/Park City Math. Ser. 21, Amer. Math. Soc., Providence, RI (2014) 7–54 MR Zbl

- [56] **D Spriano**, *Hyperbolic HHS, II: Graphs of hierarchically hyperbolic groups*, preprint (2018) arXiv 1801.01850
- [57] **P A Storm**, *The Novikov conjecture for mapping class groups as a corollary of Hamenstädt's theorem*, preprint (2005) arXiv math/0504248
- [58] **J Tao**, *Linearly bounded conjugator property for mapping class groups*, *Geom. Funct. Anal.* 23 (2013) 415–466 MR Zbl
- [59] **A J Tromba**, *On a natural algebraic affine connection on the space of almost complex structures and the curvature of Teichmüller space with respect to its Weil–Petersson metric*, *Manuscripta Math.* 56 (1986) 475–497 MR Zbl
- [60] **S A Wolpert**, *Chern forms and the Riemann tensor for the moduli space of curves*, *Invent. Math.* 85 (1986) 119–145 MR Zbl

*Department of Mathematics, University of California Riverside
Riverside, CA, United States*

*Department of Mathematics, Yale University
New Haven, CT, United States*

*Department of Mathematics, Heriot-Watt University
Edinburgh, United Kingdom*

mdurham@ucr.edu, yair.minsky@yale.edu, a.sisto@hw.ac.uk

Proposed: Mladen Bestvina

Received: 3 June 2021

Seconded: David Fisher, Benson Farb

Revised: 11 November 2021

GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITOR

András I. Stipsicz Alfréd Rényi Institute of Mathematics
stipsicz@renyi.hu

BOARD OF EDITORS

Dan Abramovich	Brown University dan_abramovich@brown.edu	Mark Gross	University of Cambridge mgross@dpms.cam.ac.uk
Ian Agol	University of California, Berkeley ianagol@math.berkeley.edu	Rob Kirby	University of California, Berkeley kirby@math.berkeley.edu
Mark Behrens	Massachusetts Institute of Technology mbehrens@math.mit.edu	Frances Kirwan	University of Oxford frances.kirwan@balliol.oxford.ac.uk
Mladen Bestvina	Imperial College, London bestvina@math.utah.edu	Bruce Kleiner	NYU, Courant Institute bkleiner@cims.nyu.edu
Martin R. Bridson	Imperial College, London m.bridson@ic.ac.uk	Urs Lang	ETH Zürich urs.lang@math.ethz.ch
Jim Bryan	University of British Columbia jbryan@math.ubc.ca	Marc Levine	Universität Duisburg-Essen marc.levine@uni-due.de
Dmitri Burago	Pennsylvania State University burago@math.psu.edu	John Lott	University of California, Berkeley lott@math.berkeley.edu
Ralph Cohen	Stanford University ralph@math.stanford.edu	Ciprian Manolescu	University of California, Los Angeles cm@math.ucla.edu
Tobias H. Colding	Massachusetts Institute of Technology colding@math.mit.edu	Haynes Miller	Massachusetts Institute of Technology hrm@math.mit.edu
Simon Donaldson	Imperial College, London s.donaldson@ic.ac.uk	Tom Mrowka	Massachusetts Institute of Technology mrowka@math.mit.edu
Yasha Eliashberg	Stanford University eliash-gt@math.stanford.edu	Walter Neumann	Columbia University neumann@math.columbia.edu
Benson Farb	University of Chicago farb@math.uchicago.edu	Jean-Pierre Otal	Université d'Orleans jean-pierre.otal@univ-orleans.fr
Steve Ferry	Rutgers University sferry@math.rutgers.edu	Peter Ozsváth	Columbia University ozsvath@math.columbia.edu
Ron Fintushel	Michigan State University ronfint@math.msu.edu	Leonid Polterovich	Tel Aviv University polterov@post.tau.ac.il
David M. Fisher	Rice University davidfisher@rice.edu	Colin Rourke	University of Warwick gt@maths.warwick.ac.uk
Mike Freedman	Microsoft Research michaelf@microsoft.com	Stefan Schwede	Universität Bonn schwede@math.uni-bonn.de
David Gabai	Princeton University gabai@princeton.edu	Peter Teichner	University of California, Berkeley teichner@math.berkeley.edu
Stavros Garoufalidis	Southern U. of Sci. and Tech., China stavros@mpim-bonn.mpg.de	Richard P. Thomas	Imperial College, London richard.thomas@imperial.ac.uk
Cameron Gordon	University of Texas gordon@math.utexas.edu	Gang Tian	Massachusetts Institute of Technology tian@math.mit.edu
Lothar Götsche	Abdus Salam Int. Centre for Th. Physics gotsche@ictp.trieste.it	Ulrike Tillmann	Oxford University tillmann@maths.ox.ac.uk
Jesper Grodal	University of Copenhagen jg@math.ku.dk	Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Misha Gromov	IHÉS and NYU, Courant Institute gromov@ihes.fr	Anna Wienhard	Universität Heidelberg wienhard@mathi.uni-heidelberg.de

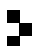
See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2023 is US \$740/year for the electronic version, and \$1030/year (+ \$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<http://msp.org/>

© 2023 Mathematical Sciences Publishers

GEOMETRY & TOPOLOGY

Volume 27 Issue 6 (pages 2049–2496) 2023

- Duality between Lagrangian and Legendrian invariants 2049
TOBIAS EKHOLM and YANKI LEKILI
- Filtering the Heegaard Floer contact invariant 2181
ÇAĞATAY KUTLUHAN, GORDANA MATIĆ, JEREMY
VAN HORN-MORRIS and ANDY WAND
- Large-scale geometry of big mapping class groups 2237
KATHRYN MANN and KASRA RAFI
- On dense totipotent free subgroups in full groups 2297
ALESSANDRO CARDERI, DAMIEN GABORIAU and
FRANÇOIS LE MAÎTRE
- The infimum of the dual volume of convex cocompact
hyperbolic 3-manifolds 2319
FILIPPO MAZZOLI
- Discrete subgroups of small critical exponent 2347
BEIBEI LIU and SHI WANG
- Stable cubulations, bicombings, and barycenters 2383
MATTHEW G DURHAM, YAIR N MINSKY and
ALESSANDRO SISTO
- Smallest noncyclic quotients of braid and mapping class groups 2479
SUDIPTA KOLAY