

A COMBINATORIAL TAKE ON HIERARCHICAL HYPERBOLICITY AND APPLICATIONS TO QUOTIENTS OF MAPPING CLASS GROUPS

JASON BEHRSTOCK, MARK HAGEN, ALEXANDRE MARTIN, AND ALESSANDRO SISTO

ABSTRACT. We give a simple combinatorial criterion, in terms of an action on a hyperbolic simplicial complex, for a group to be hierarchically hyperbolic. We apply this to show that quotients of mapping class groups by large powers of Dehn twists are hierarchically hyperbolic (and even relatively hyperbolic in the genus 2 case). In genus at least three, there are no known infinite hyperbolic quotients of mapping class groups. However, using the hierarchically hyperbolic quotients we construct, we show, under a residual finiteness assumption, that mapping class groups have many non-elementary hyperbolic quotients. Using these quotients, we relate questions of Reid and Bridson–Reid–Wilton about finite quotients of mapping class groups to residual finiteness of specific hyperbolic groups.

CONTENTS

Introduction	2
(Hierarchically) hyperbolic quotients of mapping class groups	4
Speculations	7
Largest hyperbolic quotients	8
Acknowledgements	9
1. Combinatorial HHSs	9
1.1. Basic definitions	9
1.2. On the various parts of the definition	12
1.3. Projections to links	13
1.4. Statement of Theorem 1.18	14
1.5. User’s guide and simple examples	15
1.6. Mapping class groups and blow-ups	17
2. Background on hierarchical hyperbolicity	18
2.1. Axioms	18
2.2. Useful facts	20
3. Hyperbolicity of Y_Δ	21
3.1. Preliminary lemma about hyperbolic graphs	22
3.2. Y_Δ is hyperbolic for $\mathcal{C}(\Delta)$ large	25
3.3. QI-embedding of $\mathcal{C}(\Sigma) \cap Y_\Delta$ in Y_Δ	28
4. Combinatorial HHS structures on links, and hieromorphisms	33
5. W is hierarchically hyperbolic	38
5.1. Strong bounded geodesic image	38
5.2. Proof of Theorem 1.18: verification of the HHS axioms	39
6. Hierarchical hyperbolicity from actions on hyperbolic complexes	47
7. Mapping class group quotients	52
8. Proof of Theorem 7.1	56
8.1. Outline	56
8.2. Coloring subsurfaces	57
8.3. Composite projection systems and composite rotating families	59

8.4. Setting up the induction	60
8.5. Inductive conditions	62
8.6. Verifying composite properties, and Greendlinger’s Lemma	67
8.7. Lifting	71
8.8. Supporting lemmas for Theorem 7.1	73
8.9. Checking hierarchical hyperbolicity	75
8.10. Preservation properties	79
References	81

INTRODUCTION

Hierarchically hyperbolic spaces/groups (HHS/G) were introduced in [BHS17b] as a common framework for studying the coarse geometry of mapping class groups and $\text{CAT}(0)$ cube complexes. Since then, this notion has found applications in a variety of flavors, including results of a coarse geometric [BHS17a, BHS, RST18] and of an algebraic nature [Hae16, DHS17, ABD17, AB18, ANS19], with many of these new even for well-studied examples such as mapping class groups. Also, the class of groups known to be hierarchically hyperbolic groups has expanded considerably beyond the motivating examples of mapping class groups, compact special groups, and hyperbolic groups, and there are many ways to produce new hierarchically hyperbolic groups from old: suitable graphs of groups [BHS19, BR18, RS20], graph products [BR20, BR18], and, somewhat in the spirit of the present paper, certain “small-cancellation” quotients, including, say, quotients of mapping class groups by normal subgroups generated by high powers of a pseudo-Anosov [BHS17a]. The machinery built in the present paper gives a method for producing even more new examples. Specifically, we show in Theorem 2, that quotients of mapping class groups by subgroups generated by suitable powers of (all) Dehn twists are hierarchically hyperbolic.

The main drawback of the theory has been that, despite the simplification in the definition provided by [BHS19], verifying that any particular space is an HHS requires a lot of work, and understanding of the HHS machinery. We remedy this by providing a combinatorial sufficient condition for a space/group to be hierarchically hyperbolic. This criterion is simpler to state than the definition of hierarchical hyperbolicity, and it is (in principle and in examples) easier to verify for a given space/group.

We first need a definition. Let a group H act on a (possibly disconnected) simplicial complex Y . We say that Y is a *hyperbolic H -space* if it becomes hyperbolic upon addition of finitely many H -orbits of edges (see Definition 6.2).

A simplified, but still powerful, version of the criterion is the following theorem which, informally, states that if G acts cocompactly with finite stabilizers of maximal simplices on a hyperbolic simplicial complex with hyperbolic links, then G is hierarchically hyperbolic under a geometric condition (quasi-isometric embedding in condition A) and a combinatorial condition on intersections of links (condition B).

Theorem 1. *Let the group G act cocompactly on the flag simplicial complex X , and suppose that maximal simplices have finite stabilizers. Suppose that:*

- (A) *For every simplex Δ of X , its link, $\text{Lk}(\Delta)$, is a hyperbolic $\text{Stab}_G(\text{Lk}(\Delta))$ -space quasi-isometrically embedded in $X - \bigcup_{\text{Lk}(\Sigma)=\text{Lk}(\Delta)} \Sigma$.*
- (B) *For all simplices Δ, Σ of X there exist simplices Π, Π' of $\text{Lk}(\Delta)$ such that $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \text{Lk}(\Delta \star \Pi) \star \Pi'$.*
- (C) *If the simplex Δ is not a co-dimension 1 face of a maximal simplex, then $\text{Lk}(\Delta)$ is connected.*

Then G is a hierarchically hyperbolic group.

In the statement, \star denotes the simplicial join, and we are allowing the empty simplex in the various conditions, so that in particular $X = \text{Lk}(\emptyset)$ is hyperbolic. For more explanation of the statement and comments on the various conditions, see Section 6.

A more general version is given in Theorem 1.18 below, and both Theorem 1.18 and Theorem 6.4 give further details about the hierarchically hyperbolic structure.

The setup of Theorem 1, and in particular the idea of considering a hyperbolic complex with hyperbolic links, is inspired by the curve complex $\mathcal{C}(S)$ of a surface S . Links of simplices in $\mathcal{C}(S)$ are related to curve complexes of subsurfaces of S by the following observations, which hold with some low-complexity exceptions. The link of a vertex is the curve graph of the complement of the corresponding curve. More generally, for each subsurface U of S , the curve complex $\mathcal{C}(U)$ arises as the link of a simplex in $\mathcal{C}(S)$. Conversely, the link of each simplex is either a nontrivial join (hence bounded), or is the curve graph of a subsurface.

However, curve graphs are the hyperbolic complexes that witness the hierarchical hyperbolicity of pants graphs, not mapping class groups. This is because annular curve graphs are not links of simplices in the curve graph; in fact, $\mathcal{C}(S)$ does not even contain the vertex set of $\mathcal{C}(U)$ when U is an annulus. For a discussion of the hyperbolic complex that witnesses hierarchical hyperbolicity of the mapping class group, see Section 1.6.

For the pants graph, checking the conditions of Theorem 1.18 (which are similar to those of Theorem 1) only uses hyperbolicity of curve graphs [MM99], the fact [MM00, Lemma 2.3] that subsurface projections are coarsely Lipschitz where defined (this is to check the quasi-isometric embedding condition), and arguments involving subsurfaces filled by multicurves to check the combinatorial conditions. We leave the details to the reader since pants graphs are already known to be hierarchically hyperbolic [BHS17b, Theorem G].

Some applications of Theorem 1 are discussed below, but we expect it (and the more general version, Theorem 1.18) to have many other applications. In fact, Theorem 1.18 is used in [DDLS20] to verify hierarchical hyperbolicity of surface extensions of naturally-occurring subgroups of mapping class groups, namely lattice Veech groups. Moreover, in [HMS22], Theorem 1.18 is used to show that Artin groups of extra-large type are hierarchically hyperbolic groups, and in [HRSS22], the combinatorial HHS viewpoint is used to prove that fundamental groups of graph manifolds are hierarchically hyperbolic *groups*, not merely hierarchically hyperbolic *spaces* (as was previously shown in [BHS19]).

We emphasize that the combinatorial HHS viewpoint encapsulated in Theorem 1.18 operates in tandem with the definition of hierarchical hyperbolicity from [BHS19], rather than replacing it. By this we mean the following:

- Theorem 1.18 is a good way to prove that a space/group is hierarchically hyperbolic, but many of the consequences of hierarchical hyperbolicity — for example, finite asymptotic dimension [BHS17a], control of quasiflats [BHS], “coarse rank-rigidity” and the omnibus subgroup theorem [DHS17, PS20], the Tits alternative [DHS17, DHS19], uniform exponential growth [ANS19], etc. rely on the coarse geometric machinery built on the original definition. The combinatorial viewpoint does not currently allow us to prove geometric results of this type directly. Even establishing much more basic geometric properties of a combinatorial HHS — as in Theorem 1 and its proof — takes considerable work.
- More fundamentally, the main tools of hierarchical hyperbolicity — notably, the distance formula — aren’t readily extractable from the combinatorial viewpoint on its own.
- The property of being a combinatorial HHS seems to be strictly stronger than the property of being an HHS. This can be illustrated using fairly artificial CAT(0) square complex examples to show that the natural candidate combinatorial HHS structure one

might try to build from an HHS structure does not work without additional hypotheses on the HHS. A partial converse of the form “hierarchically hyperbolic groups satisfying natural additional conditions are combinatorially hierarchically hyperbolic” is the subject of ongoing work, so we will not deal in detail with the above mentioned example here. Suffice it to say that it involves an infinite CAT(0) square complex with trivial automorphism group, which is a hierarchically hyperbolic space in view of [BHS17b] but for which the natural candidate combinatorial hierarchically hyperbolic structure involves an underlying simplicial complex in which the set of links of simplices, partially ordered by inclusion, contains arbitrarily long chains.

Nonetheless, someone wishing to establish a property of some group G that is already known for hierarchically hyperbolic groups can (with luck) now do so with no engagement whatsoever with the HHS machinery: they could instead try to build a combinatorial HHS on which G acts. A guide on how to use Theorem 1.18 for further application appears in Section 1.5 and Section 1.6.

(Hierarchically) hyperbolic quotients of mapping class groups. The main application of Theorem 1 presented in this paper is the study of certain quotients of mapping class groups. Below, we first discuss the natural quotients obtained by modding out high powers of Dehn twists. We then discuss a construction of non-elementary hyperbolic quotients, which works under the assumption that specific hyperbolic groups encountered in the construction are residually finite (and without this assumption in genus 2).

Quotients by powers of Dehn twists. We first apply Theorem 1 to quotients of mapping class groups by large powers of Dehn twists, further advancing the technology developed by Dahmani in [Dah18] to resolve Ivanov’s deep relations question [Iva06, Section 12], and the extensions of that technology from [DHS21].

We now very briefly survey the history of the study of quotients of mapping class groups by powers of Dehn twists. The study of these quotients dates to at least 1974 where it appears in Birman’s classic monograph [Bir74]. In that text, Birman notes that for the closed genus two surface the normal closure of the squares of Dehn twists is of index $6!$ in the mapping class group; she then asked whether the index is finite or infinite for arbitrary genus. Humphries later resolved this for the normal subgroups generated by squares or by cubes of Dehn twists (finite for the closed or once punctured surfaces of genus two or three; otherwise infinite), see [Hum92]. Humphries also showed that for the genus two surface with any number of punctures and any power of Dehn twists greater than 3, the corresponding quotient group was infinite. Eventually, Funar proved that as long as the genus is at least 3, then the quotient by the normal subgroups along powers of Dehn twists are infinite, as long as they are at least 13th powers [Fun99]. Funar was interested in these subgroups because of their connection with TQFT representations. See also [AF19].

Another, more recent, motivation for studying these quotients comes from the algebraic counterpart of Thurston’s Dehn filling theorem in the context of relatively hyperbolic groups [Osi07, GM08]. This algebraic version has numerous important applications, including a role in the proof of the Virtual Haken conjecture [Ago13] and in the solution of the isomorphism problem for certain relatively hyperbolic groups [DG18, DT19]. Mapping class groups are not non-trivially relatively hyperbolic except in very low complexity [AAS07, BDM09], but, in a number of ways, the subgroups generated by Dehn twists around curves in a pants decomposition play a role analogous to that of peripheral subgroups.

In [DHS21], it is proven that quotients of mapping class groups by large powers of Dehn twists are acylindrically hyperbolic, providing an analogue of the Dehn filling theorem. However, while acylindrical hyperbolicity has many interesting consequences, it only captures part of the geometry of mapping class groups. One would hope that performing Dehn fillings preserves

much more of the hierarchically hyperbolic structure of the mapping class group — indeed, in Theorem 2 below we will establish that the quotients are also hierarchically hyperbolic.

Given a surface S , we denote by DT_K the normal subgroup generated by all K -th powers of Dehn twists. Using Theorem 1, we prove:

Theorem 2. *Let S be a finite-type surface. Then there exists $K_0 > 0$ so that for any non-zero multiple K of K_0 , $MCG(S)/DT_K$ is an infinite hierarchically hyperbolic group.*

In fact, we provide an explicit HHS structure, which is described in Theorem 7.3. Corollaries of hierarchical hyperbolicity for these groups include: finiteness of the asymptotic dimension of $MCG(S)/DT_K$ [BHS17a], uniform exponential growth [ANS19], and (using the description of the HHS structure) that the maximal dimension of quasiflats is $\lfloor (3g - 3)/2 \rfloor$ and each quasi-flat of that dimension is a union of finitely many standard orthants [BHS17b, BHS]. The latter result might be useful to prove quasi-isometric rigidity, which we formulate as a question:

Question 3. *For K as in Theorem 2, is $MCG(S)/DT_K$ quasi-isometrically rigid?*

A strategy to give a positive answer to Question 3 could involve adapting arguments from [BHS, Section 5] and from [Bow15], and proving a combinatorial rigidity result whereby one shows that the automorphism group of a suitable simplicial complex coincides with the desired group.¹ Using the lifting techniques from [DHS21] that we develop further in Subsection 8.7, it might be possible to prove such a result by reducing it to combinatorial rigidity of a suitable complex on which $MCG(S)$ acts. We pointed to [Bow15] because we expect the structure of $MCG(S)/DT_K$ to more closely resemble that of the Weil-Petersson metric than that of the mapping class group itself, since, roughly speaking, we made the annular curve graphs bounded and hence irrelevant.

Hyperbolic quotients. The case of the genus 2 surface is notable in that quotients by powers of Dehn twists are not only hierarchically hyperbolic, but they are in fact relatively hyperbolic, as we will see by applying a result of Russell [Rus19] to the HHS structure on the quotients:

Theorem 4. *There exists $K_0 \geq 1$ so that for all non-zero multiples K of K_0 , the following holds. The quotient $MCG(\Sigma_2)/DT_K$ is hyperbolic relative to an infinite index subgroup commensurable to the product of two $C'(1/6)$ -groups.*

We note that the theorem can be used in conjunction with relatively hyperbolic Dehn filling [Osi07, GM08] to produce many hyperbolic quotients of $MCG(\Sigma_2)$, as stated in the following corollary, which we deduce from Theorem 4 in Remark 7.6. In said remark, we also point out a different construction relying on results of a very different nature and on a trick suggested to us by Francesco Fournier Facio.

Corollary 5. *$MCG(\Sigma_2)$ is fully residually non-elementary hyperbolic.*

Recall that a group is *fully residually \mathcal{P}* if for every finite collection of elements, there is a quotient satisfying \mathcal{P} into which the collection injects.

For higher genus, we cannot apply the relatively hyperbolic Dehn filling. However, we now outline how one could construct hyperbolic quotients using HHS machinery instead. After modding out powers of Dehn twists, we are left with a hierarchically hyperbolic group that has strictly lower complexity (in the HHS sense, the exact meaning of which is immaterial for this discussion). At the bottom of the hierarchy, we find a collection of hyperbolic groups. The idea to further reduce complexity is to repeat the previous procedure, namely modding out deep finite-index subgroups of these hyperbolic groups. We proceed inductively, reducing complexity,

¹After this paper was originally circulated, this approach has been successfully implemented for punctured spheres in [MS22].

until we are left with a hierarchically hyperbolic group of complexity 1. A general fact about hierarchically hyperbolic groups is that when the complexity is 1, the group is hyperbolic. We show that this construction works provided that all hyperbolic groups encountered at the various stages are residually finite. In Theorem 6 below we formulate this under the assumption that all hyperbolic groups are residually finite, which is much stronger than what we strictly need. We do not know whether there are enough residually finite hyperbolic groups to run our construction. Theorem 6, as well as the other theorems below, can be seen as a route towards either establishing the existence of non-residually finite hyperbolic groups or as an invitation to find a suitable class of hyperbolic groups that are residually finite. We discuss this further below.

Theorem 6. *Let S be a connected orientable surface of finite type of complexity at least 2. If all hyperbolic groups are residually finite, then $MCG(S)$ is fully residually non-elementary hyperbolic.*

As noted above in Corollary 5, our techniques prove that the mapping class group of the closed genus-2 surface is fully residually non-elementary hyperbolic without any residual finiteness assumption.

Remark 7. *In the above theorem, and the next two theorems, we condition the conclusion on residual finiteness of all hyperbolic groups. This is for simplicity in formulating the theorems. In reality, during the proofs of these theorems, one encounters specific hyperbolic groups, and it is only for these particular groups whose residual finiteness is necessary. Therefore, one can interpret these theorems as potential ways of proving the existence of a non-residually finite hyperbolic group (by exhibiting a mapping class group with no non-elementary hyperbolic quotients, say), but we prefer to see these theorems as an invitation to study the residual finiteness question for the specific hyperbolic groups encountered in the proofs.*

We use the flexibility of the construction described above to obtain the necessary multitude of hyperbolic quotients, as described in Theorem 7.1.

Remark 8. *The following was pointed out by Dawid Kielak. If $MCG(S)$ admits a non-elementary hyperbolic quotient, then $MCG(S)$ admits an affine isometric action on an L^p space with unbounded orbits [Nic13, Yu05] for some $1 < p < +\infty$, and in particular it does not have property F_{L^p} for said p . (In contrast, $SL_n(\mathbb{Z})$ does have F_{L^p} for all $1 < p < +\infty$ [BFG⁺07].) In view of Theorem 6, if some mapping class group has F_{L^p} , say for all $1 < p < +\infty$, then there is a non-residually finite hyperbolic group.*

The flexibility of our construction of quotients can also be exploited to prove further results, as we now discuss. Recall from [FM02] that $H \leq MCG(S)$ is *convex-cocompact* if some H -orbit in the Teichmüller space of S is quasiconvex. There are several characterizations of convex-cocompactness; see [KL08, Ham05, DT15]. One reason this notion is interesting is its connection with hyperbolicity of fundamental groups of surface bundles over surfaces [FM02, Ham05].

Reid posed the question of whether convex-cocompact subgroups of $MCG(S)$ are separable [Rei06, Question 3.5]. Recall that a subgroup $H < G$ is separable if for every $x \in G - H$ there exists a finite group F a surjective homomorphism $\phi : G \rightarrow F$ with $\phi(x) \notin \phi(H)$.

Note that in general $MCG(S)$ contains non-separable subgroups. In fact this is already the case for $MCG(S_{0,5})$ [LM07]. (Nonetheless, various geometrically natural subgroups, e.g. curve-stabilizers, are known to be separable in $MCG(S)$ [LM07].)

Our techniques reduce Reid's question to residual finiteness of certain hyperbolic groups, which we formulate as:

Theorem 9. *Let S be a connected orientable surface of finite type of complexity at least 2. If all hyperbolic groups are residually finite, then every convex-cocompact subgroup of $MCG(S)$ is separable.*

The proof of Theorem 9 relies on the hyperbolic quotients of $MCG(S)$ arising in the proof of Theorem 6. The assumption about residual finiteness is invoked again to apply a result of [AGM09] (namely, if all hyperbolic groups are residually finite then all hyperbolic groups are QCERF). Once again, we do not really need residual finiteness of all hyperbolic groups, just the ones encountered in our construction and in the (iterated Dehn filling) construction from [AGM09].

The next application relates to a question of Bridson–Reid–Wilton. In fact, we reduce [BRW17, Question 5.1] to the questions of residual finiteness of certain hyperbolic groups and of the congruence subgroup property for mapping class groups.

Theorem 10. *Let S be a connected orientable surface of finite type of complexity at least 2. If all hyperbolic groups are residually finite, then the following holds. Let $g, h \in MCG(S)$ be pseudo-Anosovs with no common proper power, and let $q \in \mathbb{Q}_{>0}$. Then there exists a finite group G and a homomorphism $\psi : MCG(S) \rightarrow G$ so that $\text{ord}(\psi(g))/\text{ord}(\psi(h)) = q$, where ord denotes the order.*

The property established in the theorem is called *omnipotence for pseudo-Anosovs*. In [BRW17], the authors study profinite rigidity of 3-manifold groups using the notion of $\pi_1(\Sigma)$ -congruence omnipotence, where there is the additional requirement that the finite quotients are congruence quotients. It is not known whether mapping class groups have the congruence subgroup property, see [EK95, Problem 2.10][Iva06, Conjecture on page 75], but if so, then the two notions of omnipotence for pseudo-Anosovs are equivalent.

In [BRW17] it is shown that a positive answer to their Question 5.1 implies the following. Let M be a closed hyperbolic 3-manifold with first Betti number 1. Let N be a compact 3-manifold so that $\pi_1 M$ and $\pi_1 N$ have isomorphic profinite completions. Then M and N have a common finite cyclic cover.

A heuristic discussion of the proof of Theorem 6 is given in Section 8.1. We note that the proof of Theorem 6 is an essentially self-contained part of this paper, only using the statement of Theorem 1.

Speculations. As mentioned above, Theorems 6, 9, and 10 hold provided that “sufficiently many” hyperbolic groups are residually finite, and therefore there are two natural research directions to explore. The first is to use those results to show that there exist hyperbolic groups that are not residually finite. Consider, for example, the following question:

Question 11. *Do all mapping class groups of closed oriented surface of genus at least 1 have an infinite hyperbolic quotient?*

In view of Theorem 6, a negative answer to said question shows that there exists some non-residually-finite hyperbolic group. Similarly, the same is true if either the question by Reid [Rei06, Question 3.5] or that by Bridson–Reid–Wilton [BRW17, Question 5.1] have a negative answer.

(A priori, having an infinite hyperbolic quotient is weaker than being fully residually non-elementary hyperbolic. However, in our context these properties are equivalent in view of the argument in Remark 7.6, since mapping class groups do not have two-ended quotients.)

In the other direction, one might try to find a suitable class of hyperbolic groups that can be shown to be residually finite, and are sufficient to prove the theorems above without additional assumptions. This would mirror the developments that led to the proof of the virtual Haken conjecture. Indeed, if all hyperbolic groups were residually finite, then the virtual Haken conjecture would follow from all hyperbolic groups being in fact QCERF [AGM09], the existence of quasiconvex surface subgroups [KM12], and the connection between separability and embedding into a finite cover [Sco78]. Also, although the proof does not follow exactly this template, the conjecture was eventually proven by showing that cubulated hyperbolic groups are virtually

special [Ago13], which in particular implies that they are (QCERF, whence) residually finite [HW08].

Further, provided that residual finiteness issues are resolved, it would be interesting to determine whether the hyperbolic quotients obtained via our construction can be $CAT(0)$. We believe that this is a natural question given that actions of mapping class groups on $CAT(0)$ spaces are very constrained [Bri10, KL96]. One can even ask:

Question 12. *Can an infinite hyperbolic quotient of the mapping class group of a closed oriented surface of sufficiently high genus be $CAT(0)$?*

Largest hyperbolic quotients. For the purposes of the following discussion, assume that all hyperbolic groups are residually finite. It is natural to wonder whether the hyperbolic groups we construct in the proof of Theorem 7.1 are the “largest possible”, meaning that any hyperbolic quotient of the mapping class group is a quotient of one of them.

We do not believe this to be true, because, in our quotients, too many stabilizers of subsurfaces have finite image. However, we formulate a conjecture related to this below.

Fix a closed surface S of genus at least 2, and let \mathcal{Y} be the collection of all (isotopy classes of) essential subsurfaces Y of S so that there exists a mapping class g with Y and gY disjoint and not isotopic. For example, an annulus around a non-separating curve is in \mathcal{Y} , while the annulus around the “middle curve” of the genus-2 closed surface is not.

Remark 13. *The set \mathcal{Y} coincides with the set of non- $MCG(S)$ -overlapping subsurfaces in the sense of [CMM20], as well as with the set of nondisplaceable subsurfaces in the sense of [HQR20, Theorem 1].*

For a subsurface $Y \subseteq S$, denote by $MCG(Y|S)$ the subgroup of $MCG(S)$ consisting of all mapping classes supported on Y . The conjecture is:

Conjecture 14. *Let S be a closed surface of genus at least 3. Then there are epimorphisms $\phi : MCG(S) \rightarrow G$, with G hyperbolic, such that the following holds for any essential subsurface $Y \subseteq S$: the group $\phi(MCG(Y|S))$ is finite if and only if $Y \in \mathcal{Y}$.*

The statement of the conjecture might require some adjustments. We believe the conjecture to be at least morally correct provided that no residual finiteness issues arise.

The conjecture is inspired by the observations below, which show that if $Y \in \mathcal{Y}$, then $MCG(Y|S)$ becomes virtually cyclic in every hyperbolic quotient of $MCG(S)$.

Since every infinite virtually cyclic group surjects onto $\mathbb{Z}/2\mathbb{Z}$, this implies that $MCG(Y|S)$ actually becomes finite at least if Y has genus at least 3, but we believe that with additional arguments this can also be shown in lower genus.

Let $Y \in \mathcal{Y}$, with corresponding mapping class g . Suppose that the epimorphism $\phi : MCG(S) \rightarrow G$, where G is hyperbolic, is so that $H = \phi(MCG(Y|S))$ is infinite (otherwise we are done). Then H contains an infinite order element, say h . Moreover, $\phi(g)H\phi(g)^{-1}$ is contained in the centralizer of h (notice that $gMCG(Y|S)g^{-1} = MCG(gY|S)$, and that $MCG(Y|S)$ commutes with $MCG(gY|S)$). The centralizer of h is virtually cyclic, and hence so is H .

Outline In Section 1, we introduce the notions needed to state Theorem 1.18, in particular the notion of a combinatorial HHS and state the theorem. We conclude the section with some remarks that might be of use to the reader wishing to apply the theorem, and we also include a simple example.

In Section 2 we recall the definition of HHS, and the (very few) results needed for this paper.

Sections 3–6 contain the proofs of Theorems 1.18 and 1, but only the statement of Theorem 1 is used in the subsequent sections.

In Section 3, we prove that the spaces used to define HHS projections are hyperbolic. This is crucial to prove that the candidate projections to the various hyperbolic spaces behave as expected. In Section 4, we study induced combinatorial HHS structures on links, which will enable inductive arguments. In Section 5, we complete the proof of Theorem 1.18 by checking the HHS axioms. At this stage, the hardest part of the proof will be the Uniqueness axiom.

In Section 6, we show that the conditions in Theorem 1 imply those in Theorem 1.18. At this point, we have proved Theorem 1, and we move on to the mapping class group applications.

Sections 7 and 8 focus on quotients of mapping class groups.

In Section 7, we state Theorem 7.1, about the hierarchically hyperbolic structures on our quotients of mapping class groups and deduce Theorem 2 (Theorem 7.3), Theorem 4 (Corollary 7.4), Theorem 6 (Corollary 7.7), Theorem 9 (Corollary 7.8), and Theorem 10 (Corollary 7.9).

In Section 8, we prove Theorem 7.1. Here, we improve the lifting technology of [DHS21] and combine it with Theorem 1. In fact, we expect many of the new technical lemmas in this section to be useful for other Dehn filling-type theorems, possibly even outside the hierarchically hyperbolic context. Section 8 begins with a fairly detailed heuristic outline of the proof.

Remark. This paper arose from two separate projects, which are naturally linked and we therefore merged. The results in Sections 3–6, in particular Theorems 1.18 and 1, are due to MH, AM, and AS. The remaining results are due to all four authors.

Acknowledgements. We are very grateful to François Dahmani for numerous essential discussions. A similar approach to the construction of further quotients of mapping class groups was discussed with him and was instrumental to this work, as have been his explanations of technical points in his paper [Dah18]. We thank Ian Agol, Matt Clay, Francesco Fournier Facio, Dawid Kielak, Dan Margalit, Michah Sageev, and Alex Wright for useful comments and suggestions. We also thank Carolyn Abbott, Daniel Berlyne, Thomas Ng, Alex Rasmussen, and Jacob Russell who together carefully read through an early version of the paper with a very sharp eye and provided us with extremely helpful feedback. We are also extremely grateful to the referees for a very large number of comments, and several corrections, that have enormously improved the paper.

The authors would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme *Non-positive curvature, group actions, and cohomology* where work on this paper was undertaken. This work was supported by EPSRC grant no EP/K032208/1. We thank the International Centre for Mathematical Sciences (ICMS) for their hospitality during a Research-in-Groups in which much of this work was done. Behrstock was supported by NSF grant DMS-1710890. Hagen was partially supported by EPSRC New Investigator Award EP/R042187/1. Martin was partially supported by EPSRC New Investigator Award EP/S010963/1. Sisto was partially supported by the Swiss National Science Foundation (grant #182186).

1. COMBINATORIAL HHSs

We first state Theorem 1.18 and supply tools for its proof. In Subsection 1.5, we illustrate how these tools work, via an explicit example (Section 1.5.1); it might be instructive for the reader to refer to that for motivation.

1.1. Basic definitions. Let X be a flag simplicial complex.

Definition 1.1 (Join, link, star). Given disjoint simplices Δ, Δ' of X , we let $\Delta \star \Delta'$ denote the simplex spanned by $\Delta^{(0)} \cup \Delta'^{(0)}$, if it exists. More generally, if K, L are disjoint induced subcomplexes of X such that every vertex of K is adjacent to every vertex of L , then $K \star L$ is the induced subcomplex with vertex set $K^{(0)} \cup L^{(0)}$. We refer to $K \star L$ as the *join* of K and L .

For each simplex Δ , the *link* $\text{Lk}(\Delta)$ is the union of all simplices Σ of X such that $\Sigma \cap \Delta = \emptyset$ and $\Sigma \star \Delta$ is a simplex of X . Observe that $\text{Lk}(\Delta) = \emptyset$ if Δ is a maximal simplex. Conversely, if $\text{Lk}(\Delta) = \emptyset$, then Δ is not properly contained in a simplex, i.e., Δ is maximal. More generally, if K is an induced subcomplex of X , then $\text{Lk}(K)$ is the union of all simplices Σ of X such that $\Sigma \cap K = \emptyset$ and $\Sigma \star K$ is a subcomplex of X .

The *star* of Δ is $\text{Star}(\Delta) = \text{Lk}(\Delta) \star \Delta$, i.e., the union of all simplices of X that contain Δ .

We often refer to 0–simplices as *vertices* and 1–simplices as *edges*, and make no distinction between 1–dimensional simplicial complexes and simplicial graphs. In a (not necessarily simplicial) graph or a simplicial complex Y , we use the term *open star* of a vertex v to refer to the union of $\{v\}$ with all open simplices (or open edges) of Y whose closures contain v . (By *open simplex* or *open edge*, we mean the image of the restriction to the interior of a cell of the appropriate characteristic map, which need not be open in X .) We sometimes refer to *removing the open star* of v , the result of which is an induced subcomplex of Y consisting of exactly those simplices (or edges) that do not contain v .

We emphasize that \emptyset is a simplex of X , whose link is all of X and whose star is all of X .

Definition 1.2 (X –graph, W –augmented complex). An X –graph is a graph W whose vertex set is the set of all maximal simplices of X .

For a flag complex X and an X –graph W , the W –augmented graph X^{+W} is the graph defined as follows:

- the 0–skeleton of X^{+W} is $X^{(0)}$;
- if $v, w \in X^{(0)}$ are adjacent in X , then they are adjacent in X^{+W} ;
- if two vertices in W are adjacent, then we consider σ, ρ , the associated maximal simplices of X , and in X^{+W} we connect each vertex of σ to each vertex of ρ .

We equip W with the usual path-metric, in which each edge has unit length, and do the same for X^{+W} .

We are aiming to construct a hierarchically hyperbolic structure (W, \mathfrak{S}) . The actually hierarchically hyperbolic space will be the graph W , equipped with the usual path-metric. The “curve graph” (i.e., the hyperbolic space associated to the unique \sqsubseteq –maximal element of \mathfrak{S}) will be X^{+W} .

Remark (Connectedness of W). We emphasize that, a priori, there is no assumption that W is connected. In practice, connectedness of W will be deduced using the other parts of the definition of a combinatorial HHS (Definition 1.8 below). Specifically, during the proof of Theorem 1.18, we verify that the links of X provide a hierarchically hyperbolic structure for W , and, during the part of that proof where the “uniqueness axiom” for hierarchically hyperbolic spaces (Definition 2.1.(9)) is checked, we verify that W is connected. This relies on the fact — coming from Definition 1.8 — that various auxiliary graphs related to W and X are hyperbolic when given the usual path-metric, and in particular connected. It also uses induction on the “complexity” n of X (see Definition 1.8.(1) and Definition 1.6), to say that various subgraphs W^Δ of W associated to links of simplices Δ in X are hierarchically hyperbolic, and in particular, connected.

Probably the strongest reason not to simply hypothesize connectedness of W is that, in support of the above induction, we will need to verify that for each nonempty non-maximal simplex Δ of X , the subgraph W^Δ of W spanned by maximal simplices of the form $\sigma \star \Delta$ is a $\text{Lk}(\Delta)$ –graph that combines with the complex $\text{Lk}(\Delta)$ to form a combinatorial HHS of strictly lower complexity. This is Proposition 4.9. By not including connectedness of W in the definition of an X –graph, we avoid having to verify connectedness of W^Δ when proving Proposition 4.9, and instead are able to assume it as an inductive hypothesis when proving connectedness of W later, in Theorem 1.18.

Definition 1.3 (Equivalent simplices, saturation). For Δ, Δ' simplices of X , we write $\Delta \sim \Delta'$ to mean $\text{Lk}(\Delta) = \text{Lk}(\Delta')$. We denote by $[\Delta]$ the \sim -equivalence class of Δ . Let $\text{Sat}(\Delta)$ denote the set of vertices $v \in X$ for which there exists a simplex Δ' of X such that $v \in \Delta'$ and $\Delta' \sim \Delta$, i.e.,

$$\text{Sat}(\Delta) = \left(\bigcup_{\Delta' \in [\Delta]} \Delta' \right)^{(0)}.$$

We refer to $\text{Sat}(\Delta)$ as the *saturation* of Δ . We denote by \mathfrak{S} the set of \sim -classes of **non-maximal** simplices in X .

Remark 1.4. Notice that $\text{Sat}(\Delta) \subseteq \text{Lk}(\text{Lk}(\Delta))^{(0)}$ (any vertex in $\text{Sat}(\Delta)$ is contained in a simplex all of whose vertices are connected to any vertex in $\text{Lk}(\Delta)$). Also, we have $\text{Lk}(\text{Sat}(\Delta)) = \text{Lk}(\Delta)$ (indeed, since $\Delta^{(0)} \subseteq \text{Sat}(\Delta)$, we have the inclusion “ \subseteq ”, while on the other hand any vertex connected to all vertices of Δ is connected to all elements of $\text{Sat}(\Delta)$, giving the other inclusion).

Definition 1.5 (Complement, link subgraph). Let W be an X -graph. For each simplex Δ of X , the *complement subgraph* Y_Δ is the subgraph of X^{+W} induced by the set $(X^{+W})^{(0)} - \text{Sat}(\Delta)$ of vertices.

The *augmented link* $\mathcal{C}(\Delta)$ of Δ is the induced subgraph of Y_Δ spanned by $\text{Lk}(\Delta)^{(0)}$. Note that $\mathcal{C}(\Delta) = \mathcal{C}(\Delta')$ whenever $\Delta \sim \Delta'$. We emphasize that we are taking links in X , not in X^{+W} , and then considering the subgraphs of Y_Δ induced by those links.

(The notation $\mathcal{C}(\Delta)$ is chosen since these spaces will be the underlying hyperbolic spaces in a hierarchically hyperbolic structure for W ; compare Definition 2.1. The use of “ \mathcal{C} ” was originally motivated by usage in concrete examples from [BHS17b]: curve graphs in the setting of mapping class groups, and contact graphs in the setting of $\text{CAT}(0)$ cube complexes.)

Definition 1.6. The simplicial complex X has *finite complexity* if there exists $n \in \mathbb{N}$ so that any chain $\text{Lk}(\Delta_1) \subsetneq \dots \subsetneq \text{Lk}(\Delta_i)$, where each Δ_j is a simplex of X , has length at most n ; the minimal such n is the *complexity* of X .

Remark 1.7 (Complexity versus dimension). Suppose that X has complexity n . Let Δ be a t -simplex of X . Then Δ contains a chain $\emptyset \subsetneq \Delta_0 \subsetneq \Delta_1 \subsetneq \dots \subsetneq \Delta_t$ with each Δ_i an i -simplex. Observe that $\text{Lk}(\Delta_{i+1}) \subsetneq \text{Lk}(\Delta_i)$ for all i , so $t + 2 \leq n$, i.e., $\dim X \leq n - 2$.

However, for a general simplicial complex X , the complexity cannot be bounded in terms of the dimension. Indeed, let X be the following 1-dimensional simplicial complex. Let $X^{(0)} = \{v_n\}_{n \geq 0} \sqcup \{h_n\}_{n \geq 0}$, and, for each $i \geq 0$, join v_i by an edge to h_0, \dots, h_i . Then $\text{Lk}(v_i) \subsetneq \text{Lk}(v_{i+1})$ for all i , i.e., the complexity is infinite.

Our basic object is a *combinatorial hierarchically hyperbolic space*:

Definition 1.8 (Combinatorial HHS). A *combinatorial HHS* (X, W) consists of a flag simplicial complex X and an X -graph W satisfying all of the following conditions.

- (1) X has complexity $n < +\infty$.
- (2) There is a constant δ so that for each non-maximal simplex Δ , the subgraph $\mathcal{C}(\Delta)$ is δ -hyperbolic and (δ, δ) -quasi-isometrically embedded in the complement subgraph Y_Δ , which was defined in Definition 1.5.
- (3) Let Δ and Σ be non-maximal simplices such that there exists a non-maximal simplex Γ with the following properties:
 - $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Delta)$,
 - $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Sigma)$, and
 - $\text{diam}(\mathcal{C}(\Gamma)) \geq \delta$.

Then there exists a simplex Π in the link of Σ such that $\text{Lk}(\Sigma \star \Pi) \subseteq \text{Lk}(\Delta)$ and all non-maximal simplices Γ satisfying the above three itemized properties also satisfy $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Sigma \star \Pi)$.

- (4) If v, w are distinct non-adjacent vertices of $\text{Lk}(\Delta)$, for some simplex Δ of X , contained in W -adjacent maximal simplices, then they are contained in W -adjacent simplices of the form $\Delta \star \Sigma$.

Sometimes we use the notation (X, W, δ, n) when we have to keep track of the constants.

Remark 1.9. The simplex $\Sigma \star \Pi$ in Definition 1.8.(3) is necessarily non-maximal. Indeed, its link is nonempty since it contains $\text{Lk}(\Gamma)$ for some non-maximal Γ .

1.2. On the various parts of the definition. We regard the first two conditions of Definition 1.8 as the most important ones, and the ones with solid theoretical reasons to be there.

As a side note, the quasi-isometric embedding part of Condition 2 can be viewed as an analogue of Bowditch’s fineness condition for relative hyperbolicity [Bow12, Proposition 2.1.(F5)]. We discuss this in the context of examples in Section 1.5.

While the first two conditions do not seem to be sufficient to yield an HHS, we expect that the last two conditions can be replaced with “better” ones. We do not have compelling reasons for those properties to be required; our best explanations are as follows. Our heuristic justification for Condition 3 is that it seems to be what is needed to perform arguments that in the context of the curve graph would require the use of tight geodesics, see in particular the proof of the “Uniqueness” axiom (i.e., the verification that Definition 2.1.(9) is satisfied). The heuristic justification for Condition 4 is that, without it, it might be possible to “move” between places in the link of some simplex without this being doable within the link itself.

In the interest of the reader who might need alternative conditions, or who might be interested in finding the “right” ones, we list where the last two conditions get used:

Remark 1.10 (Potentially replacing conditions (3) and (4)). Here are the only places where Condition (3) is used:

- the proof of Proposition 3.3, via Lemma 3.7,
- the proof of Lemma 4.8.2,
- the proof of Proposition 4.9 (only to check the analogous condition for a link),
- the proof of Theorem 1.18, in Section 5, in the “Consistency for nesting”, “Large links”, and “Uniqueness” (Case 1) parts.

The only uses of Condition (4) are in Section 3, in the proofs of Lemma 3.9 and Lemma 3.14, and in Section 4, via Lemma 4.1.

Here are some further comments on Definition 1.8.(3)] To understand the point of Definition 1.8.(3), we consider a more intuitive, and strictly stronger, version of the condition: one could insist that intersections of links are always links, i.e., if $\text{Lk}(\Delta) \cap \text{Lk}(\Pi)$ is nonempty, then there exists a (necessarily non-maximal) simplex Π such that $\text{Lk}(\Sigma) \cap \text{Lk}(\Delta) = \text{Lk}(\Pi)$.

This captures the right intuition. Indeed, our goal is to show that the set of equivalence classes of non-maximal simplices will give a hierarchically hyperbolic structure, where nesting is containment of links. For each Δ , the associated hyperbolic space will be $\mathcal{C}(\Delta)$, and if $\text{Lk}(\Pi) \subsetneq \text{Lk}(\Delta)$ for some non-maximal simplex Π , we need Π to correspond to a bounded subset of $\mathcal{C}(\Delta)$, because this is required by the definition of a hierarchically hyperbolic space (Definition 2.1.(2)). This much always works: $\text{Sat}(\Pi)$ has to intersect Y_Δ , so $\text{Lk}(\Pi)$ is bounded in Y_Δ and hence in $\mathcal{C}(\Delta)$ in view of Definition 1.8.(2).

Now, if Σ and Δ have intersecting links, but there is no containment between their links, then they should be transverse elements in the HHS structure (see Definition 2.1.(4)). This again requires that Σ correspond to a bounded set in $\mathcal{C}(\Delta)$. The naive hypothesis $\text{Lk}(\Sigma) \cap \text{Lk}(\Delta) = \text{Lk}(\Pi)$ would achieve this, since Π would be nested in Δ .

This hypothesis is too strong to accommodate desirable examples, including the case where X is the curve graph of a surface (so that simplices are multicurves). See Remark 6.6.

So, we define our bounded sets in $\mathcal{C}(\Delta)$ as follows. When Σ is nested in, or transverse to, Δ , we observe as above that $\text{Sat}(\Sigma) \cap Y_\Delta$ is always nonempty and bounded in Y_Δ , i.e., $\mathcal{C}(\Sigma) \cap Y_\Delta$ is “coned off” in Y_Δ .

To define the required bounded subset of $\mathcal{C}(\Delta)$ associated to Σ , which is done in Definition 1.16, we apply the coarse closest-point projection $Y_\Delta \rightarrow \mathcal{C}(\Delta)$ to $\text{Sat}(\Sigma) \cap Y_\Delta$. This projection obviously gives a uniformly bounded set if $\mathcal{C}(\Delta)$ has uniformly bounded diameter. If not, we verify in Section 3 that Y_Δ is hyperbolic, from which Definition 6.1.(2) implies that the coarse projection has the necessary properties.

The need for a condition like Definition 1.8.(3) then makes itself felt in proving hyperbolicity of Y_Δ . Recall that we only need to do this when $\mathcal{C}(\Delta)$ has diameter at least the threshold δ . The point in the proof of hyperbolicity (Lemma 3.9) where this is needed is: via Lemma 1.14, and its consequence, Lemma 3.7, if Δ is nested in some Π , then $Y_\Pi \subset Y_\Delta$. The hypothesis we have chosen seems to be a suitable way to arrange this while being weak enough to cover natural examples. As noted above, we also use it in similar ways in a few other places.

1.3. Projections to links. In this section we relate combinatorial objects to HHS objects, the connection being justified by Theorem 1.18. The reader who is not interested in the details of the hierarchical structure obtained but only wishes to use it as a simple criterion to prove the hierarchical hyperbolicity of a space/group can skip this section and go directly to Section 1.4.

Definition 1.11 (Nesting, orthogonality, transversality, complexity). Let X be a simplicial complex. Let Δ, Σ be non-maximal simplices of X . Then:

- $[\Delta] \sqsubseteq [\Sigma]$ if $\text{Lk}(\Delta) \subseteq \text{Lk}(\Sigma)$, and we say $[\Delta]$ is *nested* in $[\Sigma]$;
- $[\Delta] \perp [\Sigma]$ if $\text{Lk}(\Sigma) \subseteq \text{Lk}(\text{Lk}(\Delta))$, and we say $[\Delta]$ and $[\Sigma]$ are *orthogonal*.

If $[\Delta]$ and $[\Sigma]$ are neither \perp -related nor \sqsubseteq -related, we write $[\Delta] \pitchfork [\Sigma]$, and say $[\Delta]$ and $[\Sigma]$ are *transverse*.

Note that $[\emptyset]$ is the unique \sqsubseteq -maximal \sim -class of simplices in X and that \sqsubseteq is a partial ordering on the set of \sim -classes of simplices in X . Notice that the simplicial complex X has finite complexity if there exists $n \in \mathbb{N}$ so that any \sqsubseteq -chain has length at most n ; the minimal such n is the complexity of X .

Remark 1.12. The definition of $\Sigma \perp \Delta$ is equivalent to saying that any vertex in the link of Σ is joined by an edge to any vertex in the link of Δ .

One might be tempted to think of nesting as being equivalent to inclusion of simplices, but this only works in one direction, namely:

Remark 1.13. Let Δ, Δ' be simplices of X . If $\Delta \subseteq \Delta'$, then $[\Delta'] \sqsubseteq [\Delta]$.

Similarly, if $\text{Sat}(\Delta) \subset \text{Sat}(\Delta')$, then any vertex in $\text{Lk}(\Delta')$ is adjacent to every vertex in $\text{Sat}(\Delta)$, so $\text{Lk}(\Delta') \subset \text{Lk}(\Delta)$, i.e., $[\Delta'] \sqsubseteq [\Delta]$. Again, the converse does not hold, although Lemma 3.7 gives a partial converse.

Notice that Definition 1.8.(3) can be rephrased as follows:

- Whenever Δ and Σ are non-maximal simplices for which there exists a non-maximal simplex Γ such that $[\Gamma] \sqsubseteq [\Delta]$, $[\Gamma] \sqsubseteq [\Sigma]$, and $\text{diam}(\mathcal{C}(\Gamma)) \geq \delta$, then there exists a simplex Π in the link of Σ such that $[\Sigma \star \Pi] \sqsubseteq [\Delta]$ and all $[\Gamma]$ as above satisfy $[\Gamma] \sqsubseteq [\Sigma \star \Pi]$.

Also, note that if Π is the simplex associated to Σ and Δ , provided by Definition 1.8.(3), then since $\Pi \subset \text{Lk}(\Sigma)$, the simplex $\Pi \star \Sigma$ exists automatically and, moreover, $[\Pi \star \Sigma] \sqsubseteq [\Sigma]$.

We note the following special case of Condition (3) for later use:

Lemma 1.14. *Suppose that $[\Sigma] \sqsubseteq [\Delta]$ and $\text{diam}(\mathcal{C}(\Sigma)) \geq \delta$. Then $[\Sigma] = [\Delta \star \Pi]$ for some simplex Π of $\text{Lk}(\Delta)$.*

Proof. Definition 1.8.(3) provides a simplex Π of $\text{Lk}(\Delta)$ such that $[\Delta \star \Pi] \sqsubseteq [\Sigma]$. Since $\text{diam}(\mathcal{C}(\Sigma)) \geq \delta$, we also have (setting $\Gamma = \Sigma$) $[\Sigma] \sqsubseteq [\Delta \star \Pi]$, so $[\Sigma] = [\Delta \star \Pi]$. \square

Our next goal is to define projections from W to $\mathcal{C}([\Delta])$ for $[\Delta] \in \mathfrak{S}$. This will use the following lemma:

Lemma 1.15. *Let X be a flag simplicial complex, let Δ be a non-maximal simplex, and let Σ be a maximal simplex. Then $\Sigma \cap Y_\Delta$ is nonempty and has diameter at most 1.*

Proof. Let Z be the subcomplex of X spanned by $\text{Sat}(\Delta)$. Then for each maximal simplex Π of Z , we have $\text{Lk}(\Pi) \supseteq \text{Lk}(\Delta)$. If $\Sigma \cap Y_\Delta = \emptyset$, then $\Sigma^{(0)} \subseteq \text{Sat}(\Delta)$, so $\Sigma \subset Z$. Moreover, by maximality, we have $\text{Lk}(\Sigma) \supseteq \text{Lk}(\Delta)$. But $\text{Lk}(\Sigma) = \emptyset$, by maximality of Σ , while $\text{Lk}(\Delta) \neq \emptyset$, by non-maximality of Δ . Hence $\Sigma \cap Y_\Delta \neq \emptyset$.

Since the vertices of Σ are pairwise-adjacent in X , they are pairwise-adjacent in X^{+W} , so since Y_Δ is an induced subgraph, the vertices of $\Sigma \cap Y_\Delta$ are pairwise-adjacent in Y_Δ , as required. \square

Definition 1.16 (Projections). Let (X, W, δ, n) be a combinatorial HHS.

Fix $[\Delta] \in \mathfrak{S}$ and define a map $\pi_{[\Delta]} : W \rightarrow 2^{\mathcal{C}([\Delta])}$ as follows. Let $p : Y_\Delta \rightarrow 2^{\mathcal{C}([\Delta])}$ be the coarse closest point projection, i.e.,

$$p(x) = \{y \in \mathcal{C}([\Delta]) : d_{Y_\Delta}(x, y) \leq d_{Y_\Delta}(x, \mathcal{C}([\Delta])) + 1\}.$$

Suppose that w is a vertex of W , so w corresponds to a unique simplex Σ_w of X . Since Σ_w is maximal (by Definition 1.2), and Δ is non-maximal (by the definition of \mathfrak{S}), the graph $\Sigma_w \cap Y_\Delta$ is nonempty and has diameter at most 1, by Lemma 1.15. Define

$$\pi_{[\Delta]}(w) = p(\Sigma_w \cap Y_\Delta).$$

We have thus defined $\pi_{[\Delta]} : W^{(0)} \rightarrow 2^{\mathcal{C}([\Delta])}$. If $v, w \in W$ are joined by an edge e of W , then Σ_v, Σ_w are joined by edges in X^{+W} , and we let $\pi_{[\Delta]}(e) = \pi_{[\Delta]}(v) \cup \pi_{[\Delta]}(w)$.

Now let $[\Delta], [\Delta'] \in \mathfrak{S}$ satisfy $[\Delta] \triangleleft [\Delta']$ or $[\Delta'] \sqsubset [\Delta]$. Let

$$\rho_{[\Delta]}^{[\Delta']} = p(\text{Sat}(\Delta') \cap Y_\Delta),$$

where $p : Y_\Delta \rightarrow 2^{\mathcal{C}([\Delta])}$ is coarse closest-point projection.

Let $[\Delta] \sqsubset [\Delta']$. Let $\rho_{[\Delta]}^{[\Delta']} : \mathcal{C}([\Delta']) \rightarrow 2^{\mathcal{C}([\Delta])}$ be defined as follows. On $\mathcal{C}([\Delta']) \cap Y_\Delta$, it is the restriction of p to $\mathcal{C}([\Delta']) \cap Y_\Delta$. Otherwise, it takes the value \emptyset .

Remark 1.17 (See the future). In Lemma 3.9, we will show that Y_Δ is δ_0 -hyperbolic, for some uniform δ_0 , provided $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$. Since $\mathcal{C}([\Delta])$ is (δ, δ) -quasi-isometrically embedded, we will then have that $\text{diam}(p(\Sigma_w \cap Y_\Delta))$ is bounded in terms of δ, δ_0 , i.e., $\pi_{[\Delta]}(w)$ is a nonempty, uniformly bounded set. When $\text{diam}(\mathcal{C}(\Delta)) \leq \delta$, then the same conclusion is immediate, with no need for hyperbolicity of Y_Δ . Once we have established that either Y_Δ is hyperbolic or $\mathcal{C}(\Delta)$ is uniformly bounded, then the coarse closest-point projection will send points to uniformly bounded sets, and we will, when convenient and when we are only concerned about distances up to uniformly bounded error, think of p as a map.

1.4. Statement of Theorem 1.18. Our main theorem about combinatorial HHS is Theorem 1.18. See Section 2 for the definition of a hierarchically hyperbolic space (HHS) and a hierarchically hyperbolic group (HHG).

Theorem 1.18 (HHS structures from X -graphs). *Let (X, W) be combinatorial HHS.*

Let \mathfrak{S} be as in Definition 1.3, define nesting and orthogonality relations on \mathfrak{S} as in Definition 1.11, let the associated hyperbolic spaces be as in Definition 1.8, and define projections as in Definition 1.16.

Then (W, \mathfrak{S}) is a hierarchically hyperbolic space, and the HHS constants only depend on δ, n as in Definition 1.8.

Moreover, suppose that G is a group acting cocompactly on X . Suppose that the G -action on the set of maximal simplices of X extends to an action on W which is metrically proper and cobounded. Then (G, \mathfrak{S}) is a hierarchically hyperbolic group.

Remark 1.19. In the “moreover” part of the statement, one actually only needs something weaker than cocompactness of the G -action on X . Specifically, the exact property we need is that G acts on X with finitely many orbits of subcomplexes of the form $\text{Lk}(\Delta)$, where Δ is a non-maximal simplex of X .

As in Definition 6.1 below, we say that the group G acts on the combinatorial HHS (X, W) if G acts by simplicial automorphisms on X and the action on $W^{(0)}$ induced by the action on X extends to an action on the whole graph W (i.e., it preserves W -adjacency). In this language, we can rephrase the latter part of Theorem 1.18:

Corollary 1.20. *Let G act on the combinatorial HHS (X, W) . Suppose that the action of G on X is cocompact and the action on W is proper and cocompact. Then G is a hierarchically hyperbolic group.*

Remark 1.21. Notice that, under the assumptions of Corollary 1.20, we have that the action of G on X is acylindrical in view of [BHS17b, Theorem K].

We fix the notation of Theorem 1.18 from now on. The proof of Theorem 1.18 is in Section 5.2, after some necessary preparation.

1.5. User’s guide and simple examples. We make some remarks that could be useful for the reader interested in applying Theorem 1.18 to establish hierarchical hyperbolicity in their example of interest. First, Theorem 6.4 provides a simpler set of conditions that do not involve the X -graph W , and the reader is advised to first check whether that theorem applies in their situation. Typical obstructions to using the simplified version arise from bounded links, which are treated more flexibly in Theorem 1.18.

If not, it has to be noted that, in situations where there is a natural X to consider, the X might actually have to be changed within its quasi-isometry class to satisfy the fine geometry constraints. For example, in the amalgamated free product example just below, we see that the natural candidate, the Bass-Serre tree, may not work as our X , and we need to “blow up” the vertices of the tree to stars before proceeding.

One strategy is to build the correct “model” bottom-up, meaning starting from the hyperbolic complexes for the expected sub-HHSs (e.g., in a tree of HHSs, one might want to suitably combine the hyperbolic complexes for the various vertex spaces, see below).

It might also be useful to note that taking direct products at the level of complexes corresponds to taking joins at the level of links. Also, relative hyperbolicity corresponds to disjoint unions of cones over the hyperbolic complexes for the peripherals, as we discuss a bit more at the end of this subsection.

1.5.1. Amalgamated free product example. We now give an example of a combinatorial HHS. Let us consider an amalgamated product $G = A *_C B$ of hyperbolic groups over a common quasiconvex almost-malnormal subgroup C , so that G is hyperbolic, and hence hierarchically hyperbolic, by the Bestvina–Feighn combination theorem [BF92].

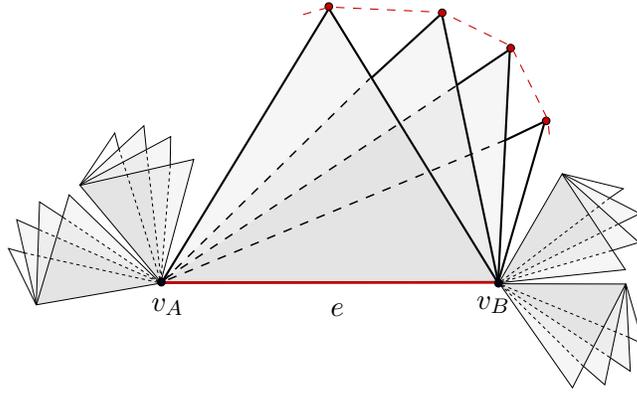


Figure 1. A portion of the complex X . The vertices in red correspond to elements of C , and all have the same link (red edge e in the picture). The link of that edge in X is a discrete set in bijection with C . However, due to the choice of W , the augmented link $\mathcal{C}(e)$ is the Cayley graph of C with respect to S_C (red dotted lines in the picture).

We will define the simplicial complex X , which will be quasi-isometric to the Bass-Serre tree for G . Notice that the Bass-Serre tree itself is not the right complex to consider since there is no link “encoding C ”, meaning that C does not act on any link in such a way that, say, C has unbounded orbits if it is infinite, as one would expect from the right complex in view of the distance formula.

Let us now construct X , as follows. The vertex set of X is $G \sqcup \{v_{gA} : gA \in G/A\} \sqcup \{v_{gB} : gB \in G/B\}$ (where G/H denotes the set of left cosets of H in G). Edges of X correspond to either containment of an element of G in a coset of A or B , or to pairs of cosets of A and B intersecting non-trivially; edges of the latter type correspond to edges of the Bass-Serre tree.

Finally, we let X be the flag complex with the 1-skeleton we just described. Roughly speaking, X is the Bass-Serre tree corresponding to the splitting, but where every edge is now contained in triangles indexed by C (see Figure 1). Notice that maximal simplices have vertex set of the form g, v_{gA}, v_{gB} , and are in bijection with G . The link of $g \in G$ is a single edge, while the vertex set of the link of v_A is in bijection with $A \sqcup A/C$.

Examples of saturations are that the saturation of $g \in G$ is gC , while the saturation of the edge with endpoints $1, A$ is $C \cup \{v_A\}$.

We now define W as any Cayley graph of G corresponding to a generating set $S_A \cup S_B \cup S_C$ with S_H a generating set of H for $H \in \{A, B, C\}$, and $S_A \cap C = S_B \cap C = \emptyset$. Then it can be checked that X^{+W} is quasi-isometric to X and to the Bass-Serre tree of G , that $\mathcal{C}(v_A)$ is the Cayley graph of A with respect to S_A with the cosets of C coned-off (and similarly for B), and that $\mathcal{C}(e)$, where e is the edge with endpoints v_A, v_B , is the Cayley graph of C with respect to S_C . The link of the edge e' joining, say, $1 \in G$ to v_A is the single vertex v_B .

1.5.2. Relative hyperbolicity. We now say a few words on relative hyperbolicity, and the analogy between the notion of a combinatorial HHS and Bowditch’s *fine graphs* [Bow12, Proposition 2.1].

First, consider infinite hyperbolic groups A, B and let $G = A \star B$. Let X be the Bass-Serre tree. The vertices are of the form v_{gA} or v_{gB} , for $g \in G$, i.e., they are indexed by left cosets of A and B . Links of vertices are discrete, and any two intersect in at most one vertex, so there is no proper containment of links beyond the fact that all links are contained in the link of \emptyset .

There is a natural G -equivariant bijection $G \rightarrow W^{(0)}$, where $W^{(0)}$ is the set of edges (maximal simplices) of X : each $g \in G$ appears in precisely one coset gA and one coset gB , and the intersection of these cosets corresponds to an edge of X . Fixing a finite generating set for G

consisting of the disjoint union of generating sets of A and B , we join $x, y \in W^{(0)}$ (edges of X) by an edge of W if and only if the corresponding elements of G are adjacent in the Cayley graph of G . So, W is a copy of the Cayley graph of G . The graph $X+W$ is naturally quasi-isometric to the coned-off Cayley graph of G , with cones over each gA, gB , so it is hyperbolic. The link of each vertex v is, upon addition of W -edges, a copy of a Cayley graph of A or B , and hence hyperbolic. One verifies that this augmented link is quasi-isometrically embedded in $X+W - \{v\}$ by constructing a Lipschitz retract from $X+W - \{v\}$ to the link (see, for instance, [HRSS22] for a more general version of such a construction in the context of Bass-Serre trees).

This is a simple example of a relatively hyperbolic group — thinking of A and B as the peripheral subgroups — where an associated hyperbolic fine graph, namely X , also functions as the underlying complex of a CHHS structure. As we saw above, obtaining CHHS structures from splittings normally requires “blowing up” the Bass-Serre tree; that it doesn’t in this example is an artifact of the example being a free product.

However, coned-off Cayley graphs for relatively hyperbolic groups do relate closely to combinatorial HHSes. Indeed, let G be hyperbolic relative to a subgroup P . Fix a Cayley graph K_0 of G associated to a finite generating set containing generators for P , and let K be obtained from K_0 by adding a vertex v_{gP} , joined by an edge to each element of gP , for each left coset gP . So, K is the standard hyperbolic fine graph witnessing relative hyperbolicity.

On the other hand, suppose that (X_P, W_P) is a combinatorial HHG structure for P , i.e., a combinatorial HHS, with a simplicial P -action on X_P such that, for simplicity, W_P with the induced P -action is equivariantly isomorphic to the Cayley graph of P with the finite generating set mentioned above. For each gP , let X'_{gP} be obtained from a copy X_{gP} of X_P by adding a vertex v_{gP} , joining it to each vertex of X_{gP} , and taking the induced flag complex. We let W_{gP} be a copy of W_P ; note that the vertices of W_{gP} correspond naturally to the elements of gP . Let X be the disjoint union of the X'_{gP} . Let W be formed from the disjoint union of the W_{gP} as follows: first, note that the vertex set of W is naturally in bijection with G , since the cosets gP partition G . Join vertices of W by an edge if the corresponding elements of G are adjacent in our Cayley graph. The pair (X, W) is a combinatorial HHS with G acting geometrically on W . For example, $X+W$ is quasi-isometric to K , so it is hyperbolic. This is a sketch of a proof that a group hyperbolic relative to combinatorially hierarchically hyperbolic groups is again combinatorially hierarchically hyperbolic, and justifies our earlier claim of an analogy with fine hyperbolic graphs.

1.6. Mapping class groups and blow-ups. This subsection contains further discussion on applications of Theorem 1.18 and is not used elsewhere in the paper.

Although $MCG(S)$ is known to be a hierarchically hyperbolic group — the index set \mathfrak{S} consists of isotopy classes of (possibly disconnected) subsurfaces, and the associated hyperbolic spaces are curve graphs [BHS19, Section 11] — one cannot apply Theorem 1.18 to the curve graph \mathcal{CS} in order to realize $MCG(S)$ as a combinatorial HHS. The reason is that annular curve graphs — projections to which need to appear one way or another in any HHS structure — do not arise as links of simplices in the curve complex of S . When we take $X = \mathcal{CS}$ and W to be the pants graph of S , then applying Theorem 1.18 yields the HHS structure on the pants graph (i.e., on Teichmüller space with the Weil-Petersson metric) as a combinatorial HHS. Note that this standard HHS structure on the Weil-Petersson metric, described in [BHS17b, Theorem G] or [Vok22, Theorem 1.1, Example 2.3], has index set consisting of the *non-annular* subsurfaces, which do correspond to links in \mathcal{CS} .

However, one can most likely modify \mathcal{CS} by something like the following “blow-up” construction to get a combinatorial HHS structure on the marking graph (which is quasi-isometric to the mapping class group). For each vertex γ of \mathcal{CS} , let $B(\gamma)$ be the graph obtained from the vertex set of the annular curve graph $\mathcal{C}\gamma$ by adding a vertex γ joined by an edge to each vertex

of $\mathcal{C}\gamma$. Let X be the flag complex on the graph obtained from $\bigsqcup_{\gamma \in \mathcal{C}S(0)} B(\gamma)$ by adding an edge joining every vertex of $B(\gamma)$ to every vertex of $B(\alpha)$ whenever γ, α are $\mathcal{C}S$ -adjacent. (So, in particular, X is quasi-isometric to $\mathcal{C}S$: just collapse each $B(\gamma)$ to the vertex γ and each subgraph $B(\gamma) \star B(\alpha)$ to the edge $[\alpha, \gamma]$.) The idea would then be to associate maximal simplices in X to markings, take W to be a suitably chosen version of the marking graph, and verify that (X, W) is a combinatorial HHS.

It's instructive to see how annular curve graphs arise as links in this setup. Let γ be a curve, and let $\gamma_1, \dots, \gamma_n$ be a maximal collection of disjoint curves with $\gamma_1 = \gamma$. For $i \geq 2$, choose a point $x_i \in \mathcal{C}\gamma_i$ (regarded as a subgraph of X). Let Σ be the simplex with vertex set $\gamma_1, \dots, \gamma_n, x_2, \dots, x_n$. Then $\text{Lk}_X(\Sigma)$ is exactly $\mathcal{C}\gamma$.

Also, let Σ be a simplex of the form $[\gamma_1, x_1], \dots, [\gamma_n, x_n]$, where each $[\gamma_i, x_i]$ joins a curve γ_i to $x_i \in \mathcal{C}\gamma_i$, and the γ_i are all disjoint. Then $\text{Lk}_X(\Sigma)$ is the union of the subgraphs $B(\alpha)$, as α varies over all the curves in the complement of $\gamma_1, \dots, \gamma_n$. In particular, if Y is a subsurface, then taking $\gamma_1, \dots, \gamma_n$ to consist of the boundary curves of Y that are essential in S , together with pants decompositions of the components of $S - Y$, we find that $\text{Lk}_X(\Sigma)$ corresponds to $\mathcal{C}Y$. So, except for some bounded links, the links in X correspond to the curve graphs of the various subsurfaces in the usual HHS structure on the marking graph, described in [BHS19, Section 11].

We expect that there is a wide range of contexts in which similar blow-up constructions can be used to construct combinatorial HHS structures. For example, we expect that for (many) right-angled Artin groups, one can “blow up” the Kim-Koberda *extension graph* [KK13, KK14] to obtain a combinatorial HHS structure, and whence an alternate proof of hierarchical hyperbolicity for these groups via Theorem 1.18.

In [HMS22], hierarchical hyperbolicity of extra-large-type Artin groups is established by constructing an appropriate version of the extension graph, blowing it up, and verifying that this blown-up graph gives the X such that (X, W) is a combinatorial HHS, where W is an appropriately-chosen Cayley graph of the Artin group.

In fact, we expect that such a blow-up construction can be used to prove a partial converse to Theorem 1.18 under some extra conditions on the hierarchically hyperbolic space.

2. BACKGROUND ON HIERARCHICAL HYPERBOLICITY

2.1. Axioms. We recall from [BHS19] the definition of a hierarchically hyperbolic space.

Definition 2.1 (Hierarchically hyperbolic space). The q -quasigeodesic space $(\mathcal{X}, d_{\mathcal{X}})$ is a *hierarchically hyperbolic space* if there exists $\delta \geq 0$, an index set \mathfrak{S} , and a set $\{\mathcal{C}U : U \in \mathfrak{S}\}$ of δ -hyperbolic spaces $(\mathcal{C}U, d_U)$, such that the following conditions are satisfied:

- (1) **(Projections.)** There is a set $\{\pi_U : \mathcal{X} \rightarrow 2^{\mathcal{C}U} \mid U \in \mathfrak{S}\}$ of *projections* sending points in \mathcal{X} to sets of diameter bounded by some $\xi \geq 0$ in the various $\mathcal{C}U \in \mathfrak{S}$. Moreover, there exists K so that for all $U \in \mathfrak{S}$, the coarse map π_U is (K, K) -coarsely Lipschitz and $\pi_U(\mathcal{X})$ is K -quasiconvex in $\mathcal{C}U$.
- (2) **(Nesting.)** \mathfrak{S} is equipped with a partial order \sqsubseteq , and either $\mathfrak{S} = \emptyset$ or \mathfrak{S} contains a unique \sqsubseteq -maximal element; when $V \sqsubseteq U$, we say V is *nested* in U . (We emphasize that $U \sqsubseteq U$ for all $U \in \mathfrak{S}$.) For each $U \in \mathfrak{S}$, we denote by \mathfrak{S}_U the set of $V \in \mathfrak{S}$ such that $V \sqsubseteq U$. Moreover, for all $V, U \in \mathfrak{S}$ with $V \sqsubseteq U$ there is a specified subset $\rho_V^U \subset \mathcal{C}U$ with $\text{diam}_{\mathcal{C}U}(\rho_V^U) \leq \xi$. There is also a *projection* $\rho_V^U : \mathcal{C}U \rightarrow 2^{\mathcal{C}V}$. (The similarity in notation is justified by viewing ρ_V^U as a coarsely constant map $\mathcal{C}U \rightarrow 2^{\mathcal{C}V}$.)
- (3) **(Orthogonality.)** \mathfrak{S} has a symmetric and anti-reflexive relation called *orthogonality*: we write $V \perp U$ when V, U are orthogonal. Also, whenever $V \sqsubseteq U$ and $U \perp A$, we require that $V \perp A$. 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 $B \in \mathfrak{S}_T - \{T\}$, so that whenever $V \perp U$ and $V \sqsubseteq T$, we have $V \sqsubseteq B$. Finally, if $V \perp U$, then V, U are not \sqsubseteq -comparable.

- (4) **(Transversality and consistency.)** If $V, U \in \mathfrak{S}$ are not orthogonal and neither is nested in the other, then we say V, U are *transverse*, denoted $V \pitchfork U$. There exists $\kappa_0 \geq 0$ such that if $V \pitchfork U$, then there are sets $\rho_U^V \subseteq \mathcal{C}W$ and $\rho_V^U \subseteq \mathcal{C}V$ each of diameter at most ξ and satisfying:

$$\min \{d_U(\pi_U(x), \rho_U^V), d_V(\pi_V(x), \rho_V^U)\} \leq \kappa_0$$

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

For $V, U \in \mathfrak{S}$ satisfying $V \sqsubseteq U$ and for all $x \in \mathcal{X}$, we have:

$$\min \{d_U(\pi_U(x), \rho_U^V), \text{diam}_{\mathcal{C}V}(\pi_V(x) \cup \rho_V^U(\pi_U(x)))\} \leq \kappa_0.$$

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

Finally, if $U \sqsubseteq V$, then $d_U(\rho_T^U, \rho_T^V) \leq \kappa_0$ whenever $T \in \mathfrak{S}$ satisfies either $V \sqsubset T$ or $V \pitchfork T$ and $T \not\sqsubseteq U$.

- (5) **(Finite complexity.)** There exists $n \geq 0$, the *complexity* of \mathcal{X} (with respect to \mathfrak{S}), so that any set of pairwise- \sqsubseteq -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 $U \in \mathfrak{S}$ and let $x, x' \in \mathcal{X}$. Let $N = \lambda d_U(\pi_U(x), \pi_U(x')) + \lambda$. Then there exists $\{T_i\}_{i=1, \dots, [N]} \subseteq \mathfrak{S}_U - \{U\}$ such that for all $T \in \mathfrak{S}_U - \{U\}$, either $T \in \mathfrak{S}_{T_i}$ for some i , or $d_T(\pi_T(x), \pi_T(x')) < E$. Also, $d_U(\pi_U(x), \rho_U^{T_i}) \leq N$ for each i .
- (7) **(Bounded geodesic image.)** There exists $E > 0$ such that for all $U \in \mathfrak{S}$, all $V \in \mathfrak{S}_U - \{U\}$, and all geodesics γ of $\mathcal{C}U$, either $\text{diam}_{\mathcal{C}V}(\rho_V^U(\gamma)) \leq E$ or $\gamma \cap \mathcal{N}_E(\rho_V^U) \neq \emptyset$.
- (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 $p_j \in \pi_{V_j}(\mathcal{X}) \subseteq \mathcal{C}V_j$. Then there exists $x \in \mathcal{X}$ so that:
- $d_{V_j}(\pi_{V_j}(x), p_j) \leq \alpha$ for all j ,
 - for each j and each $V \in \mathfrak{S}$ with $V_j \sqsubseteq V$, we have $d_V(\pi_V(x), \rho_V^{V_j}) \leq \alpha$, and
 - for each j and each $V \in \mathfrak{S}$ with $V_j \pitchfork V$, we have $d_V(\pi_V(x), \rho_V^{V_j}) \leq \alpha$.
- (9) **(Uniqueness.)** For each $\kappa \geq 0$, there exists $\theta_u = \theta_u(\kappa)$ such that if $x, y \in \mathcal{X}$ and $d_{\mathcal{X}}(x, y) \geq \theta_u$, then there exists $V \in \mathfrak{S}$ such that $d_V(\pi_V(x), \pi_V(y)) \geq \kappa$.

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} . Observe that \mathcal{X} is hierarchically hyperbolic with respect to $\mathfrak{S} = \emptyset$, i.e., hierarchically hyperbolic of complexity 0, if and only if \mathcal{X} is bounded. Similarly, \mathcal{X} is hierarchically hyperbolic of complexity 1 with respect to $\mathfrak{S} = \{\mathcal{X}\}$, if and only if \mathcal{X} is hyperbolic.

Remark 2.2. Jacob Russell has pointed out that the “ $d_W(\pi_W(x), \rho_W^{T_i}) \leq N$ ” requirement in Definition 2.1.(6) follows from the consistency and bounded geodesic image axioms, and is therefore redundant; see [Rus19, Remark 2.10].

Notation 2.3. Where it will not cause confusion, given $U \in \mathfrak{S}$, we will often suppress the projection map π_U when writing distances in $\mathcal{C}U$, i.e., given $x, y \in \mathcal{X}$ and $p \in \mathcal{C}U$ we write $d_U(x, y)$ for $d_U(\pi_U(x), \pi_U(y))$ and $d_U(x, p)$ for $d_U(\pi_U(x), p)$. Note that when we measure distance between a pair of sets (typically both of bounded diameter) we are taking the minimum distance between the two sets. Given $A \subset \mathcal{X}$ and $U \in \mathfrak{S}$ we let $\pi_U(A)$ denote $\cup_{a \in A} \pi_U(a)$.

Definition 2.4 (Hierarchically hyperbolic group). The group G is a *hierarchically hyperbolic group* (HHG) if there exists a hierarchically hyperbolic space (Z, \mathfrak{S}) such that the following hold:

- G acts metrically properly and coboundedly by isometries on the quasigeodesic space Z .
- G acts on \mathfrak{S} with finitely many orbits, and the G action preserves the relations $\sqsubseteq, \perp, \triangleleft$.
- For all $U \in \mathfrak{S}$ and $g, h \in G$, there is an isometry $g : \mathcal{C}U \rightarrow \mathcal{C}gU$ such that the isometry $(gh) : \mathcal{C}U \rightarrow \mathcal{C}ghU$ is the composition of the isometries $g : \mathcal{C}hU \rightarrow \mathcal{C}ghU$ and $h : \mathcal{C}U \rightarrow \mathcal{C}hU$.
- For all $U \in \mathfrak{S}, g \in G, z \in Z$, we have $\pi_{gU}(gz) = g(\pi_U(z))$.
- For all $U, V \in \mathfrak{S}$ such that $U \triangleleft V$ or $U \sqsubset V$, and all $g \in G$, we have $\rho_{gV}^U = g(\rho_V^U)$.

From the first bullet and Milnor-Schwarz, G is finitely generated and, when G is equipped with any word-metric, composing the projections π_U with any orbit map $G \rightarrow Z$ shows that we can take $Z = G$ in the above definition. In particular, (G, \mathfrak{S}) is an HHS. When we wish to emphasize the particular HHS structure on which G is acting, we say that (G, \mathfrak{S}) is an HHG.

Remark 2.5. In the definition, we have asked that G act *metrically properly and coboundedly*, rather than properly and cocompactly, since it is sometimes convenient to check that G is an HHG by constructing an action on an HHS (Z, \mathfrak{S}) where Z is not proper.

2.2. Useful facts. We now recall results from [BHS19] that will be useful later on. To avoid some technicalities, we will assume that, given an HHS $(\mathcal{X}, \mathfrak{S})$, the maps $\pi_U, U \in \mathfrak{S}$ are uniformly coarsely surjective, which can always be arranged (see [BHS19, Remark 1.3]).

Definition 2.6 (Consistent tuple). 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 $V, W \in \mathfrak{S}$, we have

$$\min\{d_V(b_V, \rho_V^W), d_W(b_W, \rho_W^V)\} \leq \kappa$$

whenever $V \triangleleft W$ and

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

whenever $V \sqsubset W$.

The following is [BHS19, Theorem 3.1]:

Theorem 2.7 (Realization). *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then for each $\kappa \geq 1$, there exists $\theta = \theta(\kappa)$ so that, for any κ -consistent tuple $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{\mathcal{C}U}$, there exists $x \in \mathcal{X}$ such that $d_V(x, b_V) \leq \theta$ for all $V \in \mathfrak{S}$.*

Observe that uniqueness (Definition (9)) implies that the *realization point* x for \vec{b} provided by Theorem 2.7 is coarsely unique.

The following is [BHS19, Theorem 4.5]:

Theorem 2.8 (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, K so that for all $x, y \in \mathcal{X}$,*

$$d(x, y) \asymp_{K, C} \sum_{U \in \mathfrak{S}} \{\{d_U(x, y)\}\}_s.$$

(The notation $\{\{A\}\}_B$ denotes the quantity which is A if $A \geq B$ and 0 otherwise. The notation $A \asymp_{\lambda, \lambda} B$ means $A \leq \lambda B + \lambda$ and $B \leq \lambda A + \lambda$.)

We will use the following variation, which is well-known to experts:

Theorem 2.9 (Distance Formula+ ϵ). *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. For every $\lambda \geq 1$, there exist $T \geq 2\lambda$ and $\kappa \geq 1$, depending only on the HHS constants and λ , with the following property. Let $x, y \in \mathcal{X}$, and consider for every $Y \in \mathfrak{S}$ some h_Y with $h_Y \asymp_{\lambda, \lambda} d_Y(x, y)$. Then*

$$d_X(x, y) \asymp_{\kappa, \kappa} \sum_{Y \in \mathfrak{S}} \{\{h_Y\}\}_T.$$

Proof. The proof follows from manipulating the thresholds of the usual distance formula.

Let $DF_T(x, y) = \sum_{Y \in \mathfrak{S}} \{\{d_Y(x, y)\}\}_T$, and let $H_T(x, y)$ be the sum in the statement of the theorem we are proving. Then, for $T \geq 2\lambda$, we claim that we have:

$$\frac{1}{2\lambda} DF_{\lambda(\lambda+T)}(x, y) \leq H_T(x, y) \leq 2\lambda DF_{T/\lambda-1}(x, y).$$

These inequalities, combined with the distance formula (Theorem 2.8), prove the required statement.

Let us show the second inequality. If $h_Y \geq T$ (so that h_Y contributes to $H_T(x, y)$), then we have $d_Y(x, y) \geq h_Y/\lambda - 1 \geq h_Y/(2\lambda)$ since $T \geq 2\lambda$.

Hence:

$$H_T(x, y) = \sum_{Y: h_Y \geq T} h_Y \leq 2\lambda \sum_{Y: h_Y \geq T} d_Y(x, y) \leq 2\lambda DF_{T/\lambda-1}(x, y),$$

as required.

Let us show the first inequality, with the same method. If $d_Y(x, y) \geq \lambda(\lambda + T)$ then $h_Y \geq d_Y(x, y)/\lambda - 1 \geq T$, and also $h_Y \geq d_Y(x, y)/(2\lambda)$, since T is sufficiently large. Hence:

$$DF_{\lambda(\lambda+T)}(x, y) = \sum_{Y: d_Y(x, y) \geq \lambda(\lambda+T)} d_Y(x, y) \leq 2\lambda \sum_{Y: d_Y(x, y) \geq \lambda(\lambda+T)} h_Y \leq 2\lambda H_T(x, y),$$

as required. □

The following will be used for one of the corollaries on quotients of mapping class groups in Section 8, and in the proof of Theorem 7.1.(III). This result appears as the (3) implies (2) implication of [ABD17, Theorem B].

Lemma 2.10 ([ABD17]). *Let (G, \mathfrak{S}) be an HHG, and let $Q < G$ be such that orbit maps to $\mathcal{C}S$ are quasi-isometric embeddings, where S is the \sqsubseteq -maximal element of \mathfrak{S} . Then there exists κ such that for all $Y \in \mathfrak{S} - \{S\}$ we have $d_Y(g, h) \leq \kappa$ for all $g, h \in Q$.*

Proof. A detailed proof appears in [ABD17], here is a sketch.

In view of the bounded geodesic image axiom, if $d_Y(g, h)$ is large, then ρ_S^Y lies uniformly close to a geodesic from $\pi_S(g)$ to $\pi_S(h)$. Again in view of the bounded geodesic image axiom, along such a geodesic the projection to $\mathcal{C}(Y)$ is coarsely constant far from ρ_S^Y , and similarly for a neighborhood of the geodesic. Moreover, such a geodesic stays close to the quasiconvex subset $\pi_S(Q)$ of $\mathcal{C}(S)$, and in particular we see that Q contains elements g', h' so that:

- $\pi_S(g'), \pi_S(h')$ lie within uniformly bounded distance in $\mathcal{C}(S)$,
- $d_Y(g, h)$ differs from $d_Y(g', h')$ by a uniformly bounded amount.

In view of the first item there are finitely many possible pairs (g', h') up to the Q -action. In particular, there is a bound on $d_Y(g', h')$, whence a bound on $d_Y(g, h)$ as required. □

3. HYPERBOLICITY OF Y_Δ

Let (X, W, δ, n) be a combinatorial HHS. The goal of this section is to show that Y_Δ is uniformly hyperbolic whenever Δ is a non-maximal simplex of X for which $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$. We need this for several reasons in the proof of Theorem 1.18. The most fundamental reason is: to construct an HHS structure on W , we will need the projections $\pi_\Delta : W \rightarrow \mathcal{C}(\Delta)$ from Definition 1.16 to be well-defined, coarsely Lipschitz coarse maps. This follows from the fact that $\mathcal{C}(\Delta) \rightarrow Y_\Delta$ is a uniform quasi-isometric embedding, once we show that Y_Δ is hyperbolic.

We say that a constant K is *uniform* if K depends only on δ and n .

3.1. Preliminary lemma about hyperbolic graphs. Given a graph Z , we say that a set \mathcal{V} of vertices of Z is *discrete* if no two elements of \mathcal{V} are joined by an edge of Z .

The following lemma will be used inductively to prove hyperbolicity of Y_Δ . The “moreover” part will not be used to prove hyperbolicity of Y_Δ , but rather an additional statement, namely Lemma 3.10, which will be required in Section 4.

There are various ways to prove Lemma 3.1; we chose the way that serves as a warm-up for the proof of Theorem 1.18.

Lemma 3.1. *For every δ there exists δ' with the following property. Let Z be a δ -hyperbolic graph and let \mathcal{V} be a discrete collection of vertices. For each $v \in \mathcal{V}$, let Z_v be the induced subgraph of Z with vertex set $Z^{(0)} - \{v\}$.*

Suppose that the following hold for all $v \in \mathcal{V}$:

- $\text{Lk}(v)$ is δ -hyperbolic;
- $\text{Lk}(v)$ is (δ, δ) -quasi-isometrically embedded in Z_v .

Then the induced subgraph $Z_\mathcal{V}$ of Z with vertex set $Z^{(0)} - \mathcal{V}$ is δ' -hyperbolic.

Moreover, for each λ , there exists $\mu = \mu(\lambda, \delta)$ such that the following holds.

Let $Q \subseteq Z$ be a (λ, λ) -quasi-isometrically embedded induced subgraph of Z such that Q contains the star of v whenever $v \in \mathcal{V}$ belongs to Q . Then the induced subgraph $Q_\mathcal{V}$ of Z with vertex set $Q^{(0)} - \mathcal{V} \cap Q^{(0)}$ is (μ, μ) -quasi-isometrically embedded in $Z_\mathcal{V}$.

Proof of Lemma 3.1. We will equip $Z_\mathcal{V}$ with a hierarchically hyperbolic space structure $(Z_\mathcal{V}, \mathfrak{F})$ in which \mathfrak{F} has empty orthogonality relation; from results in [BHS] we will then get that $Z_\mathcal{V}$ is δ' -hyperbolic for some δ' depending on the HHS constants, which we shall see in this case depend only on δ .

Let $\mathfrak{F} = \{S\} \cup \mathcal{V}$. The associated hyperbolic space \mathcal{CS} is Z , and for each $v \in \mathcal{V}$, the associated space $\mathcal{C}v$ is $\text{Lk}(v)$ (which is δ -hyperbolic by hypothesis). We declare $v \sqsubset S$ for all $v \in \mathcal{V}$; there is no other nontrivial nesting relation, and the orthogonality relation is empty. Therefore, two elements of \mathfrak{F} are transverse if and only if they are distinct elements of \mathcal{V} . The relation \sqsubset is easily seen to be a partial order with a unique maximal element and bounded chains (the complexity is 2).

Define the projection $\pi_S : Z_\mathcal{V} \rightarrow \mathcal{CS} = Z$ to be the inclusion. The projection $\pi_v : Z_\mathcal{V} \rightarrow \mathcal{C}v = \text{Lk}(v)$ is defined as follows: given $x \in Z_\mathcal{V}$, consider all geodesics α in Z from x to v , and let x' be the entry point of α in $\text{Lk}(v)$. Then $\pi_v(x)$ is the set of all such x' . By Claim 3.2 below, $\pi_v(x)$ has diameter bounded in terms of δ only.

(Here, we have used discreteness of \mathcal{V} to ensure that x' exists. Indeed, discreteness ensures that $\text{Lk}(v) \subset Z_v$, so that any geodesic in Z from x to v must pass through $\text{Lk}(v)$.)

Claim 3.2. *There exists $K = K(\delta)$ so that the following holds. Let $x, y \in Z_\mathcal{V}$ and let α, β be geodesics in Z starting, respectively, at x, y and intersecting $\text{Lk}(v)$ only at their other respective endpoints x', y' . If $d_{\text{Lk}(v)}(x', y') \geq K$, then any geodesic γ in Z from x to y contains v .*

Proof of Claim 3.2. Let γ be a geodesic in Z , joining x to y , and not containing v . Then each vertex of γ belongs to Z_v , i.e., γ is a path in Z_v . Moreover, since γ is a geodesic of Z , it is again a geodesic of Z_v . Fix a geodesic q of Z from x' to y' ; since $x', y' \in \text{Lk}(v)$, we have $|q| \leq 2$.

Consider the quadrilateral in Z with sides α, β, γ, q . By δ -hyperbolicity of Z , this quadrilateral is 2δ -thin. Let α' be the geodesic in Z obtained by traveling backward along α , starting from x' , for distance $\min\{10\delta + 1, |\alpha|\}$, and define β' analogously, using y' and β . By construction, α', β' are paths in Z_v , since α, β cannot contain v . The endpoints a, b of α', β' are 2δ -close to points $a', b' \in \gamma$ with $d_Z(a', b') \leq 24\delta + 4$. If $|\alpha| \leq 10\delta + 1$ then $a = \alpha'$ and if $|\beta| \leq 10\delta + 1$, we have $b = \beta'$. Otherwise, a, a' are both at least $8\delta + 1$ -far from v , so in any case a geodesic $[a, a']$ in Z must lie in Z_v .

Thus, using $\alpha', \beta', [a, a'], [b, b']$ and the part of γ between a' and b' , we construct a path of length at most $100(\delta + 1)$ that joins x' to y' and lies in Z_v .

Hence, $d_{Z_v}(x', y') \leq 100(\delta + 1)$. Since $\text{Lk}(v)$ is (δ, δ) -quasi-isometrically embedded in Z_v , we obtain $d_{\text{Lk}(v)}(x', y') \leq \delta(100(\delta + 1) + \delta)$, as required. \blacksquare

Moreover, the coarse map π_v is (K, K) -coarsely Lipschitz, by Claim 3.2. Indeed, if $x, y \in Z_{\mathcal{V}}$ are adjacent vertices, then the claim implies that their images lie at distance at most K in $\text{Lk}(v)$.

Let $\rho_S^v \subset Z$ be $\{v\}$. Define $\rho_v^S : Z \rightarrow \text{Lk}(v)$ by letting $\rho_v^S(x) = \pi_v(x)$ if $x \in Z_v$ and defining $\rho_v^S(w)$ arbitrarily when w is in the open star of v . The consistency inequality for nesting holds by definition.

For $v, w \in \mathcal{V}$ distinct vertices, let $\rho_w^v = \pi_w(v)$. Let us verify the consistency inequality for transversality. Consider distinct $v, w \in \mathcal{V}$ and $x \in Z_{\mathcal{V}}$. Suppose that $d_{\text{Lk}(v)}(\rho_v^w, \pi_v(x)) \geq K$, for K as in Claim 3.2. Then any geodesic γ from x to w passes through v . In particular, the set of entrance points in $\text{Lk}(w)$ of geodesics from x to w is contained in the set of entrance points in $\text{Lk}(w)$ for geodesics from v to w . This implies $\pi_w(x) \subseteq \pi_w(v)$, and hence the two coarsely coincide.

Large links holds again by Claim 3.2, which implies that given $x, y \in Z_{\mathcal{V}}$, the set of all $v \in \mathcal{V}$ on whose links x, y have large projections are vertices of a geodesic in X from x to y .

It remains to verify bounded geodesic image, partial realisation, and uniqueness.

Bounded geodesic image: Let γ be a geodesic in Z . If γ is disjoint from $\text{Lk}(v)$, then $\text{diam}(\rho_v^S(\gamma)) \leq K$, by Claim 3.2. This verifies bounded geodesic image.

Partial realisation: Let $p \in Z = \mathcal{CS}$ be the coordinate to realise. If $p \in Z_{\mathcal{V}}$, let $x = p$. If p lies in the open star of v , let $x \in \text{Lk}(v)$ be chosen arbitrarily. Then $d_S(\pi_S(x), p) \leq 1$, and no element of \mathfrak{F} is transverse to, or contains, S . If $p \in \mathcal{C}v = \text{Lk}(v)$ is the coordinate to realise, let $x = p$. Then $d_v(\pi_v(x), p) = 0$ and $d_S(\pi_S(x), \rho_S^v) = 1$. Also, $d_w(\rho_w^v, x)$ is bounded since π_v is coarsely Lipschitz. This verifies partial realisation. (Here, again, we used discreteness to ensure that $\text{Lk}(v) \subset Z_{\mathcal{V}}$.)

Uniqueness: Let $\kappa \geq 0$ and fix $x, y \in Z_{\mathcal{V}}$. Suppose that $d_Z(x, y) \leq \kappa$, and that

$$d_{\text{Lk}(v)}(\pi_v(x), \pi_v(y)) \leq \kappa$$

for all $v \in \mathcal{V}$. We need to bound $d_{Z_{\mathcal{V}}}(x, y)$ in terms of κ and δ .

Let η be a geodesic of $Z_{\mathcal{V}}$ from x to y . If η is a geodesic of Z , then $d_{Z_{\mathcal{V}}}(x, y) = d_Z(x, y) \leq \kappa$, and we are done. So, assume that each geodesic γ in Z from x to y passes through some $v \in \mathcal{V}$, and fix such a γ .

Let $v_1, \dots, v_n \in \mathcal{V}$ be the vertices in \mathcal{V} contained in γ . Let x_i, y_i be the entry and exit points of γ in $\text{Lk}(v_i)$, for $1 \leq i \leq n$. By discreteness, all of the x_i, y_i lie in $Z_{\mathcal{V}}$, and are in particular all distinct from all v_j .

So, we can write

$$\gamma = \gamma_0 p_1 \gamma_1 p_2 \cdots \gamma_{n-1} p_n \gamma_n,$$

where each γ_n is a (possibly trivial) geodesic in $Z_{\mathcal{V}}$ and each p_i is the path of length 2 from x_i to y_i passing through v_i . Hence

$$d_{Z_{\mathcal{V}}}(x, y) \leq \sum_{i=1}^n d_{\text{Lk}(v_i)}(x_i, y_i) + \sum_{i=1}^n |\gamma_i|.$$

Note that each term of the left sum is at most $d_{\text{Lk}(v_i)}(\pi_{v_i}(x), \pi_{v_i}(y)) + 2K$, where $K = K(\delta)$ bounds the diameter of π_{v_i} -projections of points. So, by our assumption that all projections of x, y are κ -close, we have

$$d_{Z_{\mathcal{V}}}(x, y) \leq n\kappa + n(2K - 2) + d_Z(x, y) \leq (n + 1)\kappa + n(2K - 2),$$

since the Z -geodesic γ is obtained by concatenating the γ_i along with n paths of length 2. Finally, each v_i contributes at least 1 to the length of γ , so $n \leq \kappa$. We conclude that $d_{Z_{\mathcal{V}}}(x, y) \leq \kappa^2 + \kappa(2K - 1)$. Since this depends only on κ and δ , uniqueness is verified.

Hence $(Z_{\mathcal{V}}, \mathfrak{F})$ is an HHS with empty orthogonality relation. The constants implicit in the definition of an HHS depend only on δ . Then the proof of [BHS, Corollary 2.15] (relying on Theorem 2.1 of [BHS]) implies that $Z_{\mathcal{V}}$ is a *coarse median space* of rank 1, where the constants/function in the definition of a coarse median space (see [Bow13]) depend only on the HHS constants and hence only depend on δ . Hence [Bow13, Theorem 2.1] implies that $Z_{\mathcal{V}}$ is δ' -hyperbolic, where δ' depends only on the coarse median constants and hence only on δ .

The “moreover” clause: We now prove the “moreover clause” essentially by proving that $Q_{\mathcal{V}}$ is an HHS whose structure is compatible with that of $Z_{\mathcal{V}}$. Let Q be an induced subgraph of Z as in the statement. Fix $v \in \mathcal{V}$. Let Q_v be the induced subgraph of Q spanned by $Q^{(0)} - \{v\}$. Define $q : Q_v \rightarrow 2^{\text{Lk}(v)}$ as follows. Let $x \in Q_v$ be a vertex. For each geodesic γ in Q from x to v , let x_{γ} be the entrance point of γ in $\text{Lk}(v)$. Let $q(x)$ be the set of points x_{γ} , as γ varies over the Q -geodesics from x to v .

For notational purposes, set $z = \pi_S : Z_{\mathcal{V}} \rightarrow 2^{\text{Lk}(v)}$ as above. Namely, given $x \in Z_{\mathcal{V}}$, consider all Z -geodesics α from x to v , and for each α , let x_{α} be the entrance point of α in $\text{Lk}(v)$. Let $z(x)$ be the union of the x_{α} . By Claim 3.2, $z(x)$ has diameter bounded in terms of δ , i.e., z is a (uniformly) well-defined coarse map. We now show that q is a well-defined coarse map that uniformly coarse coincides with the restriction of z to Q_v .

Fix $x \in Q$ and let γ, α be as above. By hypothesis, γ is a (λ, λ) -quasigeodesic of Z and hence ϵ -fellow-travels in Z with α , where $\epsilon = \epsilon(\lambda, \delta)$.

Let α' be the subpath of α from x to x_{α} , so that α' is a geodesic of Z avoiding v and hence a geodesic of Z_v . Let γ' be the subpath of γ from x to x_{γ} , so that γ' is a (λ, λ) -quasigeodesic of Z avoiding v .

If $|\gamma'| < 10^3\lambda\epsilon + \lambda^2 = C$, then $|\alpha'| < C$, so $(\alpha')^{-1}\gamma'$ is a path in Z_v of length $2C$ from x_{α} to x_{γ} . Since $\text{Lk}(v)$ is (δ, δ) -quasi-isometrically embedded in Z_v , this gives an upper bound on $d_{\text{Lk}(v)}(x_{\alpha}, x_{\gamma})$ in terms of $\lambda, \epsilon, \delta$.

Suppose $|\gamma'| > C$. Then $|\alpha'| > 10^3\epsilon$. Hence α', γ' contain points a, c such that $d_Z(a, c) \leq \epsilon$ and $d_Z(a, v), d_Z(c, v) \in [100\epsilon, 10^4\epsilon]$. So, any Z -geodesic η from c to a is a Z_v -geodesic of length at most ϵ . Hence, by travelling along γ' from x_{γ} to c , then along η , and then along α' from a to x_{α} , we obtain a path in Z_v of length at most $10^4\epsilon + \lambda 10^4\epsilon + \lambda$. As above, this means that $d_{\text{Lk}(v)}(x_{\alpha}, x_{\beta})$ is bounded above in terms of $\lambda, \epsilon, \delta$.

This shows that $q(x)$ is a well-defined coarse map that uniformly coarsely coincides with $z(x)$ for $x \in Q$.

Now consider the HHS structure $(Z_{\mathcal{V}}, \mathfrak{F})$ constructed in the first part of the proof. Recall that the index set is $\mathfrak{F} = \{S\} \cup \mathcal{V}$, the hyperbolic space associated to S is Z , and the hyperbolic space associated to $v \in \mathcal{V}$ is $\text{Lk}(v)$. The projection $\pi_S : Z_{\mathcal{V}} \rightarrow Z$ is the inclusion, and the projection $\pi_v : Z_{\mathcal{V}} \rightarrow \text{Lk}(v)$ is the map z discussed above.

Now, let $Q_{\mathcal{V}}$ be the induced subgraph of Q produced by removing the vertices in \mathcal{V} . By the same argument that we used for $Z_{\mathcal{V}}$, the pair $(Q_{\mathcal{V}}, \mathfrak{F})$ is an HHS, where $\mathfrak{F} = \{S\} \cup \mathcal{V}$, the hyperbolic space associated to S is Q , and the hyperbolic space for each v is $\text{Lk}(v)$. The projection $Q_{\mathcal{V}} \rightarrow Q$ is again inclusion, and $\pi_v : Q_{\mathcal{V}} \rightarrow \text{Lk}(v)$ is the map $q : Q_v \rightarrow 2^{\text{Lk}(v)}$ defined above. (We can apply the same argument we used for Z because each $\text{Lk}(v)$ is uniformly quasi-isometrically embedded in Q_v , since $\text{Lk}(v)$ is (δ, δ) -quasi-isometrically embedded in Z_v and $\text{Lk}(v) \subset Q_v \subset Z_v$.) Note that the constants for this HHS structure are different from those for $(Z_{\mathcal{V}}, \mathfrak{F})$, but they still depend only on δ and λ .

To conclude, we need to prove that $Q_{\mathcal{V}} \hookrightarrow Z_{\mathcal{V}}$ is a quasi-isometric embedding with constants depending on δ and λ . We saw above that for each $v \in \mathcal{V}$, the projections q, z to $\text{Lk}(v)$ used in the two HHS structures coarsely coincide on $Q_{\mathcal{V}}$ (constants depend on δ, λ). Moreover, the inclusions $Q_{\mathcal{V}} \rightarrow Q, Q_{\mathcal{V}} \rightarrow Z$ obviously coincide.

There exists $\lambda' = \lambda'(\lambda, \delta)$ such that the following hold for $a, b \in Q_{\mathcal{V}}$:

$$d_Q(a, b) \asymp_{\lambda', \lambda'} d_Z(a, b)$$

and, for each $v \in \mathcal{V}$,

$$d_v(q(a), q(b)) \asymp_{\lambda', \lambda'} d_v(z(a), z(b)).$$

Hence, by Theorem 2.9, there exists T_0 , depending on λ' and the HHS constants for $Z_{\mathcal{V}}$, such that the following holds. Let $T_1 \geq T_0$. Then there exists μ_1 , depending on T_1 and the HHS constants, such that

$$d_{Z_{\mathcal{V}}}(a, b) \asymp_{\mu_1, \mu_1} \{d_Q(a, b)\}_{T_1} + \sum_{v \in \mathcal{V}} \{d_v(q(a), q(b))\}_{T_1}.$$

By the distance formula, applied to the HHS structure on $Q_{\mathcal{V}}$, the following holds. For all sufficiently large T_1 , we have μ_2 , depending only on the HHS constants for $Q_{\mathcal{V}}$ (which depend only on δ, λ) such that

$$d_{Q_{\mathcal{V}}}(a, b) \asymp_{\mu_2, \mu_2} \{d_Q(a, b)\}_{T_1} + \sum_{v \in \mathcal{V}} \{d_v(q(a), q(b))\}_{T_1}.$$

Hence there exists $\mu = \mu(\delta, \lambda)$ such that $d_{Q_{\mathcal{V}}}(a, b) \asymp_{\mu, \mu} d_{Z_{\mathcal{V}}}(a, b)$, as required. \square

3.2. Y_{Δ} is hyperbolic for $\mathcal{C}(\Delta)$ large. Let Δ be a non-maximal simplex of X , and assume that $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$. Our main goal is to prove that Y_{Δ} is uniformly hyperbolic. We also have a second goal, which is to prove that if $[\Delta] \sqsubseteq [\Sigma]$ then $Y_{\Delta} \cap \mathcal{C}(\Sigma)$ is (uniformly) quasi-isometrically embedded in Y_{Δ} .

We will apply the former conclusion throughout the rest of the paper. The latter conclusion is used in the proof of Proposition 4.11.

Here is the formal statement:

Proposition 3.3. *For every δ, n there exists δ' so that the following holds. Let (X, W, δ, n) be a combinatorial HHS, and let Δ be a non-maximal simplex of X such that $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$. Then Y_{Δ} is δ' -hyperbolic.*

Moreover, let Σ be a non-maximal simplex of X with $[\Delta] \sqsubseteq [\Sigma]$. Then the inclusion $Y_{\Delta} \cap \mathcal{C}(\Sigma) \hookrightarrow Y_{\Delta}$ is a (δ', δ') -quasi-isometric embedding.

The hyperbolicity will follow from Lemma 3.9 below. The QI-embedding will follow from Lemma 3.10.

For clarity, whenever we say that a constant is uniform we mean that it depends on δ, n only.

We now fix Δ as in the statement of Proposition 3.3 and define a sequence of spaces interpolating between X^{+W} and Y_{Δ} as follows.

Definition 3.4 (Co-level). Let Σ be a non-maximal simplex. We define the *co-level* $cl[\Sigma]$ inductively as follows. First, define $cl[\emptyset] = 0$.

For $k \geq 0$, we say that $[\Sigma]$ has *colevel at least $k+1$* if there exists $[\Sigma']$ with colevel at least k and $[\Sigma] \sqsubseteq [\Sigma']$. The co-level of $[\Sigma]$, denoted $cl[\Sigma]$, is the maximal k such that $[\Sigma]$ has co-level at least k .

We will be interested in the co-level of classes $[\Sigma]$ with $[\Delta] \sqsubseteq [\Sigma]$ for our fixed simplex Δ .

Observe that for any simplices $[\Pi], [\Gamma]$, if $[\Pi] \sqsubseteq [\Sigma]$ then $cl[\Pi] > cl[\Sigma]$.

Definition 3.5. Let $\mathcal{U} = \{[\Sigma] : [\Delta] \sqsubseteq [\Sigma]\}$. For $0 \leq k \leq cl[\Delta]$, define Y_{Δ}^k to be the graph obtained from Y_{Δ} by connecting all pairs of vertices of $\text{Lk}(\Sigma) \cap Y_{\Delta}$ for all $[\Sigma] \in \mathcal{U}$ with $cl[\Sigma] > k$.

Remark 3.6. (Heuristic picture of Y_Δ^k) It might be useful to keep in mind the following rough picture of Y_Δ^k . First, Y_Δ^k contains a quasi-isometric copy of each $\mathcal{C}(\Sigma)$ for $[\Sigma]$ of co-level k , and any two of these have bounded coarse intersection. In fact, these coarse intersections are best thought of to consist of coned-off copies of certain $\mathcal{C}(\Lambda)$ for Λ of co-level strictly larger than k . Finally, Y_Δ^k does *not* contain a copy of $\mathcal{C}(\Lambda)$ for Λ of co-level strictly smaller than k , but rather a blown-up version of it similar to a Y_Λ -space for a link (this might become clearer when, in Section 4, we discuss induced combinatorial HHS structures on links). We do not formulate the proof of hyperbolicity in these terms, but essentially we will show inductively that Y_Δ^k is hyperbolic relative to the aforementioned copies of the $\mathcal{C}(\Sigma)$ for $[\Sigma]$ of co-level k , with coned-off graph quasi-isometric to Y_Δ^{k-1} .

We will need the following auxiliary lemma, which is how we use the hypothesis on the diameter of $\mathcal{C}(\Delta)$.

Lemma 3.7. *Let Σ, Δ be non-maximal simplices of X such that $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$ and $[\Delta] \sqsubseteq [\Sigma]$. Then $\text{Sat}(\Sigma) \subset \text{Sat}(\Delta)$ and hence $Y_\Delta \subset Y_\Sigma$.*

Proof. By Lemma 1.14, there exists a simplex Π of $\text{Lk}(\Sigma)$ such that $[\Delta] = [\Sigma \star \Pi]$. If Σ' is a simplex with $\text{Lk}(\Sigma) = \text{Lk}(\Sigma')$, then the simplex $\Sigma' \star \Delta$ exists, and $\text{Lk}(\Sigma' \star \Delta) = \text{Lk}(\Sigma') \cap \text{Lk}(\Delta) = \text{Lk}(\Sigma \star \Delta) = \text{Lk}(\Delta)$. By definition, any $v \in \text{Sat}(\Sigma)$ is contained in some such Σ' , and hence in $\Sigma' \star \Delta$, from which it follows from the above that $v \in \text{Sat}(\Delta)$, as required. \square

We also need a graph Z_Δ^k , which is a quasi-isometry model of Y_Δ^k . By definition, Y_Δ^k is obtained from Y_Δ by electrifying $\text{Lk}(\Sigma) \cap Y_\Delta$ when $[\Delta] \sqsubseteq [\Sigma]$ and $[\Sigma]$ has co-level at least $k + 1$. In Z_Δ^k , for the $[\Sigma]$ with co-level exactly $k + 1$, we replace the electrification by coning. The reason for working with Z_Δ^k is that it enables the use of Lemma 3.1, once we check that the links of the cone-points are hyperbolic and are quasi-isometrically embedded in the graph obtained from Z_Δ^k by removing them (i.e., in Y_Δ^{k+1}). This graph will be needed in the proof of Lemma 3.9, and also in the proof of Lemma 3.10.

Definition 3.8. Let Z_Δ^k be the graph obtained from Y_Δ by

- connecting all pairs of vertices of $\text{Lk}(\Sigma) \cap Y_\Delta$ for all $[\Sigma] \in \mathcal{U}$ with $cl[\Sigma] > k + 1$,
- adding a vertex $v_{[\Sigma]}$ for each $[\Sigma] \in \mathcal{U}$ with $cl[\Sigma] = k + 1$, and connecting any such $v_{[\Sigma]}$ to all vertices of $\text{Lk}(\Sigma) \cap Y_\Delta$.

We now inductively prove hyperbolicity of the Y_Δ^k , in the lemma below.

Lemma 3.9. *The following hold provided $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$.*

- (1) Y_Δ^0 is uniformly quasi-isometric to X^{+W} .
- (2) $Y_\Delta^{cl[\Delta]} = Y_\Delta$.
- (3) Y_Δ^k is hyperbolic for $0 \leq k \leq cl[\Delta]$, with uniform hyperbolicity constant.

In particular, Y_Δ is uniformly hyperbolic.

Proof. Assertions (1) and (2) hold by construction.

For (3), we prove hyperbolicity by induction on k . The case $k = 0$ holds by item (1) and Definition 1.8. Suppose that Y_Δ^k is hyperbolic for some $k \geq 1$ and $k < cl[\Delta]$.

Since Z_Δ^k is uniformly quasi-isometric to Y_Δ^k , we have that Z_Δ^k is uniformly hyperbolic. Moreover, Y_Δ^{k+1} is obtained from Z_Δ^k by removing the open star of all vertices $v_{[\Sigma]}$ described above (by construction, these vertices form a discrete set). In view of Lemma 3.1, it suffices to show that the link $\text{Lk}(v_{[\Sigma]})$ of any $v_{[\Sigma]}$ in Z_Δ^k is hyperbolic and quasi-isometrically embedded in Y_Δ^{k+1} .

Hyperbolicity: The vertex set of $\text{Lk}(v_{[\Sigma]})$ is by definition $\text{Lk}(\Sigma) \cap Y_\Delta$.

We first construct a Lipschitz map $\psi : \text{Lk}(v_{[\Sigma]}) \rightarrow \mathcal{C}(\Sigma)$. Since $\text{Lk}(v_{[\Sigma]}) = \text{Lk}(\Sigma) \cap Y_\Delta$, whose vertex set is contained in $\mathcal{C}(\Sigma)$, we first define ψ to be the inclusion on vertices.

Next, let e be an edge of $\text{Lk}(v_{[\Sigma]})$. We claim that the endpoints of e are joined by a path of uniformly bounded length in $\mathcal{C}(\Sigma)$, which suffices to extend ψ to a (uniformly) Lipschitz map.

We check the claim by considering the various possibilities for e . By the definition of Z_{Δ}^k , one of the following holds:

- e is an edge of Y_{Δ} with endpoints in $\text{Lk}(\Sigma)$, and hence is already an edge of $\mathcal{C}(\Sigma)$. In this case, the uniformly bounded path is e itself.
- There exists $\Sigma' \in \mathcal{U}$ with $cl[\Sigma'] > cl[\Sigma]$, and the endpoints of e are in $\text{Lk}(\Sigma')$. We claim that $\text{Lk}(\Sigma') \cap \text{Lk}(\Sigma)$ has uniformly bounded diameter in $\mathcal{C}(\Sigma)$. Indeed, first note that since $cl[\Sigma'] > cl[\Sigma]$, we have $[\Sigma] \not\sqsupseteq [\Sigma']$. In particular, $\text{Sat}(\Sigma')$ contains a vertex α not contained in $\text{Sat}(\Sigma)$, i.e., $\alpha \in Y_{\Sigma}$. Now, every vertex of $\text{Lk}(\Sigma) \cap \text{Lk}(\Sigma')$ is joined by an edge of Y_{Σ} to α , so $diam_{Y_{\Sigma}}(\text{Lk}(\Sigma) \cap \text{Lk}(\Sigma')) \leq 2$. Since $\mathcal{C}(\Sigma) \hookrightarrow Y_{\Sigma}$ is a uniform quasi-isometric embedding, we see that $\text{Lk}(\Sigma) \cap \text{Lk}(\Sigma')$ is uniformly bounded in $\mathcal{C}(\Sigma)$, which provides the desired path.

Having thus shown that ψ is uniformly Lipschitz, we now show it is a quasi-isometry. To do this, we will define a map $\phi : \mathcal{C}(\Sigma) \rightarrow \text{Lk}(v_{[\Sigma]})$ and verify that it is a uniformly Lipschitz quasi-inverse for ψ .

First, let ϕ be the identity on $\text{Lk}(\Sigma) \cap Y_{\Delta} \subset \mathcal{C}(\Sigma)$. Fix an arbitrary vertex $w \in \text{Lk}(\Delta) \subseteq \mathcal{C}(\Sigma)$, and complete the definition of ϕ on vertices by sending every $v \in \text{Lk}(\Sigma) - Y_{\Delta}$ to w .

Second, we check that ϕ is a quasi-inverse for ψ . On vertices of $\text{Lk}(\Sigma) \cap Y_{\Delta}$, this is clear; the two maps are literally inverses. On the other hand, if $v \in \text{Lk}(\Sigma) - Y_{\Delta}$, then by definition $\phi(v) = w \in \text{Lk}(\Delta)$. Now, since $v \notin Y_{\Delta}$, then $v \in \text{Sat}(\Delta)$, so v and w are joined by an edge of X , since $w \in \text{Lk}(\Delta)$. Now, since $\text{Lk}(\Delta) \subset \text{Lk}(\Sigma)$, we have $w \in \mathcal{C}(\Sigma)$, and since $v \in \mathcal{C}(\Sigma)$, this edge is in $\mathcal{C}(\Sigma)$, and hence $\psi(\phi(v)) = w$ is distance 1 in $\mathcal{C}(\Sigma)$ from v . This establishes that ψ and ϕ are quasi-inverses.

We are left to show that ϕ is coarsely Lipschitz. This reduces to showing that if $v \in \text{Lk}(\Sigma) - Y_{\Delta}$ is connected to $v' \in \text{Lk}(\Sigma) \cap Y_{\Delta}$, then v' and w lie within bounded distance.

First note, as above, that $v \in \text{Lk}(\Sigma) - Y_{\Delta}$ is equivalent to $v \in \text{Lk}(\Sigma) \cap \text{Sat}(\Delta)$. Since $\text{Lk}(\Delta) \subset \text{Lk}(\Sigma)$, every vertex of $\text{Lk}(\Delta)$ is connected in X to every vertex of Σ . Together with the fact that $v \in \text{Sat}(\Delta)$, this implies that $\text{Lk}(\Delta) \subset \text{Lk}(\Sigma \star v)$. In other words, $[\Delta] \sqsubseteq [\Sigma \star v]$, i.e., $\Sigma \star v \in \mathcal{U}$. Also, clearly $[\Sigma \star v] \not\sqsubseteq [\Sigma]$, so that $cl[\Sigma \star v] > cl([\Sigma])$. From the definition of Z_{Δ}^k , it now follows that any pair of vertices in $\text{Lk}(\Sigma \star v) \cap Y_{\Delta}$ are connected in $\text{Lk}(v_{[\Sigma]})$.

Now we have two cases, reflecting that the edge of $\mathcal{C}(\Sigma)$ joining v, v' is of one of two types:

- If v and v' are connected in X , then $v' \in \text{Lk}(\Sigma \star v) \cap Y_{\Delta}$. On the other hand, $w \in \text{Lk}(\Delta) \subset \text{Lk}(\Sigma \star v) \cap Y_{\Delta}$, so by the preceding discussion, v' and w are adjacent in $\text{Lk}(v_{[\Sigma]})$, as required.
- The other possibility is that v, v' are not connected in X , and are therefore joined by a W -edge of $\mathcal{C}(\Sigma)$. By Definition 1.8.(4), there exist maximal simplices x, x' of X , connected in W , and, moreover, we have simplices Π, Π' , maximal in $\text{Lk}(\Sigma)$, such that $x = \Sigma \star \Pi, x' = \Sigma \star \Pi'$, and $v \in \Pi, v' \in \Pi'$.

We claim that there exists $v'' \in \Pi \cap Y_{\Delta}$. Indeed, if not, then $\Pi \subset \text{Sat}(\Delta)$. Hence, for any $t \in \text{Lk}(\Delta)^{(0)} \subset \text{Lk}(\Sigma)$, we have a simplex $\Pi \star t \subset \text{Lk}(\Sigma)$, contradicting maximality of Π . Hence there is a vertex $v'' \in \Pi \cap Y_{\Delta}$.

Since $v'' \in \Pi$, we have $v'' \in \text{Lk}(\Sigma)$. Since $v \in \Pi$, there is an edge of X joining v, v'' , so $v'' \in \text{Lk}(\Sigma) \cap \text{Lk}(v) = \text{Lk}(\Sigma \star v)$. Finally, we have chosen $v'' \in Y_{\Delta}$. So $v'' \in \text{Lk}(\Sigma \star v) \cap Y_{\Delta}$.

Now, recall that, since w and v'' are both in $\text{Lk}(\Sigma \star v) \cap Y_{\Delta}$, they are joined by an edge of $\text{Lk}(v_{[\Sigma]})$. On the other hand, v' and v'' are connected by a W -edge, since x and x' are W -adjacent. Since $v', v'' \in Y_{\Delta}$, we see that v' is connected to v'' in $\text{Lk}(v_{[\Sigma]})$. Hence v' and w are at distance at most 2 in $\text{Lk}(v_{[\Sigma]})$, as required.

This completes the proof that ψ is a uniform quasi-isometry, so by uniform hyperbolicity of $\mathcal{C}(\Sigma)$, which comes from Definition 1.8, we get the required hyperbolicity of $\text{Lk}(v_{[\Sigma]})$.

Quasi-isometric embedding of $\text{Lk}(v_{[\Sigma]})$ in Y_{Δ}^{k+1} : To show that $\text{Lk}(v_{[\Sigma]})$ is quasi-isometrically embedded in Y_{Δ}^{k+1} it suffices to find a coarsely Lipschitz extension $\hat{\psi} : Y_{\Delta}^{k+1} \rightarrow Y_{\Sigma}$ of ψ . Indeed, given such a $\hat{\psi}$, consider the uniformly coarsely commutative diagram:

$$\begin{array}{ccc} \text{Lk}(v_{[\Sigma]}) & \xrightarrow{\psi} & \mathcal{C}([\Sigma]) \\ \downarrow & & \downarrow \\ Y_{\Delta}^{k+1} & \xrightarrow{\hat{\psi}} & Y_{\Sigma} \end{array}$$

where the inclusion on the right is a uniform quasi-isometric embedding (by Definition 1.8), the map ψ is a uniform quasi-isometric embedding, and the left inclusion is necessarily Lipschitz. If we show that $\hat{\psi}$ is uniformly coarsely Lipschitz, the map $\text{Lk}(v_{[\Sigma]}) \rightarrow Y_{\Delta}^{k+1}$ will then be forced to be a uniform quasi-isometric embedding.

It is now that we use the assumption that $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$, which gives us access to Definition 1.8.(2) and its consequences. Specifically, by Lemma 3.7, we have $\text{Sat}(\Sigma) \subset \text{Sat}(\Delta)$. It follows that the vertex set of Y_{Δ}^{k+1} is contained in the vertex set of Y_{Σ} , and at the level of vertex sets we declare $\hat{\psi}$ to be the inclusion.

The argument for why $\hat{\psi}$ gives a coarsely Lipschitz map is now similar to the argument for ψ (it is still true that if $\text{Lk}(\Sigma') \cap Y_{\Sigma}$ is not bounded, then $[\Sigma]$ is nested into $[\Sigma']$). This completes the proof that Y_{Δ} is uniformly hyperbolic. \square

3.3. QI-embedding of $\mathcal{C}(\Sigma) \cap Y_{\Delta}$ in Y_{Δ} . We now turn to the QI-embedding part of Proposition 3.3, which we restate for convenience:

Lemma 3.10. *The inclusion $\mathcal{C}(\Sigma) \cap Y_{\Delta} \rightarrow Y_{\Delta}$ is a uniform quasi-isometric embedding, where Σ is any non-maximal simplex of X with $[\Delta] \sqsubseteq [\Sigma]$.*

The proof is similar in spirit to that of hyperbolicity, but a bit more technical. Namely, we will construct a sequence of spaces embedded in the various hyperbolic spaces Y_{Δ}^k and prove inductively that they are quasi-isometrically embedded, with the last space in the sequence being $\mathcal{C}(\Sigma) \cap Y_{\Delta}$. As in the proof of hyperbolicity, we will also need intermediate spaces with suitable "cone points" in order to apply Lemma 3.1.

Before the proof, we need some preliminary discussion to construct all the relevant spaces.

3.3.1. Preliminaries on electrifying and coning off links inside $\mathcal{C}(\Sigma) \cap Y_{\Delta}$. Recall that for $0 \leq k \leq \text{cl}[\Delta]$, we have defined graphs Y_{Δ}^k and Z_{Δ}^k in Definition 3.5 and Definition 3.8 respectively.

Fix $[\Sigma]$ as in Lemma 3.10.

Definition 3.11. For $0 \leq k \leq \text{cl}[\Delta]$, defined Q_{Δ}^k to be the induced subgraph of Y_{Δ}^k spanned by $\text{Lk}(\Sigma) \cap Y_{\Delta}^k$. Define R_{Δ}^k to be the induced subgraph of Z_{Δ}^k spanned by the vertex set of $\text{Lk}(\Sigma) \cap Y_{\Delta}^k$ and the cone-points $v_{[\Pi]}$ in Z_{Δ}^k associated to $[\Pi]$ for which $[\Pi] \sqsubseteq [\Sigma]$.

Remark 3.12. By Definition 3.5 and Definition 3.11, the graph Q_{Δ}^k is obtained from $\mathcal{C}(\Sigma) \cap Y_{\Delta}$ by electrifying each subgraph of the form $\text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_{\Delta}$, where $[\Pi]$ is such that $[\Delta] \sqsubseteq [\Pi]$ and $\text{cl}[\Pi] > k$. In particular, Q_{Δ}^{k-1} is obtained from Q_{Δ}^k by electrifying the subgraphs $\text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_{\Delta}$ with $[\Delta] \sqsubseteq [\Pi]$ and $\text{cl}[\Pi] = k$.

Meanwhile, the graph R_{Δ}^{k-1} is obtained from Q_{Δ}^k by coning off $\text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_{\Delta}$ for those $[\Pi]$ with $[\Delta] \sqsubseteq [\Pi]$, $\text{cl}[\Pi] = k$, and the additional property that $\text{Lk}(\Pi) \subset \text{Lk}(\Sigma)$. So, if $v_{[\Pi]}$ is

a cone-point of Z_Δ^{k-1} that lies in R_Δ^{k-1} , then the entire star of $v_{[\Pi]}$ in Z_Δ^{k-1} lies in R_Δ^{k-1} , which is a property we will need to apply Lemma 3.1.

But, if $cl[\Pi] = k$ and $[\Delta] \sqsubseteq [\Pi]$ but $[\Pi] \not\sqsubseteq [\Sigma]$, then $\text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_\Delta$ is electrified in Q_Δ^{k-1} but not coned off in R_Δ^{k-1} . Hence, it is not immediately clear that Q_Δ^{k-1} is quasi-isometric to R_Δ^{k-1} , but we will show later that it is.

We relate the above spaces in the following commutative diagram, which we call the *main diagram* (where one can roughly think of the bottom row as obtained by taking intersections with $\mathcal{C}(\Sigma)$):

$$\begin{array}{ccccc}
 Y_\Delta^k & \xrightarrow{\ell_k} & Z_\Delta^{k-1} & \xrightarrow{\eta} & Y_\Delta^{k-1} \\
 \uparrow i_k & & \uparrow j_{k-1} & & \uparrow i_{k-1} \\
 Q_\Delta^k & \xrightarrow{\ell_k} & R_\Delta^{k-1} & \xrightarrow{\chi} & Q_\Delta^{k-1}
 \end{array}$$

The maps are as follows:

- i_k and i_{k-1} are inclusions; by Definition 3.11, these are graph homomorphisms and hence 1-Lipschitz;
- j_{k-1} is an inclusion, and again, by Definition 3.11, it is 1-lipschitz;
- ℓ_k are inclusions of induced subgraphs, as explained in Remark 3.12 and again 1-lipschitz;
- η is the inclusion on the image of ℓ_k , and sends each $v_{[\Pi]}$ to an arbitrary vertex of $\text{Lk}(\Pi) \cap Y_\Delta$ — this is again 1-lipschitz since $v_{[\Pi]}$ is adjacent to every vertex of $\text{Lk}(\Pi) \cap Y_\Delta$, and this subgraph is electrified in Q_Δ^{k-1} — in fact, as noted in Section 3.2, η is a uniform quasi-isometry;
- $\chi = i_{k-1}^{-1} \circ \eta \circ j_{k-1}$. Note that if $v_{[\Pi]} \in j_{k-1}(R_\Delta^{k-1})$, then $\text{Lk}(\Pi) \cap Y_\Delta \subset \text{Lk}(\Sigma) \cap Y_\Delta$, so $\eta(v_{[\Pi]}) \in \text{Lk}(\Sigma)$, whence χ is well-defined, and it is 1-Lipschitz for the same reason η is.

The next lemma about the Q_Δ^k supports the base case and the final step of an induction in the proof of Lemma 3.10.

Lemma 3.13. *The graphs Q_Δ^k have the following properties:*

- (1) Q_Δ^k has uniformly bounded diameter for $k < cl[\Sigma]$.
- (2) $Q_\Delta^{cl[\Delta]} = \mathcal{C}(\Sigma) \cap Y_\Delta$.

Proof. The second assertion holds by definition: when $k = cl[\Delta]$, we start with $\mathcal{C}(\Sigma) \cap Y_\Delta$ and add no further edges.

The first assertion follows from the definition since the vertex set of Q_Δ^k is the vertex set of $\text{Lk}(\Sigma) \cap Y_\Delta$, and $[\Delta] \sqsubseteq [\Sigma]$, so when $k < cl[\Sigma]$, the definition says that any two vertices are joined by an edge. \square

3.3.2. Y_Δ^k -bounded intersections of links. The following lemma will help us to control subsets $\text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_\Delta$ that are electrified in Q_Δ^{k-1} , and coned off in Z_Δ^{k-1} , but not coned off in R_Δ^{k-1} .

Lemma 3.14. *Suppose that $0 \leq k \leq cl[\Delta]$. Let Π be a non-maximal simplex of X such that $[\Delta] \sqsubseteq [\Pi]$, and $cl[\Pi] = k$, but $[\Pi] \not\sqsubseteq [\Sigma]$.*

Then for all $x, y \in \mathcal{C}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_\Delta$, there exists a path γ in Y_Δ^k , with all vertices in $\mathcal{C}(\Pi)$, such that

- γ joins x to y ;
- γ has length at most $\delta(\delta + 2) + 4$.

In particular, $\text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_\Delta$ has bounded diameter in the graph metric on Y_Δ^k .

Proof. Fix a vertex $w_0 \in \text{Lk}(\Delta) \subset \text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_\Delta$. Fix $x' \in \text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_\Delta$.

Suppose that there exists $v \in \text{Lk}(\Pi) \cap \text{Sat}(\Delta)$ and that x' is adjacent in $\mathcal{C}(\Pi)$ to v . There is a simplex $v \star \Pi$ of X . Now, since $v \in \text{Sat}(\Delta)$, we have $\text{Lk}(\Delta) \subset \text{Lk}(v) \cap \text{Lk}(\Pi)$, so $v \star \Pi$ is non-maximal and $[\Delta] \sqsubseteq [v \star \Pi] \subsetneq [\Pi]$. Hence $cl[v \star \Pi] > k$, and it follows that any two vertices of $\text{Lk}(v \star \Pi)$ are joined by an edge in Y_Δ^k .

There are two possibilities to consider, according to the type of edge e joining x' to v :

- If e is an edge of X , then $x' \in \text{Lk}(v) \cap \text{Lk}(\Pi)$, so $x' \in \text{Lk}(v \star \Pi)$, and thus Y_Δ^k has an edge α joining x' to w_0 .
- If e is a W -edge (and there is no X -edge), then since $x', v \in \text{Lk}(\Pi)$, Definition 1.8.(4) provides maximal simplices σ, τ of $\text{Lk}(\Pi)$ such that $x' \in \sigma \star \Pi$, and $v \in \tau \star \Pi$, and $\sigma \star \Pi, \tau \star \Pi$ are W -adjacent. Since $\tau \star \Pi$ is maximal, there exists $v' \in \tau \star \Pi - \text{Sat}(\Delta)$ (since saturations cannot contain maximal simplices). Moreover, since $\mathcal{C}(\Delta)$ has diameter at least δ , Lemma 3.7 implies that $\Pi \subset \text{Sat}(\Delta)$, so $v' \in \tau$. Thus x' is joined to v' by a W -edge, and hence by an edge α_0 of Y_Δ^k . And since $v' \in \text{Lk}(v \star \Pi)$, there is an edge α_1 of Y_Δ^k from v' to w_0 . So x' is joined to w_0 by the length-2 path $\alpha = \alpha_0 \alpha_1$ in Y_Δ^k whose vertices are in $\text{Lk}(\Pi)$.

Now let x, y be as in the statement. Since $[\Pi] \not\sqsubseteq [\Sigma]$, we have $\text{Sat}(\Sigma) \not\subset \text{Sat}(\Pi)$. Hence there exists $u \in \text{Sat}(\Sigma) \cap Y_\Pi$, so since $x, y \in \text{Lk}(\Pi)$, there is a path of length 2 in Y_Π (consisting of X -edges) from x to y . Since $\mathcal{C}(\Pi) \rightarrow Y_\Pi$ is a (δ, δ) -quasi-isometric embedding, we get a path $\gamma' \subset \mathcal{C}(\Pi)$ from x to y with $|\gamma'| \leq \delta(\delta + 2)$. If $\gamma' \subset Y_\Delta$, we are done, taking $\gamma = \gamma'$.

Hence suppose that γ' passes through $\text{Sat}(\Delta)$. Let v_0, v_1 be the first and last vertices of $\text{Sat}(\Delta)$ along γ' (it is possible that $v_0 = v_1$). Hence we have vertices x, x', v_0, v_1, y', y in γ' , in that order, such that x, x' and y', y are joined by subgeodesics of γ' lying in Y_Δ , and $v_0, v_1 \in \text{Sat}(\Delta)$, and the subgeodesics of γ' from x' to v_0 and from v_1 to y' are edges of $\mathcal{C}(\Pi)$. Then the first part of the proof shows that x' is joined to w_0 by a path α_0 in Y_Δ^k , and w_0 is joined to y' by a path α_1 in Y_Δ^k , such that $|\alpha_0|, |\alpha_1| \leq 2$, and all vertices in both paths are in $\text{Lk}(\Pi)$.

Hence x is joined to y by a path γ obtained by concatenating the subgeodesic of γ' from x to x' , the path $\alpha_0 \alpha_1$, and the subgeodesic of γ' from y' to y . So, $|\gamma| \leq \delta(\delta + 2) + 4$, as claimed. \square

3.3.3. Proof of Lemma 3.10. To prove that $\mathcal{C}(\Sigma) \cap Y_\Delta \hookrightarrow Y_\Delta$ is a quasi-isometric embedding, we actually prove the following lemma:

Lemma 3.15. *For all $0 \leq k \leq cl[\Delta]$, the map $i_k : Q_\Delta^k \rightarrow Y_\Delta^k$ is a uniform quasi-isometric embedding.*

Proof. First observe that the lemma holds for $0 \leq k < cl[\Sigma]$ because Q_Δ^k has uniformly bounded diameter in this case, by Lemma 3.13. So we only have to prove the lemma for $cl[\Sigma] \leq k \leq cl[\Delta]$.

Let $\ell = cl[\Delta] - cl[\Sigma]$. We will argue by induction on ℓ , i.e., for all $[\Sigma]$ with $cl[\Delta] - cl[\Sigma] = \ell$ and $[\Delta] \sqsubseteq [\Sigma]$, we will prove the claim in the lemma for all k .

Base case $\ell = 0$: When $\ell = 0$, the fact that $[\Delta] \sqsubseteq [\Sigma]$ implies that $[\Delta] = [\Sigma]$. In this case we have $Y_\Delta^{cl(\Delta)} = Y_\Delta$ by Lemma 3.9 and $Q_\Delta^{cl(\Delta)} = \mathcal{C}(\Delta) \cap Y_\Delta = \mathcal{C}(\Delta)$ by Lemma 3.13. Hence, by Definition 1.8.(2), $i_{cl[\Delta]}$ is a (δ, δ) -quasi-isometric embedding. Since we only need consider $k = cl[\Delta] = cl[\Sigma]$ in this case, we are done.

Inductive step from $\ell - 1$ to ℓ : Let $\ell \geq 1$ and suppose that the lemma holds for any $[\Sigma']$ such that $[\Delta] \sqsubseteq [\Sigma']$ and $cl[\Delta] - cl[\Sigma'] < \ell$. Suppose that $cl[\Delta] - cl[\Sigma] = \ell$.

We now argue by induction on k that i_k is a uniform quasi-isometric embedding. We saw above that this holds in the base cases $k < cl[\Sigma]$. So, suppose $cl[\Sigma] \leq k \leq cl[\Delta]$ and suppose, by induction, that i_{k-1} is a (δ_0, δ_0) -quasi-isometric embedding, where δ_0 is a uniform constant. Recall the main diagram:

$$\begin{array}{ccccc}
 Y_{\Delta}^k & \xrightarrow{\ell_k} & Z_{\Delta}^{k-1} & \xrightarrow{\eta} & Y_{\Delta}^{k-1} \\
 \uparrow i_k & & \uparrow j_{k-1} & & \uparrow i_{k-1} \\
 Q_{\Delta}^k & \xrightarrow{\ell_k} & R_{\Delta}^{k-1} & \xrightarrow{\chi} & Q_{\Delta}^{k-1},
 \end{array}$$

in which all maps are 1-Lipschitz, η is a uniform quasi-isometry, and i_{k-1} is a uniform quasi-isometric embedding. We show below, in Claim 3.16 and the surrounding discussion, that χ is a uniform quasi-isometry. Inspecting the main diagram then shows that j_{k-1} is a uniform quasi-isometric embedding.

From here, we conclude as follows. By construction, $j_{k-1}(R_{\Delta}^{k-1})$ — which is isomorphic to R_{Δ}^{k-1} via j_{k-1} — is an induced subgraph of Z_{Δ}^{k-1} that contains the full star of any cone-vertex of Z_{Δ}^{k-1} whenever it contains the cone-vertex. Indeed, if $v_{[\Pi]}$ is a cone-vertex in $j_{k-1}(R_{\Delta}^{k-1})$, then by Definition 3.11, $\text{Lk}(\Pi) \cap Y_{\Delta} \subset \text{Lk}(\Sigma) \cap Y_{\Delta}$. It then follows immediately from Lemma 3.1 that i_k is a uniform quasi-isometric embedding. The fact that j_{k-1} is a quasi-isometric embedding is needed in order to satisfy the hypothesis of Lemma 3.1, which requires a quasi-isometric embedding at the level of coned-off graphs. The hypotheses involving the "ambient" graphs Y_{Δ}^k and Z_{Δ}^{k-1} were already checked during the proof of Lemma 3.9.

Proving that χ is a quasi-isometry: It remains to show that χ is a uniform quasi-isometry. We have already seen that χ is 1-Lipschitz, so it remains to construct a uniformly coarsely Lipschitz quasi-inverse $\bar{\chi}$.

On the subgraph $\mathcal{C}(\Sigma) \cap Y_{\Delta}$, which is a subgraph of both Q_{Δ}^{k-1} and R_{Δ}^{k-1} , and which contains all vertices of Q_{Δ}^{k-1} , we define $\bar{\chi}$ to be the identity. To conclude that χ is a uniform quasi-isometry, it therefore suffices to bound the distance in R_{Δ}^{k-1} between any two vertices $x, y \in Q_{\Delta}^{k-1}$ that are joined by an edge in Q_{Δ}^{k-1} that is not an edge in $\mathcal{C}(\Sigma) \cap Y_{\Delta}$.

First, if there exists Π such that $[\Delta] \sqsubseteq [\Pi]$ and $cl[\Pi] > k$ and $x, y \in \text{Lk}(\Pi)$, then x, y are joined by an edge in Q_{Δ}^k and hence in R_{Δ}^{k-1} , as required.

Second, if there exists Π such that $[\Delta] \sqsubseteq [\Pi] \sqsubseteq [\Sigma]$ and $cl[\Pi] = k$ and $x, y \in \text{Lk}(\Pi)$, then R_{Δ}^{k-1} contains a cone-point $v_{[\Pi]}$ connected by edges to both x and y , and we are again done.

The remaining case is where there exists Π such that $[\Delta] \sqsubseteq [\Pi]$, and $cl[\Pi] = k$, but $[\Pi] \not\sqsubseteq [\Sigma]$, and $x, y \in \text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_{\Delta}$. This case is dealt with in the following claim, which is the whole reason we are arguing by induction on ℓ and k , and not merely on k .

Claim 3.16. *For x, y as above, there exists a uniform constant C such that $d_{R^{k-1}}(x, y) \leq C$.*

Proof of Claim 3.16. We will produce a path in Q_{Δ}^k , of uniformly bounded length, joining x to y . Since Q_{Δ}^k is a subgraph of R_{Δ}^{k-1} , this suffices.

First, observe that $k < cl[\Delta]$. Indeed, if $k = cl[\Delta]$, then $[\Pi] = [\Delta]$ since $cl[\Pi] = k$ and $[\Delta] \sqsubseteq [\Pi]$. But since $[\Delta] \sqsubseteq [\Sigma]$, this implies $[\Pi] \sqsubseteq [\Sigma]$, contradicting our choice of x, y (this case was handled above).

We next reduce to the case where $y \in \text{Lk}(\Delta)$ (which is contained in $\text{Lk}(\Pi) \cap \text{Lk}(\Sigma) \cap Y_{\Delta}$ since $[\Delta] \sqsubseteq [\Sigma], [\Delta] \sqsubseteq [\Pi]$). Suppose that x can be joined to some $x' \in \text{Lk}(\Delta)$ by a path in Q_{Δ}^k of length bounded by a uniform constant C_0 , and that the same is true with x replaced by y and x' replaced by some $y' \in \text{Lk}(\Delta)$. Now, since $k < cl[\Delta]$, any two vertices of $\text{Lk}(\Delta)$ are connected by an edge of Y_{Δ}^k , and hence x', y' are connected by an edge of Q_{Δ}^k . Thus x, y are connected by a path of length at most $2C_0 + 1$ in Q_{Δ}^k . Since this bound is uniform, we are done. Hence we can and shall assume $y \in \text{Lk}(\Delta)$.

Next, since $[\Pi] \not\sqsubseteq [\Sigma]$ and $cl[\Sigma] \leq k = cl[\Pi]$, there is no nesting relation between $[\Pi]$ and $[\Sigma]$. As in the proof of Lemma 3.14, but with Π and Σ switching roles, this implies that there

exists $v \in \text{Sat}(\Pi) - \text{Sat}(\Sigma) = Y_\Sigma \cap \text{Sat}(\Pi)$. Since $x, y \in \text{Lk}(\Pi)$, we obtain a path of length 2 that joins x to y (passing through v) and lies in Y_Σ . Hence, by Definition 1.8.(2), there exists a geodesic α of $\mathcal{C}(\Sigma)$ that joins x to y and satisfies $|\alpha| \leq \delta(\delta + 2)$.

If $\alpha \subset Y_\Delta$, then since $\mathcal{C}(\Sigma) \cap Y_\Delta$ is contained in Q_Δ^k , and α has uniformly bounded length, we are done. Therefore, we can assume that α contains a vertex u of $\text{Sat}(\Delta)$. Since $y \in \text{Lk}(\Delta)$, there is an edge f of X from u to y , so since α is a geodesic, this edge must be the terminal edge of α . Hence we have $\alpha = \alpha_0 \cdot e \cdot f$, where α_0 is a path in $\mathcal{C}(\Sigma) \cap Y_\Delta \subset Q_\Delta^k$ of length at most $\delta(\delta + 2)$ and e is an edge of $\mathcal{C}(\Sigma) \subset Q_\Delta^k$.

Let x' be the vertex of α_0 immediately preceding u , so that e joins x' to u .

Note that X contains a simplex $u \star \Sigma$, since $u \in \mathcal{C}(\Sigma)$. Since $u \in \text{Sat}(\Delta)$ and $[\Delta] \sqsubseteq [\Sigma]$, we have $\text{Lk}(\Delta) \subset \text{Lk}(u) \cap \text{Lk}(\Sigma) = \text{Lk}(u \star \Sigma)$, i.e., $[\Delta] \sqsubseteq [u \star \Sigma] \subsetneq [\Sigma]$. So, $cl[\Delta] \geq cl[u \star \Sigma] > cl[\Sigma]$, and $y \in \text{Lk}(\Delta)$, so $y \in \text{Lk}(u \star \Sigma)$.

Now we analyse two cases, according to whether or not e is an X -edge.

If e is an X -edge, then $x' \in \text{Lk}(u \star \Sigma)$. For later purposes, note that the following argument gives in particular that whenever $x' \in \text{Lk}(u \star \Sigma)$ there exists a path of uniformly bounded length in Q_Δ^k from x' to y .

Notice that $cl[\Delta] - cl[u \star \Sigma] < cl[\Delta] - cl[\Sigma] = \ell$, so by our induction hypothesis (from the ‘‘outer’’ induction, on ℓ), the inclusion $\mathcal{Q} \hookrightarrow Y_\Delta^k$ is a (δ_1, δ_1) -quasi-isometric embedding, where δ_1 is uniform and \mathcal{Q} is the induced subgraph of Y_Δ^k spanned by the vertex set of $\mathcal{C}(u \star \Sigma) \cap Y_\Delta^k$. Note that $\mathcal{Q} \subset Q_\Delta^k$ since $[u \star \Sigma] \sqsubseteq [\Sigma]$.

Now, since $x, y \in \text{Lk}(\Sigma) \cap \text{Lk}(\Pi) \cap Y_\Delta$ and $[\Pi] \not\sqsubseteq [\Delta]$, Lemma 3.14 provides a path β in Y_Δ^k , of length at most $\delta(\delta + 2) + 4$, that joins x to y . By traveling backward along α_0 , and then along β , we get a uniformly bounded path in Y_Δ^k from x' to y . Thus \mathcal{Q} contains a path of uniformly bounded length (in terms of δ_1) from x' to y . So, since \mathcal{Q} is a subgraph of Q_Δ^k , we have a path in Q_Δ^k that has uniformly bounded length and joins x' to y . Prepending α_0 to this path completes the proof in this case.

Finally, suppose that e is a W -edge, and that x' is not joined to u by an X -edge. By Definition 1.8.(4), there exist maximal simplices σ, τ of $\text{Lk}(\Sigma)$ such that $x' \in \sigma \star \Sigma$ and $u \in \tau \star \Sigma$, and every vertex of σ is joined to every vertex of τ by an edge of $\mathcal{C}(\Sigma)$. Since saturations cannot contain maximal simplices, there exists $u' \in (\tau \star \Sigma)^{(0)} \cap Y_\Delta$. Note that u' cannot lie in $\text{Sat}(\Sigma)$, because $\text{Sat}(\Sigma) \subset \text{Sat}(\Delta)$ by Lemma 3.7. Hence $u' \in \tau \subset \text{Lk}(\Sigma)$.

It follows that the concatenation of α_0 with the edge of $\mathcal{C}(\Sigma)$ from x' to u' is a uniformly bounded path in $\mathcal{C}(\Sigma) \cap Y_\Delta$, so it suffices to find a uniformly bounded path from u' to y . But u' is joined to u by an X -edge, so that, as observed above, the previous argument applies with u' replacing x' to bound the distance in Q_Δ^k from u' to y , and hence from x to y . This proves the claim. \blacksquare

Conclusion: We have shown that i_k is a quasi-isometric embedding. The constants increased by a uniform amount (arising from Lemma 3.1) in the inductive step from $k - 1$ to k , but there are boundedly many such steps. The constants also depended in a uniform way on the constants obtained at stage $\ell - 1$, but there are boundedly many possible values of ℓ (namely the complexity). Hence the constants are uniform. \square

Finally:

Proof of Lemma 3.10. Apply Lemma 3.15 to the given $[\Sigma]$ in the case where $k = cl[\Delta]$; this shows that i_k is a uniform quasi-isometric embedding. But in view of Lemma 3.13, i_k is the inclusion $\mathcal{C}(\Sigma) \cap Y_\Delta \hookrightarrow Y_\Delta$, as required. \square

The proof of Proposition 3.3 is now complete: Lemma 3.9 established uniform hyperbolicity of Y_Δ , and Lemma 3.10 established that $\mathcal{C}(\Sigma) \cap Y_\Delta \rightarrow Y_\Delta$ is a uniform quasi-isometric embedding when $[\Delta] \sqsubseteq [\Sigma]$. \square

We remind the reader that both conclusions of the proposition require that $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$, but as we shall see below, we only need the proposition in that case. Indeed, morally, we need hyperbolicity in many places for there to be a uniformly coarsely lipschitz projection $Y_\Delta \rightarrow \mathcal{C}(\Delta)$, but we can define this arbitrarily when $\mathcal{C}(\Delta)$ is bounded. We need the QI-embedding statement once, in the proof of Proposition 4.11, which obviously holds when the simplex $\Sigma \star \Delta$ in the statement of the proposition has uniformly bounded augmented link.

4. COMBINATORIAL HHS STRUCTURES ON LINKS, AND HOMOMORPHISMS

We continue to let (X, W, δ, n) be a combinatorial HHS. In this section, we discuss induced combinatorial HHS structures on links of simplices. The goal is to show, roughly speaking, that the inclusion of a link in X is compatible with the combinatorial HHS structures. This will allow us to perform inductive arguments, since the complexity of the HHS structure on a link is strictly smaller than that of X .

The main results of this section are Propositions 4.9, describing the combinatorial HHS structure on links, and 4.11, stating exactly what we mean by the inclusion being compatible with the combinatorial HHS structures.

Fix a non-maximal simplex Δ of X . Let $\mathcal{C}_0(\Delta)$ be the subgraph of Y_Δ obtained from $\text{Lk}(\Delta)$ by adding an edge from v to w if the following holds: there exist simplices x, y of X such that $\Delta \star x, \Delta \star y$ are maximal simplices of X that respectively contain v and w and are connected in W . (Note that when $[\Delta] = [\emptyset]$, we have $\text{Lk}(\Delta) = X$, and $\mathcal{C}_0(\Delta)$ is obtained from X by joining v, w whenever they are respectively contained in maximal simplices x, y of X with x, y adjacent in W . To see this, note that $[\Delta] = [\emptyset]$ implies that $\Delta = \emptyset$, since $\text{Lk}(\emptyset) = X$. Hence the maximal simplices of the form $\Delta \star x$ of X are exactly the maximal simplices x of X .)

The advantage of this definition for us is that it has better inheritance properties than $\mathcal{C}(\Delta)$. However, and this is the only use of Definition 1.8.4, it defines the same object:

Lemma 4.1. *Given a combinatorial HHS (X, W) and a non-maximal simplex Δ of X , we have $\mathcal{C}_0(\Delta) = \mathcal{C}(\Delta)$.*

Proof. Clearly, $\mathcal{C}_0(\Delta)$ is a subgraph of $\mathcal{C}(\Delta)$, and the vertex sets are the same. So, it suffices to show that any edge e of $\mathcal{C}(\Delta)$ is also an edge of $\mathcal{C}_0(\Delta)$. If e is an edge of X , then this is clear. Otherwise, e is an edge coming from W , meaning that the endpoints of e are contained in W -adjacent maximal simplices. Definition 1.8.4 provides W -adjacent maximal simplices containing the endpoints of the form required to yield an edge of $\mathcal{C}_0(\Delta)$. \square

In the rest of this section we will most often use the notation $\mathcal{C}_0(\Delta)$, except in the proof of Proposition 4.11, where we need to use results from Section 3.

Definition 4.2 (Induced $\text{Lk}(\Delta)$ -graph). Let Δ be a non-maximal simplex of X . The *induced $\text{Lk}(\Delta)$ -graph* is the graph W^Δ defined as follows.

The vertex set of W^Δ is the set of maximal simplices Σ of $\text{Lk}(\Delta)$. Notice that any such simplex Σ has the property that $\Delta \star \Sigma$ is a maximal simplex of X , and hence a vertex of W .

Two simplices Σ, Σ' of $\text{Lk}(\Delta)$ are connected in W^Δ if and only if $\Delta \star \Sigma$ and $\Delta \star \Sigma'$ are connected in W . By definition, W^Δ is a $\text{Lk}(\Delta)$ -graph in the sense of Definition 1.2.

Remark 4.3 (The induced map $W^\Delta \rightarrow W$). There is a natural injective graph morphism $W^\Delta \hookrightarrow W$ whose image is an induced subgraph. Indeed, the maximal simplex Σ of $\text{Lk}(\Delta)$ is sent to the maximal simplex $\Delta \star \Sigma$ of X . Hence we can view W^Δ as a subgraph of W .

Definition 4.4 (The induced map $\text{Lk}(\Delta)^{+W^\Delta} \rightarrow X^{+W}$). There is an induced simplicial map $\iota_\Delta : \text{Lk}(\Delta)^{+W^\Delta} \rightarrow X^{+W}$ defined as follows.

For each $v \in \text{Lk}(\Delta)^{(0)}$, let $\iota_\Delta(v) = v$, where v is viewed as a vertex of X . If $v, w \in \text{Lk}(\Delta)^{(0)}$ are joined by an edge e of $\text{Lk}(\Delta)$, then e is an edge of X (and hence X^{+W}) and we let $\iota_\Delta(e) = e$. If σ, σ' are maximal simplices of $\text{Lk}(\Delta)$ corresponding to vertices of W^Δ , then by definition, $\sigma \star \Delta$ and $\sigma' \star \Delta$ are maximal simplices of X corresponding to vertices of W . If σ, σ' are adjacent in W^Δ , then $\sigma \star \Delta$ and $\sigma' \star \Delta$ are adjacent in W , by Definition 4.2. Thus $\text{Lk}(\Delta)^{+W^\Delta}$ contains the 1-skeleton D of the join $\sigma \star \sigma'$ and X^{+W} contains the 1-skeleton D' of the join $(\sigma \star \sigma') \star \Delta$, and we define ι_Δ on D to be the inclusion $D \rightarrow D'$.

Remark 4.5. Chasing the definitions, one sees that $\iota_\Delta(\text{Lk}(\Delta)^{+W^\Delta}) = \mathcal{C}_0(\Delta)$. In fact, the two graphs have the same vertex set, and in either case two vertices are connected by an edge if and only if they are contained in maximal simplices of the form $\Delta \star \Pi$ that are connected in W .

Notation. We now revisit the notions associated to a combinatorial HHS in the context of $\text{Lk}(\Delta)$. For each non-maximal simplex Σ of $\text{Lk}(\Delta)$, let $\text{Lk}^\Delta(\Sigma)$ be the link of Σ in $\text{Lk}(\Delta)$, let \sim_Δ be the resulting equivalence relation (as in Definition 1.3) on the simplices of $\text{Lk}(\Delta)$, and let $\text{Sat}^\Delta(\Sigma)$ be the union of the vertex sets of simplices of $\text{Lk}(\Delta)$ in the \sim_Δ -class of Σ . Let Y_Σ^Δ be the induced subgraph of $\text{Lk}(\Delta)^{+W^\Delta}$ induced by the vertices of $\text{Lk}(\Delta) - \text{Sat}^\Delta(\Sigma)$.

Let $\mathcal{C}_0^\Delta([\Sigma])$ be the subgraph of $\text{Lk}(\Delta)^{+W^\Delta}$ obtained from $\text{Lk}^\Delta(\Sigma)$ by adding edges connecting v, w if v, w are vertices of maximal simplices of $\text{Lk}(\Delta)$ of the form $\Sigma \star x, \Sigma \star y$ that are connected in W^Δ .

Let \mathfrak{S}_Δ be the set of \sim_Δ -classes of non-maximal simplices in $\text{Lk}(\Delta)$. The relation \sim_Δ and the accompanying graphs $\text{Lk}^\Delta(\Sigma), \text{Sat}^\Delta(\Sigma)$ (where Σ is a simplex of $\text{Lk}(\Delta)$) allow us to define relations $\sqsubseteq, \perp, \pitchfork$ on \mathfrak{S}_Δ as in Definition 1.11.

The \perp relation on \mathfrak{S}_Δ warrants some extra comment. Let Σ, Σ' be non-maximal simplices of $\text{Lk}(\Delta)$. Following Definition 1.11, we declare Σ, Σ' to be orthogonal in \mathfrak{S}_Δ if and only if $\text{Lk}^\Delta(\Sigma') \subseteq \text{Lk}^\Delta(\text{Lk}^\Delta(\Sigma))$. Note that the left side is $\text{Lk}(\Sigma') \cap \text{Lk}(\Delta)$, and the right side is $\text{Lk}(\text{Lk}(\Sigma) \cap \text{Lk}(\Delta)) \cap \text{Lk}(\Delta)$.

Lemma 4.6 (Induced map on index sets). *Let Σ be a non-maximal (possibly empty) simplex of $\text{Lk}(\Delta)$. The assignment $\Sigma \mapsto \Sigma \star \Delta$ induces an injective map $\iota^* : \mathfrak{S}_\Delta \rightarrow \mathfrak{S}$ that preserves the $\sqsubseteq, \perp, \pitchfork$ relations.*

Proof. We first show that $\Sigma \sim_\Delta \Sigma'$ if and only if $\Sigma \star \Delta \sim \Sigma' \star \Delta$, showing that ι^* is well-defined and injective.

Recall that $\text{Lk}(\Sigma \star \Delta) = \text{Lk}(\Sigma) \cap \text{Lk}(\Delta) = \text{Lk}^\Delta(\Sigma)$, and similarly for Σ' , so that we have $\text{Lk}(\Sigma \star \Delta) = \text{Lk}(\Sigma' \star \Delta)$ if and only if $\text{Lk}^\Delta(\Sigma) = \text{Lk}^\Delta(\Sigma')$, as required.

Preservation of nesting and non-nesting: For the same reason as above, we have $\text{Lk}(\Sigma \star \Delta) \subseteq \text{Lk}(\Sigma' \star \Delta)$ if and only if $\text{Lk}^\Delta(\Sigma) \subseteq \text{Lk}^\Delta(\Sigma')$, that is, $[\Sigma] \sqsubseteq [\Sigma']$ in \mathfrak{S}^Δ if and only if $[\Sigma \star \Delta] \sqsubseteq [\Sigma' \star \Delta]$ in \mathfrak{S} .

Preservation of orthogonality and non-orthogonality: Suppose that Σ, Σ' are simplices of $\text{Lk}(\Delta)$ satisfying $[\Sigma] \perp [\Sigma']$. Then by definition, we have

$$\text{Lk}(\Sigma') \cap \text{Lk}(\Delta) \subseteq \text{Lk}(\text{Lk}(\Sigma) \cap \text{Lk}(\Delta)) \cap \text{Lk}(\Delta).$$

In other words, $\text{Lk}(\Sigma' \star \Delta) \subseteq \text{Lk}(\text{Lk}(\Sigma \star \Delta)) \cap \text{Lk}(\Delta)$. Hence $[\Sigma' \star \Delta] \perp [\Sigma \star \Delta]$ in \mathfrak{S} , as required.

Conversely, if $[\Sigma' \star \Delta] \perp [\Sigma \star \Delta]$ in \mathfrak{S} , then by definition $\text{Lk}(\Sigma') \cap \text{Lk}(\Delta) \subseteq \text{Lk}(\text{Lk}(\Sigma) \cap \text{Lk}(\Delta))$, but since the left-hand side is contained in $\text{Lk}(\Delta)$ we also have

$$\text{Lk}(\Sigma') \cap \text{Lk}(\Delta) \subseteq \text{Lk}(\text{Lk}(\Sigma) \cap \text{Lk}(\Delta)) \cap \text{Lk}(\Delta),$$

that is $[\Sigma] \perp [\Sigma']$ in \mathfrak{S}_Δ . □

Corollary 4.7. *If Δ is non-empty, then the complexity of $\text{Lk}(\Delta)$ is strictly less than the complexity n of (X, W) .*

Proof. Any \sqsubseteq -chain in \mathfrak{S}_Δ maps via ι^* to a \sqsubseteq -chain in \mathfrak{S} of the same length, by Lemma 4.6. Let Σ be the \sqsubseteq -maximal element of such a chain in \mathfrak{S}_Δ . Then $[\Sigma \star \Delta]$ is properly nested in $[\emptyset]$, which lies in $\mathfrak{S} - \iota^*(\mathfrak{S}_\Delta)$, whence the complexity of \mathfrak{S}_Δ is strictly less than the complexity n of \mathfrak{S} . \square

We are aiming to show that $(\text{Lk}(\Delta), W^\Delta, \delta', m)$ inherits a combinatorial HHS structure from (X, W, δ, n) , where δ' is uniform, and $m < n$. The preceding statements established what we need at the level of index sets. Now we study the spaces $\mathcal{C}_0^\Delta(\Sigma)$.

Lemma 4.8. *Let Σ be a simplex of $\text{Lk}(\Delta)$. Then:*

- (1) ι_Δ restricts to an isomorphism of graphs from $\mathcal{C}_0^\Delta([\Sigma])$ to $\mathcal{C}_0([\Sigma \star \Delta])$.
- (2) Suppose that $\text{diam}(\mathcal{C}_0(\Sigma \star \Delta)) \geq \delta$. Then $\iota_\Delta(Y_\Sigma^\Delta) = Y_{\Sigma \star \Delta} \cap \mathcal{C}_0(\Delta)$, and ι_Δ restricts to a graph isomorphism from Y_Σ^Δ to $Y_{\Sigma \star \Delta} \cap \mathcal{C}_0(\Delta)$.

Proof. Proof of item (1). Let v be a vertex of $\mathcal{C}^\Delta([\Sigma])$, that is, v is a vertex of $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$. But then v is a vertex of $\text{Lk}(\Delta \star \Sigma)$, so that $\iota_\Delta(v) = v$ is a vertex of $\mathcal{C}([\Sigma \star \Delta])$. Notice also that, for similar reasons, every vertex of $\text{Lk}(\Delta \star \Sigma)$ is in the image of ι_Δ , and hence ι_Δ restricts to a bijection between the vertex sets of $\mathcal{C}^\Delta([\Sigma])$ and $\mathcal{C}([\Sigma \star \Delta])$. We now have to show that v, w are connected by an edge in $\mathcal{C}_0^\Delta([\Sigma])$ if and only if they are connected by an edge in $\mathcal{C}_0([\Sigma \star \Delta])$. Clearly, such v, w are connected by an edge of $\text{Lk}(\Delta)$ if and only if they are connected by an edge of X . Hence, suppose that v, w are connected by one of the additional edges coming from W^Δ , meaning that v, w are vertices of maximal simplices $\Sigma \star x, \Sigma \star y$ of $\text{Lk}(\Delta)$ that are connected in W^Δ . But then, by definition of W^Δ , $\Delta \star \Sigma \star x$ and $\Delta \star \Sigma \star y$ are connected in W , showing that v, w are connected in $\mathcal{C}_0([\Sigma \star \Delta])$. The other implication is similar.

Proof of item (2). Let us first check the statement at the level of vertex sets.

\subseteq : Let $x \in Y_\Sigma^\Delta$ be a vertex, which is to say that $x \in \text{Lk}(\Delta)^{(0)} - \text{Sat}^\Delta(\Sigma)$. Since $x \in \text{Lk}(\Delta)$, we have $x \in \text{Lk}(\Delta)^{+W^\Delta}$. Clearly $x \in \mathcal{C}_0(\Delta)$, so we claim that $x \in Y_{\Sigma \star \Delta}$, i.e., that $x \in X^{(0)} - \text{Sat}(\Sigma \star \Delta)$. Suppose to the contrary that there exists a simplex Π' so that $x \in \Pi'$ and $\text{Lk}(\Pi') = \text{Lk}(\Sigma) \cap \text{Lk}(\Delta)$.

Then we know that $[\Sigma \star \Delta] \sqsubseteq [\Delta \star x]$, since

$$\text{Lk}(\Sigma \star \Delta) = \text{Lk}(\Delta \star x) = \text{Lk}(\Delta) \cap \text{Lk}(x) \supseteq \text{Lk}(\Delta) \cap \text{Lk}(\Pi').$$

Hence, by Lemma 1.14 there exists a simplex Π'' in $\text{Lk}(\Delta \star x)$ with $[\Delta \star x \star \Pi''] = [\Sigma \star \Delta]$. In particular, $\text{Lk}^\Delta(\Pi'' \star x) = \text{Lk}^\Delta(\Sigma)$, showing $x \in \text{Sat}^\Delta(\Sigma)$, a contradiction.

\supseteq : This is similar, but easier. Let $x \in Y_{\Sigma \star \Delta} \cap \mathcal{C}_0(\Delta)$ be a vertex. We have to show that $x \in \text{Lk}(\Delta)^{(0)} - \text{Sat}^\Delta(\Sigma)$. If we had $x \in \text{Sat}^\Delta(\Sigma)$ then there would exist a simplex Π of $\text{Lk}(\Delta)$ so that $x \in \Pi$ and $\text{Lk}^\Delta(\Pi) = \text{Lk}^\Delta(\Sigma)$. But then $\text{Lk}(\Delta \star \Pi) = \text{Lk}(\Delta \star \Sigma)$, and we would have $x \in \text{Sat}(\Sigma \star \Delta)$, a contradiction.

Finally, we are left with arguing that the edge sets are also the same. Observe that two vertices x, y of Y_Σ^Δ are connected by an edge if and only if they are connected by an edge in $\text{Lk}(\Delta)^{+W^\Delta}$. By Remark 4.5, we have $\text{Lk}(\Delta)^{+W^\Delta} = \mathcal{C}_0(\Delta)$. Also, since $Y_{\Sigma \star \Delta} \cap \mathcal{C}_0(\Delta)$ is an induced subgraph of $\mathcal{C}_0(\Delta)$, we have that vertices in $Y_{\Sigma \star \Delta} \cap \mathcal{C}_0(\Delta)$ are connected if and only if they are connected in $\mathcal{C}_0(\Delta)$. To sum up, two vertices x, y of Y_Σ^Δ are connected by an edge if and only if they are connected by an edge in $\mathcal{C}_0(\Delta)$, if and only if they are connected in $Y_{\Sigma \star \Delta} \cap \mathcal{C}_0(\Delta)$, as required. \square

Proposition 4.9 (Combinatorial HHS structure on links). *For each δ there exists δ'' so that the following holds. Given a combinatorial HHS (X, W, δ, n) and a non-maximal, nonempty simplex Δ of X , there is $m < n$ so that $(\text{Lk}(\Delta), W^\Delta, \delta'', m)$ is a combinatorial HHS.*

Proof. The required bound on complexity of \mathfrak{S}_Δ comes from Corollary 4.7.

We will first check Condition 4, which (as in the proof of Lemma 4.1) guarantees that $\mathcal{C}_0^\Delta(\Sigma) = \mathcal{C}^\Delta(\Sigma)$ for every simplex Σ of $\text{Lk}(\Delta)$.

Let $v, w \in \text{Lk}^\Delta(\Sigma)$, and suppose that they are contained in W^Δ -adjacent simplices Σ_1, Σ_2 of $\text{Lk}(\Delta)$. By definition of the edges of W^Δ , we have that $\Delta \star \Sigma_1$ and $\Delta \star \Sigma_2$ are W -adjacent (and clearly still contain v, w respectively). Moreover, $v, w \in \text{Lk}(\Delta \star \Sigma)$, so that by Condition 4 for (X, W) we have that v, w are contained in W -adjacent maximal simplices $\Delta \star \Sigma \star \Gamma_1, \Delta \star \Sigma \star \Gamma_2$. This means that v, w are contained in the W^Δ -adjacent maximal simplices $\Sigma \star \Gamma_1, \Sigma \star \Gamma_2$ of $\text{Lk}(\Delta)$, as required.

Next, we produce δ' so that each $\mathcal{C}_0^\Delta([\Sigma])$ is δ' -hyperbolic, and $\mathcal{C}_0^\Delta([\Sigma]) \hookrightarrow Y_\Sigma^\Delta$ is a (δ', δ') -quasi-isometric embedding.

By Lemma 4.8.1, $\mathcal{C}_0^\Delta([\Sigma])$ is isometric to $\mathcal{C}_0([\Sigma \star \Delta])$. Hence $\mathcal{C}_0^\Delta([\Sigma])$ is δ -hyperbolic.

We now verify the quasi-isometric embedding part. It suffices to consider the case where $\text{diam}(\mathcal{C}_0^\Delta([\Sigma])) \geq \delta$. Hence, in view of Lemma 4.8.2, we can consider the following commutative diagram, where the horizontal arrows are restriction of ι_Δ and the vertical arrows are inclusions:

$$\begin{array}{ccc} Y_\Sigma^\Delta & \longrightarrow & Y_{\Sigma \star \Delta} \\ \uparrow & & \uparrow \\ \mathcal{C}_0^\Delta([\Sigma]) & \xrightarrow{\cong} & \mathcal{C}([\Sigma \star \Delta]) \end{array}$$

the right vertical arrow is a (δ, δ) -quasi-isometric embedding by Definition 1.8, and the top horizontal arrow is a 1-Lipschitz map, and the left vertical arrow is also 1-Lipschitz. Since the bottom horizontal arrow is an isometry, there exists δ'' , depending only on δ , so that the left vertical arrow is a (δ'', δ'') -quasi-isometric embedding. (For instance, $\delta'' = \max\{\delta, 1\}$ suffices.)

It remains to verify condition (3) from Definition 1.8, which is going to easily follow from the same condition for (X, W) . Let Σ, Ω be non-maximal simplices of $\text{Lk}(\Delta)$, and suppose that there exists some non-maximal simplex Γ of $\text{Lk}(\Delta)$ whose \sim_Δ -class is nested into those of Σ and Ω , and $\text{diam}(\mathcal{C}_0^\Delta([\Gamma])) \geq \delta$. Then $[\Gamma \star \Delta] \subseteq [\Sigma \star \Delta], [\Omega \star \Delta]$ by Lemma 4.6, and $\text{diam}(\mathcal{C}_0([\Delta \star \Gamma])) \geq \delta$ by Lemma 4.8.1. By Definition 1.8.(3), applied to the combinatorial HHS (X, W) , there exists a non-maximal simplex Π of $\text{Lk}(\Sigma \star \Delta) = \text{Lk}(\Sigma) \cap \text{Lk}(\Delta)$ such that $[\Pi \star (\Sigma \star \Delta)] \subseteq [\Omega \star \Delta]$. Hence Π is a simplex of $\text{Lk}(\Delta)$ (so $[\Pi] \in \mathfrak{S}_\Delta$), and Π is a simplex of $\text{Lk}^\Delta(\Sigma)$ (as required by Definition 1.8.(3)). Moreover, by the preceding nesting relation, we have $\text{Lk}^\Delta(\Pi \star \Sigma) \subseteq \text{Lk}^\Delta(\Omega)$, i.e., $\Pi \star \Sigma$ is nested in Ω , with respect to the nesting in \mathfrak{S}_Δ .

To check the remaining clause of Definition 1.8.(3), suppose that Γ is a non-maximal simplex of $\text{Lk}(\Delta)$ with $\mathcal{C}_0^\Delta(\Gamma) \cong \mathcal{C}_0(\Gamma \star \Delta)$ having diameter at least δ . Then $[\Gamma \star \Delta] \subseteq [(\Sigma \star \Delta) \star \Pi]$, by Definition 1.8.(3), applied to the combinatorial HHS (X, W) . Hence, as above, we see that $[\Gamma] \subseteq [\Sigma \star \Pi]$, where equivalence classes and nesting are taken in \mathfrak{S}_Δ . This completes the proof. \square

From Proposition 4.9 and the first part of Proposition 3.3, we immediately obtain:

Corollary 4.10. *Let (X, W, δ, n) be a combinatorial HHS. Then there exists δ' , depending only on δ and n , so that the following holds. Let Δ be a non-maximal simplex of X and let Σ be a non-maximal simplex of $\text{Lk}(\Delta)$. Then Y_Σ^Δ is δ' -hyperbolic provided $\text{diam}(\mathcal{C}^\Delta([\Sigma])) \geq \delta'$.*

Compatibility of structures. For convenience, we now recall some more notation from Section 1, and what it yields in the case of $\text{Lk}(\Delta)$ and W^Δ . For each non-maximal simplex Σ of $\text{Lk}(\Delta)$, the projection $\pi_\Sigma^\Delta : W^\Delta \rightarrow \mathcal{C}^\Delta([\Sigma])$ is defined as follows: each vertex $w \in W^\Delta$ is a maximal simplex of $\text{Lk}(\Delta)$, and in particular w is not contained in $\text{Sat}^\Delta(\Sigma)$, since $\text{Lk}^\Delta(\Sigma) \neq \emptyset$. Thus

$w \cap Y_{\Sigma}^{\Delta}$ is a complete graph, and we take $\pi_{[\Sigma]}^{\Delta}(w)$ to be the closest-point projection of $w \cap Y_{\Sigma}^{\Delta}$ on $\mathcal{C}^{\Delta}([\Sigma])$. When Σ, Σ' are non-orthogonal simplices of $\text{Lk}(\Delta)$, $\rho_{[\Sigma]}'$ is defined in Definition 1.16.

Now we are ready to prove the second main proposition of this section:

Proposition 4.11 (Inclusions induce hieromorphisms). *Let $X, W, \delta, n, \Delta, \Sigma$ be as in Proposition 4.9. Then the map $\Sigma \mapsto \Sigma \star \Delta$ induces an isometry $\mathcal{C}^{\Delta}([\Sigma]) \rightarrow \mathcal{C}([\Sigma \star \Delta])$ so that the following diagrams uniformly coarsely commute whenever Σ, Σ' are simplices of $\text{Lk}(\Delta)$ for which $[\Sigma] \pitchfork [\Sigma']$, $[\Sigma] \sqsubset [\Sigma']$ or $[\Sigma'] \sqsubset [\Sigma]$ in \mathfrak{S}_{Δ} :*

$$\begin{array}{ccc} W^{\Delta} \hookrightarrow W & & \mathcal{C}^{\Delta}([\Sigma]) \xrightarrow{\cong} \mathcal{C}([\Sigma \star \Delta]) \\ \pi_{[\Sigma]}^{\Delta} \downarrow & & \rho_{[\Sigma]}' \downarrow \\ \mathcal{C}^{\Delta}([\Sigma]) \xrightarrow{\cong} \mathcal{C}([\Sigma \star \Delta]) & & \mathcal{C}([\Sigma'] \star \Delta) \end{array}$$

Proof. First of all, let us recall that in view of Lemma 4.1 there is no difference between \mathcal{C}_0 -spaces and \mathcal{C} -spaces. This allows us to use results from this section and from Section 3. Second, we may assume that $\text{diam}(\mathcal{C}(\Sigma' \star \Delta)) \geq \delta$, for otherwise uniform coarse commutation of both diagrams is immediate, since the maps are all uniformly coarsely constant in this case.

Hence, by Proposition 3.3, there exists $\delta' = \delta'(\delta)$ so that $Y_{\Sigma \star \Delta}$ is δ' -hyperbolic. By (uniformly) enlarging δ' , we have that Y_{Σ}^{Δ} is δ' -hyperbolic for each non-maximal simplex Σ of $\text{Lk}(\Delta)$, by Corollary 4.10 and Proposition 3.3.

In view of Lemma 4.8.2, whenever $\text{diam}(\mathcal{C}_0^{\Delta}(\Sigma)) \geq \delta$, the ‘‘moreover’’ clause of Proposition 3.3 says that there exists a uniform ϵ such that $Y_{\Sigma}^{\Delta} \hookrightarrow Y_{\Sigma \star \Delta}$ is an (ϵ, ϵ) -quasi-isometric embedding.

Now, as in the proof of Proposition 4.9, provided that $\text{diam}(\mathcal{C}_0^{\Delta}(\Sigma)) \geq \delta$, we have a commutative diagram

$$\begin{array}{ccc} Y_{\Sigma}^{\Delta} \hookrightarrow Y_{\Sigma \star \Delta} & & \\ \uparrow & & \uparrow \\ \mathcal{C}^{\Delta}([\Sigma]) \xrightarrow{\cong} \mathcal{C}([\Sigma \star \Delta]) & & \end{array}$$

where all arrows are (ϵ', ϵ') -quasi-isometric embeddings, where ϵ' is uniform. Let $p : Y_{\Sigma \star \Delta} \rightarrow \mathcal{C}([\Sigma \star \Delta])$ be coarse closest-point projection, and let $p^{\Delta} : Y_{\Sigma}^{\Delta} \rightarrow \mathcal{C}^{\Delta}([\Sigma])$ be the coarse closest-point projection.

For convenience, in the following argument we identify Y_{Σ}^{Δ} and $\mathcal{C}^{\Delta}([\Sigma])$ with their images under ι_{Δ} . Let $a \in Y_{\Sigma}^{\Delta}$. Note that $p(a) \in \mathcal{C}(\Sigma \star \Delta) = \mathcal{C}_{\Sigma}^{\Delta} \subset Y_{\Sigma}^{\Delta}$.

Let γ be a Y_{Σ}^{Δ} -geodesic from a to $p^{\Delta}(a)$. So, γ is a uniform quasi-geodesic of $Y_{\Sigma \star \Delta}$, and hence contains a point b at uniformly bounded distance C (in Y_{Σ}^{Δ}) from $p(a)$, since γ fellow-travels with a geodesic of $Y_{\Sigma \star \Delta}$ that starts at a and ends at $p^{\Delta}(a) \in \mathcal{C}(\Sigma \star \Delta)$. Now, travelling along γ from a to b and then travelling along a Y_{Σ}^{Δ} -geodesic from b to $p(a)$ gives a path from a to $p(a)$ of length at most $d_{Y_{\Sigma}^{\Delta}}(a, b) + C \geq |\gamma| - 1$. Hence $d_{Y_{\Sigma}^{\Delta}}(b, p^{\Delta}(a)) \leq C + 1$, so $d_{Y_{\Sigma}^{\Delta}}(p^{\Delta}(a), p(a)) \leq 2C + 1$. Hence we have a uniform bound on $d_{Y_{\Sigma \star \Delta}}(p(a), p^{\Delta}(a))$, and thus on $d_{\mathcal{C}(\Sigma \star \Delta)}(p(a), p^{\Delta}(a))$. Thus p^{Δ} uniformly coarsely coincides with the restriction of p to Y_{Σ}^{Δ} .

Since $\pi_{[\Sigma \star \Delta]}^{\Delta}, \pi_{[\Sigma]}^{\Delta}$, and $\rho_{[\Sigma]}'$ and $\rho_{[\Sigma \star \Delta]}^{\Sigma' \star \Delta}$ were all defined in terms of p and p_{Δ} (and since when the target of the projection under consideration has bounded diameter the statement holds automatically), checking that the diagrams in the statement coarsely commute is now straightforward; we give the details below.

Let w be a maximal simplex of $\text{Lk}(\Delta)$, i.e., a vertex of W^Δ . So, the image of w under $W^\Delta \hookrightarrow W$ is the maximal simplex $w \star \Delta$ of W . Let Σ be a non-maximal simplex of $\text{Lk}(\Delta)$.

Observe that $(w \star \Delta) \cap Y_{\Sigma \star \Delta} = w \cap Y_\Sigma^\Delta$. Indeed, if a is a vertex of $(w \star \Delta) \cap Y_{\Sigma \star \Delta}$, and each vertex of Δ belongs to $\Sigma \star \Delta$ and hence to $\text{Sat}(\Sigma \star \Delta)$, we must have $a \in w$. Since $w \subset \text{Lk}(\Delta)$, we have by Lemma 4.8.(2) that $a \in Y_\Sigma^\Delta$. Conversely, if $a \in w \cap Y_\Sigma^\Delta$, then by the same lemma, $a \in \text{Lk}(\Delta) \cap Y_{\Sigma \star \Delta}$, and obviously $a \in w \star \Delta$.

Now, by definition,

$$\pi_{[\Sigma]}^\Delta(w) = p^\Delta(w \cap Y_\Sigma^\Delta)$$

and

$$\pi_{[\Sigma \star \Delta]}(w \star \Delta) = p((w \star \Delta) \cap Y_{\Sigma \star \Delta}) = p(w \cap Y_\Sigma^\Delta),$$

so by uniform closeness of p^Δ and p , established earlier in the proof, the sets $\pi_{[\Sigma]}^\Delta(w)$ and $\pi_{[\Sigma \star \Delta]}(w \star \Delta)$ uniformly coarsely coincide. Hence the diagram in the statement uniformly coarsely commutes.

Suppose that Σ, Σ' are non-maximal simplices of $\text{Lk}(\Delta)$ such that $[\Sigma] \pitchfork [\Sigma']$ (resp. $[\Sigma] \sqsubset [\Sigma']$). Then $[\Sigma \star \Delta] \pitchfork [\Sigma' \star \Delta]$ (resp. $[\Sigma \star \Delta] \sqsubset [\Sigma' \star \Delta]$).

By definition, $\rho_{[\Sigma]}^{[\Sigma]} = \pi_{[\Sigma]}^\Delta(Y_{\Sigma'}^\Delta \cap \text{Sat}^\Delta(\Sigma))$ and $\rho_{[\Sigma' \star \Delta]}^{[\Sigma \star \Delta]} = \pi_{[\Sigma' \star \Delta]}(Y_{\Sigma' \star \Delta} \cap \text{Sat}(\Sigma \star \Delta))$.

If a is a vertex of $\text{Sat}^\Delta(\Sigma)$, then $a \in \Pi$, for some simplex Π of $\text{Lk}(\Delta)$ with $\text{Lk}(\Pi) \cap \text{Lk}(\Delta) = \text{Lk}(\Sigma) \cap \text{Lk}(\Delta)$. So, $\Pi \star \Delta$, and hence a , is contained in $\text{Sat}(\Sigma \star \Delta)$. Moreover, if $a \in Y_{\Sigma'}^\Delta$, then $a \in Y_{\Sigma' \star \Delta} \cap \text{Lk}(\Delta)$. Thus, $Y_{\Sigma'}^\Delta \cap \text{Sat}^\Delta(\Sigma)$ is contained in $Y_{\Sigma' \star \Delta} \cap \text{Sat}(\Sigma \star \Delta)$. Combining this with the above discussion relating p^Δ to p , we have that $\rho_{[\Sigma]}^{[\Sigma]}$ (computed in the combinatorial HHS structure on $\text{Lk}(\Delta)$) is coarsely contained in, and hence coarsely coincides with, $\rho_{[\Sigma' \star \Delta]}^{[\Sigma \star \Delta]}$, as required.

Now suppose that $[\Sigma] \sqsubset [\Sigma']$. The map $\rho_{[\Sigma]}^{[\Sigma']}$ is defined to be the restriction of $p^\Delta : Y_\Sigma^\Delta \rightarrow \mathcal{C}(\Sigma \star \Delta)$ to $\mathcal{C}(\Sigma' \star \Delta) \cap Y_\Sigma^\Delta$, and \emptyset otherwise. Meanwhile, $\rho_{[\Sigma \star \Delta]}^{[\Sigma' \star \Delta]}$ is the restriction of p to $\mathcal{C}(\Sigma' \star \Delta) \cap Y_{\Sigma \star \Delta}$ and \emptyset otherwise. So, these maps coarsely agree on $\mathcal{C}(\Sigma' \star \Delta) \cap Y_\Sigma^\Delta$. Since $\mathcal{C}(\Sigma' \star \Delta) \cap \text{Sat}^\Delta(\Sigma)$ is contained in $\mathcal{C}(\Sigma' \star \Delta) \cap \text{Sat}(\Sigma \star \Delta)$, we have that $\rho_{[\Sigma \star \Delta]}^{[\Sigma' \star \Delta]}(a) = \emptyset$ whenever $\rho_{[\Sigma]}^{[\Sigma']}$ is defined and equal to \emptyset . This completes the proof that the required diagrams coarsely commute. \square

5. W IS HIERARCHICALLY HYPERBOLIC

Now we prove Theorem 1.18. Throughout this section, (X, W, δ, n) is a combinatorial HHS.

5.1. Strong bounded geodesic image. The following is the key lemma.

Lemma 5.1 (Strong BGI). *For each δ, δ' , there exists C such that the following holds. Let Z be a δ -hyperbolic graph and let V be a nonempty subgraph of Z . Let L_V be an induced subgraph of Z disjoint from V with the property that $z \in L_V$ implies that z is adjacent to each vertex of V . Let Z_V be the induced subgraph of Z with vertex set $Z^{(0)} - V^{(0)}$, and suppose Z_V is δ' -hyperbolic.*

Suppose that L_V is δ -hyperbolic and (δ, δ) -quasi-isometrically embedded (hence quasiconvex with constant depending only on δ, δ') in Z_V , and let $\pi : Z_V \rightarrow L_V$ be the coarse closest-point projection.

Let $x, y \in Z^{(0)}$ and let γ be a geodesic in Z from x to y . Suppose that $\gamma \cap V = \emptyset$. Then $d_{L_V}(\pi(x), \pi(y)) \leq C$.

Proof. Let x, y, γ be as in the statement. Since $\gamma \cap V = \emptyset$, we have $\gamma \subset Z_V$, and γ is necessarily a geodesic of Z_V . Consider a geodesic quadrilateral formed by geodesics $[x, p], [p, q], [q, y]$ and

γ in Z_V , where $p \in \pi(x), q \in \pi(y)$. Since L_V is quasiconvex in Z_V , there exists a constant $K = K(\delta, \delta')$ such that $[p, q]$ lies in the K -neighbourhood of L_V , as measured in Z_V (and hence in Z).

The quadrilateral is $2\delta'$ -thin in Z_V . Let $x' \in [x, p]$ and suppose that x' is $2\delta'$ -close to some point of $[p, q]$. Then x' is $(2\delta' + K)$ -close to L_V , which is possible only if $d_{Z_V}(x', p) \leq 2(\delta' + K)$, for otherwise we contradict that $p \in \pi(x)$. Similarly, the $2\delta'$ -neighbourhood of $[y, q]$ only intersects $[p, q]$ at points within $2(\delta' + K)$ of q . But $[p, q]$ lies in the union of the $2\delta'$ -neighbourhoods of the other three sides of the quadrilateral. So, either $d_{Z_V}(p, q) \leq 5(\delta' + K)$, and hence $d_{L_V}(\pi(x), \pi(y)) \leq C$ for C uniformly bounded, and we are done, or: a non-empty subpath of $[p, q]$, consisting of all but initial and terminal length- $5(K + \delta')$ segments, lies $2\delta'$ -close to γ . Hence we can assume that there exist $a, b \in \gamma$ such that $d_{Z_V}(a, \pi(x)) \leq 10K\delta'$ and $d_{Z_V}(b, \pi(y)) \leq 10K\delta'$.

Hence, for any $v \in V$, we have $d_Z(a, v) \leq 10K\delta' + 1$ and $d_Z(b, v) \leq 10K\delta' + 1$. So, since γ is a geodesic of Z , the subpath of γ from a to b has length at most $20K\delta' + 2$, i.e., $d_{Z_V}(a, b) \leq 20K\delta' + 2$. Hence $d_{Z_V}(\pi(x), \pi(y)) \leq 40K\delta' + 2$. Since L_V is (δ, δ) -quasi-isometrically embedded in Z_V , this gives $d_{L_V}(\pi(x), \pi(y)) \leq 40K\delta\delta' + \delta^2 + 2\delta$, and we take the latter value to be C . \square

Now fix $[\Delta] \in \mathfrak{S}$. By assumption, $Z = X^{+W}$ is δ -hyperbolic. The subgraph $V = \text{Sat}(\Delta)$ has the property that $Z_V = Y_\Delta$. Moreover, each vertex of $L_V = \mathcal{C}([\Delta])$ is joined by an edge of Z to each vertex in V , and L_V is δ -hyperbolic and (δ, δ) -quasi-isometrically embedded in Z_V , because of Definition 1.8. Finally, Proposition 3.3 ensures that Z_V is δ' -hyperbolic.

Hence, by Lemma 5.1, we have:

Lemma 5.2. *There exists $C = C(\delta, \delta')$ so that the following holds. Let $[\Delta] \in \mathfrak{S}$ and let $x, y \in Y_\Delta$. If $d_{\mathcal{C}([\Delta])}(x, y) > C$, then any geodesic γ in X^{+W} from x to y intersects $\text{Sat}(\Delta)$.*

The preceding lemma will be used repeatedly in the next subsection.

5.2. Proof of Theorem 1.18: verification of the HHS axioms. Before proceeding to the proof of Theorem 1.18, we remind the reader of the roles of the different spaces from Definition 1.8:

Remark 5.3 (Connectedness, and the roles of the different spaces). Recall that, given a combinatorial HHS (X, W) , Definition 1.8 does not assert that either X or W is connected. In the case of X , the reader should bear in mind that this is by design: the objects about which we are making geometric claims and/or assumptions are W and the various Y_Δ and $\mathcal{C}(\Delta)$, including X^{+W} .

The complex X does not function as a metric object, but rather as a “database of links” recording the index set (and nesting, orthogonality, and transversality relations) for an HHS structure on W . We never refer to metric properties of X , only combinatorial properties.

The spaces $\mathcal{C}(\Delta)$ are connected by assumption: Definition 1.8.(2) asserts that these are hyperbolic, and in particular they are geodesic metric spaces (not extended metric spaces).

What we will have to prove is connectedness of W . Indeed, our goal is to prove that W — a graph with the usual graph metric — is an HHS, and, in particular, by the requirement in Definition 2.1 that HHSes are quasigeodesic metric spaces, this means proving that W is connected. The proof of Theorem 1.18 is by induction on the complexity of the combinatorial HHS (X, W) . In the base cases, we will show that W is either a single point, or $W = X^{+W} = \mathcal{C}(\emptyset)$, which is connected by assumption.

In the inductive step, we will assume that, for each non-empty, non-maximal simplex $(\text{Lk}(\Delta), W^\Delta)$, we have an HHS structure $(W^\Delta, \mathfrak{S}_\Delta)$. Built into this is the assumption that W^Δ is connected since, by Definition 2.1, a graph whose graph-metric admits an HHS structure is a geodesic metric space, i.e., it is connected. We will verify, as part of the proof that (X, \mathfrak{S}) satisfies the

“Uniqueness” axiom (Definition 2.1.(9)), that any two vertices of W are at finite distance in the graph metric, which is to say that W is connected. We will do this by building paths.

We are now ready for:

Proof of Theorem 1.18. Fix (X, W, δ, n) as in the statement.

For “Large Links” and “Uniqueness” we will have to proceed by induction on complexity, and assume that Theorem 1.18 holds for all the links of nonempty non-maximal simplices of the simplicial complex X .

Base cases: First we explain the base cases $n \leq 1$. First note that \emptyset is always a simplex of X . So the complexity $n = 0$ only if \emptyset is a maximal simplex, which means that $X = \emptyset$. By Definition 1.2, this implies that W is a single point, and (W, \emptyset) is a hierarchically hyperbolic space in this case, as required.

When the complexity $n = 1$, the only non-maximal simplex is the empty set, so X is a discrete set of vertices (0-simplices have to be maximal). So, by Definition 1.2, $W = X^{+W}$. On the other hand, by the definition of a combinatorial HHS (Definition 1.8.(2)) X^{+W} is hyperbolic, which is to say that W is hyperbolic. In particular, $(W, \{\emptyset\})$ is an HHS, where the projection from W to $\mathcal{C}(\emptyset) = X^{+W}$ is the identity.

Inductive step: We refer the reader to Definition 2.1. The proof will consist of verifying that (W, \mathfrak{S}) satisfies each requirement in that definition.

Inductive hypothesis: By Proposition 4.9, there is a constant $\delta'' = \delta''(\delta, n)$ such that $(\text{Lk}(\Delta), W^\Delta, \delta'', m)$ is a combinatorial HHS whenever Δ be a nonempty non-maximal simplex of X , where $m < n$ is the complexity of $\text{Lk}(\Delta)$. Our inductive hypothesis is therefore that for each such Δ , the pair $(W^\Delta, \mathfrak{S}_\Delta)$ is an HHS, with the constants in Definition 2.1 depending only on δ'' and m (and hence bounded in terms of δ and n). Moreover, we assume that the projections π_\bullet and ρ_\bullet are as given in Definition 1.16, with $\text{Lk}(\Delta)$ playing the role of X , and W^Δ playing the role of W , in that definition. We assume that the relations are as in Definition 1.11, with $\text{Lk}(\Delta)$ playing the role of X .

Verifying the HHS axioms: We now proceed to verify the axioms from Definition 2.1.

First, the underlying space, W , is a geodesic extended metric space, since it is a graph. Definition 1.8 requires W to be a quasigeodesic space, which in our situation means that we have to show that W is connected. This is done below, when we verify the uniqueness axiom (Definition 2.1.(9)). We will not require connectedness of W for the earlier parts of the proof. (The simplest argument that we are aware of for connectedness requires less precision, but it otherwise follows the same outline.)

Second, let \mathfrak{S} be the associated index set, with relations described in Definition 1.11; for each $[\Delta] \in \mathfrak{S}$, the associated δ -hyperbolic space is $\mathcal{C}([\Delta])$, and the various projections are as described in Definition 1.16.

We now verify the enumerated axioms from Definition 2.1:

Projections: The projections $\pi_{[\Delta]}$ are coarsely Lipschitz coarse maps because that is true of coarse closest point projections to quasiconvex subsets of hyperbolic spaces, if $\mathcal{C}([\Delta])$ has diameter at least δ (hyperbolicity of Y_Δ holds by Proposition 3.3 in this case, and $\mathcal{C}([\Delta])$ is quasi-isometrically embedded in Y_Δ by assumption). If $\text{diam}(\mathcal{C}([\Delta])) \leq \delta$, then the claim is obvious.

Let v be a vertex of $\mathcal{C}(\Delta)$. Then v is contained in some maximal simplex w of X , so $\pi_{[\Delta]}(w)$ contains v . So, $\mathcal{C}(\Delta) = \cup_w \pi_{[\Delta]}(w)$, and in particular $\pi_{[\Delta]}$ has quasiconvex image.

Nesting: The relation \sqsubseteq is defined in Definition 1.11, and it is clearly a partial order. The maximal element is the equivalence class of \emptyset . The bounded sets and coarse maps ρ_\bullet are defined in Definition 1.16. If $[\Delta'] \sqsubset [\Delta]$, then by definition $\text{Lk}(\Delta') \subsetneq \text{Lk}(\Delta) \subseteq Y_\Delta$. Notice that some vertex v of Δ' is contained in Y_Δ . Indeed, if we had $\Delta' \subseteq \text{Sat}(\Delta)$, then we would have $\text{Lk}(\Delta) = \text{Lk}(\text{Sat}(\Delta)) \subseteq \text{Lk}(\Delta')$ (see Remark 1.4), contradicting $\text{Lk}(\Delta') \subsetneq \text{Lk}(\Delta)$. Since

$v \in \text{Sat}(\Delta')$, $\rho_{[\Delta]}^{[\Delta']}$ is non-empty. Moreover, $\text{Sat}(\Delta')$ has diameter at most 2 in Y_Δ because any vertex of $\text{Sat}(\Delta')$ is connected to any vertex of $\text{Lk}(\Delta') \subseteq \text{Lk}(\Delta) \subseteq Y_\Delta$ (notice that $\text{Lk}(\Delta')$ is non-empty because, by definition, Δ' is non-maximal). Hence $\rho_{[\Delta]}^{[\Delta']}$ $\subset \mathcal{C}([\Delta])$ is uniformly bounded.

Orthogonality: We defined \perp in Definition 1.16. It is symmetric because if $\text{Lk}(\Delta') \subseteq \text{Lk}(\text{Lk}(\Delta))$ then $\text{Lk}(\text{Lk}(\Delta')) \supseteq \text{Lk}(\text{Lk}(\text{Lk}(\Delta)))$, and then:

Claim 5.4. $\text{Lk}(\Delta) = \text{Lk}(\text{Lk}(\text{Lk}(\Delta)))$.

Proof of Claim 5.4. For any subcomplex Σ , chasing the definitions we see that $\Sigma \subseteq \text{Lk}(\text{Lk}(\Sigma))$. For $\Sigma = \text{Lk}(\Delta)$, we get the inclusion “ \subseteq ”. For $\Sigma = \Delta$, and applying Lk to both sides, we get the inclusion “ \supseteq ”. \blacksquare

Anti-reflexivity of \perp follows from the fact that $\text{Lk}(\Delta)$ is always disjoint from $\text{Lk}(\text{Lk}(\Delta))$, and $\text{Lk}(\Delta)$ is non-empty for any non-maximal simplex (recall that we are excluding maximal simplices).

Next, suppose that $[\Delta] \sqsubseteq [\Delta']$ and $[\Delta'] \perp [\Delta'']$. Then by the definitions, $\text{Lk}(\Delta) \subseteq \text{Lk}(\Delta') \subseteq \text{Lk}(\text{Lk}(\Delta''))$, so $[\Delta] \perp [\Delta'']$. Also, if $[\Delta] \perp [\Delta']$, then $\text{Lk}(\Delta)$ (which is non-empty) is contained in $\text{Lk}(\text{Lk}(\Delta'))$, which is disjoint from $\text{Lk}(\Delta')$, so that $[\Delta] \not\sqsubseteq [\Delta']$.

Now fix $[\Delta]$ and $[\Delta_1] \sqsubseteq [\Delta]$, and suppose that $A = \{[\Delta_2] : [\Delta_2] \sqsubseteq [\Delta], [\Delta_2] \perp [\Delta_1]\}$ is non-empty. We need to find $[\Delta']$ with $[\Delta'] \not\sqsubseteq [\Delta]$ and $[\Delta_2] \sqsubseteq [\Delta']$ for all $[\Delta_2] \in A$.

By definition, $[\Delta_2] \in A$ has $\text{Lk}(\Delta_2) \subseteq \text{Lk}(\Delta)$ and $\text{Lk}(\Delta_2) \subseteq \text{Lk}(\text{Lk}(\Delta_1))$. Consider $B = \bigcup_{[\Delta_2] \in A} \text{Lk}(\Delta_2)$; we are looking for a simplex strictly containing Δ whose link contains B . We have $\text{Lk}(B) = \bigcap_{[\Delta_2] \in A} \text{Lk}(\text{Lk}(\Delta_2))$, and also $\text{Lk}(\Delta_1) \subseteq \text{Lk}(\text{Lk}(\Delta_2))$ for each $[\Delta_2] \in A$, so that any vertex v in $\text{Lk}(\Delta_1)$ (which exists) has:

- $v \in \text{Lk}(\Delta)$, because $\text{Lk}(\Delta_1) \subseteq \text{Lk}(\Delta)$, since $[\Delta_1] \sqsubseteq [\Delta]$. In particular, the simplex $\Delta' = v \star \Delta$ is well-defined.
- $v \in \text{Lk}(B)$, so that $B \subseteq \text{Lk}(v)$. So, for any $[\Delta_2] \in A$, $\text{Lk}(\Delta_2) \subseteq \text{Lk}(\Delta') = \text{Lk}(v) \cap \text{Lk}(\Delta)$ (since $\text{Lk}(\Delta_2) \subseteq \text{Lk}(\Delta)$), i.e., $[\Delta_2] \sqsubseteq [\Delta']$.

Since $\Delta \not\sqsubseteq \Delta'$, we have $[\Delta'] \not\sqsubseteq [\Delta]$ (by Remark 1.13). Notice that the nesting is indeed proper because $v \in \text{Lk}(\Delta) - \text{Lk}(\Delta')$.

Transversality and consistency: Recall that $\rho_{[\Delta]}^{[\Delta']}$ is defined in Definition 1.16. If $[\Delta] \pitchfork [\Delta']$ then:

- $\text{Sat}(\Delta') \cap Y_\Delta \neq \emptyset$, for otherwise we would have $\text{Sat}(\Delta') \subseteq \text{Sat}(\Delta)$, and hence $\text{Lk}(\Delta) = \text{Lk}(\text{Sat}(\Delta)) \subseteq \text{Lk}(\text{Sat}(\Delta')) = \text{Lk}(\Delta')$ and $[\Delta] \sqsubseteq [\Delta']$ (we used Remark 1.4).
- $\text{Lk}(\Delta') \cap Y_\Delta \neq \emptyset$, for otherwise we would have $\text{Lk}(\Delta')^{(0)} \subseteq \text{Sat}(\Delta)$, and $[\Delta] \perp [\Delta']$ since $\text{Sat}(\Delta) \subseteq \text{Lk}(\text{Lk}(\Delta))$ by Remark 1.4. So, $\text{Sat}(\Delta') \cap Y_\Delta$ has diameter at most 2 in Y_Δ , and hence $\rho_{[\Delta]}^{[\Delta']}$ is uniformly bounded.

Now, suppose that $d_{[\Delta]}(w, \rho_{[\Delta]}^{[\Delta']}) \geq C$, for some vertex w of W (corresponding to a maximal simplex Σ_w of X), where C is as in Lemma 5.2. Then there is a geodesic γ in X^{+W} from $\text{Sat}(\Delta') - \text{Sat}(\Delta)$ to w' intersecting $\text{Sat}(\Delta)$, where $w' \in \Sigma_w - \text{Sat}(\Delta)$. The minimal subgeodesic of γ from w' to $\text{Sat}(\Delta)$ does not intersect $\text{Sat}(\Delta')$. By Lemma 5.2 we must have $d_{[\Delta]}(w, \rho_{[\Delta]}^{[\Delta']}) < C$.

Consistency for nesting: Now suppose that $[\Delta'] \not\sqsubseteq [\Delta]$. Let $w \in W$.

Choose a vertex $w' \in X$ as follows. Suppose $w \in W^{(0)}$ corresponds to a maximal simplex Σ_w . If every vertex of Σ_w is in $\text{Sat}(\Delta) \cup \text{Sat}(\Delta')$, then since $\text{Lk}(\Delta') \subset \text{Lk}(\Delta)$, X would contain $\Sigma_w \star \text{Lk}(\Delta')$, contradicting maximality. So w' can be chosen in $\Sigma_w - (\text{Sat}(\Delta) \cup \text{Sat}(\Delta'))$. So, $w' \in Y_\Delta$ and $w' \in Y_{\Delta'}$.

Let $p_\Delta : Y_\Delta \rightarrow \mathcal{C}([\Delta])$ and $p_{\Delta'} : Y_{\Delta'} \rightarrow \mathcal{C}([\Delta'])$ be closest-point projections.

If $\mathcal{C}([\Delta'])$ has diameter at most δ , then we are done, so assume that $\mathcal{C}([\Delta'])$ has diameter more than δ . Then by Lemma 1.14, we have $[\Delta'] = [\Pi \star \Delta]$ for some simplex Π of $\text{Lk}(\Delta)$.

Let α be a geodesic in the δ' -hyperbolic space Y_Δ joining w' to $p_\Delta(w') \in \text{Lk}(\Delta)$.

First suppose that $\alpha \cap \text{Sat}(\Delta') = \emptyset$. In this case we use the following lemma:

Lemma 5.5. *Let Σ be a simplex of $\text{Lk}(\Delta)$. Then $Y_{\Sigma \star \Delta} \subset Y_\Sigma \cap Y_\Delta$.*

Proof. Suppose that $x \in \text{Sat}(\Delta)$, i.e., there exists a non-maximal simplex Π of X such that $\text{Lk}(\Pi) = \text{Lk}(\Delta)$ and $x \in \Pi$. Since Σ is a simplex of $\text{Lk}(\Delta)$, there is a simplex $\Sigma \star \Pi$ in X , and $\text{Lk}(\Sigma \star \Pi) = \text{Lk}(\Sigma) \cap \text{Lk}(\Pi) = \text{Lk}(\Sigma) \cap \text{Lk}(\Delta) = \text{Lk}(\Sigma \star \Delta)$. So $x \in \Pi \star \Sigma \sim_X \Delta \star \Sigma$, i.e., $x \in \text{Sat}(\Sigma \star \Delta)$, as required. Hence $Y_{\Sigma \star \Delta} \subset Y_\Delta$. Similarly, $Y_{\Sigma \star \Delta} \subset Y_\Sigma$. \square

By Lemma 5.5, $Y_{\Delta'} \subset Y_\Delta$. Since α is a geodesic of Y_Δ and is entirely contained in $Y_{\Delta'}$, we have that α is also a geodesic of $Y_{\Delta'}$. Now, $Y_{\Delta'}$ is δ' -hyperbolic, and is obtained from Y_Δ by deleting $\text{Sat}(\Delta') - \text{Sat}(\Delta)$ (and any edges with at least one endpoint in $\text{Sat}(\Delta')$). Now, $\mathcal{C}([\Delta'])$ is a subgraph of $Y_{\Delta'}$, each of whose vertices is adjacent in Y_Δ to each vertex of $\text{Sat}(\Delta') - \text{Sat}(\Delta)$. Moreover, $\mathcal{C}([\Delta'])$ is δ -hyperbolic and (δ, δ) -quasi-isometrically embedded in $Y_{\Delta'}$. Hence, by Lemma 5.1, w' and $p(w')$ have C -close $p_{\Delta'}$ -images on $\mathcal{C}([\Delta'])$, which is to say that $\text{diam}(\pi_{[\Delta']}(w) \cup \rho_{[\Delta']}^{\Delta}(\pi_{[\Delta]}(w)))$ is uniformly bounded.

Next, suppose that α passes through some $v \in \text{Sat}(\Delta') - \text{Sat}(\Delta)$. Since α is a geodesic from w' to $p_\Delta(w')$, we have from the definition of p_Δ that

$$d_{Y_\Delta}(w', p_\Delta(w')) = d_{Y_\Delta}(w', v) + d_{Y_\Delta}(v, p_\Delta(w')) \leq d_{Y_\Delta}(w, \text{Lk}(\Delta)) + 1.$$

Since $v \in \text{Sat}(\Delta')$, for any $\ell \in \text{Lk}(\Delta') \subset \text{Lk}(\Delta)$, we have that v is adjacent to ℓ . Fixing such an ℓ (which exists because Δ' is non-maximal), we have $d_{Y_\Delta}(w', \ell) \leq d_{Y_\Delta}(w', v) + 1$. Hence

$$d_{Y_\Delta}(w', \text{Lk}(\Delta)) \leq d_{Y_\Delta}(w', \ell) \leq d_{Y_\Delta}(w', v) + 1 \leq d_{Y_\Delta}(w', \text{Lk}(\Delta)) - d_{Y_\Delta}(v, p_\Delta(w')) + 2.$$

This implies that $d_{Y_\Delta}(v, p_\Delta(w')) \leq 2$, which provides a uniform bound, in terms of δ , on $d_{Y_\Delta}(p_\Delta(w'), p_\Delta(v))$. This in turn gives a uniform bound on $d_{[\Delta]}(p_\Delta(w'), p_\Delta(v))$, in view of Definition 1.8.(2). Finally, $p_\Delta(w') \in \pi_{[\Delta]}(w)$ and $v \in \text{Sat}(\Delta') \cap Y_\Delta$, so by Definition 1.16, we have shown that $d_{[\Delta]}(\pi_{[\Delta]}(w), \rho_{[\Delta]}^{\Delta'})$ is uniformly bounded, as required.

Finally, we need to check the following. Suppose that $[\Delta'] \sqsubset [\Delta]$, and $[\Delta] \sqsubset [\Sigma]$ or $[\Delta] \pitchfork [\Sigma]$, and $[\Sigma] \not\sqsubset [\Delta']$. We claim that $d_{[\Sigma]}(\rho_{[\Sigma]}^{\Delta}, \rho_{[\Sigma]}^{\Delta'})$ is uniformly bounded, as required by item 4 of the definition of an HHS (Definition 2.1). By definition, $\rho_{[\Sigma]}^{\Delta} = p_\Sigma(\text{Sat}(\Delta) \cap Y_\Sigma)$ and $\rho_{[\Sigma]}^{\Delta'} = p_\Sigma(\text{Sat}(\Delta') \cap Y_\Sigma)$. Since $[\Delta'] \pitchfork [\Sigma]$ or $[\Delta'] \sqsubset [\Sigma]$, there exists $v \in \text{Lk}(\Delta') \cap Y_\Sigma$. Now, v is adjacent in Y_Σ to each vertex in $\text{Sat}(\Delta') \cap Y_\Sigma$, and to each vertex in $\text{Sat}(\Delta) \cap Y_\Sigma$. So, $d_{Y_\Sigma}(\text{Sat}(\Delta), \text{Sat}(\Delta')) \leq 2$, so, since p_Σ is uniformly coarsely Lipschitz, $d_{[\Sigma]}(\rho_{[\Sigma]}^{\Delta}, \rho_{[\Sigma]}^{\Delta'})$ is uniformly bounded.

Finite complexity: This follows from Definition 1.8.(1).

Bounded geodesic image: Let $[\Delta] \sqsubset [\Delta']$. By definition of nesting, $\mathcal{C}([\Delta])$ is (properly) contained in $\mathcal{C}([\Delta'])$. Let E be a constant to be determined, and suppose that γ is a geodesic in $\mathcal{C}([\Delta'])$ that is disjoint from the E -neighborhood of $\mathcal{C}([\Delta])$. If γ contains a vertex v of $\text{Sat}(\Delta)$, then γ is joined by an edge of X , and hence of X^{+W} , to some vertex $w \in \text{Lk}(\Delta) \subset \text{Lk}(\Delta')$. Hence γ passes through the 1-neighborhood in $\mathcal{C}([\Delta'])$ of $\text{Lk}(\Delta)$; by choosing $E > 1$, this is impossible. Hence γ can be regarded as a geodesic in Y_Δ which is far from the quasiconvex subset $\mathcal{C}([\Delta])$ of the δ' -hyperbolic space Y_Δ . Hence, provided E is chosen sufficiently large in terms of δ' and δ , the projections of the endpoints of γ to $\mathcal{C}([\Delta])$ are E -close to each other, as required.

Large links: Let $[\Delta] \in \mathfrak{S}$. Let $x, y \in W$. We need to produce a constant E , depending on δ and n but independent of $[\Delta]$, such that there exist $[\Sigma_1], \dots, [\Sigma_N] \sqsubset [\Delta]$ with the property

that any $[\Sigma'] \sqsubset [\Delta]$ with $d_{[\Sigma']}(x, y) > E$ satisfies $[\Sigma'] \sqsubseteq [\Sigma_i]$ for some i . Moreover, we need to bound N by a uniform linear function of $d_{[\Delta]}(x, y)$, and also bound $d_{[\Delta]}(x, \rho_{\Delta}^{[\Sigma_i]})$ above by N .

First, suppose that Δ is nonempty. Then by our induction hypothesis, $(W^\Delta, \mathfrak{S}^\Delta)$ is an HHS, and the projections for W^Δ coarsely coincide with those for (X, W) , as stated in Proposition 4.11.

The coordinates $(\pi_Y(x))_{Y \in \mathfrak{S}^\Delta}$ are consistent (with uniform constants), as we checked above (“Transversality and consistency” and “Consistency for nesting”). Hence, using our induction hypothesis, we can apply the realisation theorem for hierarchically hyperbolic spaces (Theorem 2.7) to obtain a vertex $x' \in W^\Delta$ whose projections coarsely coincide with $\pi_Y(x)$ for all $Y \in \mathfrak{S}^\Delta$. We can similarly construct $y' \in W^\Delta$ starting from the coordinates of y .

In view of Lemma 1.14, every $[\Sigma]$ so that $[\Sigma] \sqsubseteq [\Delta]$ and $\mathcal{C}(\Sigma)$ has diameter at least δ is in the image of the map ι^* from Lemma 4.6. Hence, it is easily seen that Large Links for $[\Delta], x, y$ follows from Large Links for $[\emptyset], x', y'$ in $(W^\Delta, \mathfrak{S}^\Delta)$ (up to increasing the constants).

Now we handle the case where $\Delta = \emptyset$. Fix a geodesic α in X^+W from x' to y' , where x', y' are vertices of x, y . For every non-maximal simplex $[\Sigma]$, by Lemma 5.2 we have that either $d_{[\Sigma]}(x, y)$ is uniformly bounded by some E , or α contains a vertex v in $\text{Sat}(\Sigma)$. Notice that $v \in \text{Sat}(\Sigma)$ implies $[\Sigma] \sqsubseteq [v]$. Hence it suffices to let $\{\Sigma_i\}$ be the collection of all vertices of α .

Partial realization: Let $[\Delta_1], \dots, [\Delta_k] \in \mathfrak{S}$ be pairwise-orthogonal, i.e., $\text{Lk}(\Delta_i) \subset \text{Lk}(\text{Lk}(\Delta_j))$ for all $i \neq j$. For each i , let $p_i \in \text{Lk}(\Delta_i)$. Since $p_i \in \text{Lk}(\text{Lk}(\Delta_j))$, we have that p_1, \dots, p_k are pairwise-adjacent vertices of X . Hence there exists a maximal simplex Π of X containing p_1, \dots, p_k . Since Π is a maximal simplex, it corresponds to a vertex w of W . For each i , we have $\Pi \cap \mathcal{C}([\Delta_i]) \ni p_i$, so $p_i \in \pi_{[\Delta_i]}(w)$.

Next, suppose that $[\Sigma] \in \mathfrak{S}$ satisfies $[\Delta_i] \pitchfork [\Sigma]$ or $[\Delta_i] \sqsubset [\Sigma]$ for some i . Let $p : Y_\Sigma \rightarrow \mathcal{C}([\Sigma])$ be the coarse closest-point projection, so that $\rho_{[\Sigma]}^{[\Delta_i]}$ is by definition $p(\text{Sat}(\Delta_i) \cap Y_\Sigma)$. Now, if $[\Delta_i] \sqsubset [\Sigma]$, then $\text{Lk}(\Delta_i) \sqsubset \text{Lk}(\Sigma)$, so $p_i \in \text{Lk}(\Sigma)$, whence $p_i \in Y_\Sigma$. Hence $p(p_i) \subset \pi_{[\Sigma]}(w)$. On the other hand, there exists $v \in \text{Sat}(\Delta_i)$ such that $v \notin \text{Sat}(\Sigma)$, so $p(v) \subset \rho_{[\Sigma]}^{[\Delta_i]}$. Since p_i, v are adjacent in Y_Σ , we have that $p(p_i)$, and hence $\pi_{[\Sigma]}(w)$, lies uniformly close to $\rho_{[\Sigma]}^{[\Delta_i]}$. When $[\Delta_i] \pitchfork [\Sigma]$, we again have some $v \in \text{Sat}(\Delta_i) \cap Y_\Sigma$, so v is necessarily adjacent to p_i . Hence v lies at distance at most 2 in Y_Σ from every vertex of $\Pi \cap Y_\Sigma$. Since $\Pi \cap Y_\Sigma$ contains at least one vertex, $d_{[\Sigma]}(p(v), p(\Pi \cap Y_\Sigma))$ is uniformly bounded, i.e., $d_{[\Sigma]}(\rho_{[\Sigma]}^{[\Delta_i]}, \pi_{[\Sigma]}(w))$ is uniformly bounded, as required.

Uniqueness axiom and connectedness of W : In this part of the proof, we again use our induction assumption: links of nonempty non-maximal simplices carry a natural combinatorial HHS structure of strictly lower complexity, as in Proposition 4.9, which is furthermore compatible with the structure for X as explained in Proposition 4.11. These combinatorial HHS structures have their associated constants (from Definition 1.8) bounded in terms of δ and n , and in particular, independently of the link in question.

Moreover, we are assuming by induction on complexity that Theorem 1.18 holds for all links of nonempty non-maximal simplices of X , i.e., for each nonempty non-maximal $[\Delta]$, we have an HHS structure $(W^\Delta, \mathfrak{S}_\Delta)$ where all of the constants from Definition 2.1 depend only on the above combinatorial HHS constants, and are hence bounded uniformly in terms of δ, n . In particular, the function θ_u from Definition 2.1.(9), can be taken to be the same for all HHS structures $(W^\Delta, \mathfrak{S}_\Delta)$ as Δ varies over nonempty non-maximal simplices of X .

Let x, y be vertices in W , i.e., maximal simplices in X . To prove connectedness amounts to proving that $d_W(x, y)$ is finite. To prove uniqueness requires a strict strengthening of this, namely that $d_W(x, y)$ is bounded above by a fixed function of $\sup_{[\Delta] \in \mathfrak{S}} d_{[\Delta]}(x, y)$. To do this, we will construct a path in W from x to y , and moreover, bound the length of this path in terms of $\sup_{[\Delta] \in \mathfrak{S}} d_{[\Delta]}(x, y)$.

First:

Claim 5.6. *There exists a uniform constant E such that the set of $[\Delta] \in \mathfrak{S}$ such that $d_{[\Delta]}(x, y) > E$ is finite. Hence*

$$\kappa = \kappa(x, y) = \sup_{[\Delta] \in \mathfrak{S}} d_{[\Delta]}(x, y) < \infty.$$

Proof of Claim 5.6. By the large link axiom, which we have just verified above, there exists a constant E such that the following holds. Let α be a geodesic in X^{+W} from x' to y' , where $x' \in x, y' \in y$. Then any $[\Delta] \in \mathfrak{S} - \{[\emptyset]\}$ for which $d_{[\Delta]}(x, y) > E$ satisfies $[\Delta] \sqsubseteq [v]$, where v is one of the $d_{[\emptyset]}(x', y') + 1$ vertices of α .

By our induction hypothesis, for each such v , the pair (W^v, \mathfrak{S}_v) is an HHS (with uniform constants and uniqueness functions), so by applying Theorem 2.7 (realisation theorem) exactly as in the verification of the large link axiom (“Large links”), we obtain $x_v, y_v \in W^v$ whose projections to each $\mathcal{C}([\Delta]), [\Delta] \sqsubseteq [v]$, uniformly coarsely coincide with $\pi_{[\Delta]}(x)$ and $\pi_{[\Delta]}(y)$ respectively. Now, up to uniformly enlarging E , it follows, by, for example, an application of the distance formula (Theorem 2.8) in the HHS (W^v, \mathfrak{S}_v) , that there are finitely many $[\Delta] \sqsubseteq [v]$ with $d_{[\Delta]}(x, y) > E$. Since there are finitely many choices for v , this proves the first assertion, from which the second assertion follows immediately. \blacksquare

Hence, to simultaneously prove connectedness of W and the uniqueness axiom, it suffices to prove that, for any $\kappa \geq 0$ and any vertices $x, y \in W$ with $d_{[\Delta]}(x, y) \leq \kappa$ for all $[\Delta] \in \mathfrak{S}$, there is a path in W that joins x to y and has length bounded in terms of κ, δ, n .

Therefore, fix $\kappa \geq 0$. Let x, y be vertices of W (that is, maximal simplices of X) so that $d_{[\Delta]}(x, y) \leq \kappa$ for all $[\Delta] \in \mathfrak{S}$. We need θ , depending on κ, δ, n only, so that x and y can be joined by a path in W of length at most θ , which is to say that $d_W(x, y) \leq \theta$ and x, y lie in the same path-component of W .

To achieve this, we will prove by induction on k that for every κ there exists $\theta(\kappa, \delta, n, k)$ so that whenever $d_{X+W}(x, y) \leq k$ and $d_{[\Delta]}(x, y) \leq \kappa$ for all $[\Delta] \in \mathfrak{S} - \{[\emptyset]\}$, the vertices x and y can be joined by a path in W of length at most θ , and in particular, $d_W(x, y) \leq \theta$.

Note that we are inducting on distance in $\mathcal{C}(\emptyset) = X^{+W}$, which is connected by Definition 1.8.(2).

Base case $k = 0$: First suppose that $k = d_{X+W}(x, y) = 0$. Then x and y share a vertex v , and they are of the form $x' \star v, y' \star v$, for some simplices x', y' in the link of v , and necessarily x', y' are maximal simplices in $\text{Lk}(v)$. Moreover,

$$d_{\mathcal{C}^v([\Delta])}(x', y') \leq \kappa'(\kappa, \delta, n)$$

for any non-maximal simplex Δ of $\text{Lk}(v)$, by Proposition 4.11.

By induction on complexity, in W^v there is path $x' = x'_1, \dots, x'_l = y'$ of maximal simplices of $\text{Lk}(v)$ with $l \leq \theta(\kappa, \delta, n)$. Hence, in W we have a path $x = x'_1 \star v, \dots, x'_l \star v = y$ (notice that these are indeed maximal simplices of X), showing $d_W(x, y) \leq l$.

Inductive step: Now suppose that $k = d_{X+W}(x, y) > 0$, and consider a geodesic γ of the δ -hyperbolic space X^{+W} such that γ is of minimal length among geodesics connecting a vertex u of x to a vertex of y . (Recall that hyperbolicity, and in particular the property of being a geodesic space, holds for X^{+W} by Definition 1.8.(2).) We consider two cases:

Case 1. Suppose that one of the following holds:

- (1) there is a vertex t on $\gamma - \{u\}$ that lies in $\text{Sat}(\sigma)$ for some non-maximal simplex σ with $[\sigma] \sqsubseteq [u]$ and $\text{diam}(\mathcal{C}(\sigma)) \geq \delta$.
- (2) u is connected in X to the second vertex t of γ .

Consider the closest t to y satisfying either condition.

Any σ as above has the property that $[\sigma] \sqsubseteq [t]$ since $t \in \text{Sat}(\sigma)$. Hence, in view of Definition 1.8.(3), there exists a simplex τ in the link of u so that $[\tau \star u] \sqsubseteq [t]$ and any $[\omega]$ with $\text{diam}(\mathcal{C}(\omega)) \geq \delta$ nested into both $[u]$ and $[t]$ is nested into $[\tau \star u]$. For later purposes, we can and will pick $\tau = t$ if t is the second vertex of γ and u, t are connected by an edge of X . Set $\tau' = \tau \star u$.

Before proceeding, observe that in either of the two itemized situations, τ' is a non-maximal simplex. Indeed, in the first case, where τ' comes from Definition 1.8.(3), this is because of Remark 1.9. In the second case, where $\tau' = t \star u$, we can assume that τ' is non-maximal by the following argument. If τ' is maximal, then $d_{X+W}(\tau', y) \leq k-1$, so by the inductive hypothesis, $d_W(\tau', y) \leq \theta(\kappa, \delta, n, k-1)$, and in particular y, τ' lie in the same component of W .

Also, $d_{X+W}(\tau', x) = 0$, so $d_W(x, \tau') \leq \theta(\kappa, \delta, n, 0)$, and in particular x, τ' are in the same component of W .

Transitivity of the binary relation “are connected by a path in W ” on the vertex set of W shows that x, y are in the same component of W , so $d_W(x, y)$ is finite and the triangle inequality bounds $d_W(x, y)$, as required. So, for the remainder of the argument, we can and shall assume that τ' is non-maximal.

Claim 5.7. *There exists $C = C(\delta, n, \kappa)$ so that there is a maximal simplex z'' in $\text{Lk}(\tau')$ with the property that, for $z = z'' \star \tau'$, we have $d_{\mathcal{C}([\Delta])}(z, y) \leq C$ for all non-maximal, non-empty simplices Δ of X .*

Proof. Let us consider $W^{\tau'}$, which is an HHS by Proposition 4.9 and induction on n . Moreover, the projections to the various hyperbolic spaces for $W^{\tau'}$ can be computed from those for W , by Proposition 4.11.

We can now apply the realization theorem (Theorem 2.7) to the coordinates of y (which we verified to be consistent when we checked consistency) and find a maximal simplex z'' in $\text{Lk}(\tau')$ with the property that $d_{\mathcal{C}([\Delta])}(z'' \star \tau', y) \leq C$ for all $[\Delta]$ nested into $[\tau']$. Consider now some $[\Delta]$ not nested into $[\tau']$, and $\Delta \neq \emptyset$. Moreover, we can assume $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$. There are a few cases.

If $[\Delta]$ is not nested into $[u]$, then u is not in $\text{Sat}(\Delta)$, and the projections of y and z to $\mathcal{C}([\Delta])$ coarsely coincide with the projection of u , as required.

Suppose that $[\Delta]$ is nested into $[u]$. Let γ' be the final subgeodesic of γ starting at t . If γ' does not intersect $\text{Sat}(\Delta)$, then the projections of y and z to $\mathcal{C}([\Delta])$ coarsely coincide and we are done. Hence, suppose that γ' intersects $\text{Sat}(\Delta)$. Since we chose t as close as possible to y , the intersection must consist of t only. Hence, t lies in $\text{Sat}(\Delta)$, which implies that $[\Delta]$ is nested into $[t]$. Since $[\Delta]$ is nested into both $[u]$ and $[t]$, it is nested into τ' , and in this case the bound C exists by construction as stated above. \blacksquare

Let z' be a (necessarily) maximal simplex in $\text{Lk}(u)$ so that $z'' \star \tau' = z' \star u$ (so that $z = z' \star u$), and write $x = x' \star u$.

Now consider the HHS (W^u, \mathfrak{S}_u) , which exists by induction. By assumption, for all non-maximal simplices Δ , we have $d_{[\Delta]}(x, y) \leq \kappa$, and by Claim 5.7, $d_{[\Delta]}(y, z) \leq C$, from which the triangle inequality gives $d_{[\Delta]}(x, z) \leq \kappa + C$. Hence $d_{\mathcal{C}^u(\Delta)}(x', z') \leq \kappa + C$ whenever $[\Delta] \sqsubseteq [u]$. So, by the uniqueness axiom in the HHS (W^u, \mathfrak{S}_u) , the distance between x' and z' in W^u can be bounded in terms of δ, n, κ .

In other words, in W^u there is a path $x' = x'_1, \dots, x'_l = z'$ of maximal simplices of $\text{Lk}(u)$ with $l \leq \theta(\kappa, \delta, n)$. Hence, there is a path $x = x'_1 \star u, \dots, x'_l \star u = z$ in W .

Claim 5.8. *Either $d_{X+W}(z, y) < d_{X+W}(x, y)$ or $d_{X+W}(z, y) = d_{X+W}(x, y)$ and the first edge of γ is not in X , while the first edge of some geodesic of minimal length from z to y is in X .*

Proof. Since $[\tau']$ is nested into $[t]$, we have in particular that any vertex of z not in τ' is connected in X to t . That is, z is within distance 1 of t in X^{+W} , and more precisely either t is

a vertex of z or t is connected in X to a vertex of z . Hence we have $d_{X+w}(z, y) < d_{X+w}(x, y)$ unless t is the second vertex of γ and *not a vertex of z* . But then in this case the first edge of γ is not an edge of X : if that edge e was in X , since t is an endpoint of e we would have $t = \tau$. This is a contradiction: recall that $\tau \star u = \tau'$, so $t = \tau$ implies $t \in \tau' \subset z'' \star \tau' = z$, contradicting that t is not a vertex of z .

On the other hand, there is a path of length $d_{X+w}(x, y)$ from z to y starting with an edge of X . If this path is not a geodesic, then $d_{X+w}(z, y) < d_{X+w}(x, y)$, and otherwise the other possibility holds. \blacksquare

If $d_{X+w}(z, y) < d_{X+w}(x, y) = k$, then by induction on k , there is a path in W , of length bounded by a constant depending only on κ, k, δ, n , that joins z to y . Otherwise, by Claim 5.8, $d_{X+w}(z, y) = k$, and some minimal-length X^{+W} -geodesic from z to y starts with an edge e of X . Now, e is contained in some maximal simplex z_1 of X , and z, z_1 share a vertex, namely the initial vertex of e . Noting that $d_{X+w}(z_1, y) < d_{X+w}(x, y)$, we argue as above to uniformly bound the distance in W from z_1 to y , and apply the base case to produce a path in W from z to z_1 and a bound on $d_W(z, z_1)$. Hence, by concatenating paths, we again see that z, y are in the same component of W and at distance bounded in terms of κ, δ, n, k .

Finally, we earlier exhibited a path of bounded length in W from x to z , so by concatenating again, we see that x, y are in the same component of W , and obtain a bound on $d_W(x, y)$, as required.

Case 2. For each non-maximal simplex σ with $[\sigma] \sqsubseteq [u]$ and $\text{diam}(\mathcal{C}(\sigma)) \geq \delta$, $\gamma - \{u\}$ does not intersect $\text{Sat}(\sigma)$. Moreover, the first edge of γ is not an edge of X .

Let v be the second vertex of γ . In this case, there exist maximal simplices p and q that are joined by an edge in W , p contains u and q contains v .

Claim 5.9. *There is $C = C(\delta, n, \kappa)$ so that for every $[\Delta]$ nested into $[u]$ we have $d_{\mathcal{C}([\Delta])}(p, x) \leq C$.*

Proof. Consider any $[\Delta]$ nested into $[u]$, and we can assume $\text{diam}(\mathcal{C}(\Delta)) \geq \delta$. The projection of p to $\mathcal{C}(\Delta)$ coarsely coincides with the projection of q , since we checked that projections are coarsely Lipschitz and p, q are adjacent in W , and in turn the projection of q coarsely coincides with that of v , since the latter is well-defined by the hypothesis about γ avoiding saturations. For the same reason plus Lemma 5.2, the projection of v coarsely coincides with that of y . Finally, by hypothesis the projection of y coarsely coincides with that of x , and hence we are done. \blacksquare

Write $x = x' \star u$ and $p = p' \star u$. Since W^u is an HHS, in W^u there is a path $x' = x'_1, \dots, x'_l = p'$ of length l of maximal simplices of $\text{Lk}(u)$ with $l \leq \theta(\kappa, \delta, n)$. Hence, there is a path $x = x'_1 \star u, \dots, x'_l \star u = p$ in W . In particular, there is a path of length $l + 1$ from x to q . Clearly, $d_{X+w}(q, y) < d_{X+w}(u, y)$, and we are done.

Conclusion: We have shown that the combinatorial HHS (X, W) gives rise to a hierarchically hyperbolic space (W, \mathfrak{S}) . It remains to prove the statement about group actions. Let G act on X . By hypothesis, there are finitely many G -orbits of links of non-maximal simplices. Since \mathfrak{S} corresponds G -equivariantly and bijectively to this set of links, the action of G on \mathfrak{S} is cofinite, as required by Definition 2.4.

The action on X also induces an action on the set of maximal simplices of X , and hence on the vertex set of W . By hypothesis, this action extends to a metrically proper, cobounded action on W , as required by Definition 2.4.

Let Δ be a non-maximal simplex, i.e., $[\Delta] \in \mathfrak{S}$. For any $g \in G$, the automorphism $g : X \rightarrow X$ restricts to an isomorphism $g : \text{Lk}(\Delta) \rightarrow \text{Lk}(g\Delta)$. In particular, the G -action on \mathfrak{S} induced by the action on X preserves $\sqsubseteq, \perp, \flat$.

Moreover, if $v, w \in \text{Lk}(\Delta)$ are contained in W -adjacent maximal simplices σ, τ , then gv, gw are contained in W -adjacent maximal simplices $g\sigma, g\tau$, so the isomorphism $g : \text{Lk}(\Delta) \rightarrow \text{Lk}(g\Delta)$ extends to an isometry $g : \mathcal{C}(\Delta) \rightarrow \mathcal{C}(g\Delta)$.

Similarly, since $g\text{Sat}(\Delta) = \text{Sat}(g\Delta)$, we get an isometry $g : Y_\Delta \rightarrow Y_{g\Delta}$. Because $\pi_{[\Delta]}$ was defined in terms of coarse closest-point projection $Y_\Delta \rightarrow \mathcal{C}(\Delta)$, we have $\pi_{g[\Delta]}(gw) = g(\pi_{[\Delta]}(w))$ for all $w \in W$. Similarly, if $[\Sigma] \sqsubset [\Delta]$ or $[\Sigma] \flat [\Delta]$, then $\rho_{g[\Delta]}^{g[\Sigma]} = g(\rho_{[\Delta]}^{[\Sigma]})$. Thus the action of G on the HHS (X, W) has all the properties listed in Definition 2.4, so (G, \mathfrak{S}) is an HHG. \square

6. HIERARCHICAL HYPERBOLICITY FROM ACTIONS ON HYPERBOLIC COMPLEXES

A motivating application of Theorem 1.18 is to groups acting on hyperbolic simplicial complexes with finite stabilizers of maximal simplices. In this setup, at the cost of reducing the generality, one can even formulate conditions that imply those in Theorem 1.18 and do not need to refer to an X -graph; Theorem 6.4 achieves just that. We start with two preliminary definitions:

Definition 6.1. We say that the group G *acts on* the combinatorial HHS (X, W) if it acts on X in such a way that the action on maximal simplices extends to W . We say that the action on (X, W) is proper (resp. cocompact) if the action on W is proper (resp. cocompact).

Definition 6.2 (Hyperbolic H -space). Let H be a group acting on a simplicial complex Y . We say that Y is a *hyperbolic H -space* if there is a graph Y' , consisting of $Y^{(1)}$ together with a set of additional edges \mathcal{E} called the *additional edge set*, such that Y' is hyperbolic, the action of H on $Y^{(1)}$ extends to an action of H on Y' , and \mathcal{E} contains finitely many H -orbits of edges.

Remark 6.3. Varying the additional edge set in the definition above does not change the quasi-isometry type. In particular, if Y is already connected, then it is a hyperbolic H -space if and only if it is already hyperbolic.

Now we can state the theorem:

Theorem 6.4. *Let G be a group acting cocompactly on the flag simplicial complex X , and suppose that the stabilizer of each maximal simplex is finite. Suppose, in addition, that for all non-maximal simplices Δ, Σ of X :*

(A) $\text{Lk}(\Delta)$ is a hyperbolic $\text{Stab}_G(\text{Lk}(\Delta))$ -space quasi-isometrically embedded in

$$\mathcal{E} \cup \left(X - \bigcup_{\text{Lk}(\Sigma) = \text{Lk}(\Delta)} \Sigma \right),$$

where \mathcal{E} is the additional edge set,

(B) $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \text{Lk}(\Delta \star \Pi) \star \Pi'$ for some simplices Π, Π' of $\text{Lk}(\Delta)$,

(C) $\text{Lk}(\Delta)$ is connected unless Δ is a codimension-1 face of a maximal simplex.

Then G acts properly and cocompactly on a combinatorial HHS (X, W) . In particular, G is a hierarchically hyperbolic group. Moreover, W can be chosen with the properties that

- any two maximal simplices of X corresponding to W -adjacent vertices share a codimension-1 face;
- $\mathcal{C}(\Delta)$ contains finitely many $\text{Stab}_G(\text{Lk}(\Delta))$ -orbits of edges for each simplex Δ which is a codimension-1 face of a maximal simplex.

Remark 6.5 (Metric in Condition (A)). For clarity, in item A, the quasi-isometric embedding statement refers to the restriction to $\text{Lk}(\Delta)$ of the hyperbolic metric built into the assumption that it is a hyperbolic $\text{Stab}_G(\text{Lk}(\Delta))$ -space; recall that all such metrics are naturally quasi-isometric to each other. Since $\text{Lk}(\Delta)$ is contained in $X - \bigcup_{\text{Lk}(\Sigma)=\text{Lk}(\Delta)} \Sigma$, it makes sense to add the additional edges to the latter graph. Finally, notice that, in view of condition C, for Δ not codimension-1 in a maximal simplex we could have more simply stated that $\text{Lk}(\Delta)$ is hyperbolic and quasi-isometrically embedded into $X - \bigcup_{\text{Lk}(\Sigma)=\text{Lk}(\Delta)} \Sigma$.

Condition (A) is playing the role of Definition 1.8.(2): adding \mathcal{E} amounts to augmenting the links, and $X - \bigcup_{\text{Lk}(\Sigma)=\text{Lk}(\Delta)} \Sigma$ is playing the role of Y_Δ .

Remark 6.6. Condition (B) holds for curve graphs. Since we do not need this fact explicitly, we will only mention that it can be proven using arguments similar to those used in Claim 8.24, and we leave the details to the interested reader. An interesting case to keep in mind, which shows the reason behind the “ $\star\Pi$ ”, is when Δ and Σ are each a curve, and these curves fill the complement of some other curve. The intersection of the links consists of the latter curve only.

Remark 6.7 (Improper actions). We will only use the condition on stabilizers of maximal simplices being finite to get properness of the action of G on W . Dropping that condition, we can still conclude the following:

- G acts coboundedly by HHS automorphisms on W , i.e., the HHS structure (W, \mathfrak{S}) , and the G -action on W , satisfy everything from Definition 2.4 except for the requirement that the action of G on W is proper;
- $\text{Cay}(G, S \cup \{H_i\})$ is equivariantly quasi-isometric to W , where the H_i are stabilizers of representatives of orbits of maximal simplices and S is a finite subset of G so that $S \cup \{H_i\}$ is a generating set.

Proof of Theorem 6.4. We will construct W as in the statement, and it will then follow from Theorem 1.18 that G is a hierarchically hyperbolic group.

Construction of W : A simplex Δ of X is *almost-maximal* if Δ is a codimension-1 face of a maximal simplex of X .

Let $\{\text{Lk}(\Delta_1), \dots, \text{Lk}(\Delta_k)\}$ contain exactly one element of each G -orbit of links of almost-maximal simplices. (Note that this is not quite the same as taking a list of G -orbit representatives of almost-maximal simplices and then taking links, since multiple G -distinct almost-maximal simplices can have the same link.)

For each $i \leq k$, there is an additional edge set $\mathcal{E}^i = \{e_1^i, \dots, e_{\ell(i)}^i\}$ of edges such that, by adding edges ge_j^i to $\text{Lk}(\Delta_i)$ (for $j \leq \ell(i)$ and $g \in \text{Stab}_G(\text{Lk}(\Delta_i))$), we obtain a hyperbolic graph; this set of additional edges exists by condition (A). For each i, j , let v_j^i, w_j^i denote the endpoints of e_j^i .

We are free to replace the hyperbolic graph obtained from $\text{Lk}(\Delta_i)$ in this way by any other $\text{Stab}_G(\text{Lk}(\Delta_i))$ -equivariantly quasi-isometric graph with finitely many $\text{Stab}_G(\text{Lk}(\Delta_i))$ -orbits of “additional” edges. For later use we make the following specific choice for the additional edges:

Remark 6.8. By hypothesis, $\text{Lk}(\Delta_i)$ (with the metric obtained by adding the edges e_j^i and their translates) is quasi-isometrically embedded in $\mathcal{E}^i \cup (X - \text{Sat}(\Delta_i))$. Suppose that $v, w \in \text{Lk}(\Delta_i)$ lie at distance at most 2 in $\mathcal{E}^i \cup X - \text{Sat}(\Delta_i)$. Then v, w lie at uniformly bounded distance (denoted B) in $\text{Lk}(\Delta_i)$ (with the extra edges).

Notice that $\text{Stab}(\text{Lk}(\Delta_i))$, and in fact even $\text{Stab}(\Delta_i)$, acts with finitely many orbits of vertices on $\text{Lk}(\Delta_i)$, since vertices v', w' in $\text{Lk}(\Delta_i)$ are in the same $\text{Stab}(\Delta_i)$ -orbit if $\Delta_i \star v', \Delta_i \star w'$ are in the same G -orbit of simplices with a marked vertex.

Hence $\text{Stab}(\text{Lk}(\Delta_i))$ acts cocompactly on $\text{Lk}(\Delta_i) \cup \text{Stab}_G(\text{Lk}(\Delta_i)) \cdot \{e_1^i, \dots, e_{\ell(i)}^i\}$. Moreover, $\text{Lk}(\Delta_i) \cup \text{Stab}_G(\text{Lk}(\Delta_i)) \cdot \{e_1^i, \dots, e_{\ell(i)}^i\}$ is locally finite since $\text{Stab}_G(\text{Lk}(\Delta_i))$ acts on $\text{Lk}(\Delta_i)$ with

finite stabilizers of vertices (since Δ_i is almost-maximal, and hence $\Delta_i \star v$ is maximal for any vertex v in $\text{Lk}(\Delta_i)$). In particular, there are finitely many orbits of paths of length at most B . Hence, by adding finitely many more $\text{Stab}_G(\text{Lk}(\Delta_i))$ -orbits of edges to $\text{Lk}(\Delta_i)$, we can and shall assume that any such v, w are joined by a $\text{Stab}_G(\text{Lk}(\Delta_i))$ -translate of an edge in $\{e_j^i\}$. This assumption will simplify matters in the proof of Claim 6.11 below.

Now, define an X -graph W as follows:

- The vertex set of W is the set of maximal simplices of X .
- If x, y are maximal simplices of X , we join x, y by an edge if there exists $i \leq k, j \leq \ell(i)$ and $g \in G$ such that $x = g(\Delta'_i \star v_j^i)$ and $y = g(\Delta'_i \star w_j^i)$, where Δ'_i is an almost-maximal simplex with $\text{Lk}(\Delta'_i) = \text{Lk}(\Delta_i)$. By construction, W -adjacent maximal simplices must intersect in a common almost-maximal simplex, which verifies the first item in the “moreover” clause of Theorem 6.4.

The G -action on W : The G -action on the set of maximal simplices of X induces a G -action on W by graph automorphisms, by the construction of W . Moreover, since G acts cocompactly on X , there are finitely many G -orbits of vertices in W . Every edge has the form $\{g(\Delta_i \star v_j^i), g(\Delta_i \star w_j^i)\}$, where Δ_i is one of finitely many almost-maximal simplices and (v_j^i, w_j^i) is one of finitely many pairs of vertices, so W has finitely many G -orbits of edges. Hence G acts cocompactly on W .

Since we have added finitely many orbits of edges to each $\text{Lk}(\Delta_i)$, and stabilizers of maximal simplices are finite, it follows that W is locally finite. Since G acts on W with finite vertex stabilizers (because maximal simplices in X have finite stabilizers by hypothesis), it then follows that G acts properly on W .

Hence, to complete the proof, it suffices to verify that (X, W) is a combinatorial HHS.

(X, W) is a combinatorial HHS: First, W is an X -graph by construction.

Also, observe that condition B is stronger than Definition 1.8.(3). In fact, any Γ whose link is not a non-trivial join (e.g., if $\text{diam}(\mathcal{C}_0(\Gamma)) \geq 3$) and so that $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$ cannot intersect Π' , and so we have $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Delta \star \Pi)$. (Moreover, $\text{Lk}(\Delta \star \Pi) \subseteq \text{Lk}(\Sigma)$ by definition, so $[\Delta \star \Pi] \sqsubseteq [\Sigma]$, as required.)

By definition, X is a flag complex. The remaining parts of Definition 1.8 are checked in the following claims.

Claim 6.9. *X has finite complexity.*

Proof. If we had $\Pi' = \emptyset$ in condition B, then inclusion of links would yield reverse inclusion of representative simplices, and the proof would be straightforward from finite dimension. Since Π' could be non-empty (which happens even in curve graphs), we need some understanding of join structures on links.

For a simplex Δ , let Θ_Δ be any simplex in $\text{Lk}(\Delta)$ so that $\text{Lk}(\Delta)$ is a join of Θ_Δ and some sub-complex, and Θ_Δ is maximal with this property. For convenience, for a simplex Δ we let $\#\Delta = |\Delta^{(0)}|$. Define $c(\Delta) = (\#\Gamma_\Delta, -\#\Theta_\Delta)$, where Γ_Δ is any simplex with $\text{Lk}(\Gamma_\Delta) = \text{Lk}(\Delta \star \Theta_\Delta)$ with the maximal number of vertices among all choices.

We claim that if $\text{Lk}(\Delta) \subsetneq \text{Lk}(\Sigma)$, then $c(\Sigma) < c(\Delta)$ in the lexicographic order. Since G acts on X cocompactly, $\dim X < \infty$, so this readily implies finite complexity.

First, let us show $\text{Lk}(\Delta \star \Theta_\Delta) \subseteq \text{Lk}(\Sigma \star \Theta_\Sigma)$. Notice that $\Theta_\Sigma \cap \text{Lk}(\Delta) \subseteq \Theta_\Delta$, since if we had a vertex v in $\Theta_\Sigma \cap \text{Lk}(\Delta) - \Theta_\Delta$, then we could add it to Θ_Δ to form a larger simplex with the property that $\text{Lk}(\Delta)$ is a join of the simplex and some sub-complex, contradicting maximality of Θ_Δ . Consider now a vertex $v \in \text{Lk}(\Delta \star \Theta_\Delta)$, that is $v \in \text{Lk}(\Delta) - \Theta_\Delta$. Then $v \in \text{Lk}(\Sigma)$, and it cannot lie in Θ_Σ , so $v \in \text{Lk}(\Sigma \star \Theta_\Sigma)$, as required.

If $\text{Lk}(\Delta \star \Theta_\Delta) \subsetneq \text{Lk}(\Sigma \star \Theta_\Sigma)$, then condition B implies that we can write $\text{Lk}(\Gamma_\Delta) = \text{Lk}(\Gamma_\Sigma \star \Pi) \star \Pi'$. By maximality of Θ_Δ , we have $\Pi' = \emptyset$, so we must have $\Pi \neq \emptyset$. Hence, $\#\Gamma_\Sigma < \#\Gamma_\Sigma \star \Pi \leq \#\Gamma_\Delta$, and hence $c(\Sigma) < c(\Delta)$ (regardless of the number of vertices of $\Theta_\Delta, \Theta_\Sigma$).

If $\text{Lk}(\Delta \star \Theta_\Delta) = \text{Lk}(\Sigma \star \Theta_\Sigma)$, then $\#\Gamma_\Delta = \#\Gamma_\Sigma$, so we have to show $\#\Theta_\Delta < \#\Theta_\Sigma$.

Let $v \in \Theta_\Delta$ be a vertex. Then $v \in \text{Lk}(\Delta)$, and in particular it lies either in $\text{Lk}(\Sigma \star \Theta_\Sigma)$ or in Θ_Σ . However, in the present situation, the former cannot occur since $v \notin \text{Lk}(\Delta \star \Theta_\Delta)$. Hence, $\Theta_\Delta \subseteq \Theta_\Sigma$. If we had equality, then we would have

$$\text{Lk}(\Delta) = \text{Lk}(\Delta \star \Theta_\Delta) \star \Theta_\Delta = \text{Lk}(\Sigma \star \Theta_\Sigma) \star \Theta_\Sigma = \text{Lk}(\Sigma),$$

but we are assuming $\text{Lk}(\Delta) \subsetneq \text{Lk}(\Sigma)$. Hence, $\#\Theta_\Delta < \#\Theta_\Sigma$, as required. \blacksquare

Claim 6.10. *Let Δ be a simplex, let $v, w \in \text{Lk}(\Delta)$ be vertices, and let $\Sigma \star v, \Sigma \star w$ be maximal simplices of X . Then there exists a simplex Π in $\text{Lk}(\Delta)$ so $\Delta \star \Pi \star v, \Delta \star \Pi \star w$ are maximal simplices of X .*

Proof. If $v = w$, then we can just take any maximal simplex $\Pi \star v$ in $\text{Lk}(\Delta)$, so suppose that this is not the case.

Write $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \text{Lk}(\Sigma \star \Pi_0) \star \Pi'_0$, as in condition B. Note that Σ is almost maximal since, for instance, $\Sigma \star v$ is maximal. Hence, $\text{Lk}(\Sigma)$ is discrete, and thus $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$ is discrete.

We claim that $\Pi'_0 = \emptyset$. If not, then since $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$ is discrete, Π'_0 is a single vertex and $\text{Lk}(\Sigma \star \Pi_0) = \emptyset$. Thus contradicts that $\text{Lk}(\Sigma) \cap \text{Lk}(\Delta)$ contains two distinct vertices, namely v and w . Hence $\Pi'_0 = \emptyset$.

Next, observe that $\Pi_0 = \emptyset$. Indeed, if not, then $\Sigma \star \Pi_0$ is a simplex whose link contains v and w , so $\Pi_0 \star v \star \Sigma$ is a simplex properly containing the $\Sigma \star v$, which was assumed to be maximal, a contradiction.

Hence $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \text{Lk}(\Sigma)$, which is to say that $\text{Lk}(\Sigma) \subseteq \text{Lk}(\Delta)$.

Now, write $\text{Lk}(\Sigma) = \text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \text{Lk}(\Delta \star \Pi) \star \Pi'$, again using condition B. Since $\text{Lk}(\Sigma)$ is discrete, we have $\Pi' = \emptyset$. But then Π is the simplex we wanted (notice that $\text{Lk}(\Delta \star \Pi)$ is discrete, so $\Delta \star \Pi \star v, \Delta \star \Pi \star w$ are maximal simplices). \blacksquare

Claim 6.11. *Let Δ be a simplex of X and let σ_v, σ_w be W -adjacent maximal simplices of X respectively containing vertices $v, w \in \text{Lk}(\Delta)$ that are distinct and non-adjacent in X . Then there exist simplices Σ_v, Σ_w of $\text{Lk}(\Delta)$ such that $\Sigma_v \star \Delta$ and $\Sigma_w \star \Delta$ are maximal simplices respectively containing v, w , and $\Sigma_v \star \Delta, \Sigma_w \star \Delta$ are adjacent in W .*

Proof of Claim 6.11. Let v, w be as in the statement. Then there exists Δ_i such that $\sigma_v = g(\Delta_i \star v_j^i)$ and $\sigma_w = g(\Delta_i \star w_j^i)$ for some $g \in G, j \leq \ell(i)$.

Notice that since v, w are distinct and not adjacent, neither v nor w is contained in $g\Delta_i$. (Indeed, if $v \in g\Delta_i$, then $v \in g\Delta_i \star gw_j^i = g(\Delta_i \star w_j^i)$, and hence v is adjacent in X to w , contradicting our hypothesis.)

Hence, $gv_j^i = v, gw_j^i = w$. Since σ_v, σ_w share the almost-maximal face $g\Delta_i$, Claim (6.10) provides a simplex Π of $\text{Lk}(\Delta)$ such that $\Delta \star \Pi \star v$ and $\Delta \star \Pi \star w$ are maximal simplices of X .

Suppose that $g\Delta_i$ is not contained in $\text{Sat}(\Delta \star \Pi)$. Then there is a path of length 2 in X from v to w that avoids $\text{Sat}(\Delta \star \Pi)$. So, $\Delta \star \Pi \star v$ and $\Delta \star \Pi \star w$ are W -adjacent, because of how we added extra edges to the link of the almost-maximal simplex $\Delta \star \Pi$ — see Remark 6.8. Hence we are done, with $\Sigma_v = \Pi \star v$ and $\Sigma_w = \Pi \star w$.

Otherwise, suppose that $g\Delta_i \subseteq \text{Sat}(\Delta \star \Pi)$. So, $[\Delta \star \Pi] \subseteq [g\Delta_i]$. Hence, there exist simplices Π' and Π'' such that $\text{Lk}(\Delta \star \Pi) = \text{Lk}(g\Delta_i \star \Pi') \star \Pi''$. But $\Delta \star \Pi$, being almost-maximal, has discrete link, so $\Pi'' = \emptyset$. Thus there exists Π' (a simplex of $\text{Lk}(g\Delta_i)$) with $[\Delta \star \Pi] = [g\Delta_i \star \Pi']$. Since $g\Delta_i$ is almost-maximal and $g\Delta_i \star \Pi'$ is necessarily non-maximal, we have $\Pi' = \emptyset$. So, $[\Delta \star \Pi] = [g\Delta_i]$. By definition, this means that $\text{Lk}(\Delta \star \Pi) = \text{Lk}(g\Delta_i)$. So, the extra edges

added to $\text{Lk}(\Delta \star \Pi)$ — which are determined by the link of $\Delta \star \Pi$, not the simplex itself — are exactly the edges determined by $g\Delta_i$, so $\Delta \star \Pi \star v, \Delta \star \Pi \star w$ are W -adjacent. Again, we are done, with $\Sigma_v = \Pi \star v$ and $\Sigma_w = \Pi \star w$. \blacksquare

Claim 6.11 says that (X, W) satisfies condition (4) from Definition 1.8. We now verify condition (2), which has two parts, the second of which verifies the second item in the “moreover” clause of the theorem:

Claim 6.12. *There exists δ such that $\mathcal{C}(\Delta)$ is δ -hyperbolic for each non-maximal simplex Δ of X , and it is moreover obtained from $\text{Lk}(\Delta)$ by adding finitely many $\text{Stab}_G(\text{Lk}(\Delta))$ -orbits of edges.*

Proof of Claim 6.12. Fix Δ . We first prove the statement about finitely many $\text{Stab}_G(\text{Lk}(\Delta))$ -orbits of edges. When $\text{Lk}(\Delta)$ is in the same G -orbit as $\text{Lk}(\Delta_i)$ for some i , we have by construction that vertices of $\text{Lk}(\Delta)$ are joined by a W -edge only if they are joined by a $\text{Stab}_G(\text{Lk}(\Delta))$ -translate of one of the finitely many additional edges in \mathcal{E}^i . And, by discreteness of $\text{Lk}(\Delta)$ in this case, no two vertices are joined by an X -edge. Hence there are finitely many $\text{Stab}_G(\text{Lk}(\Delta))$ -orbits of edges in $\mathcal{C}(\Delta)$ in this case.

In general, if Δ is not (up to equivalence of links) almost maximal, we argue as follows. By Claim 6.11, vertices $v, w \in \text{Lk}(\Delta)$ are $\mathcal{C}(\Delta)$ adjacent only if they are either adjacent in X (i.e. in $\text{Lk}(\Delta)$), or belong to W -adjacent maximal simplices of the form $\Delta \star \Sigma_v, \Delta \star \Sigma_w$. In the latter case, the proof of Claim 6.11 shows that there is a simplex Π of $\text{Lk}(\Delta)$ such that $\Delta \star \Sigma_v = \Delta \star \Pi \star v$ and $\Delta \star \Sigma_w = \Delta \star \Pi \star w$, and $[\Delta \star \Pi] = [g\Delta_i]$ for some $i \leq k$, and some $g \in G$, and moreover there exists $j \leq \ell(i)$ such that $gv_i^j = v, gw_i^j = w$.

Recall that $\text{Stab}_G(\Delta) \leq \text{Stab}_G(\text{Lk}(\Delta))$ acts cocompactly on $\text{Lk}(\Delta)$. In particular, gv_i^j belongs to one of finitely many $\text{Stab}_G(\text{Lk}(\Delta))$ -orbits of maximal simplices $\Pi \star v$ of $\text{Lk}(\Delta)$. Since W is locally finite, it follows that the edge of W from $\Delta \star \Pi \star v$ to $\Delta \star \Pi \star w$ belongs to one of finitely many $\text{Stab}_G(\text{Lk}(\Delta))$ -orbits, and hence the edge from v to w in $\mathcal{C}(\Delta)$ also belongs to one of finitely many $\text{Stab}_G(\text{Lk}(\Delta))$ -orbits.

Hence it remains to show that $\mathcal{C}(\Delta)$ is hyperbolic.

By the proof of Lemma 4.1 and the fact that we have already verified condition (4) from Definition 1.8, we know that $\mathcal{C}_0(\Delta) = \mathcal{C}(\Delta)$ for any simplex Δ , so it suffices to show that $\mathcal{C}_0(\Delta)$ is hyperbolic. If Δ is almost-maximal, then the choice of edges implies that $\mathcal{C}_0(\Delta) = \mathcal{C}(\Delta)$ is hyperbolic.

Otherwise, $\text{Lk}(\Delta)$ is connected, by condition (C). Since $\text{Lk}(\Delta)$ can be made hyperbolic by adding finitely many $\text{Stab}_G(\text{Lk}(\Delta))$ -orbits of edges, by condition (A), $\text{Lk}(\Delta)$ is quasi-isometric to a hyperbolic graph and is therefore hyperbolic. So, it suffices to show that $\mathcal{C}(\Delta)$ is quasi-isometric to $\text{Lk}(\Delta)$.

Now, the inclusion $\text{Lk}(\Delta) \rightarrow \mathcal{C}(\Delta)$ is Lipschitz and bijective on vertex sets, so we need to show that the inverse map on vertex sets is coarsely Lipschitz.

Suppose that $v, w \in \mathcal{C}(\Delta)$ are adjacent. If v, w are adjacent in $\text{Lk}(\Delta)$, then we are done. So, suppose that v, w belong to W -adjacent maximal simplices x, y . Then $x \cap y = \Sigma$ is an almost-maximal simplex, because of how the edges in W were defined. Now, if $\Sigma \subseteq \text{Sat}(\Delta)$, then $\text{Lk}(\Delta) \subseteq \text{Lk}(\Sigma)$. But since Σ is almost-maximal, $\text{Lk}(\Sigma)$ contains no X -edges, whence $\text{Lk}(\Delta)$ also contains no X -edges. But this means that Δ is almost-maximal, a contradiction. So Σ contains a vertex u of $X - \text{Sat}(\Delta)$, so v, u, w is a path of length 2 in $X - \text{Sat}(\Delta)$ from v to w . Since $\text{Lk}(\Delta)$ is quasi-isometrically embedded in $X - \text{Sat}(\Delta)$, by condition (A) and the fact that $\text{Lk}(\Delta)$ is connected (see Remark 6.3), this means that v, w lie at uniformly bounded distance in $\text{Lk}(\Delta)$, as required. \blacksquare

The final claim is:

Claim 6.13. *There exists δ such that $\mathcal{C}(\Delta)$ is (δ, δ) -quasi-isometrically embedded in Y_Δ , for all non-maximal simplices Δ of X .*

Proof of Claim 6.13. It suffices to show the claim for a fixed Δ .

Let Z_Δ be obtained from $X - \text{Sat}(\Delta)$ by connecting $\mathcal{C}(\Delta)$ -adjacent vertices of $\text{Lk}(\Delta)$. By the second part of Claim 6.12 and condition (A), $\mathcal{C}(\Delta)$ is quasi-isometrically embedded into Z_Δ .

We now show that $Z_\Delta \hookrightarrow Y_\Delta$ is a uniform quasi-isometric embedding.

Since Z_Δ is a subgraph of Y_Δ and the inclusion $Z_\Delta \hookrightarrow Y_\Delta$ is bijective on vertices, it suffices to show that if e is an edge of Y_Δ that is not an edge of Z_Δ , then the endpoints v, w of e are uniformly close in Z_Δ . Any such v, w are contained in maximal simplices x, y of W . Now, by construction of W , we have that $x \cap y$ is an almost-maximal simplex. If $(x \cap y)^{(0)} \not\subseteq \text{Sat}(\Delta)$, then v, w are joined by a path of length 2 in $X - \text{Sat}(\Delta) \subset Z_\Delta$, and we are done. Otherwise, $(x \cap y)^{(0)} \subset \text{Sat}(\Delta)$. Hence $\text{Lk}(\Delta) \subseteq \text{Lk}(x \cap y)$ so, since $x \cap y$ is almost-maximal, we have that $[\Delta] = [x \cap y]$. Hence v, w are joined by an edge of $\mathcal{C}(\Delta)$ (coming from the hyperbolic G -graph structure on $\text{Lk}(\Delta) = \text{Lk}(x \cap y)$). Thus v, w are adjacent in Z_Δ , as required. \blacksquare

This completes the proof that Definition 1.8.(2) holds for (X, W) , and hence completes the proof. \square

7. MAPPING CLASS GROUP QUOTIENTS

We now state Theorem 7.1, describing hierarchically hyperbolic quotients of mapping class groups. In this section we discuss the various consequences of the theorem, while the proof is given in Section 8.

In this section, and the next section, we abuse language slightly: we often use the same notation for a simplicial graph and for the flag complex determined by the graph. In particular, we do not distinguish between the curve graph $\mathcal{C}(S)$ of a surface S , and the curve complex $\mathcal{C}(S)$. For example, when talking about the metric on $\mathcal{C}(S)$, we mean the graph metric on the 1-skeleton; when talking about a combinatorial HHS (X, W) with $X = \mathcal{C}(S)$, we mean the full curve complex.

Recall that the complexity of a connected orientable surface S of finite type is $\xi(S) = 3\text{Gen}(S) + p(S) - 3$, where $\text{Gen}(S)$ is the genus and $p(S)$ is the number of punctures.

Theorem 7.1. *Let S be a connected orientable surface of finite type of complexity at least 2. Let $F \subseteq \text{MCG}(S)$ be any finite set, and let $Q < \text{MCG}(S)$ be a convex-cocompact subgroup. If all hyperbolic groups are residually finite, then the following holds.*

For all $-1 \leq i < \xi(S)$ there exists a normal subgroup $N_i \triangleleft \text{MCG}(S)$ such that the quotient $\phi: \text{MCG}(S) \rightarrow \text{MCG}(S)/N_i = G_i$ has the following properties:

- (I) (**Large injectivity radius.**) $\phi|_F$ is injective.
- (II) (**Explicit HHS structure.**) The action of \bar{G}_i on $\mathcal{C}(S)/N_i$ satisfies the hypotheses of Theorem 6.4, so that \bar{G}_i is a hierarchically hyperbolic group. More precisely, \bar{G}_i acts properly and cocompactly on a combinatorial HHS $(\mathcal{C}(S)/N_i, W)$, and the corresponding HHS structure $(\bar{G}_i, \mathfrak{S}_{N_i})$ satisfies:
 - the map $b: \mathfrak{S}_{N_i} \rightarrow \mathfrak{S}^{\geq 1}/N_i$ from Definition 7.2 below is well-defined and a bijection, where $\mathfrak{S}^{\geq 1}$ is the set of isotopy classes subsurfaces of S without annular or thrice-punctured sphere components;
 - two elements $U, V \in \mathfrak{S}_{N_i}$ are nested (resp. orthogonal) if and only if $b(U), b(V)$ have representatives in $\mathfrak{S}^{\geq 1}$ that are nested (resp. disjoint);
 - there exists B so that for any element U of \mathfrak{S}_{N_i} such that $b(U)$ has a representative of complexity at most i , we have that $\mathcal{C}(U)^{(0)}$ is finite and $\text{diam}(\mathcal{C}(U)) \leq B$.
- (III) (**Convex-cocompact injects.**) $\phi|_Q$ is injective and the orbit maps of Q to $\mathcal{C}(S)/N_i$ are quasi-isometric embeddings; in particular \bar{G}_i is infinite.

The existence of a bijection b with the stated properties is sufficient for many applications; in practice we will use the following construction (the existence of lifts of simplices of $\mathcal{C}(S)/N_i$ to $\mathcal{C}(S)$ is part of the claim in the theorem that b is well-defined):

Definition 7.2. Given any non-maximal simplex Δ of $\mathcal{C}(S)/N_i$, consider a lift $\hat{\Delta}$ to $\mathcal{C}(S)$. The vertex set of the link of $\hat{\Delta}$ in $\mathcal{C}(S)$ consists of all curves (regarded as vertices of $\mathcal{C}(S)$) contained in a subsurface that we denote $S_{\hat{\Delta}}$. Define $b([\Delta]) = [S_{\hat{\Delta}}]_{N_i}$, where $[\cdot]_{N_i}$ denotes the N_i -orbit.

The groups \bar{G}_i will be constructed inductively, with \bar{G}_{i+1} being a quotient of \bar{G}_i , and \bar{G}_{-1} being a quotient by suitable powers of Dehn twists. In particular, N_{-1} is a normal subgroup generated by normal powers of Dehn twists, and $N_i < N_{i+1}$ for all i . Roughly speaking, for each i , we have to ensure that subgroups of the mapping class group coming from subsurfaces of complexity i become finite, and to do so we mod out finite-index subgroups of those subgroups. More information on all this is provided at the beginning of Section 8.

Remark. The residual finiteness hypothesis will be applied to particular hyperbolic groups which arise in the proof.

Theorem 7.1 will be proven in Section 8.9. We now establish the corollaries stated in the introduction. Additionally, in the case of the closed genus 2 surface, we prove Theorem 7.1 without needing to assume residual finiteness for hyperbolic groups.

For Theorem 7.1, the case of $i = -1$ warrants extra focus, since there the mapping classes being quotiented are those supported in annuli, namely Dehn twists. Accordingly, in this case we provide a more explicit description of the kernel N_{-1} , which we state as Theorem 7.3. (In the general case, we prefer to keep the statement of Theorem 7.1 more concise rather than adding a more detailed description of N_i .) Also, note that in this case there is no residual finiteness assumption required.

Theorem 7.3. *Let S be a connected orientable surface of finite type of complexity at least 2. Given $K \geq 1$, denote by DT_K the normal subgroup generated by all K -th powers of Dehn twists. There exists $K_0 \geq 1$ so that, for any multiple K of K_0 , $MCG(S)/DT_K$ is an infinite hierarchically hyperbolic group. More precisely, given F, Q as in Theorem 7.1, all conclusions of Theorem 7.1 hold with $i = -1$ and $N_{-1} = DT_K$ for any sufficiently large multiple of K_0 .*

Proof. The proof follows verbatim the proof of Theorem 7.1, applied in the case $i = -1$, with the following modifications:

- We do not need the choice of H in Lemma 8.1, only the coloring of the subsurfaces described in the lemma, and the following fact: there is $K'_0 > 0$ so that any K'_0 -th power of an element of $MCG(S)$ preserves the coloring.
- The choice of $N = N_{-1}$ made in Notation 8.11 below can be replaced by choosing $\Gamma_Y^\theta < \langle \tau_Y \rangle$ to be $\langle \tau_Y^K \rangle$ for a suitably large multiple K of K'_0 .

Here, and in Corollary 7.4, we take K_0 to be such a suitably large multiple of K'_0 . □

In the special case of Theorem 7.3 for the the genus-two closed surface, an even stronger conclusion holds which we obtain in the following.

Corollary 7.4. *There exists $K_0 \geq 1$ so that for all non-zero multiples K of K_0 , the following holds. The quotient $MCG(\Sigma_2)/DT_K$ is hyperbolic relative to an infinite index subgroup commensurable to the product of two $C'(1/6)$ -groups, where DT_K denotes the normal subgroup generated by all K -th powers of Dehn twists.*

Proof. The peripheral subgroup of the relatively hyperbolic structure will be the image H in G_K of the stabilizer \hat{H} of a fixed curve γ that cuts Σ_2 into two $\Sigma_{1,1}$ subsurfaces. Note that \hat{H} is virtually a central extension by a Dehn twist of a product of virtually free groups, which

are isomorphic to the mapping class group of $\Sigma_{1,1}$. By [DHS21, Proposition 5.8], provided K is sufficiently large, we have the following. The subgroup H arises from \widehat{H} as the quotient by the subgroup generated by K -th powers of Dehn twists around γ and curves contained in one of the $\Sigma_{1,1}$. In particular, H is commensurable to the product of two groups, each of which is the quotient of a free group by K -th powers of certain elements, and this finite collection of elements is independent of K . In particular, up to increasing K , H is commensurable to the product of two finitely presented $C'(1/6)$ -groups.

Note that H has infinite index, for example because G_K is acylindrically hyperbolic [DHS21, Theorem 3.1], and thus cannot be commensurable to a product of infinite groups.

We are left to check relative hyperbolicity, for which we use [Rus19]. Recall that the index set of the HHS structure is $\mathfrak{S}^{\geq 1}/DT_K$, with elements being nested if and only if they have nested representatives, and similarly for orthogonality. The only orthogonal pairs in $\mathfrak{S}^{\geq 1}$ are pairs Y_i, W_i of surfaces homeomorphic to $\Sigma_{1,1}$. The set of surfaces $\{Y_i \cup W_i\}$ satisfies the following two properties:

- Whenever $U, V \in \mathfrak{S}^{\geq 1}$ satisfy $U \perp V$, there exists i so that $U, V \subseteq Y_i \cup W_i$,
- For $i \neq j$, there is no $U \in \mathfrak{S}^{\geq 1}$ so that $U \subseteq Y_i \cup W_i$ and $U \subseteq Y_j \cup W_j$.

Therefore, the analogous properties hold for $\{Y_i \cup W_i\}/DT_K \subseteq \mathfrak{S}^{\geq 1}/DT_K$, that is, $(G_K, \mathfrak{S}^{\geq 1}/DT_K)$ has *isolated orthogonality* in the sense of [Rus19, Definition 4.1]. Hence, by [Rus19, Theorem 4.3], G_K is hyperbolic relative to H . □

Remark 7.5. We believe that that quotients of $MCG(\Sigma_2)$ by suitable large powers of Dehn twists around *separating* curves are hyperbolic relative to subgroups which are virtually a direct product of free groups. However, we cannot use quotients of curve graphs to witness this, since, roughly, those quotients are expected to have an HHS structure that still has annular curve graphs corresponding to non-separating curves, and these annular curve graphs are not “visible” in the curve graph of Σ_2 .

Remark 7.6. Corollary 7.4 implies that $MCG(\Sigma_2)$ is fully residually non-elementary hyperbolic; we now explain why, and then provide a different argument for this fact relying on results of a very different nature dating back to [Pic81, Pic85].

Let K_0 be as in Theorem 7.3. Up to replacing K_0 with a multiple, we can assume this constant is large enough to satisfy the hypothesis of [DHS21, Theorem 3.1, Proposition 6.8]. Let $F \subseteq MCG(\Sigma_2)$ be finite. Using Theorem 7.3, choose a non-zero multiple K of K_0 so that $\phi|_F$ is injective, where $\phi : MCG(\Sigma_2) \rightarrow MCG(\Sigma_2)/DT_K = G_K$ is the quotient map.

The peripheral subgroup of the relatively hyperbolic structure on G_K from Corollary 7.4 is residually finite by residual finiteness of $C'(1/6)$ -groups, which follows from applying [Wis04, Theorem 1.2], [Ago13, Theorem 1.1], and [HW08, Theorem 4.4].

We can then apply the relatively hyperbolic Dehn filling theorem [Osi07, GM08] to construct a non-elementary hyperbolic quotient of G_K where F embeds, as we wanted.

The simpler argument is based on the following observation that was pointed out to us by Francesco Fournier Facio. Suppose that the residually finite group G has a non-elementary hyperbolic quotient H . Then G is fully residually non-elementary hyperbolic. To see this, fix a finite subset F of G , and consider a finite quotient Q of G where F embeds. Then G also maps to $H \times Q$, the image is non-elementary hyperbolic, and F embeds.

This observation can be applied to $MCG(\Sigma_2)$. In fact, mapping class groups are residually finite [Gro75], and a non-elementary hyperbolic quotient of $MCG(\Sigma_2)$ can be constructed as follows, as pointed out to us by Ian Agol. First, $MCG(\Sigma_2)$ maps onto $MCG(\Sigma_{0,6})$ by modding out by the hyperelliptic involution (see e.g. [BB01, Proposition 3.3]). In turn, $MCG(\Sigma_{0,6})$ maps onto the fundamental group of a finite-volume complex hyperbolic orbifold, see [Thu98, Theorem 0.2] or references therein. Such fundamental group has a non-elementary hyperbolic quotient, say by the relatively hyperbolic Dehn filling theorem.

As a side remark, we note that *pure* mapping class groups of punctured spheres with at least 4 punctures map onto the non-elementary hyperbolic group $PMCG(\Sigma_{0,4})$ via repeated use of the Birman exact sequence. It would be interesting to know if this could be promoted to the full mapping class groups of punctured spheres.

In the proof of Corollary 7.7 we explain how to construct hyperbolic quotients of the mapping class group; we then employ this construction in the remaining corollaries.

Corollary 7.7. *Let S be a connected orientable surface of finite type of complexity at least 2. If all hyperbolic groups are residually finite, then $MCG(S)$ is fully residually non-elementary hyperbolic.*

Proof. Let $F \subseteq MCG(S)$ be finite, and let Q be any convex-cocompact subgroup of $MCG(S)$ isomorphic to the free group on 2 generators (for instance, by [Fuj15], any sufficiently high powers of a pair of independent pseudo Anosovs will yield such a Q).

By Theorem 2, there is a hierarchically hyperbolic quotient \bar{G} of $MCG(S)$ such that: $MCG(S) \rightarrow \bar{G}$ is injective on F and Q ; Q quasi-isometrically embeds into \bar{G} ; and, all hyperbolic spaces in the HHS structure are bounded, except for the one space associated to the \sqsubseteq -maximal domain (the quotient of $\mathcal{C}(S)$). Thus, by [BHS, Corollary 2.15], \bar{G} is hyperbolic, and since Q embeds, \bar{G} is non-elementary. \square

Corollary 7.8. *Let S be a connected orientable surface of finite type of complexity at least 2. If all hyperbolic groups are residually finite, then every convex-cocompact subgroup of $MCG(S)$ is separable.*

Proof. We will show below that all *torsion-free* convex-cocompact subgroups are separable. This is sufficient, by the following argument. Let Q be convex-cocompact and let \hat{Q} be a finite index torsion-free subgroup of Q , which exists by, say, intersecting Q with a torsion-free finite-index subgroup of $MCG(S)$. Then \hat{Q} is closed in the profinite topology, which implies that its cosets are also closed. So, Q is closed since it is a finite union of closed sets. This reduces the claim to the case where Q is torsion-free, which we now address.

Let Q be a convex-cocompact torsion-free subgroup, and let $g \in MCG(S) - Q$. We consider two cases:

Non-pseudo-Anosov case. First suppose that g is reducible or periodic (that is, it acts with bounded orbits on $\mathcal{C}(S)$). Construct a hyperbolic quotient \bar{G} of $MCG(S)$ as in the proof of Corollary 7.7, for $F = \{1, g\}$ and our given Q . The image \bar{Q} of Q is quasi-convex in \bar{G} , and the image \bar{g} of g is non-trivial. Since g has bounded orbits in $\mathcal{C}(S)$, we have that \bar{g} has bounded orbits in the quotient of $\mathcal{C}(S)$. Hence, \bar{g} has finite order, and in particular it is not in \bar{Q} , since \bar{Q} is also torsion-free. In view of the fact that we are assuming that all hyperbolic groups are residually finite, by [AGM09, Theorem 0.1] we can find a finite quotient of \bar{G} in which the image of \bar{g} is not in the image of \bar{Q} .

Pseudo-Anosov case. For sufficiently large $n > 0$, the subgroup $\langle g^n, Q \rangle$ is convex-cocompact and is naturally isomorphic to $\langle g^n \rangle * Q$ (see e.g. [RST19, Theorem M]). Construct a hyperbolic quotient \bar{G} of $MCG(S)$, as in the proof of Corollary 7.7, except using the convex-cocompact subgroup $\langle g^n, Q \rangle$. Since $\langle g^n, Q \rangle$ is quasi-isometrically embedded in \bar{G} , so is Q . Moreover, $\bar{g} \notin \bar{Q}$ since $\langle g^n \rangle * Q \rightarrow \bar{G}$ is a quasi-isometric embedding. We conclude as above using [AGM09, Theorem 0.1]. \square

Corollary 7.9. *Let S be a connected orientable surface of finite type of complexity at least 2. If all hyperbolic groups are residually finite, then the following holds. Let $g, h \in MCG(S)$ be pseudo-Anosovs with no common proper power, and let $q \in \mathbb{Q}_{>0}$. Then there exists a finite group G and a homomorphism $\psi : MCG(S) \rightarrow G$ so that $\text{ord}(\psi(g))/\text{ord}(\psi(h)) = q$, where ord denotes the order.*

Proof. For sufficiently large n , the elements g^n, h^n freely generate a convex-cocompact free subgroup Q . Construct \bar{G} as in the proof of Corollary 7.7. Hence we have a hyperbolic quotient \bar{G} of $MCG(S)$ where the images \bar{g}, \bar{h} of g and h have infinite order and have no common proper power.

We can now quotient \bar{G} by suitable (large) powers of \bar{g} and \bar{h} to find a further hyperbolic quotient \tilde{G} where the images of g and h satisfy the condition on the orders as in the statement. Using residual finiteness of \tilde{G} we finally find the finite quotient of $MCG(S)$ that we were looking for. \square

8. PROOF OF THEOREM 7.1

8.1. Outline. We start with a rough outline of the proof of Theorem 7.1, in which we will take successive quotients of $MCG(S)$.

8.1.1. First quotient: Dehn twists. We start by describing the first quotient, which is the quotient of $MCG(S)$ by the normal subgroup generated by suitable powers of Dehn twists. In this outline, we denote this normal subgroup by N . We will check hierarchical hyperbolicity of $MCG(S)/N$ by considering its action on $\mathcal{C}(S)/N$ and applying Theorem 6.4. It was already proven in [DHS21] that $\mathcal{C}(S)/N$ is hyperbolic; here we will further develop the technology from [DHS21] to gain additional information about $\mathcal{C}(S)/N$.

The key tool will be lifts. In particular, the way that hyperbolicity of $\mathcal{C}(S)/N$ is proven in [DHS21] is by showing that geodesic triangles in $\mathcal{C}(S)/N$ can be lifted to geodesic triangles in $\mathcal{C}(S)$ (compare with Proposition 8.31 below). An important tool for doing such lifting is a version of “Greendlinger’s Lemma.” Roughly, this provides us with a normal form in which every term contributes a large projection to some domain and forces the lift to travel near some specific vertex in the curve graph. The tool through which we obtain our Greendlinger’s Lemmas, Lemma 8.10, is that of a *composite rotating family* in the sense of [Dah18].

The key generalization we provide here is that, rather than lifting triangles, here we lift more general objects, namely *generalized m -gons*. A generalized m -gon is roughly a concatenation of simplices and geodesics in links; we formalize this idea in Definition 8.12.

We will show that generalized m -gons can be lifted, provided that m is not too large. This will be the main tool to reduce various statements about links in $\mathcal{C}(S)/N$ to statements about links in $\mathcal{C}(S)$, which can be verified by curves-on-surfaces considerations.

To prove hierarchical hyperbolicity of $MCG(S)/N$ we use this type of argument repeatedly; see Subsections 8.8 and 8.9.

The condition of Theorem 6.4 which requires the most new work to check is the quasi-isometric embedding requirement in Theorem 6.4.(A). In order to check that condition, we use certain concatenations of geodesics in links which we call *approach paths* below. These will give rise to the most complicated generalized m -gons that we will consider.

Finally, choosing sufficiently large powers of Dehn twists allows one to make sure that a given finite set embeds in the quotient. Moreover, this allows one to preserve the contracting directions, which are characterized by having bounded projections to all proper subsurfaces; these are the “convex-cocompact” directions. This is essentially because the version of Greendlinger’s Lemma mentioned above says that nontrivial elements of N create large subsurface projections.

8.1.2. Further quotients. So far, we found that the quotient of $MCG(S)$ by suitable powers of Dehn twists is hierarchically hyperbolic. To pass to further hierarchically hyperbolic quotients, and eventually to a hyperbolic quotient, we use a similar method, approximately speaking.

At any given stage, we have a hierarchically hyperbolic group, and at the “bottom level of the hierarchy” we have hyperbolic groups (in the case of the first quotients, we had \mathbb{Z} subgroups). In each of the hyperbolic groups, we take sufficiently deep finite-index subgroups, and quotient

by those. In these later stages, in order to find the appropriate finite-index subgroups, we use the hypothesis that hyperbolic groups are residually finite.

Once again, in this setting we establish a composite rotating family, a Greendlinger's Lemma, and the ability to lift. However, this approximate description hides several technical difficulties, some quite serious, as we now explain.

As mentioned above, we use the combinatorial/geometric structure of $\mathcal{C}(S)$ to prove various properties of the quotient via lifting. This means that either one has to make sure that all those properties pass to the quotients of the curve graphs, or to take lifts to $\mathcal{C}(S)$ for the further quotients as well. We choose the second option.

The ability to first lift to the previous quotient of $\mathcal{C}(S)$, and from there to $\mathcal{C}(S)$ is an essential aspect of our induction. Accordingly, we establish this as Proposition 8.13.(VI)-(VII) below.

More generally, we collect all the properties that are required for the inductive hypothesis in Proposition 8.13; a couple of them follow from the others, but we found it helpful to have all of them collected in a single place.

Even given the ability to lift, just checking that the aforementioned sufficiently deep finite-index subgroups define a composite rotating family requires a significant amount of work, since relatively straightforward lifting arguments are not sufficient. This is one of a few places where we found it efficient to insist that the kernels of our quotients are contained in a carefully chosen finite-index subgroup of the mapping class group (as required by Proposition 8.13.(IV)), see Lemma 8.1. The subgroup we use is contained in the one constructed by Bestvina–Bromberg–Fujiwara [BBF15, Section 5]. We do not think that choosing this subgroup so specifically is strictly needed, but the strategies we are aware of to get around using it are significantly more complicated than using the subgroup.

8.1.3. Structure of the section. We now explain how the rest of the section is organized. In Subsection 8.2, we construct the ambient finite-index subgroup of $MCG(S)$. Then, after recalling the definitions of composite projection system and composite rotating family in Subsection 8.3, we set up the case of the first quotient in Subsection 8.4. In Subsection 8.5, we set up the induction in Proposition 8.13. After that, we can describe the composite projection system and composite rotating family for the further quotients, in Subsection 8.6. From that point on (and **only** from that point on), the proofs in the case of Dehn twists and in the case of the further quotients are largely the same, and are done together, with the occasional digression where the two cases are treated differently.

8.1.4. Warning to the reader. Some of the combinatorial properties of quotients of curve graphs that we will encounter correspond to topological properties that can be stated in terms of subsurfaces, curves, etc. While this might help with intuition and to motivate why they are relevant, we emphasize that very often it will not be straightforward at all to relate properties of quotients of curve graphs to topological properties, and we will have to rely on the combinatorial HHS viewpoint, taking advantage of topological arguments only after lifting to the curve graph. In particular, in all of our statements we can only use combinatorial, rather than topological, language.

8.2. Coloring subsurfaces. Throughout the proof of Theorem 7.1, we will use a strengthened form of the coloring of the subsurfaces of S constructed in [BBF15].

Lemma 8.1 (Enhanced BBF subgroups). *There exists a finite coloring $\mathfrak{S} = \mathfrak{S}_1 \sqcup \dots \sqcup \mathfrak{S}_t$ of the collection \mathfrak{S} of all subsurfaces of S , so that distinct elements with the same color overlap.*

Moreover, for every integer $q > 0$ there exists a finite-index torsion-free normal subgroup H_q of $MCG(S)$ so that

- (1) *the coloring is H_q -invariant, meaning $H_q \mathfrak{S}_j = \mathfrak{S}_j$ for all j ;*

- (2) for every subsurface Y with at least one component which is not an annulus or a pair of pants, and any curve γ on Y , there exists a curve α on Y such that no H_q -translate of α is disjoint from, or isotopic to, γ ;
- (3) for every Dehn twist $\tau \in MCG(S)$ we have $H_q \cap \langle \tau \rangle \leq \langle \tau^q \rangle$.

Proof. Bestvina–Bromberg–Fujiwara constructed a coloring of \mathfrak{S} with the property that distinct elements with the same color overlap [BBF15, Proposition 5.8]. Additionally, the coloring they construct has the property that there is a normal, torsion-free finite-index subgroup H^0 of $MCG(S)$ so that the colors are exactly the H^0 -orbits of the induced action on \mathfrak{S} . Thus, this coloring satisfies item (1) as well as the requirements before it.

(Strictly speaking, Bestvina–Bromberg–Fujiwara produced a coloring of the set of *connected* subsurfaces by H^0 -orbits with no two subsurfaces of the same color being disjoint. This extends to all of \mathfrak{S} : just color each disconnected subsurface by its H^0 -orbit. If Y is disconnected, then Y and gY cannot be disjoint, for $g \in H^0$, because then each component of Y would be disjoint from its g -translate.)

Starting from the BBF coloring, we will pass to increasingly deep finite-index subgroups of $MCG(S)$. We begin by noting that each of the enumerated properties in the statement is stable under passing to further finite-index subgroups.

We now arrange for the coloring to satisfy item (2). Consider a subsurface Y as in the statement, and any curve γ on Y . Let g be a mapping class that is supported on Y and does not stabilise γ . By taking g to be, for example, a partial pseudo-Anosov supported on Y , we can pass to a positive power and assume that our g with the preceding property also satisfies $g \in H^0$.

By [LM07, Theorem 1.4], $\text{Stab}(\gamma)$ is separable in $MCG(S)$. Hence H^0 has a finite-index subgroup $H^{Y,\gamma}$ such that $\text{Stab}(\gamma) \cap H^0 \leq H^{Y,\gamma}$ but $g \notin H^{Y,\gamma}$. It follows that $H^{Y,\gamma}g \cap \text{Stab}(\gamma) = \emptyset$.

Let $\alpha = g\gamma$. Then for any $h \in H^{Y,\gamma}$, we have that $h\alpha = hg\gamma$ is in the same H^0 -orbit as γ . On the other hand, $h\alpha \neq \gamma$, for otherwise we would have $hg \in H^{Y,\gamma} \cap \text{Stab}(\gamma)$, which is impossible. Thus $h\alpha$ and γ must intersect, by the defining property of H^0 , applied to annular subsurfaces.

The above paragraph holds for any fixed γ and Y . Since there are only finitely many H^0 -orbits of pairs (Y, γ) , we can conclude by taking a finite intersection of the $H^{Y,\gamma}$, where (Y, γ) varies over orbit-representatives. Denote this intersection by H^1 and we now have a subgroup which satisfies item (2), as desired.

Let us now fix q and arrange item (3). We will construct a finite-index normal subgroup $H^2 < MCG(S)$ so that for every Dehn twist $\tau \in MCG(S)$, we have $H^2 \cap \langle \tau \rangle < \langle \tau^q \rangle$. Then, setting $H_q = H^1 \cap H^2$, we will have that H_q is a finite-index subgroup of H^1 , which ensures that properties 1 and 2 hold as well.

Since there are finitely many conjugacy classes of Dehn twists, it suffices to show that for any given Dehn twist τ there is a normal finite-index subgroup H^τ so that $H^\tau \cap \langle \tau \rangle < \langle \tau^q \rangle$ (so that we can take a finite intersection of the H^τ for all τ in a complete list of conjugacy representatives).

Fix a Dehn twist τ . There are two cases:

- Suppose that τ is a Dehn twist around a non-separating curve. Let $S = S_{g,p}$, where g is the genus and p is the number of punctures.

If $p = 0$, then the usual action of $MCG(S)$ on $H_1(S, \mathbb{Z}/q\mathbb{Z})$ gives a homomorphism $\Psi : MCG(S) \rightarrow Sp(2g, \mathbb{Z}/q\mathbb{Z})$. By [FM12, Proposition 6.3], $\Psi(\tau)$ has order q . We let $H^\tau = \ker(\Psi)$.

Suppose $p \geq 1$. Let $S' = S_{g,p-1}$, so that S is obtained from S' by removing a point x . Let $F : PMCG(S) \rightarrow PMCG(S')$ denote the surjection in the Birman exact sequence, where $PMCG(S) \leq MCG(S)$ is the finite index subgroup fixing each puncture. Then

$F(\tau)$ is a Dehn twist τ' in S' around a non-separating curve. By induction, $MCG(S')$ has a finite-index subgroup $H^{\tau'}$ such that $H^{\tau'} \cap \langle \tau' \rangle < \langle (\tau')^q \rangle$. The subgroup $H^\tau = F^{-1}(H^{\tau'})$, which has finite index in $PMCG(S)$ and hence in $MCG(S)$, has the desired property.

- Now suppose τ is a Dehn twist around a separating curve. It suffices to consider the case that q is a power of a prime p , since in general we can take the intersection of the finite-index subgroups dictated by the prime factorization of q . By [Par09, Theorem 1.2], τ lies in a finite-index normal subgroup H^3 of $MCG(S)$ which is residually p . In particular, there is a finite quotient H^3/N of H^3 so that τ maps to an element of order q' with $q|q'$, and $[H^3 : N]$ is a power of p . Since the intersection of finitely many subgroups of index a power of p also has index a power of p , we can take N to be characteristic in H^3 , whence normal in $MCG(S)$, and set $H^\tau = N$.

This completes the proof of the lemma. \square

8.3. Composite projection systems and composite rotating families. Now we recall two definitions from [Dah18] that we will need below. Specifically, the following combines [Dah18, Definitions 1.1-1.2]. The reader can keep in mind as \mathbb{Y}_* the family of curves on a surface, which can be split into finitely many families with the property that curves in each family pairwise intersect by a result of [BBF15]. Also, $\text{Act}(Y)$ is the set of curves that intersect the curve Y , and d_Y denotes the distance between subsurface projections. (This was the motivating example for this notion in [Dah18].)

Definition 8.2 (Composite projection system). Let \mathbb{Y}_* be a countable set equipped with a finite coloring $\mathbb{Y}_* = \sqcup_{j=1}^m \mathbb{Y}_j$. For each $Y \in \mathbb{Y}_*$, let $j(Y)$ denote the value j for which $Y \in \mathbb{Y}_j$.

A *composite projection system* on a countable set \mathbb{Y}_* is the data consisting of: a constant $\theta > 0$; a family of subsets, one for each $Y \in \mathbb{Y}_*$ denoted $\text{Act}(Y) \subset \mathbb{Y}_*$ (called the *active set* for Y) such that $\mathbb{Y}_{j(Y)} \subset \text{Act}(Y)$; and a family of functions $d_Y : (\text{Act}(Y) \setminus \{Y\}) \times (\text{Act}(Y) \setminus \{Y\}) \rightarrow \mathbb{R}_+$, satisfying the following whenever all quantities are defined:

- (CPS1) (symmetry) $d_Y(X, Z) = d_Y(Z, X)$ for all X, Y, Z ;
- (CPS2) (triangle inequality) $d_Y(X, Z) + d_Y(Z, W) \geq d_Y(X, W)$ for all X, Y, Z, W ;
- (CPS3) (Behrstock inequality) $\min\{d_Y(X, Z), d_Z(X, Y)\} \leq \theta$ for all X, Y, Z ;
- (CPS4) (properness) $\{Y \in \mathbb{Y}_j, d_Y(X, Z) > \theta\}$ is finite for all X, Z ;
- (CPS5) (separation) $d_Y(Z, Z) \leq \theta$ for all Z, Y .

The map Act is required to satisfy three further properties:

- (CPS6) (symmetry in action) $X \in \text{Act}(Y)$ if and only if $Y \in \text{Act}(X)$;
- (CPS7) (closeness in inaction) if $X \notin \text{Act}(Z)$, for all $Y \in \text{Act}(X) \cap \text{Act}(Z)$, $d_Y(X, Z) \leq \theta$;
- (CPS8) (finite filling) for all $\mathcal{Z} \subset \mathbb{Y}_*$, there is a finite collection of elements X_j in \mathcal{Z} such that $\cup_j \text{Act}(X_j)$ covers $\cup_{X \in \mathcal{Z}} \text{Act}(X)$.

An *automorphism* of a composite projection system is a bijection $g : \mathbb{Y}_* \rightarrow \mathbb{Y}_*$ such that:

- g preserves each \mathbb{Y}_j ;
- for all $Y \in \mathbb{Y}_*$, we have $\text{Act}(gY) = g(\text{Act}(Y))$;
- for all Y and all $X, Z \in \text{Act}(Y)$, we have $d_{g(Y)}(g(X), g(Z)) = d_Y(X, Z)$.

The following is a variant of the notion introduced in [Dah18, Definition 2.1].

Definition 8.3 ((Θ_P, Θ_{Rot}) -Composite rotating family). Let Θ_P and Θ_{Rot} be constants. A (Θ_P, Θ_{Rot}) -*composite rotating family* on a composite projection system endowed with an action of a group G by automorphisms is a family of subgroups $\Gamma_Y, Y \in \mathbb{Y}_*$ such that

- (CRF1) for all $X \in \mathbb{Y}_*$, $\Gamma_X < G_X = \text{Stab}_G(X)$, is an infinite group of rotations around X , with proper isotropy, meaning that for all $R > 0$ and $Y \in \text{Act}(X)$ the set $F_Y^X(R) = \{\gamma \in \Gamma_X : d_X(Y, \gamma Y) < R\}$ is finite.

- (CRF2) for all $g \in G$, and all $X \in \mathbb{Y}_*$, one has $\Gamma_{gX} = g\Gamma_X g^{-1}$,
 (CRF3) if $X \notin \text{Act}(Z)$ then Γ_X and Γ_Z commute,
 (CRF4) for all j , for all $X, Y, Z \in \mathbb{Y}_j$, if $d_Y(X, Z) \leq \Theta_P$ then for all $g \in \Gamma_Y \setminus \{1\}$, $d_Y(X, gZ) \geq \Theta_{Rot}$.

Remark 8.4. Dahmani's original definition of a composite rotating family, [Dah18, Definition 2.1], doesn't use the metric d_Y from the composite projection system, but rather a perturbation which differs from d_Y by a bounded amount.

Our definition above relies on two constants Θ_P and Θ_{Rot} . In [Dah18, Definition 2.1], instead the constants Θ_P and Θ_{Rot} are fixed depending only on the constant θ from the composite projection system. For suitably chosen values of Θ_P and Θ_{Rot} a (Θ_P, Θ_{Rot}) -composite rotating family is a composite rotating family in the sense of Dahmani's [Dah18, Definition 2.1].

Specifically, in the notation of [Dah18, § 1.2.1] we can choose $\Theta_P = c_* + 21m\kappa + \kappa$ and $\Theta_{Rot} > 2c_* + 2\Theta_P + 20(\kappa + \Theta) + \kappa$, where all these constants are functions of the constant θ from the composite projection system. These two relations differ from the analogous ones in [Dah18, § 1.2.1] by an additional term of κ to take into account the perturbation of the metric by at most κ , and $\kappa \geq \theta$.

Accordingly, with a slight abuse of terminology we define:

Definition 8.5 (Composite rotating family). A *composite rotating family* is a (Θ_P, Θ_{Rot}) -composite rotating family with constants as in Remark 8.4.

Remark. The definition of proper isotropy given above is taken from [DHS21, Definition 2.3], and it is weaker than the corresponding definition from [Dah18, §1.2.2], which requires that the finite set from Definition 8.5.(CRF1) is independent of $Y \in \text{Act}(X)$, that is, $F_Y^X(R) = F^X(R)$. We now explain why, despite this, all results from [Dah18] still apply to a (Θ_P, Θ_{Rot}) -composite rotating family defined as above (that is, using proper isotropy from [DHS21]).

- First of all, proper isotropy is only needed in [Dah18] to the extent that it is needed to prove [Dah18, Lemma 1.4]; there are no other uses of $F^X(R)$ (simply denoted $F(R)$ in [Dah18]). In said lemma, the only relevant value of R is $R = 10\kappa$ for κ as in Remark 8.4. Hence, it suffices to show that for a (Θ_P, Θ_{Rot}) -composite rotating family, we can take $F^X(10\kappa) = \{1\}$.
- Spelling out the above, we have to show that for all $Y \in \text{Act}(X)$ and $\gamma \in \Gamma_X - \{1\}$ we have $d_X(Y, \gamma Y) > 10\kappa$; this would follow directly from Definition 8.5.(CRF4) except that we have to consider $Y \in \text{Act}(X)$ but not necessarily in $\mathbb{Y}_{j(X)}$. However, given such Y , we use that Γ_Y is infinite, as stipulated in Definition 8.5.(CRF1), to produce $\alpha \in \Gamma_Y - F_X^Y(10\kappa)$. By definition, $d_Y(X, \alpha X) \geq 10\kappa > \theta$. Hence, by Definition 8.2.(CPS3), we have $d_X(Y, \alpha X) \leq \theta \leq \Theta_P$. Since G preserves each \mathbb{Y}_i , we have $X, \alpha X \in \mathbb{Y}_{j(X)}$.
- Now let $\gamma \in \Gamma_X - \{1\}$. Then by Definition 8.5.(CRF4), we have $d_X(\alpha X, \gamma \alpha X) \geq \Theta_{rot}$. By equivariance, we have $d_X(\gamma \alpha X, \gamma Y) \leq \theta$, so by Definition 8.2.(CPS2), we get $d_X(Y, \gamma Y) \geq \Theta_{rot} - 2\theta > 10\kappa$, in view of our choice of Θ_{rot} .

This shows that a (Θ_P, Θ_{Rot}) -composite rotating family defined as above satisfies all of the statements about composite rotating families established in [Dah18] and in [DHS21], and hence we can freely apply both sets of results below. It also explains why the results from [Dah18] used in [DHS21] apply in the setting of the latter paper, where Θ_{rot} is always assumed to be sufficiently large.

8.4. Setting up the induction. Theorem 7.1 is proven by induction on i .

8.4.1. The annular case $i = -1$. The base case, where $i = -1$, will be verified almost identically to the inductive step, but the notation used in the proof has a slightly different meaning in the base case.

Notation 8.6. When $i = -1$, we will use the following notation:

- $N_{-2} = \{1\}$,
- $X = \mathcal{C}(S)$,
- $\mathbb{Y}_* = (\mathcal{C}(S))^{(0)}$
- d_Y denotes the distance in the annular curve graph $\mathcal{C}(Y)$ of Y .

Remark 8.7. The collection \mathbb{Y}_* above provides a composite projection system as shown in [Dah18, §3], where for all $W, Y, Z \in \mathbb{Y}_*$ the distance between Y and Z as measured in W is $d_W(\pi_W(Y), \pi_W(Z))$, where π_W is the (annular) subsurface projection.

In the case $i = -1$, we have the following:

Lemma 8.8 (Composite rotating family, annular case). *For every $\theta > 0$ there exists $\theta_0 > 0$ so that the following holds.*

Let τ_1, \dots, τ_k be a complete list of conjugacy representatives of Dehn twists in $MCG(S)$. For each $j \leq k$, let θ_j be a positive multiple of θ_0 . For $Y \in \mathbb{Y}_$, let $\Gamma_Y^\theta = \langle \tau_Y^{\theta_j(Y)} \rangle$, where τ_Y is the Dehn twist around Y , and $j(Y)$ has the property that τ_Y is conjugate to $\tau_{j(Y)}$.*

Then the subgroups Γ_Y^θ form a composite rotating family on the composite projection system \mathbb{Y}_ . Moreover,*

$$\min\{d_{\mathcal{C}(Y)}(x, \gamma x) : Y \in \mathbb{Y}_*, \gamma \in \Gamma_Y^\theta - \{1\}, x \in \mathcal{C}(Y)\} > \theta.$$

Finally, set $N = \langle\langle \Gamma_Y^\theta \rangle\rangle$. If Σ is a simplex of $\mathcal{C}(S)$ so that the complement of the multicurve $\Sigma^{(0)}$ has one complexity-1 component Y while all others are pairs of pants, then $\text{Stab}(\Sigma)/(N \cap \text{Stab}(\Sigma))$ is an infinite hyperbolic group acting with finite point-stabilizers on $\text{Lk}(\Sigma)/(N \cap \text{Stab}(\Sigma))$.

Remark 8.9. We remind the reader of the convention for (combinatorial) hierarchically hyperbolic spaces: when writing distances in a hyperbolic space/link, we suppress the π_\bullet notation for projection maps. So, e.g. $d_{\mathcal{C}(Y)}(x, y)$ means $d_{\mathcal{C}(Y)}(\pi_Y(x), \pi_Y(y))$. See Notation 2.3.

Proof. Except the last conclusion, the proof is identical to the one give by Dahmani in [Dah18, Section 3], with two minor changes: first, while in Dahmani's case all the Dehn twists are raised to the same power, ours are allowed to vary; and, second our coloring was chosen in a more specific way.

To verify the last conclusion, we may thus apply the versions of results from [DHS21] which allow for variable powers of Dehn twists and our particular coloring, as noted above.

For the last conclusion, with a suitable choice of powers of Dehn twists we can apply [DHS21, Proposition 5.8]. This says that $N \cap \text{Stab}(\Sigma)$ is generated by the powers of Dehn twists supported in Y and those supported around the curves of $\Sigma^{(0)}$. More precisely, [DHS21, Proposition 5.8] applies to stabilizers of vertices in the curve graph, but we can apply it inductively passing to subsurfaces. But then $\text{Stab}(\Sigma)/(N \cap \text{Stab}(\Sigma))$ is virtually a quotient \hat{G} of $MCG(Y)$ by powers of Dehn twists, which is a hyperbolic group by [DHS21, Theorem 6.8.(1)].

We now argue that the action of $\text{Stab}(\Sigma)/(N \cap \text{Stab}(\Sigma))$ on $\text{Lk}(\Sigma)/(N \cap \text{Stab}(\Sigma))$ has finite stabilizers. These stabilizers are quotients of stabilizers of vertices of $\text{Lk}(\Sigma)$ since every element of the stabilizer is the image in the quotient by an element in the stabilizer of an orbit, but this representative can be multiplied by an element of $N \cap \text{Stab}(\Sigma)$ to ensure that it lies in a vertex stabilizer. Quotients of stabilizers of vertices are virtually generated by commuting Dehn twists, and therefore become finite after quotienting by $N \cap \text{Stab}(\Sigma)$. \square

The following lemma is crucial to show that we can lift various objects from quotients of curve graphs to curve graphs. Roughly, it says that for N the normal subgroup generated by large powers of Dehn twists, two vertices being in the same N -orbit can be witnessed by a large

annular projection that can be “shortened” with a suitable power of a Dehn twist which lies in N . This kind of “shortening” lemmas are often referred to as “Greendlinger lemmas” in analogy with the the Greendlinger lemma for classical small cancellation groups which allows one to shorten a word representing the identity using a relator.

Lemma 8.10 (“Greendlinger lemma”, annular case). *There exists a diverging function \mathfrak{T} so that the following holds for $\theta > 0$. Let $N = N_{-1} = \langle\langle\{\Gamma_Y^\theta\}\rangle\rangle$. Then there is a well-ordered set \mathfrak{C} , and an assignment $\gamma \in N \mapsto c(\gamma) \in \mathfrak{C}$, with $c(1)$ the minimal element of \mathfrak{C} .*

Moreover, for all $\gamma \in N - \{1\}$ and all simplices Δ of X , there is $Y \in \mathbb{Y}_$ and $\gamma_Y \in \Gamma_Y^\theta$ so that $c(\gamma_Y \gamma) < c(\gamma)$ and either*

- $\Delta \subseteq \text{Fix}(\Gamma_Y^\theta)$, or $\gamma \Delta \subseteq \text{Fix}(\Gamma_Y^\theta)$, or
- $d_Y(\Delta, \gamma \Delta) > \mathfrak{T}(\theta)$ (and the quantity is defined).

Proof. We note that the proof of this result is a minor variation of the proof of [DHS21, Corollary 4.6]. The key difference is that here we must consider simplices rather than just single vertices.

The proof in the present annular case is identical, verbatim, to the proof in the general case (Lemma 8.28), so we postpone the proof until then. \square

Lemma 8.8 will be used in combination with Lemma 8.1 to choose a finite-index subgroup H of $MCG(S)$, and an H -invariant coloring of the subsurfaces, as follows.

Notation 8.11. Below $MCG(S)$ is given its usual HHS structure as in [BHS19, Section 11]. Let C be the constant from Definition 2.1.(7)(bounded geodesic image) for $MCG(S)$, see [BHS19, Section 11] or [MM00, Theorem 3.1]. Let κ be given by Lemma 2.10 applied to the subgroup $Q < MCG(S)$ from Theorem 7.1. Fix some $\theta > 0$ such that:

- $\mathfrak{T}(\theta) > (6\text{Gen}(S) + 2p(S) - 3)C$,
- $\mathfrak{T}(\theta) > \max\{d_{C(Y)}(fx, gx) : Y \in \mathbb{Y}_*, f, g \in F, x \in X\}$,
- $\mathfrak{T}(\theta) > \kappa + C$.

Let θ_0 be as in Lemma 8.8, for the given θ . Fix from now on $H = H_{\theta_0}$ as in Lemma 8.1. For $Y \in \mathbb{Y}_*$, let $\Gamma_Y^\theta = H \cap \langle \tau_Y \rangle < \langle \tau_Y^{\theta_0} \rangle$; notice that Lemma 8.8 applies to Γ_Y^θ by the construction of H . For each Y , let $\Gamma_Y = \Gamma_Y^\theta$. Set $N = N_{-1} = \langle\langle\{\Gamma_Y\}\rangle\rangle$.

At this point, we also fix an H -invariant coloring $\mathfrak{S} = \bigsqcup_j \mathfrak{S}_j$ as in Lemma 8.1, which will remain the same at all stages of the induction.

We emphasize that H will remain the same at all subsequent stages of the induction, even though we have thus far only defined Γ_Y^θ in cases where Y is an annulus.

8.5. Inductive conditions. We now set the notation for $i \geq -1$. Suppose we have constructed $\bar{G}_i = MCG(S)/N_i$ for a given F and Q , at all complexities up to i . We will make further assumptions for the inductive step. To state these we need the following definition.

Definition 8.12 (Generalized m -gon). A *generalized m -gon* in a simplicial graph is a sequence $\tau_0, \dots, \tau_{m-1}$ so that:

- Each τ_j is either a simplex (**type S**), together with non-empty sub-simplices τ_j^\pm , or a geodesic in $\text{Lk}(\Delta_j)$ for some (possibly empty) simplex Δ_j (**type G**) with endpoints τ_j^\pm .
- $\tau_j^+ = \tau_{j+1}^-$ (indices are taken modulo m).

(The second bullet implies that $\tau_j \cap \tau_{j+1}$ is non-empty.)

The main inductive statement is the following proposition, which is a more precise version of Theorem 7.1; many of the additional points are required to inductively obtain composite projection systems. We denote pointwise stabilizers by PStab . Also, given a simplicial map $q : Y \rightarrow Z$ of simplicial complexes, a *lift* of an ordered simplex Δ of Z (that is, a simplex with an ordering on its vertices) is an ordered simplex Σ of Y so that $q(\Sigma) = \Delta$, and the map is

order-preserving at the level of vertices. We will denote ordering on vertices on the subscripts, e.g., (v_0, \dots, v_k) if the vertex set of the simplex consists of the v_j . A lift of a simplex is a lift of the simplex with any order on its vertices. A *lift* of a generalized m -gon $\tau = \tau_0, \dots, \tau_{m-1}$ in Z is an m -gon $\tau'_0, \dots, \tau'_{m-1}$ in Y with $q(\tau'_j) = \tau_j$, τ'_j is **type S/type G** if and only if τ_j is, and if τ_j is a geodesic in $\text{Lk}(\Delta_j)$ then τ'_j is a geodesic in $\text{Lk}(\Sigma_j)$ for some lift Σ_j of Δ_j .

Proposition 8.13. *Let S , F , and Q be as in Theorem 7.1, with S having genus $\text{Gen}(S)$ with $p(S)$ punctures. For $-1 \leq i \leq 3\text{Gen}(S) + p(S) - 4$, there exists a quotient $\phi : \text{MCG}(S) \rightarrow \text{MCG}(S)/N_i = \bar{G}_i$ such that properties Theorem 7.1.(I)-(III) hold. Moreover, the following additional properties also hold, where $q : \mathcal{C}(S) \rightarrow \mathcal{C}(S)/N_i$ is the quotient map:*

- (IV) $N_i < H$, where $H < \text{MCG}(S)$ is as in Notation 8.11, and $N_{-1} < N_i$.
- (V) For all distinct $f, g \in F$ either $f^{-1}g$ has finite order or there exists a vertex x of $\mathcal{C}(S)/N_i$ so that $f(x) \neq g(x)$.
- (VI) For $m \leq \max\{4, 6\text{Gen}(S) + 2p(S) - 3\}$, any generalized m -gon in $\mathcal{C}(S)/N_i$ can be lifted to $\mathcal{C}(S)$.
- (VII) For every ordered simplex Δ of $\mathcal{C}(S)/N_i$ there is a unique N_i -orbit of lifts $\hat{\Delta}$ in $\mathcal{C}(S)$, and for any such lift we have $q(\text{Sat}(\hat{\Delta})) = \text{Sat}(\Delta)$.
- (VIII) There exists C_i with the following property. Let Δ be a simplex of $\mathcal{C}(S)/N_i$, and let v_0, \dots, v_k be a geodesic of $\text{Lk}(\Delta)$. Suppose that for some simplex Σ of $\mathcal{C}(S)/N_i$ we have that $d_{\mathcal{C}(\Sigma)}(v_0, v_k)$ is defined and at least C_i . Then there exists i so that $\Delta^{(0)} \cup \{v_i\}$ is contained in $\text{Sat}(\Sigma)$.
- (IX) If an element $b([\Delta]) \in \mathfrak{S}^{\geq 1}/N_i$, for $[\Delta] \in \mathfrak{S}_{N_i}$, has a representative of complexity $i + 1$, then $\text{PStab}(\text{Sat}(\Delta))$ is hyperbolic and acts properly and cocompactly on $\mathcal{C}(\Delta)$.

Remark 8.14. The combinatorial HHS $(\mathcal{C}(S)/N_i, W)$ from Theorem 7.1.(II) is obtained by applying Theorem 6.4 to the action of \bar{G}_i on $\mathcal{C}(S)/N_i$; see Section 8.9.

8.5.1. *Guide to the proof of Proposition 8.13 (hence Theorem 7.1).*

Convention 8.15. From now and until the end of the Section we assume that either:

- $i = -1$, with the notation from Notation 8.6,
- $i > -1$, and Proposition holds with i replaced by $i - 1$.

Remark 8.16. For $i = 0$, the composite projection system is empty and thus the quotient we are taking is trivial in the sense that we are just quotienting by the trivial subgroup. The reason this is empty boils down to the fact that the complexity 0 subsurfaces are thrice punctured spheres and thus for each the curve graph is empty. Thus, since the composite projection system is empty, statements involving the set \mathbb{Y}_* are all vacuously true.

Here is a list of where properties (I)–(VII) are verified:

- Items (I) and (V) are proven together in Subsection 8.10.
- Item (II) is the content of Subsection 8.9.
- Item (III) is also proven in Subsection 8.10.
- Item (IV) is simply a restatement of the assumptions we made in Notation 8.6, for $i = -1$, and Notation 8.30, for $i > -1$.
- Item (VI) and the first part of Item (VII) hold by Proposition 8.31, with the second part of Item (VII) being Lemma 8.42.

The remaining properties follow from properties (I)–(VII), as shown in the following lemmas.

Lemma 8.17. *Assume properties (I)–(VII) hold for our current i . Then property (VIII) holds for our current i .*

Proof. By property (II), we are working in a combinatorial HHS with underlying simplicial complex $\mathcal{C}(S)/N_i$.

If some vertex of $\Delta^{(0)}$ was not in $\text{Sat}(\Sigma)$, then there would be a path of length 2 in Y_Σ from v_0 to v_k and thus $d_{\mathcal{C}(\Sigma)}(v_0, v_k)$ would be uniformly bounded. This is a contradiction when C is sufficiently large. Thus all vertices in $\Delta^{(0)}$ are contained in $\text{Sat}(\Sigma)$, which yields that $[\Sigma] \subseteq [\Delta]$.

Using the combinatorial HHS structure of $\text{Lk}(\Delta)$ (Propositions 4.9 and 4.11) we have that the geodesic must intersect $\text{Sat}(\Sigma)$ by Lemma 5.2, as required. \square

Lemma 8.18. *Assume properties (I)–(VII) for our current i . Then property (IX) holds for our current i .*

Proof. Recall that PStab denotes pointwise stabilizers.

We use the combinatorial HHS $(\mathcal{C}(S)/N_i, W)$, which exists by item (II). We will check that $\text{PStab}(\text{Sat}(\Delta))$ is hierarchically hyperbolic using the technology of Section 4.

Recall the graph W^Δ defined in Definition 4.2, which has vertex set the set of maximal simplices of $\text{Lk}(\Delta)$, and has the property that $(\text{Lk}(\Delta), W^\Delta)$ is a combinatorial HHS. It is readily seen that $\text{PStab}(\text{Sat}(\Delta))$ acts on W^Δ . We now check that this action is proper and cocompact.

Properness: Let Σ be a vertex of W^Δ , that is, a maximal simplex of $\text{Lk}(\Delta)$. Then $\text{Stab}_{\text{PStab}(\text{Sat}(\Delta))}(\Sigma) \leq \text{Stab}_{\bar{G}_i}(\Delta \star \Sigma)$, and the latter group is finite, since \bar{G}_i acts properly on W again by item (II). Hence $\text{PStab}(\text{Sat}(\Delta))$ acts properly on W^Δ .

Cocompactness for $i = -1$: We check that there are finitely many $\text{PStab}(\text{Sat}(\Delta))$ -orbits of edges in W^Δ . The proof is different depending on i .

For $i = -1$ (that is, \bar{G}_i is a quotient of $MCG(S)$ by powers of Dehn twists), no simplex of $\mathcal{C}(S)/N_i$ corresponds (via the map b from Definition 7.2) to a subsurface of complexity $i + 1$, since $\mathfrak{S}^{\geq 1}$ does not contain any complexity-0 surfaces. So the lemma holds vacuously for $i = -1$.

Cocompactness for $i = 0$: The argument for cocompactness when $i = 0$ is more complicated. In this case, Remark 8.16 says that $\bar{G}_0 = \bar{G}_{-1}$ and $\mathcal{C}(S)/N_0 = \mathcal{C}(S)/N_{-1}$. Our simplex Δ has the property that $b([\Delta])$ is an orbit of complexity-1 subsurfaces.

We will use the following claim about $\bar{G}_{-1} = \bar{G}_0$ and $N_{-1} = N_0$ later as well, so we state it separately, together with the assumptions we need to prove it.

Claim 8.19. *Assume that for $i = -1$ the following items from Proposition 8.13 hold:*

- item (IV);
- item (VI);
- the part of (VII) about uniqueness of N_{-1} -orbits.

Let Δ be an almost-maximal simplex of $\mathcal{C}(S)/N_{-1}$, and let Γ_Δ be the quotient of $\text{Stab}(\text{Lk}(\Delta))$ by the kernel of the action on $\text{Lk}(\Delta)$. Then Γ_Δ has finite vertex-stabilisers in $\text{Lk}(\Delta)$.

Proof. Fix a vertex v of $\text{Lk}(\Delta)$. The simplex $\Delta \star v$ is maximal because of the assumption that Δ is almost-maximal. By the third bullet point, we can fix, once and for all, lifts $\hat{\Delta}$ and \hat{v} of Δ and v such that $\mathcal{C}(S)$ contains the simplex $\hat{\Delta} \star \hat{v}$.

The subsurface Y : The 0-skeleton of our fixed lift $\hat{\Delta}$ is a multicurve in S whose complement is a subsurface we denote by Y' . By assumption, $\Delta \star v$ is a maximal simplex of $\mathcal{C}(S)/N_{-1}$, so the simplex $\hat{\Delta} \star \hat{v}$ of $\mathcal{C}(S)$ is maximal. Indeed, if it was not maximal, we could extend it to a bigger simplex, and distinct vertices in a common simplex of $\mathcal{C}(S)$ are not in the same N_{-1} -orbit, by the first bullet point, resulting in a simplex in $\mathcal{C}(S)/N_{-1}$ strictly containing $\Delta \star v$. Hence the vertex set of $\hat{\Delta} \star \hat{v}$ is a pants decomposition of S . Therefore, $\{\hat{v}\}$ is a pants decomposition of a

complexity-1 component Y of Y' . The curves in ∂Y belong to $\widehat{\Delta}$, and the remaining curves of $\widehat{\Delta}$ form a pants decomposition of the complement of Y .

Keeping track of groups: Recall that $\phi : MCG(S) \rightarrow \bar{G}_{-1}$ denotes the quotient map. Let $\omega : \text{Stab}(\text{Lk}(\Delta)) \rightarrow \text{Sym}(\text{Lk}(\Delta))$ be the action on $\text{Lk}(\Delta)$, and we identify Γ_Δ with $\text{im}(\omega)$.

We now argue that we have

$$q(\text{Lk}(\widehat{\Delta})) = \text{Lk}(\Delta).$$

Indeed, given a vertex w of $\text{Lk}(\Delta)$, we can lift $\Delta \star w$ to a simplex of the form $\widehat{\Delta} \star \widehat{w}$, and hence $q(\widehat{w}) = w$. This shows the containment “ \supseteq ”, the other containment follows from a use of the first bullet point similar to the argument above.

Since $q(\text{Lk}(\widehat{\Delta})) = \text{Lk}(\Delta)$, the homomorphism ϕ restricts on $\text{Stab}_{MCG(S)}(\text{Lk}(\widehat{\Delta}))$ to a homomorphism to $\text{Stab}(\text{Lk}(\Delta))$.

Next, observe that $\text{Lk}(\widehat{\Delta})$ corresponds to the set of curves in Y , and hence $\text{Stab}_{MCG(S)}(\text{Lk}(\widehat{\Delta}))$ is contained in $\text{Stab}_{MCG(S)}(\partial Y)$. Let Λ_0 be the kernel of the action of $\text{Stab}_{MCG(S)}(\text{Lk}(\widehat{\Delta}))$ on the set of boundary curves of Y , so that Λ_0 has finite index in $\text{Stab}_{MCG(S)}(\partial Y)$. We summarize this with the following diagram,

$$\Lambda_0 \hookrightarrow \text{Stab}_{MCG(S)}(\text{Lk}(\widehat{\Delta})) \xrightarrow{\phi} \text{Stab}(\text{Lk}(\Delta)) \xrightarrow{\omega} \Gamma_\Delta,$$

where the rightmost arrow is surjective and the leftmost arrow has finite-index normal image.

Reducing to a claim about lifting elements: We need to show that $\omega(\text{Stab}(\text{Lk}(\Delta)) \cap \text{Stab}(v))$ is finite. To do this, we will show that for each $g \in \text{Stab}(\text{Lk}(\Delta)) \cap \text{Stab}(v)$, there exists $\widehat{g} \in \text{Stab}_{MCG(S)}(\text{Lk}(\widehat{\Delta})) \cap \text{Stab}_{MCG(S)}(\widehat{v})$ such that $\phi(\widehat{g}) = g$. We first explain why the latter statement suffices to prove the former.

First, $\Lambda_0 \cap \text{Stab}_{MCG(S)}(\widehat{v})$ has finite index in $\text{Stab}_{MCG(S)}(\text{Lk}(\widehat{\Delta})) \cap \text{Stab}_{MCG(S)}(\widehat{v})$, so we can fix (independently of g) finitely many elements $\widehat{g}_1, \dots, \widehat{g}_k \in \text{Stab}_{MCG(S)}(\text{Lk}(\widehat{\Delta})) \cap \text{Stab}_{MCG(S)}(\widehat{v})$ representing all cosets of $\Lambda_0 \cap \text{Stab}_{MCG(S)}(\widehat{v})$.

So, for some j , we have $\widehat{g} = \widehat{g}_j \widehat{h}$, where $\widehat{h} \in \Lambda_0 \cap \text{Stab}_{MCG(S)}(\widehat{v})$. Now, \widehat{h} stabilises each component of ∂Y , and stabilises the curve \widehat{v} , so \widehat{h} has the same action on $\mathcal{C}(Y)$ as does its restriction to Y (which is defined since \widehat{h} stabilises all boundary components). Hence, $\phi(\widehat{h})$ acts in the same way as $\phi(\tau^\ell)$ on $\text{Lk}(\Delta)$, i.e. $\omega \circ \phi(\widehat{h}) = \omega \circ \phi(\tau^\ell)$ (by a simple argument using $q(\text{Lk}(\widehat{\Delta})) = \text{Lk}(\Delta)$).

Since τ has a positive power contained in N_{-1} — see Notation 8.11 — we conclude that there are finitely many possibilities for $\omega \circ \phi(\widehat{h})$. Since $\omega(g) = \omega \circ \phi(\widehat{g}_i) \cdot \omega \circ \phi(\widehat{h})$, there are therefore only finitely many possibilities for g , i.e. the stabiliser of v in Γ_Δ is finite, as required.

So, it remains to produce the lift \widehat{g} of g .

Auxiliary vertices in $\text{Lk}(\Delta)$: There exists a vertex $w \in \text{Lk}(\Delta) - \{v\}$. Indeed, by the first bullet point, N_{-1} is contained in the subgroup $H < MCG(S)$ from Notation 8.11 (where H is in particular a subgroup provided by Lemma 8.1). By Lemma 8.1.(2), we have the following. Recall that \widehat{v} is a curve on the complexity-1 subsurface Y , so by the lemma, there is a curve \widehat{w} on Y such that \widehat{v} intersects (and differs from) every H -translate of \widehat{v} . In particular, \widehat{v} and \widehat{w} are not in the same N_{-1} -orbit, so the image w of \widehat{w} in $\mathcal{C}(S)/N_{-1}$ is different from v . On the other hand, $\widehat{w} \in \text{Lk}(\widehat{\Delta})$, so $w = q(\widehat{w}) \in \text{Lk}(\Delta)$.

Lifting elements by lifting 4-gons: Fix $w \in \text{Lk}(\Delta) - \{v\}$ and recall that $g \in \text{Stab}(\text{Lk}(\Delta)) \cap \text{Stab}(v)$. We have a generalised 4-gon $\Delta \star v = \Delta \star gv, gv \star g\Delta, g\Delta \star w, w \star \Delta$. (The first equality $gv = v$ is because $g \in \text{Stab}(v)$ and the simplex $g\Delta \star w$ exists because $w \in \text{Lk}(\Delta) = g\text{Lk}(\Delta) = \text{Lk}(g\Delta)$ where the second equality holds since $g \in \text{Stab}(\text{Lk}(\Delta))$.) Using the second bullet point, and the the third bullet point, we can lift this generalised 4-gon to $\mathcal{C}(S)$ in such a way that our lift contains $\widehat{\Delta}$ and \widehat{v} as the lifts of Δ and v .

We claim that there exists $\hat{g} \in MCG(S)$ such that $\phi(\hat{g}) = g$ and such that $\hat{g}(\hat{\Delta} \star \hat{v})$ is the lift of $g\Delta \star gv = g\Delta \star v$ appearing in the lift of our generalised 4-gon. Indeed, for any $\hat{g}_0 \in \phi^{-1}(g)$, the simplex $\hat{g}_0(\hat{\Delta} \star \hat{v})$ is a lift of $g\Delta \star v$. Moreover, by the third bullet point, there exists $n_0 \in N_{-1}$ such that $n_0\hat{g}_0(\hat{\Delta} \star \hat{v})$ is the lift of $g\Delta \star v$ appearing in the lift of the generalised 4-gon, and we take $\hat{g} = n_0\hat{g}_0$.

The lifted 4-gon tells us that $\hat{g}(\hat{\Delta} \star \hat{v})$ and $\hat{\Delta} \star \hat{v}$ have a common vertex which is a lift of $v = gv$. The unique lift of v in $\hat{\Delta} \star \hat{v}$ is \hat{v} , so $\hat{v} \in \hat{g}(\hat{\Delta} \star \hat{v})$. We cannot have $\hat{v} \in \hat{g}\hat{\Delta}$, since that would imply $v \in g\Delta$, contradicting that $v = gv \in \text{Lk}(g\Delta)$. Hence $\hat{v} = \hat{g}\hat{v}$, i.e. $\hat{g} \in \text{Stab}_{MCG(S)}(\hat{v})$.

On the other hand, $\hat{\Delta}^{(0)}$ and $\hat{g}\hat{\Delta}^{(0)}$ are both multicurves on S , both disjoint from the curves \hat{v} and \hat{w} , as can be seen by examining the lifted generalised 4-gon. By definition, Y is the complexity-1 component of the complement of $\hat{\Delta}^{(0)}$, so $\hat{g}Y$ is the complexity-1 component of the complement of $\hat{g}\hat{\Delta}^{(0)}$.

Now, since $v \neq w$, the lifts \hat{v} and \hat{w} of v and w have to be distinct. Since \hat{v} and \hat{w} are both in $\text{Lk}(\hat{\Delta})$, both of these vertices, regarded as curves, are in Y (and they are not boundary curves of Y because the boundary curves of Y all belong to $\hat{\Delta}$). Since Y has complexity 1, and \hat{v}, \hat{w} are distinct, these curves fill Y .

On the other hand, \hat{v} and $\hat{g}\hat{w}$ are in the link of $\hat{g}\hat{\Delta}$ (by considering the lifted 4-gon), so, as curves, they belong to $\hat{g}Y$. Indeed, they are intersecting curves in the complement of $\hat{g}\hat{\Delta}^{(0)}$, so they must belong to a common component of that complement, and since they fill a complexity-1 subsurface, the component they belong to must be $\hat{g}Y$. Hence Y is a subsurface of $\hat{g}Y$, i.e. $Y = \hat{g}Y$. Since \hat{g} stabilises Y , it must stabilise ∂Y .

We have produced \hat{g} with $\phi(\hat{g}) = g$ and $\hat{g} \in \text{Stab}_{MCG(S)}(\partial Y) \cap \text{Stab}_{MCG(S)}(\hat{v})$. As explained above, this completes the proof of the claim. \blacksquare

Now we complete the proof of cocompactness for $i = 0$.

Denote by Λ_Δ the quotient of $\text{PStab}(\text{Sat}(\Delta))$ by the kernel of its action on $\text{Lk}(\Delta)$. Both Λ_Δ and the group Γ_Δ from Claim 8.19 can be thought of as subgroups of $\text{Sym}(\text{Lk}(\Delta))$, where we have $\Lambda_\Delta < \Gamma_\Delta$. We now claim that Λ_Δ is in fact a finite-index subgroup of Γ_Δ . This suffices to prove cocompactness since, by item (II) and the ‘‘moreover’’ part of Theorem 6.4, there are finitely many $\text{Stab}(\text{Lk}(\Delta))$ -orbits of edges in $\mathcal{C}(\Delta)$. Hence, the same is true for $\text{PStab}(\text{Sat}(\Delta))$.

To show the claim, we use the fact if two groups $\Lambda < \Gamma$ act faithfully on a non-empty set, with Γ having finite stabilisers and Λ having finitely many orbits, then Λ has finite index in Γ .

Fix a lift $\hat{\Delta}$ of Δ to $\mathcal{C}(S)$, and let Y be the complexity-1 component of the complement in S of the multicurve $\hat{\Delta}^{(0)}$, which exists since $b([\Delta])$ is an orbit of complexity-1 subsurfaces (recall Definition 7.2).

Since Y has complexity 1, the simplex $\hat{\Delta}$ is almost-maximal in $\mathcal{C}(S)$, so Δ is almost-maximal in $\mathcal{C}(S)/N_{-1}$. Thus Claim 8.19 implies that Γ_Δ has finite stabilisers.

On the other hand, the kernel of the action of $\text{Stab}_{MCG}(Y)$ on the complement of Y is contained in $\text{PStab}(\text{Sat}(\hat{\Delta}))$. This kernel acts with finitely many orbits of curves on Y . Since, by item (VII), we have $q(\text{Sat}(\hat{\Delta})) = \text{Sat}(\Delta)$ (which implies that $\text{PStab}(\text{Sat}(\hat{\Delta}))$ maps to $\text{PStab}(\text{Sat}(\Delta))$), we conclude that the same holds for $\text{PStab}(\text{Sat}(\Delta))$. This completes the proof of cocompactness for $i = 0$.

Cocompactness for $i > 0$: Now suppose $i > 0$. Now the simplex Δ of $\mathcal{C}(S)/N_i$ corresponds to a subsurface of complexity $i + 1 > 1$.

Suppose that σ, τ are W^Δ -adjacent maximal simplices of W^Δ . By Lemma 4.1, this occurs if and only if $\sigma \star \Delta$ and $\tau \star \Delta$ are adjacent in W . By the first item in the ‘‘moreover’’ clause of Theorem 6.4, there exists a common codimension-1 face $\eta \star \Delta$ of $\tau \star \Delta$ and $\sigma \star \Delta$. Since $i + 1 > 1$, the simplex Δ is not almost-maximal, so $\eta \neq \emptyset$, and therefore $[\eta \star \Delta] \subsetneq [\Delta]$.

From Theorem 7.1.(II), first and second bullet points, $b([\eta \star \Delta])$ is an orbit of subsurfaces of complexity at most i , so by the third bullet point of the same statement, $\mathcal{C}([\eta \star \Delta])$ is finite. Hence, fixing η , there are only finitely many possibilities for $\tau \star \Delta$ and $\sigma \star \Delta$ and thus finitely many possibilities for the W -edge joining them.

Thus, so show that the action of $\text{PStab}(\text{Sat}(\Delta))$ on W^Δ has finitely many orbits of edges, it suffices to show that the simplex η of $\text{Lk}(\Delta)$ belongs to one of finitely many $\text{PStab}(\text{Sat}(\Delta))$ orbits.

We argue as in the previous case that $\text{PStab}(\text{Sat}(\Delta))$ acts with finitely many orbits of simplices on $\text{Lk}(\Delta)$. By assumption, we can fix a lift $\hat{\Delta}$ of Δ to $\mathcal{C}(S)$. By (VII), $q(\text{Sat}(\hat{\Delta})) = \text{Sat}(\Delta)$, so it suffices to show that $\text{PStab}(\text{Sat}(\hat{\Delta}))$ acts with finitely many orbits of simplices on $\text{Lk}(\hat{\Delta})$. Now, let Y be the union of the non-pants components of the complement in S of the multicurve $\hat{\Delta}^{(0)}$. Then $\text{Lk}(\hat{\Delta})^{(0)}$ is the vertex set of $\mathcal{C}(Y)$, and $\hat{\Delta}^{(0)}$ consists of the boundary curves of Y plus a pants decomposition of $S - Y$. Now, $\text{Stab}_{\text{MCG}(S)}(Y)$ acts on Y with finitely many orbits of multicurves, i.e. it acts with finitely many orbits of simplices on $\text{Lk}(\hat{\Delta})$. Hence the same is true for the finite index subgroup Λ of $\text{Stab}_{\text{MCG}(S)}(Y)$ that stabilises each boundary curve of Y . But since any $\hat{g} \in \Lambda$ can be restricted to Y , there exists $\hat{h} \in \text{MCG}(S)$ that acts as the identity on $S - Y$ and as \hat{g} on Y . Now, $\hat{h} \in \text{PStab}(\text{Sat}(\hat{\Delta}))$, and it follows that the latter group also has finitely many orbits of simplices in $\text{Lk}(\hat{\Delta})$, as required.

Conclusion: Hence, $\text{PStab}(\text{Sat}(\Delta))$ acts properly and cocompactly on the combinatorial HHS $(\text{Lk}(\Delta), W^\Delta)$, and it is therefore a hierarchically hyperbolic group, by Theorem 1.18.

To show that W^Δ and hence $\text{PStab}(\text{Sat}(\Delta))$ are hyperbolic, it suffices to show that this HHS structure on W^Δ has at most one unbounded hyperbolic space, namely $\mathcal{C}(\Delta) = \mathcal{C}_0^\Delta([\emptyset])$ (see Lemma 4.8.(1) for the latter equality). Indeed, let Σ be any non-empty, non-maximal simplex of $\text{Lk}(\Delta)$. By Lemma 4.8.(1) and Lemma 4.1 (saying roughly that links are the same for X as for $\text{Lk}(\Delta)$), we have to bound the diameter of $\mathcal{C}([\Delta \star \Sigma])$. But since Σ is nonempty, we have $[\Delta \star \Sigma] \sqsubset [\Delta]$. This implies that $b([\Delta \star \Sigma])$ has a representative U of complexity at most i , since $b([\Delta])$ has a representative Y with U properly nested in Y , so that the complexity of U is strictly lower than that of Y (recall that b preserves nesting, see item (II)).

We can now apply the third bullet point of item (II) (recall that we are hypothesising this property for the current i) to conclude that $\mathcal{C}([\Delta \star \Sigma])$ is uniformly bounded. Hence W^Δ and $\text{PStab}(\text{Sat}(\Delta))$ are hyperbolic. \square

8.6. Verifying composite properties, and Greendlinger's Lemma. In this subsection only, we are working in the case $i > -1$, i.e., we are doing the inductive step laid out in Convention 8.15. The base case, $i = -1$, was already handled when we verified that the subgroups Γ_Y used to form N_{-1} form a composite rotating family, etc. Many of the additional points in Proposition 8.13 were introduced to enable arguments in this subsection.

Using our inductive hypothesis, we consider the combinatorial hierarchically hyperbolic space $(\mathcal{C}(S)/N_{i-1}, W)$. We let $X = \mathcal{C}(S)/N_{i-1}$.

We let \mathbb{Y}_* denote the collection of equivalence classes $[\Delta]$ of simplices Δ of X so that:

- $b([\Delta]) = N_{i-1}Y$ for some subsurface Y of S of complexity i (recall that b is the bijection between equivalence classes of simplices and N_{i-1} -orbits of subsurfaces),
- $\mathcal{C}([\Delta])$ is unbounded.

In view of (IV), the coloring of \mathfrak{S} descends to a coloring of \mathfrak{S}/N_{i-1} , so that \mathbb{Y}_* is partitioned into finitely many countable families.

Before showing that \mathbb{Y}_* defines a composite projection system, we need a preliminary lemma.

Lemma 8.20. *Let $[\Delta] \in \mathbb{Y}_*$. Then $\text{Lk}(\text{Lk}(\Delta))^{(0)} = \text{Sat}(\Delta)$. In particular, $[\Sigma] \in \mathbb{Y}_*$ satisfies $[\Delta] \perp [\Sigma]$ if and only if $\text{Lk}(\Sigma)^{(0)} \subseteq \text{Sat}(\Delta)$.*

Proof. The containment $\text{Sat}(\Delta) \subseteq \text{Lk}(\text{Lk}(\Delta))$ holds for any simplex in any complex (by definition of the saturation), so it suffices to prove the other containment.

Since $\mathcal{C}(\Delta)$ is unbounded, we can find vertices $v, w \in \text{Lk}(\Delta)$ that are not connected by any path in $\text{Lk}(\Delta)$ of length less than 3. Let x be a vertex of $\text{Lk}(\text{Lk}(\Delta))$. We have a generalized 4-gon $\Delta \star v, v \star x, x \star w, w \star \Delta$, which we can lift to $\mathcal{C}(S)$. We denote by $\hat{\Delta}, \hat{v}$, etc. the various lifts. In terms of curves, the multicurve $\Delta^{(0)}$ is disjoint from the curves \hat{v} and \hat{w} . Moreover, from the hypothesis that v and w are sufficiently far in $\text{Lk}(\Delta)$, we see that $\hat{v} \cup \hat{w}$ fills the (necessarily unique, since $\text{Lk}(\Delta)$ is unbounded) component of the complement of $\Delta^{(0)}$ which is not a pair of pants. Since \hat{x} is disjoint from \hat{v} and \hat{w} , we then see that $\hat{x} \in \text{Sat}(\hat{\Delta})$ (that is, it is part of a pants decomposition of the complement of the surface filled by $\hat{v} \cup \hat{w}$, as is $\Delta^{(0)}$). Since $q(\text{Sat}(\hat{\Delta})) = \text{Sat}(\Delta)$, we have $x \in \text{Sat}(\Delta)$, as required. \square

Lemma 8.21. \mathbb{Y}_* defines a composite projection system.

Proof. We will use the HHS structure associated to (X, W) via Theorem 1.18 and the inductive hypothesis that (X, W) is a combinatorial HHS.

First, since all the elements of \mathbb{Y}_* are associated to equivalence classes of domains in \mathfrak{S} of the same complexity, no nesting can occur and thus any pair of these associated domains is either transverse or orthogonal.

Accordingly, to each $Y \in \mathbb{Y}_*$, we define the set $\text{Act}(Y)$ to be the set of all elements of $\mathbb{Y}_* - \{Y\}$ which are associated to domains transverse to Y . We will almost always work with $\text{Act}(Y) \setminus \{Y\}$. The symmetry in action axiom (Definition 8.2.(CPS6)) is immediate.

Recall that projections in a combinatorial HHS are defined in Definition 1.16. Consider $Y \in \mathbb{Y}_*$ and $W, Z \in \text{Act}(Y) \setminus \{Y\}$, and define $d_Y(W, Z) = d_Y(\rho_Y^W, \rho_Y^Z)$.

Symmetry (Definition 8.2.(CPS1)) and the triangle inequality (Definition 8.2.(CPS2)) follow immediately from the fact that d_Y is a distance function. The Behrstock inequality (Definition 8.2.(CPS3)) follows from Definition 2.1.(8),(4). Properness (Definition 8.2.(CPS4)) follows from the distance formula (Theorem 2.8) and Definition 2.1.(8).

The separation axiom (Definition 8.2.(CPS5)) holds trivially since d_Y is a distance function.

The closeness in inaction axiom (Definition 8.2.(CPS7)) holds, since if $W \notin \text{Act}(Z)$ and $Y \in \text{Act}(Z) \cap \text{Act}(W)$, then W and Z are orthogonal and Y is transverse to both of them. This implies that $d_Y(W, Z)$ is uniformly bounded, by [DHS17, Lemma 1.5].

The finite filling axiom (Definition 8.2.(CPS8)) turns out to be the hardest. We first record the following which allows us to verify the active sets in this axiom via links and saturations:

Claim 8.22. Let $\mathcal{W} \subseteq \mathbb{Y}_*$. Then, for $[\Sigma] \in \mathbb{Y}_*$, we have $[\Sigma] \in \bigcup_{w \in \mathcal{W}} \text{Act}(w)$ if and only if $\text{Lk}(\Sigma) \not\subseteq \bigcap_{[\Delta] \in \mathcal{W}} \text{Sat}(\Delta)$.

Proof. Passing to the complements, we show that $[\Sigma] \in \bigcap_{w \in \mathcal{W}} \text{Act}(w)^c$ if and only if $\text{Lk}(\Sigma) \subseteq \bigcap_{[\Delta] \in \mathcal{W}} \text{Sat}(\Delta)$.

This is just because we have $[\Sigma] \notin \text{Act}([\Delta]) \iff [\Sigma] \perp [\Delta]$, and in view of Lemma 8.20 this is equivalent to $\text{Lk}(\Sigma) \subseteq \text{Sat}(\Delta)$. \blacksquare

In view of the claim it suffices to show that for any $\mathcal{W} \subseteq \mathbb{Y}_*$ there exists a finite subcollection $\{[\Delta_1], \dots, [\Delta_n]\} \subseteq \mathcal{W}$ so that $\bigcap_j \text{Sat}(\Delta_j) = \bigcap_{[\Delta] \in \mathcal{W}} \text{Sat}(\Delta)$. This readily follows from the following claim.

Claim 8.23. There does not exist an infinite sequence $[\Delta_1], [\Delta_2], \dots$ in \mathbb{Y}_* so that

$$\bigcap_{j \leq n+1} \text{Sat}(\Delta_j) \subsetneq \bigcap_{j \leq n} \text{Sat}(\Delta_j)$$

for all n .

Proof. First, we show that for every $[\Delta] \in \mathbb{Y}_*$ there are simplices $\Pi(\Delta)$ and $\Pi'(\Delta)$ so that we have $\text{Sat}(\Delta) = \text{Lk}(\Pi(\Delta)) \star \Pi'(\Delta)$.

We record the following observation about curve graphs:

Claim 8.24. *Let Σ be a simplex in $\mathcal{C}(S)$. Then there exist simplices Σ', Σ'' of $\mathcal{C}(S)$ such that $\text{Sat}(\Sigma) = \text{Lk}(\Sigma') \star \Sigma''$.*

Proof. First of all, the vertex set of $\text{Lk}(\Sigma)$ consists of all essential curves in the complement of Σ , that is, it coincides with the vertex set of the curve graph of the (possibly disconnected) subsurface Y of S consisting of all components of the complement of Σ which are not pairs of pants. Moreover, Σ is the simplex whose vertex set consists of

- all curves of S isotopic to a boundary component of Y in S (note that two boundary components of Y can be isotopic in S), and
- a pants decomposition of the union Z of the non-annular components of the complement of Y .

We let Σ'' be the simplex with vertex set the curves as in the first item.

Any simplex with the same link as Σ admits the same description, and moreover any essential curve in Z can be completed to a pants decomposition. This implies that $\text{Sat}(\Sigma)$ consists of the vertex set of Σ'' together with all essential curves of Z . We can then pick Σ' to be any simplex whose link has vertex set the set of essential curves of Z ; namely, Σ' has vertex set consisting of all curves of S isotopic to a boundary component of Z in S , together with a pants decomposition of Y . ■

Consider any lift $\widehat{\Delta}$ of Δ to $\mathcal{C}(S)$. Claim 8.24 implies that $\text{Sat}(\widehat{\Delta}) = \text{Lk}(\Delta') \star \Delta''$ for some simplices Δ', Δ'' of $\mathcal{C}(S)$.

Also notice that $q(\text{Lk}(\Delta')) = \text{Lk}(q(\Delta'))$. Indeed, the containment “ \subseteq ” is clear, and the other containment follows from the fact that for any vertex v in $\text{Lk}(q(\Delta'))$ we can lift $v \star q(\Delta')$ to a simplex containing Δ' (by existence and uniqueness of the orbit of lifts of simplices, item (VII)). In particular, we get that $\text{Sat}(\Delta) = q(\text{Sat}(\widehat{\Delta})) = \text{Lk}(q(\Delta')) \star q(\Delta'')$, and we are done.

Now suppose, for a contradiction, that we have a sequence $[\Delta_1], [\Delta_2], \dots$ as in the statement of Claim 8.23. By induction, suppose we proved that $I_n = \bigcap_{j \leq n} \text{Sat}(\Delta_j) = \text{Lk}(\Pi_n) \star \Pi'_n$ for some simplices Π_n, Π'_n with $\Pi_n \subseteq \Pi_{n+1}$. The base case $n = 1$ is given by the argument above.

Then $I_{n+1} = I_n \cap \text{Lk}(\Pi(\Delta_{n+1})) \star \Pi'(\Delta_{n+1}) = \text{Lk}(\Pi_n) \star \Pi'_n \cap \text{Lk}(\Pi(\Delta_{n+1})) \star \Pi'(\Delta_{n+1})$ is readily seen to be of the required form in view of condition B of Theorem 6.4, which holds for X by inductive hypothesis (Theorem 7.1.II). Since $I_{n+1} \subsetneq I_n$, for each n we have either $\Pi'_{n+1} \subsetneq \Pi'_n$ or, if not, $\Pi_n \subsetneq \Pi_{n+1}$. This is easily seen to imply that the Π_n have arbitrarily many vertices, contradicting the finite dimensionality of X . This proves the claim. ■

This completes the proof of Lemma 8.21. □

Lemma 8.25. *For every $\theta > 0$ the following holds. For any $Y = [\Delta] \in \mathbb{Y}_*$ we can choose a finite-index subgroup $\Gamma_Y^\theta < \text{PStab}(\text{Sat}(\Delta))$ contained in $H/N_{i-1} < \bar{G}_{i-1} = \text{MCG}(S)/N_{i-1}$ so that the subgroups Γ_Y^θ form a composite rotating family on the composite projection system \mathbb{Y}_* . Moreover,*

$$\min\{d_Y(x, \gamma x) : Y \in \mathbb{Y}_*, \gamma \in \Gamma_Y - \{1\}, x \in \mathcal{C}(Y)\} > \theta.$$

Proof. In Lemma 8.21 we established that \mathbb{Y}_* is a composite projection system.

Definition 8.3.(CRF1): Roughly, this says that Γ_Y^θ needs to be an infinite group acting on \mathbb{Y}_* fixing Y and, for any R , having only finitely many elements that move some point of $\mathcal{C}(Y)$ at most distance R .

Any sufficiently deep finite-index normal subgroup Γ_Y^θ satisfies this by inductive hypothesis (IX) (which guarantees that Γ_Y^θ acts properly on an unbounded graph, so that it must be infinite).

Definition 8.3.(CRF2): It suffices to choose the Γ_Y^θ as follows. For each $Y = [\Delta]$ in a given set of representatives of \bar{G}_{i-1} -orbits in \mathbb{Y}_* , choose a normal subgroup Γ_Y^θ of $\text{PStab}(\text{Sat}(\Delta))$, and extend the choice to all Y .

Definition 8.3.(CRF3): For this, we use that Γ_Y^θ is contained in H/N_{i-1} . We first make a preliminary claim.

Claim 8.26. *Let $[\Delta] \perp [\Sigma]$, and let $g \in \text{PStab}(\text{Sat}(\Delta))$, $h \in \text{PStab}(\text{Sat}(\Sigma))$. Then $gh(x) = hg(x)$ for all $x \in \text{Sat}(\Delta) \cup \text{Lk}(\Delta)$, and $gh(x) = hg(x) = g(x)$ for all $x \in \text{Lk}(\Delta)$.*

Proof. First, we prove that $h(\text{Sat}(\Delta)) \subseteq \text{Sat}(\Delta)$. The same argument also shows $g(\text{Sat}(\Sigma)) \subseteq \text{Sat}(\Sigma)$.

Let $v \in \text{Sat}(\Delta)$, that is, v is a vertex of a simplex Δ' with $\text{Lk}(\Delta') = \text{Lk}(\Delta)$. We have to show that $hv \in \text{Sat}(\Delta)$. We have that hv is a vertex of $h\Delta'$, so we have $hv \in \text{Sat}(\Delta)$ provided that $\text{Lk}(h\Delta') = \text{Lk}(\Delta')$. But $\text{Lk}(\Delta') = \text{Lk}(\Delta) \subseteq \text{Sat}(\Sigma)$ by Lemma 8.20, and since $h \in \text{PStab}(\text{Sat}(\Sigma))$, we have that h fixes $\text{Lk}(\Delta')$. Hence, $\text{Lk}(h\Delta') = \text{Lk}(\Delta')$, as required.

Now let $x \in \text{Sat}(\Delta)$. Then $h(x) \in \text{Sat}(\Delta)$, so that $g(h(x)) = h(x)$. On the other hand, $g(x) = x$, so that $h(g(x)) = h(x)$. This shows that $gh(x) = hg(x)$ for all $x \in \text{Sat}(\Delta)$. Similarly, $gh(x) = hg(x) = g(x)$ for all $x \in \text{Sat}(\Sigma)$. To conclude, just notice that $\text{Sat}(\Delta) \cup \text{Lk}(\Delta) \subseteq \text{Sat}(\Delta) \cup \text{Sat}(\Sigma)$ since $\text{Lk}(\Delta) \subseteq \text{Sat}(\Sigma)$ by Lemma 8.20. \blacksquare

Suppose that $X \notin \text{Act}(Z)$, where $X = [\Sigma]$, $Z = [\Delta]$; we have to show that Γ_X^θ commutes with Γ_Z^θ . Then $X \perp Z$, since there is no nesting relation between distinct elements of \mathbb{Y}_* . For $g \in \text{PStab}(\text{Sat}(\Delta))$, $h \in \text{PStab}(\text{Sat}(\Sigma))$, by the claim above we have that gh and hg act in the same way on $\text{Sat}(\Delta) \cup \text{Lk}(\Delta)$, so we are done once we prove the following:

Claim 8.27. *Let $a \in H/N_{i-1} \leq \bar{G}_{i-1}$ act trivially on $\text{Sat}(\Delta) \cup \text{Lk}(\Delta)$. Then a is the identity.*

Proof. Let $\phi : \text{MCG}(S) \rightarrow \bar{G}_{i-1}$ be the quotient map. Let P be a maximal simplex of X containing Δ . The vertex set of P is contained in $\text{Sat}(\Delta) \cup \text{Lk}(\Delta)$, so a fixes P pointwise. We can lift P to a simplex \hat{P} of $\mathcal{C}(S)$, and since there is a unique N_{i-1} -orbit of such lifts, there is $\hat{a} \in H$ so that $\phi(\hat{a}) = a$ and \hat{a} fixes \hat{P} pointwise.

We claim that \hat{P} is a maximal simplex of $\mathcal{C}(S)$. Indeed, suppose to the contrary that \hat{P} is properly contained in a simplex \hat{P}' . Then, since X is simplicial, \hat{P}' projects to a simplex P' of X strictly containing P , a contradiction. So, \hat{P} is maximal.

In terms of curves, \hat{a} fixes the pants decomposition $\hat{P}^{(0)}$, so that \hat{a} is a product of powers of Dehn twists around the curves of $\hat{P}^{(0)}$. Since $N_{-1} \cap \langle \tau \rangle = H \cap \langle \tau \rangle$ and $N_{-1} < N_{i-1} < H$ (by induction hypothesis (IV)), we see that $\hat{a} \in N_{i-1}$, so that a is the identity, as required. \blacksquare

Definition 8.3.(CRF4): Since Γ_Y^θ acts properly on $\mathcal{C}(Y)$, and there are finitely many \bar{G}_{i-1} -orbits in \mathbb{Y}_* , choosing sufficiently deep subgroups for orbit representatives as in the proof of (CRF2) ensures the required property. The ‘‘moreover’’ part follows similarly from $\text{PStab}(\text{Sat}(\Delta))$ acting properly and cocompactly on $\mathcal{C}(\Delta)$ (inductive hypothesis (IX)), so that it suffices to choose deep subgroups for orbit representatives. \square

Lemma 8.28 (‘‘Greendlinger’’). *There exists a diverging function \mathfrak{T} so that the following holds. Let $\theta > 0$ and let $\Gamma_Y^\theta = \Gamma_Y$ be as in Lemma 8.25.*

Then, for $N = \langle\langle \{\Gamma_Y^\theta\} \rangle\rangle$, the following holds. There is a well-ordered set \mathfrak{C} , and an assignment $N \ni \gamma \mapsto c(\gamma) \in \mathfrak{C}$, with $c(1)$ minimal in \mathfrak{C} and the following additional properties.

For all $\gamma \in N - \{1\}$ and all simplices Δ of X , there is $Y \in \mathbb{Y}_$ and $\gamma_Y \in \Gamma_Y^\theta$ so that $c(\gamma_Y \gamma) < c(\gamma)$ and either*

- $\Delta \subseteq \text{Fix}(\Gamma_Y)$, or $\gamma \Delta \subseteq \text{Fix}(\Gamma_Y)$, or
- $d_Y(\Delta, \gamma \Delta) > \mathfrak{T}(\theta)$ (and the quantity is defined).

Remark 8.29. For $Y \in \mathbb{Y}_*$, $\text{Sat}(Y)$ is fixed pointwise by Γ_Y , so that $\pi_Y(\Delta)$ is defined for any simplex Δ of X not contained in $\text{Fix}(\Gamma_Y)$.

Proof of Lemma 8.28. Independently of θ , we now choose, for each simplex Δ of X and each \mathbb{Y}_k , some $Y_k^\Delta \in \mathbb{Y}_k$. We make the choice as follows.

First, we choose H/N_{i-1} -orbit-representatives Δ_j of simplices, and make the choice for each of them. Then, for any other simplex Δ , we choose a group element $g \in H/N_{i-1}$ mapping the suitable orbit representative, say Δ_j , to the given simplex Δ , and set $Y_k^\Delta = gY_k^{\Delta_j}$ (since H/N_{i-1} preserves the colors, this is a well-defined element of \mathbb{Y}_k). Set

$$D = \max_{j,k} \sup \{d_Y(\Delta_j, Y_k^{\Delta_j}) : Y \in \mathbb{Y}_k, \Delta_j \notin \text{Sat}(\Delta_j)\},$$

which is finite since there are finitely many colors and finitely many orbits (since H/N_{i-1} has finite index in \bar{G}_{i-1} , and \bar{G}_{i-1} acts cocompactly on X), and the distance formula takes finite values.

We refer the reader to [DHS21, Theorem 4.1] (which follows by combining Dahmani's construction in [Dah18, Section 2.4.2] with [Dah18, Prop. 2.13, Lemma 2.16, Lemma 2.17]), where the subgroup N in Lemma 8.28 is described as an increasing union of subgroups $N = \bigcup_\alpha N_\alpha$, where α varies over countable ordinals.

For $\gamma \in N$, we denote by $\alpha(\gamma)$ the smallest ordinal such that γ is conjugate into $N_{\alpha(\gamma)}$. Two properties of $\alpha(\gamma)$ observed in [DHS21] are that $\alpha(\gamma)$ is never a limit ordinal and $\alpha(\gamma) = 0$ if and only if $\gamma = 1$. Moreover, N_α for α a successor ordinal has a certain amalgamated product decomposition used in [DHS21, Definition 4.3] to define $n(\gamma)$, which is the length of the cyclic normal form in $N_{\alpha(\gamma)}$ for the conjugacy class of γ . Following [DHS21], we consider the complexity $c(\gamma)$ given by $(\alpha(\gamma), n(\gamma))$, ordered lexicographically.

After conjugation by a suitable element h and replacing Δ with $h^{-1}\Delta$, one can assume that $\gamma \in N_{\alpha(\gamma)}$.

Let $Y = v$ and $\gamma_Y = \gamma_v$ be as in [DHS21, Proposition 4.5, fourth and second bullet] with $\nu = Y_{j(\gamma)}^\Delta$, so that $\gamma_Y\gamma$ has shorter cyclic normal form. We argue that this implies that $c(\gamma_Y\gamma) < c(\gamma)$. Indeed, if the length of the new normal form is still greater than 2, we have $\alpha(\gamma_Y\gamma) = \alpha(\gamma)$ but the second coordinates satisfy $n(\gamma_Y\gamma) < n(\gamma)$. If not, the length is reduced to 1 (or 0) and $\gamma_Y\gamma$ is actually conjugate into $N_{\alpha(\gamma)-1}$, so that $\alpha(\gamma_Y\gamma) < \alpha(\gamma)$. In either case we have $c(\gamma_Y\gamma) < c(\gamma)$, as required.

We are left to show that one of the alternatives applies. If $\Delta, \gamma\Delta \notin \text{Fix}(\Gamma_Y)$ (otherwise we are done), then we have that

$$d_Y(\Delta, \gamma\Delta) \geq d_Y(Y_{j(\gamma)}^\Delta, \gamma Y_{j(\gamma)}^\Delta) - 2D,$$

as required; we can take $\mathfrak{T}(\theta) = \theta - 2D$. □

Notation 8.30. We now choose θ sufficiently large to ensure that

- $\mathfrak{T}(\theta) > (6\text{Gen}(S) + 2p(S) - 3)C_{i-1}$, for C_{i-1} as in Proposition 8.13.(VIII),
- $\mathfrak{T}(\theta) > \max\{d_{\mathcal{C}(Y)}(fx, gx) : Y \in \mathbb{Y}^*, f, g \in F, x \in X\}$,
- $\mathfrak{T}(\theta) > \kappa + C_{i-1}$, for κ as in Lemma 2.10 applied to the image of Q in \bar{G}_{i-1} .

We set $\Gamma_Y = \Gamma_Y^\theta$ as in Lemma 8.28, and set $N = \langle\langle \{\Gamma_Y\} \rangle\rangle$. We emphasize that N is the kernel of the quotient $\bar{G}_{i-1} \rightarrow \bar{G}_i$, and is not the same as N_i .

8.7. Lifting. In this subsection, we are back in the setting of Convention 8.15. In other words, either $i = -1$ (base case) or $i > -1$ (inductive step). The goal in this section is to lift generalized m -gons from X/N to X . This will typically be used in conjunction with our inductive hypothesis (see Proposition 8.13), to lift all the way back up to $\mathcal{C}(S)$.

(For clarity, in the proposition we recall parts of the definitions of lifts.)

Proposition 8.31 (Lifting generalized m -gons through q). *For every $m \leq 6\text{Gen}(S) + 2p(S) - 3$ the following hold.*

- (1) *For every simplex Δ of X/N , together with an order (v_0, \dots, v_k) on the vertices there exists a unique N -orbit of simplices Σ in X , each with an order (w_0, \dots, w_k) on its vertices, such that Σ is a lift of the ordered simplex Δ (meaning that $q(w_j) = v_j$ for $0 \leq j \leq k$).*
- (2) *Given a simplex Δ in X/N , a lift Σ of Δ in X , and a geodesic γ in the link of Δ , we have that γ can be lifted to a geodesic in the link of Σ .*
- (3) *Any generalized m -gon $\tau = \tau_0, \dots, \tau_{m-1}$ in X/N can be lifted to a generalized m -gon $\tau'_0, \dots, \tau'_{m-1}$ in X so that τ'_j is **type S/type G** if and only if τ_j is, and if τ_j is a geodesic in $\text{Lk}(\Delta_j)$ then τ'_j is a geodesic in $\text{Lk}(\Sigma_j)$ for some lift Σ_j of Δ_j .*

Proof. Consider a generalized m -gon $\tau = \tau_0, \dots, \tau_{m-1}$. If τ_j is of **type G**, and Δ_j is the corresponding simplex, let d_j be the number of vertices of Δ_j . If τ_j is **type S**, let d_j be the number of vertices of τ_j minus 2. Finally, set $d(\tau) = \max_j \{d_j\}$.

Good lifts: We say that τ_j of **type S** has a *good lift* if it can be lifted, and that τ_j of **type G** has a *good lift* if it can be lifted to a geodesic in the link of a lift of the corresponding simplex.

Structure of the proof: The proofs of the 3 items are interlaced, and more specifically we will prove the following claims, for any $d \geq 2$:

- (a) If Item 1 holds whenever Δ has at most d vertices, then Item 2 holds whenever Δ has at most $d - 2$ vertices.
- (b) If Item 1 holds whenever Δ has at most d vertices, then Item 3 holds whenever $d(\tau) \leq d - 2$.
- (c) If Item 3 holds whenever $d(\tau) \leq d - 2$ and Item 1 holds whenever Δ has at most d vertices, then Item 1 holds whenever Δ has at most $d + 1$ vertices.

(a) – (c) feed into an induction that proves all 3 items, with the base case specified below and corresponding to Item 1 for Δ with at most two vertices (from which by (a) and (b) we can deduce Items 2 and 3 for Δ with at most 2 vertices, which in turn by (c) yield Item 1 for Δ with at most 3 vertices, etc.).

Base case: The base cases are the following:

- For each vertex \bar{v} of X/N , there is a unique N -orbit of vertices $v \in X$ such that $q(v) = \bar{v}$, just because of how X/N is defined.
- For each edge \bar{e} of X/N , with endpoints \bar{v}, \bar{w} , the fact that N acts simplicially on X (in such a way that stabilizers of simplices fix them pointwise) implies that there is a unique orbit of edges e of X , with endpoints v, w , such that $q(v) = \bar{v}, q(w) = \bar{w}$.

We now prove (a) – (c), keeping (b) last since it is the hardest.

Proof of (a): Let v_0, v_1, \dots be the sequence of vertices along γ . Then $\Delta \star v_0$ is a simplex of X/N , which can be lifted to a simplex $\Sigma' \star v'_0$ of X . By the uniqueness clause of Item 1, up to applying an element of N , we can assume $\Sigma' = \Sigma$, so that we lifted the first vertex of the geodesic in the appropriate link. Suppose that we lifted the sequence of vertices v_0, \dots, v_k to v'_0, \dots, v'_k . Since $\Delta \star v_k \star v_{k+1}$ is a simplex of X/N , it can be lifted to X , and similarly to the argument for v_0 , the lift can be chosen to be of the form $\Sigma \star v'_k \star v'_{k+1}$, so that we can lift v_k to v'_k . Inductively, we can lift the whole geodesic γ .

Proof of (c): Consider a simplex Δ with at most $d + 1$ vertices. By the base case of the induction, we can assume that Δ has at least 3 vertices, so $\Delta = \Delta' \star v \star w$, for some non-empty simplex Δ' . We can think of Δ as a generalized 3-gon $\Delta' \star v, v \star w, \Delta' \star w$, with all sides of **type S**, and each $d_j \leq d - 2$. Hence, the 3-gon can be lifted, which also provides a lift of Δ .

For the uniqueness part, suppose that Δ , endowed with an order on the vertices, has two lifts not in the same N -orbit, and consider an arbitrary codimension-1 face Δ'' , with opposite vertex u . We know that Δ'' has a unique orbit of lifts, so that we see that there exist two

lifts of Δ of the form $\Sigma \star u_1, \Sigma \star u_2$ that are not in the same N -orbit, but u_1, u_2 are (as two lifts of Δ can be made to coincide on the lifts of Δ'' that they contain). But this means that there exists a vertex t in Σ so that $t \star u_1, t \star u_2$ are not in the same N -orbit. However, their endpoints are, and hence either the projections of the edges to X/N yield loops in X/N , or the two projections together form a bigon. In both cases we contradict Item 3 (for $m = 1$ or $m = 2$), since the loops/bigon cannot be lifted to X , which is a simplicial graph. For later purposes, we note that we also just proved:

Lemma 8.32. *X/N is a simplicial graph.*

Proof of (b): By (a) we know that Item 2 holds whenever Δ has at most $d - 2$ vertices, and from this we see that if we have a generalized m -gon $\tau = \tau_0, \dots, \tau_{m-1}$ with $d(\tau) \leq d - 2$ then we can lift it to an “open” generalized m -gon $\eta = \eta_0, \dots, \eta_{m-1}$, which is defined in the same way as a generalized m -gon except that we do not require $\eta_{m-1}^+ = \eta_0^-$. We call η_0^- and η_{m-1}^+ the initial and terminal marking. However, we still have that η_{m-1}^+ and η_0^- are in the same N -orbit, since they are both lifts of τ_0^- , say $\eta_{m-1}^+ = g\eta_0^-$, for some $g \in N$. Assume that among all possible choices of lifts and elements $g \in N$ with $g\eta_0^- = \eta_{m-1}^+$, we picked one that minimizes the complexity $c(g)$ from Lemma 8.28 or Lemma 8.10. If $g\eta_0^- = \eta_0^-$, we are done.

Otherwise, we use Lemma 8.28 or Lemma 8.10 (depending on whether $i > -1$ or not) to change the lift; let $Y \in \mathbb{Y}_*$ and $\gamma_Y \in \Gamma_Y$ be as in the lemma for $\gamma = g$ and $\Delta = \eta_0^-$. In particular, the “complexities” satisfy $c(\gamma_Y g) < c(g)$.

First, suppose $\eta_0^- \subseteq \text{Fix}(\Gamma_Y)$. Then we can apply γ_Y to all the lifts, contradicting minimality of $c(g)$ since $(\gamma_Y g)\gamma_Y \eta_0^- = \gamma_Y g\eta_0^- = \gamma_Y \eta_{m-1}^+$ ($\gamma_Y \eta_0^-$ and $\gamma_Y \eta_{m-1}^+$ being the new initial and terminal markings).

Second, suppose $g\eta_0^- = \eta_{m-1}^+ \subseteq \text{Fix}(\Gamma_Y)$. Then $\eta_{m-1}^+ = \gamma_Y g\eta_0^-$, again contradicting minimality of $c(g)$ (without even changing the lifts).

Lastly, suppose $d_Y(\eta_0^-, \eta_{m-1}^+) > \mathfrak{T}(\theta)$. Assume that some η_k^+ is contained in $\text{Fix}(\Gamma_Y)$. In this case, we can replace the lifts η_j for $j > k$ with the lifts $\gamma_Y \eta_j$, and get a new open generalized m -gon, $\gamma_Y g$ maps the initial marking to the terminal marking, contradicting minimality of $c(g)$.

Suppose instead that for each k there is $x_k \in \eta_k^+$ and $x_{-1} \in \eta_0^-$ so that $x_k \notin \text{Fix}(\Gamma_Y)$. Then we have $\mathfrak{T}(\theta) < d_Y(x_{-1}, gx_{-1}) \leq \sum d_Y(x_j, x_{j+1})$. By the choice of θ in Notation 8.30 large compared to the BGI constant, we see that there must be a type G η_k , with corresponding simplex Δ_k , and $v_k \in \eta_k$ so that $\Delta_k \cup \{v_k\}$ is contained $\text{Fix}(\Gamma_Y)$ (here we are using Proposition 7.1.(VIII) together with Remark 8.29). We can replace the lifts η_j for $j > k$ with the lifts $\gamma_Y \eta_j$, as well as replacing the terminal path of η_k starting at v_k also with $\gamma_Y \eta_k$. Then, we conclude as before. \square

8.8. Supporting lemmas for Theorem 7.1. Recall that $X = \mathcal{C}(S)/N_{i-1}$, and N is the kernel of the quotient $\bar{G}_{i-1} \rightarrow \bar{G}_i$, so that the map $\mathcal{C}(S) \rightarrow \mathcal{C}(S)/N_i = X/N$ factors as $\mathcal{C}(S) \rightarrow X \rightarrow X/N$.

Convention 8.33. Since generalized m -gons can be lifted from X/N to X , and from X to $\mathcal{C}(S)$, they can be lifted from X/N to $\mathcal{C}(S)$. All the lifts in this subsection are of the latter type. When we cite Proposition 8.31 in this subsection, we will always use it together with lifting from X to $\mathcal{C}(S)$.

Lemma 8.34. *For every simplex Δ of X/N , and any simplex Σ of $\mathcal{C}(S)$ that is a lift of Δ we have that $q(\text{Lk}(\Sigma)) = \text{Lk}(\Delta)$. Moreover, if $i = -1$, then $\text{Lk}(\Delta) = \text{Lk}(\Sigma)/(N \cap \text{Stab}(\Sigma))$.*

Proof. Fix a lift Σ of Δ . Given a vertex v of $\text{Lk}(\Delta)$, we can lift $\Delta \star v$ to a simplex in $\mathcal{C}(S)$, and since there is a unique orbit of lifts of Δ , we can choose the lift to contain Σ , and therefore be of the form $\Sigma \star \hat{v}$. Then $q(\hat{v}) = v$, and similarly we can show that edges of $\text{Lk}(\Delta)$ arise from edges in the link of Σ .

Let us now prove the moreover part. We have to show that if two vertices v, w of $\text{Lk}(\Sigma)$ are N -translates, then they are $(N \cap \text{Stab}(\Sigma))$ -translates. This is because v, w being N -translates implies that $\Sigma \star v, \Sigma \star w$ are lifts of the same simplex, and in particular the simplices are in the same N -orbit. This implies that there exists $h \in N$ that stabilizes Σ and maps v to w , as required. \square

Remark 8.35 (Connected links). A consequence of Lemma 8.34 is that all simplices of X/N have connected links except co-dimension 1 faces in maximal simplices (since this holds in $\mathcal{C}(S)$).

Definition 8.36 (Approach path). An *approach path* in X/N is a sequence of paths $\gamma_1, \dots, \gamma_m$ and simplices $\Delta_1, \dots, \Delta_{m+1}$, so that

- the endpoint of γ_j is the starting point of γ_{j+1} ,
- γ_j is a geodesic in the link of a (possibly empty) simplex Δ_j ,
- the endpoint of γ_m is in the link of Δ_{m+1} ,
- Δ_j is a proper sub-simplex of Δ_{j+1} .

We say that the approach path starts (resp. ends) at x if x is the starting point of γ_0 (resp. endpoint of γ_m). We call Δ_{m+1} the *terminal simplex*, and resulting path the concatenation of the γ_j .

Remark 8.37. m as above is bounded by $3\text{Gen}(S) + p(S) - 3$ (that is, the complexity of S), which equals the maximal number of vertices of a simplex in X/N by Proposition 8.31.1.

Lemma 8.38. *Given a vertex $x \in X/N$ and a simplex Δ of X/N so that $x \notin \text{Sat}(\Delta)$, there exists an approach path that starts at x and has terminal simplex Δ' so that $[\Delta'] = [\Delta]$.*

Proof. Consider x and Δ as in the statement, and pick any geodesic γ'_1 in $X/N = \text{Lk}(\emptyset)$ that intersects $\text{Sat}(\Delta)$ only at its endpoint v_1 ; notice that $[\Delta] \sqsubseteq [v_1]$. Consider the subgeodesic γ_1 of γ'_1 obtained removing the last edge, and set $\Delta_2 = v_1$. If $[\Delta_2] = [\Delta]$, we are done.

Otherwise, inductively, suppose that we have an approach path $\gamma_1, \dots, \gamma_j$ starting at x and terminating at $x_j \in \text{Lk}(\Delta_{j+1})$ and $[\Delta] \sqsubseteq [\Delta_{j+1}]$. In particular, Δ_{j+1} has connected link (see Remark 8.35), so that we can consider a shortest geodesic γ'_{j+1} in $\text{Lk}(\Delta_{j+1})$ to $\text{Lk}(\Delta)$. If γ'_{j+1} does not intersect $\text{Sat}(\Delta)$, we conclude by setting $\gamma_{j+1} = \gamma'_{j+1}$, and otherwise we can find an initial subgeodesic γ_{j+1} of γ'_{j+1} that intersects $\text{Sat}(\Delta)$ only at its endpoint v_{j+1} . We set $\Delta_{j+2} = \Delta_j \star v_{j+1}$, notice that we are done if $[\Delta_{j+2}] = [\Delta]$, and otherwise reapply the inductive procedure. This terminates by Remark 8.37. \square

Lemma 8.39. *Let Δ be a simplex of X/N , and endow $\text{Lk}(\Delta)$ with any metric induced by adding finitely many $\text{Stab}(\text{Lk}(\Delta))$ -orbits of edges, as in Definition 6.2. Let $\rho : X/N - \text{Sat}(\Delta) \rightarrow \text{Lk}(\Delta)$ map the vertex x to the endpoint of an arbitrary approach path that starts at x and has terminal simplex equivalent to Δ . Then ρ is coarsely Lipschitz, and ρ restricts to the identity on $\text{Lk}(\Delta)$.*

Proof. The fact that ρ restricts to the identity on $\text{Lk}(\Delta)$ is immediate from the definition of ρ .

It is enough to show that adjacent vertices v, w map uniformly close under ρ . To this end, form a $(2m+3)$ -gon, where $m \leq 3\text{Gen}(S) + p(S) - 3$, using approach paths starting at v and w , a single edge from v to w , and simplices $\Delta \star \rho(v), \Delta \star \rho(w)$.

Consider a lift of this $(2m+3)$ -gon (notice that $2m+3 \leq 6\text{Gen}(S) + 2p(S) - 3$), which contains a lift Σ of Δ . Notice that all sides of **type G** of the lifted $(2m+3)$ -gon are geodesics in links of simplices of $\mathcal{C}(S)$ with strictly fewer vertices than Σ , and in particular they are geodesics in links of simplices not equivalent to Σ . Moreover, no vertex on a side of **type G** is in $\text{Sat}(\Sigma)$, since the image of a vertex of $\text{Sat}(\Sigma)$ is in $\text{Sat}(\Delta)$ (this is because such vertex is contained in a simplex Σ' in $\text{Sat}(\Sigma)$ equivalent to Σ , and using Lemma 8.34 we see that $q(\Sigma')$ is equivalent to Δ). In particular, by the Bounded Geodesic Image Theorem [MM00, Theorem 3.1] (recall that we are in $\mathcal{C}(S)$) all sides have bounded subsurface projection to $\text{Lk}(\Sigma)$, providing a bound

on the distance between the lifts of $\rho(v)$ and $\rho(w)$ and hence on the distance between $\rho(v)$ and $\rho(w)$. This concludes the proof. \square

8.9. Checking hierarchical hyperbolicity. We now have all the tools to prove the main conclusion of Theorem 7.1, namely hierarchical hyperbolicity of the quotient group. We apply Convention 8.33 about lifting to $\mathcal{C}(S)$ rather than to X in this subsection as well, with exceptions clearly marked.

Proof of Theorem 7.1.(II). We check that the action of $\bar{G}_i = \bar{G}_{i-1}/N$ on X/N satisfies the hypotheses of Theorem 6.4. That is, we take N_i to be the kernel of the map $MCG(S) \rightarrow \bar{G}_i$.

First, X/N is simplicial by Lemma 8.32, and \bar{G}_{i-1}/N acts on X/N by simplicial automorphisms, and the action is cocompact since the action of $MCG(S)$ on $\mathcal{C}(S)$ is cocompact.

In view of Proposition 8.31.1, any maximal simplex Δ of X/N is the projection of some maximal simplex Σ of $\mathcal{C}(S)$, which represents a unique N_i -orbit. Since Σ is maximal, $\text{Stab}_{MCG(S)}(\Sigma)$ contains a finite-index abelian subgroup A generated by powers by Dehn twists around the curves corresponding to the vertices of Σ . Now, if $\bar{g} \in \text{Stab}_{MCG(S)/N_i}(\Delta)$, let g represent the left coset \bar{g} of N_i , so $g\Sigma = h\Sigma$ for some $h \in N_i$. In other words, $\bar{g} = \bar{g}'$ for some $g' \in \text{Stab}_{MCG(S)}(\Sigma)$, i.e., $\text{Stab}_{MCG(S)/N_i}(\Delta)$ is contained in the image of $\text{Stab}_{MCG(S)}(\Sigma)$. Now, A has finite image in $MCG(S)/N_i$, so since g' represents one of finitely many cosets of A in $\text{Stab}_{MCG(S)}(\Sigma)$, we see that $\text{Stab}_{MCG(S)/N_i}(\Delta)$ is finite, as required. It remains to check conditions (A)–(C) from Theorem 6.4.

Proof of (A): Let Δ be a non-maximal simplex of X/N . If Δ is not almost-maximal, then lift Δ to a simplex Σ of $\mathcal{C}(S)$, note that $\text{Lk}(\Sigma)$ is hyperbolic (since it is either a non-trivial join or the curve graph of a surface of complexity at least 2), and deduce that $\text{Lk}(\Delta)$ is hyperbolic since we can lift triangles in $\text{Lk}(\Delta)$ by Proposition 8.31.3.

Almost-maximal Δ case: If Δ is almost-maximal, we divide into cases according to the value of i . For $i = -1$, the group \bar{G}_{-1} is obtained as a quotient by powers of Dehn twists. The quotient \bar{G}_0 is equal to \bar{G}_{-1} (see Remark 8.16). For $i > -1$, we are taking further proper quotients.

The case $i = -1$: Let $i = -1$. Let Γ_Δ denote the image of $\text{Stab}(\text{Lk}(\Delta))$ in the group of permutations of $\text{Lk}(\Delta)$. We will show that Γ_Δ acts with finitely many orbits and with finite stabilisers on $\text{Lk}(\Delta)$, and that Γ_Δ is a hyperbolic group. Together, these facts imply that $\text{Lk}(\Delta)$ is a hyperbolic $\text{Stab}(\text{Lk}(\Delta))$ -space, as required.

We first show that Γ_Δ acts with finite stabilisers. Since we are in the case $i = -1$ and Δ is almost-maximal, this will follow from Claim 8.19 once we check that the three itemised assumptions in that claim hold in the present situation. The first assumption is that Proposition 8.13.(IV) holds for $i = -1$, but this is immediate from the choice of N_{-1} in Notation 8.11. The second assumption is that Proposition 8.13.(VI) holds; this is because of Proposition 8.31, which also implies that the first part of Proposition 8.13.(VII) holds. So, we can apply Claim 8.19 and conclude that Γ_Δ acts on $\text{Lk}(\Delta)$ with finite stabilisers.

Next, we verify that Γ_Δ acts with finitely many orbits. Using Proposition 8.31, we have a lift $\hat{\Delta}$ of Δ to $\mathcal{C}(S)$. By Lemma 8.34, we have that $q(\text{Lk}(\hat{\Delta})) = \text{Lk}(\Delta)$, and therefore, since $\text{Stab}_{MCG(S)}(\hat{\Delta}) \leq \text{Stab}_{MCG(S)}(\text{Lk}(\hat{\Delta}))$, we have $\phi(\text{Stab}_{MCG(S)}(\hat{\Delta})) \leq \text{Stab}(\text{Lk}(\Delta))$. Hence, letting L denote the image of $\phi(\text{Stab}_{MCG(S)}(\hat{\Delta}))$ in the permutation group of $\text{Lk}(\Delta)$, we have $L \leq \Gamma_\Delta$.

Since Δ is almost-maximal in $\mathcal{C}(S)/N_{-1}$, the simplex $\hat{\Delta}$ is almost-maximal in $\mathcal{C}(S)$, so its 0-skeleton is a multicurve whose complement has a single non-pants component, a complexity-1 subsurface of S denoted Y . The vertices in $\text{Lk}(\hat{\Delta})$ are the curves on Y , and the subgroup

$\text{Stab}_{MCG(S)}(\widehat{\Delta})$ acts on Y , with finitely many orbits of curves. Hence $\phi(\text{Stab}_{MCG(S)}(\widehat{\Delta}))$ acts on $\text{Lk}(\Delta)$ with finitely many orbits. In other words, L , and hence Γ_Δ , has finitely many orbits in $\text{Lk}(\Delta)$.

It remains to show that Γ_Δ is hyperbolic. Now, we have shown that $L < \Gamma_\Delta$ acts with finitely many orbits on $\text{Lk}(\Delta)$, and that Γ_Δ , and hence also L , acts on $\text{Lk}(\Delta)$ with finite stabilisers. By construction, the actions of L and Γ_Δ are faithful. So L has finite index in Γ_Δ . Hence, to show hyperbolicity of Γ_Δ , we just have to show hyperbolicity of L .

Now, Lemma 8.8 implies that $\phi(\text{Stab}_{MCG(S)}(\widehat{\Delta}))$ is hyperbolic, so to get hyperbolicity of L , it is enough to show that the kernel of the action of $\phi(\text{Stab}_{MCG(S)}(\widehat{\Delta}))$ on $\text{Lk}(\Delta)$ is finite. But, by Lemma 8.34, $\text{Lk}(\Delta) = \text{Lk}(\widehat{\Delta})/(N \cap \text{Stab}_{MCG(S)}(\widehat{\Delta}))$, and so by Lemma 8.8, $\text{Stab}_{MCG(S)}(\widehat{\Delta})/(N \cap \text{Stab}_{MCG(S)}(\widehat{\Delta}))$ acts on $\text{Lk}(\Delta)$ with finite point-stabilisers, as needed.

The case $i = 0$: For $i = 0$, we have that $N_0 = N_{-1}$, and $G_0 = G_{-1}$. As before, Δ is an almost-maximal simplex; all that has changed is the relationship between i and the complexity of the subsurface Y obtained by lifting Δ to $\mathcal{C}(S)$: that complexity is still 1, but now $1 = i + 1$. However, the exact same argument as in the case $i = -1$ shows that Γ_Δ is a hyperbolic group acting with finitely many orbits and finite stabilisers on $\text{Lk}(\Delta)$, so again $\text{Lk}(\Delta)$ is a hyperbolic $\text{Stab}(\text{Lk}(\Delta))$ -space, as required.

The case $i = 1$: Suppose that $i = 1$. Since Δ is almost-maximal, the associated subsurface Y obtained above as the non-pants component of the complement of the multicurve corresponding to a lift $\widehat{\Delta}$ of Δ still has complexity 1.

Let Δ' be the image of $\widehat{\Delta}$ under the quotient $\mathcal{C}(S) \rightarrow \mathcal{C}(S)/N_0 = X$. Then Δ' is a lift of Δ to X ; this is one of the exceptions to the convention on lifting all the way to $\mathcal{C}(S)$.

Now, by induction, we can assume that the action of \bar{G}_0 on $\mathcal{C}(S)/N_0$ satisfies all the conclusions of Proposition 8.13, and in particular item (IX), which says in particular that $\text{PStab}(\text{Sat}(\Delta'))$ acts with finitely many orbits of vertices on $\text{Lk}(\Delta')$, since Y has complexity 1 and $\widehat{\Delta}$ is also a lift of Δ' to $\mathcal{C}(S)$.

Now, the kernel N of $\bar{G}_0 \rightarrow \bar{G}_1$ contains a finite-index subgroup of $\text{PStab}(\text{Sat}(\Delta'))$. Indeed, $[\Delta'] \in \mathbb{Y}$ since Y has complexity 1 (see the beginning of Section 8.6 for the definition of \mathbb{Y} , which we are applying in the case where $X = \mathcal{C}(S)/N_0$). By the definition of N in Notation 8.30, N therefore contains a finite-index subgroup $\Gamma_{[\Delta']}^\theta$ of $\text{PStab}(\text{Sat}(\Delta'))$. By the previous part of the argument, it follows that $N \cap \text{Stab}(\text{Lk}(\Delta'))$ acts with finitely many orbits of vertices on $\text{Lk}(\Delta')$, and therefore the image of $\text{Lk}(\Delta')$ under the quotient map $X \rightarrow X/N$ is finite. Since the conclusion of Lemma 8.34 about links mapping to links also applies to $X \rightarrow X/N$, we conclude that $\text{Lk}(\Delta)$ is finite. In particular, it is a hyperbolic $\text{Stab}(\text{Lk}(\Delta))$ -space.

The case $i > 1$: Now suppose that $i > 1$. Since Δ is almost-maximal, the subsurface Y has complexity $1 \leq i - 1$, so $\text{Lk}(\Delta)$ has finite vertex set by the fact that the third bullet of Theorem 7.1.(II) holds inductively.

This completes the proof of the $\text{Stab}(\text{Lk}(\Delta))$ -hyperbolicity clause of Theorem 6.4.(A) about $\text{Stab}(\text{Lk}(\Delta))$ -hyperbolicity. The quasi-isometric embedding clause follows from the existence of the coarse retraction ρ provided by Lemma 8.39.

Proof of (C): This is Remark 8.35.

Proof of (B): Let Σ, Δ be simplices of X/N . Recall that we need to show that there exist simplices Π, Π' of $\text{Lk}(\Delta)$ such that

(i) $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \text{Lk}(\Delta \star \Pi) \star \Pi'$.

We will in fact prove the following:

Claim 8.40. *For each Δ, Σ one of the following holds:*

- *There exists a vertex v in $\text{Lk}(\Delta)$ so that $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) \subseteq \text{Star}(v)$, or*
- $\text{Lk}(\Delta) \subseteq \text{Lk}(\Sigma)$.

Induction on co-level. We now show how to conclude the proof given the claim.

First, notice that (i) holds whenever Δ is maximal.

Consider some pair Δ, Σ , and suppose that (i) holds for any pair Δ', Σ' for which Δ' has strictly more vertices than Δ . If the second bullet holds, then we can set $\Pi = \Pi' = \emptyset$, and (i) holds for Δ, Σ . If the first bullet holds, then consider the simplices Π_0, Π'_0 obtained from (i) applied to $\Delta \star v, \Sigma$. If $v \in \text{Lk}(\Sigma)$, then we can set $\Pi = \Pi_0 \star v, \Pi' = \Pi'_0 \star v$. If not, we set $\Pi = \Pi_0 \star v, \Pi' = \Pi'_0$. In either case, (i) holds, we spell out the first case, the other being very similar:

$$\begin{aligned} \text{Lk}(\Delta) \cap \text{Lk}(\Sigma) &= (\text{Lk}(\Delta) \cap \text{Star}(v)) \cap \text{Lk}(\Sigma) = (\text{Lk}(\Delta \star v) \star v) \cap \text{Lk}(\Sigma) = \\ &= (\text{Lk}(\Delta \star v) \cap \text{Lk}(\Sigma)) \star v = \text{Lk}(\Delta \star v \star \Pi_0) \star \Pi'_0 \star v. \end{aligned}$$

It remains to prove the claim:

Proof of Claim 8.40. If Δ is maximal, the second bullet holds, so we assume that this is not the case. Also, if $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \emptyset$, then we can take as v any vertex in $\text{Lk}(\Delta)$.

So, from now on we assume $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) \neq \emptyset$.

Let Λ be a maximal simplex of $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$. Let $\hat{\Delta} \star \hat{\Lambda}$ be a lift of the simplex $\Delta \star \Lambda$, provided by Proposition 8.31.

In terms of curves, $(\hat{\Lambda})^{(0)} \cup (\hat{\Delta})^{(0)}$ is a multicurve, and as such it can be completed to a pants decomposition $M_{\Lambda, \Pi} = (\hat{\Lambda})^{(0)} \cup (\hat{\Delta})^{(0)} \cup (\hat{\Pi})^{(0)}$.

Reducing to $\Pi = \emptyset$. Let $v \in \text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$. We claim that $v \in \text{Lk}(\Delta \star \Pi)$. We can lift the generalized 4-gon $\Delta \star \Lambda, \Lambda \star \Sigma, \Sigma \star v, v \star \Delta$ to $\mathcal{C}(S)$, obtaining the lifts $\hat{\Delta}$, etc., and in fact we can assume $\hat{\Delta}, \hat{\Lambda}$ coincide with the previously chosen lifts. If we had $v \notin \text{Lk}(\Delta \star \Pi)$, then we would also have $\hat{v} \notin \text{Lk}(\hat{\Delta} \star \hat{\Pi})$, and also $\hat{v} \notin \hat{\Lambda}$. But $\hat{v} \in \text{Lk}(\hat{\Delta})$, so $\hat{v} \notin \text{Lk}(\hat{\Pi})$. Moreover, by maximality of Λ , we also have $\hat{v} \notin \text{Lk}(\hat{\Lambda})$. Hence, as a curve, \hat{v} intersects both $(\hat{\Lambda})^{(0)}$ and $(\hat{\Pi})^{(0)}$ (but not $(\hat{\Delta})^{(0)}$ nor $(\hat{\Sigma})^{(0)}$). We now make a multicurve σ from \hat{v} , which we assume to be in minimal position with respect to $M_{\Lambda, \Pi}$, by considering the boundary of a regular neighborhood of $\hat{v} \cup (\hat{\Lambda})^{(0)}$. Notice that σ is disjoint from $(\hat{\Sigma})^{(0)}, (\hat{\Lambda})^{(0)}$, and $(\hat{\Delta})^{(0)}$. Also, at least one component σ_0 of σ is not parallel to $(\hat{\Lambda})^{(0)}$ (since it intersects $(\hat{\Pi})^{(0)}$ non-trivially), and this contradicts the maximality of Λ . This proves that $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) \subseteq \text{Lk}(\Delta \star \Pi)$. If Π is non-empty, the first bullet holds with v any vertex of Π . Hence, we now assume $\Pi = \emptyset$.

The case $\Pi = \emptyset$. We now know that $(\hat{\Lambda})^{(0)} \cup (\hat{\Delta})^{(0)}$ is a pants decomposition, and that in fact this holds for any maximal simplex Λ of $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$. If all such maximal simplices share a vertex v , then the first bullet holds for this v .

Otherwise, for each vertex w of a fixed maximal simplex Λ in $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$, we can find another such simplex Θ not sharing w with Λ . We claim that for each curve δ in $(\hat{\Delta})^{(0)}$ which is the boundary curve of a pair of pants not all of whose boundary curves are in $(\hat{\Delta})^{(0)}$, Σ has a lift avoiding that curve. In fact, we can lift a generalized 4-gon $\Delta \star \Lambda, \Lambda \star \Sigma, \Sigma \star \Theta, \Theta \star \Delta$, where Θ and Λ do not share a vertex corresponding to one of the boundary curves of a pair of pants as above. Then, the lift of Σ will not intersect δ , for otherwise it will have to intersect some curve either in the lift of Λ or in the lift of Θ .

Let now Δ_0 be the sub-simplex of Δ consisting of all vertices $q(\delta)$ for δ as above. Then each vertex of Σ is in $\text{Star}(\Delta_0)$. Hence, Δ_0, Λ , and Σ are contained in a common simplex, which

can be lifted to a simplex in $\mathcal{C}(S)$. We can also arrange the corresponding lifts of Δ_0 and Λ so that the lift of Λ is $\widehat{\Lambda}$ and the lift of Δ_0 has vertex set consisting of all curves δ as above.

This gives a lift $\widehat{\Sigma}'$ of Σ which is contained in the union of the pairs of pants in the complement of $(\widehat{\Lambda})^{(0)} \cup (\widehat{\Delta})^{(0)}$, all of whose boundary curves are in $\widehat{\Delta}^{(0)}$. This implies that $\text{Lk}(\widehat{\Delta}) \subseteq \text{Lk}(\widehat{\Sigma}')$, and hence $\text{Lk}(\Delta) \subseteq \text{Lk}(\Sigma)$, that is, the second bullet.

This concludes the proof of Claim 8.40. \blacksquare

We now prove the statement about the index set \mathfrak{S}_{N_i} of the HHS structure on \bar{G}_i . Recall that \mathfrak{S}_{N_i} is the set of equivalence classes of non-maximal simplices. We first recall the bijection b .

Given any non-maximal simplex Δ of X/N , consider a lift $\widehat{\Delta}$ to $\mathcal{C}(S)$. The vertex set of the link of $\widehat{\Delta}$ in $\mathcal{C}(S)$ consists of all curves (regarded as vertices of $\mathcal{C}(S)$) contained in a subsurface that we denote $S_{\widehat{\Delta}}$. Define $b([\Delta]) = [S_{\widehat{\Delta}}]_{N_i}$, where $[\cdot]_{N_i}$ denotes the N_i -orbit.

Notice that choosing a different lift yields a subsurface in the same N_i -orbit, since all lifts of Δ are in the same N_i -orbit.

We now complete the proof that b is well-defined. What is left to prove is that equivalent simplices yield the same orbit.

Claim 8.41. *No link is a join of a non-empty simplex and some subcomplex. In particular, no link consists of a single non-empty simplex.*

Proof of Claim 8.41. Consider the link of Δ and some vertex $v \in \text{Lk}(\Delta)$, and let us show that there is a vertex w of $\text{Lk}(\Delta)$ not connected to, or equal to, v . Consider a lift $\widehat{\Delta} \star \widehat{v}$ of $\Delta \star v$. By Lemma 8.1.(2) there exists a vertex \widehat{w} of $\text{Lk}(\widehat{\Delta})$ so that \widehat{v} and \widehat{w} are not in the same H -orbit, and no H -translate of \widehat{w} is adjacent to \widehat{v} . Since $N_i < H$, this means that the image w of \widehat{w} , which is in $\text{Lk}(\Delta)$, is distinct from v and not connected to v , as required. \blacksquare

In view of condition B, we now see that $\text{Lk}(\Sigma) \subseteq \text{Lk}(\Delta)$ if and only if there exists a simplex Π in $\text{Lk}(\Delta)$ so that $\text{Lk}(\Sigma) = \text{Lk}(\Delta \star \Pi)$. This also holds in $\mathcal{C}(S)$, and hence we get that $[\Delta] \sqsubseteq [\Sigma]$ if and only if there are $S_{\widehat{\Delta}}, S_{\widehat{\Sigma}}$ as above with $S_{\widehat{\Delta}}$ nested into $S_{\widehat{\Sigma}}$. This implies that if $\text{Lk}(\Delta) = \text{Lk}(\Sigma)$ then $S_{\widehat{\Delta}} \subseteq S_{\widehat{\Sigma}} \subseteq gS_{\widehat{\Delta}}$ for some $g \in N_i$, and since $S_{\widehat{\Delta}}, gS_{\widehat{\Delta}}$ have the same complexity they need to coincide, showing $S_{\widehat{\Delta}} = S_{\widehat{\Sigma}}$, as we wanted.

Notice that we also showed that $[\Delta] \sqsubseteq [\Sigma]$ if and only if the corresponding orbits contain nested representatives. We are only left to show the analogous statement for orthogonality/disjointness, which we will reduce to the nesting statement.

Consider a non-maximal simplex Δ of X/N , and a lift $\widehat{\Delta}$. Then $\widehat{\Delta}$ contains the boundary multicurve $(\widehat{\Delta}_0)^{(0)}$ of $S_{\widehat{\Delta}}$. Denote $\Delta_0 = q(\widehat{\Delta}_0)$, and pick any maximal simplex Λ in $\text{Lk}(\Delta)$. We claim that a simplex Σ satisfies $[\Delta] \perp [\Sigma]$ if and only if $[\Sigma] \sqsubseteq [\Delta_0 \star \Lambda]$. This implies the orthogonality/disjointness statement, in view of the nesting statement.

Fix a lift $\widehat{\Lambda}$ of Λ contained in $\text{Lk}(\widehat{\Delta})$, and notice that the vertex set forms a pants decomposition of $S_{\widehat{\Delta}}$. If $[\Sigma] \sqsubseteq [\Delta_0 \star \Lambda]$, then there is a lift $\widehat{\Sigma}$ so that $\text{Lk}(\widehat{\Sigma}) \subseteq \text{Lk}(\widehat{\Delta}_0 \star \widehat{\Lambda})$. Any curve disjoint from $(\widehat{\Delta}_0 \star \widehat{\Lambda})^{(0)}$ is disjoint from curves in $S_{\widehat{\Lambda}}$, so that $\text{Lk}(\widehat{\Sigma})$ and $\text{Lk}(\widehat{\Delta})$ form a join. The same then holds for Δ, Σ , showing $[\Delta] \perp [\Sigma]$, as required.

Suppose now that $[\Delta] \perp [\Sigma]$. Our goal is to show that any vertex $v \in \text{Lk}(\Sigma)$ lies in $\text{Lk}(\Delta_0 \star \Lambda)$. For later use, note that there is a well-defined simplex $\Lambda \star v$ in view of the definition of orthogonality, since $\Lambda \subset \text{Lk}(\Delta)$.

We claim that we can find another maximal simplex Θ in $\text{Lk}(\Delta)$ so that any lift of Θ contained in $\text{Lk}(\widehat{\Delta})$ has vertex set which, together with $(\widehat{\Lambda})^{(0)}$ fills $S_{\widehat{\Delta}}$. We will do so by showing that we can find Θ in $\text{Lk}(\Delta)$ such that for every vertex v of Λ there exists a vertex w of Θ such that v and w are not connected. In particular, as curves, any lifts of v and w intersect, from which we see that the vertices of any lift of Θ contained in $\text{Lk}(\widehat{\Delta})$, as curves, intersect all curves of $(\widehat{\Lambda})^{(0)}$.

Since $(\widehat{\Lambda})^{(0)}$ is a pants decomposition of $S_{\widehat{\Delta}}$, this yields the required filling statement. Now, by Claim 8.41, $\text{Lk}(\Delta)$ does not consist of the maximal simplex Λ only, so that there exists a vertex w_1 in $\text{Lk}(\Delta)$ which is not connected to some vertex of Λ (recall that Λ is a maximal simplex of $\text{Lk}(\Delta)$). If w_1 is not connected to any vertex of Λ , we can simply complete w_1 to a maximal simplex Θ of $\text{Lk}(\Delta)$. If not, let $\Lambda_2 = \Lambda \cap \text{Lk}(w_1)$. Again by Claim 8.41, used as above, there exists $v_2 \text{Lk}(\Delta \star w_1)$ which is not connected to some vertex of Λ^2 . If it is not connected to any vertex, we can take Θ with vertex set containing w_1, w_2 , and otherwise we continue for finitely many steps.

Consider now any vertex $v \in \text{Lk}(\Sigma)$, which we need to show is in $\text{Lk}(\Delta_0 \star \Lambda)$. We can lift a generalized 4-gon $\Delta_0 \star \Lambda, \Lambda \star v, v \star \Theta, \Theta \star \Delta_0$, with the lift of $\Delta_0 \star \Lambda$ being $\widehat{\Delta}_0 \star \widehat{\Lambda}$. We then see that, as a curve, the lift \widehat{v} is disjoint from $S_{\widehat{\Delta}}$ (since $\widehat{\Theta}$ and $\widehat{\Lambda}$ fill $S_{\widehat{\Delta}}$), so that v lies either in $\text{Lk}(\Delta_0 \star \Lambda)$ or in Δ_0 .

We are left to argue that v does not lie in Δ_0 . Suppose that this was the case. Again by 8.41, there is a vertex w in $\text{Lk}(\Sigma)$ which is not connected to v , and in particular any lift \widehat{w} intersects \widehat{v} . Applying the above argument with w replacing v , we would find a lift \widehat{w} that cuts $S_{\widehat{\Delta}}$ since it intersects its boundary, a contradiction.

Finally, we complete the proof of Theorem 7.1.(II) by proving that if $b([\Delta]) = N_i Y$ for Y of complexity at most i , then $\mathcal{C}(\Delta)$ has finite vertex set. This suffices to get a uniform diameter bound since there are finitely many orbits of simplices.

This is the only place in this subsection where lifts are taken to be in X rather than in $\mathcal{C}(S)$. First, if $b([\Delta]) = N_i Y$ for Y of complexity less than i , then the link of any lift $\widehat{\Delta}$ of Δ in X has finitely many vertices, by the inductive hypothesis. Since the conclusion of Lemma 8.34 about links mapping to links applies to the map $X \rightarrow X/N$ as well, with the same proof, we conclude that the link of Δ has finite vertex set, as required.

Suppose now $b([\Delta]) = N_i Y$ for Y of complexity i . Consider a lift $\widehat{\Delta}$ of Δ in X . By inductive hypothesis (specifically, Proposition 8.13.(IX)), $\text{PStab}(\text{Sat}(\widehat{\Delta}))$ acts cocompactly on $\text{Lk}(\widehat{\Delta})$. This readily implies that $\text{PStab}(\text{Sat}(\Delta))$ acts cocompactly on $\text{Lk}(\Delta)$, by Lemma 8.34 (used as above). Now, $\text{PStab}(\text{Sat}(\Delta))$ is a quotient of $\text{PStab}(\text{Sat}(\widehat{\Delta}))$ by a finite-index subgroup by construction of the composite rotating family, and we are done. \square

We conclude this subsection with the proof Proposition 8.13.(VII), which we now have the tools to prove:

Lemma 8.42. *Let Δ be a non-maximal simplex of X/N , and consider a lift $\widehat{\Delta}$ to $\mathcal{C}(S)$. Then $q(\text{Sat}(\widehat{\Delta})) = \text{Sat}(\Delta)$.*

Proof. We use that the bijection b defined in Definition 7.2 is well-defined. Let $v \in \text{Sat}(\Delta)$. Then v is a vertex of some Δ' with the same link as Δ .

Since b is well-defined, Δ' must have a lift $\widehat{\Delta}'$ such that the vertex set of $\text{Lk}(\widehat{\Delta}')$ consists of all curves contained $S_{\widehat{\Delta}}$. That is, $\widehat{\Delta}'$ has the same link as $\widehat{\Delta}$, so that $\widehat{\Delta}' \subseteq \text{Sat}(\widehat{\Delta})$ and $v \in \Delta' = q(\widehat{\Delta}') \subseteq q(\text{Sat}(\widehat{\Delta}))$. We just proved $\text{Sat}(\Delta) \subseteq q(\text{Sat}(\widehat{\Delta}))$.

Let $v \in q(\text{Sat}(\widehat{\Delta}))$. Then $v = q(\widehat{v})$ for some vertex \widehat{v} of a simplex $\widehat{\Delta}'$ with the same link as $\widehat{\Delta}$. This implies that $\Delta' = q(\widehat{\Delta}')$ has the same link as Δ , by Lemma 8.34. Hence $v \in \text{Sat}(\Delta)$, so $q(\text{Sat}(\widehat{\Delta})) \subseteq \text{Sat}(\Delta)$. \square

8.10. Preservation properties.

Lemma 8.43. *Let $\gamma \in N$ and $x \in X^{(0)}$. Then either $\gamma x = x$ or there exists $Y \in \mathbb{Y}_*$ so that $d_Y(x, \gamma x) > \mathfrak{T}(\theta)$.*

Proof. Since \mathfrak{C} is well-ordered, we can argue by (transfinite) induction on $c(\gamma)$. The statement holds for $\gamma = 1$, that is, for the minimal element of \mathfrak{C} .

Suppose $\gamma \neq 1$ and suppose $\gamma x \neq x$. Let γ_Y be as in Lemma 8.28/Lemma 8.10. If the second conclusion of Lemma 8.28/Lemma 8.10 applies, then we are done.

Otherwise, suppose $x \in \text{Fix}(\Gamma_Y)$. Then, $\gamma_Y \gamma x \neq x$ (since $\gamma_Y^{-1} x = x$), and $c(\gamma_Y \gamma) < c(\gamma)$, so by induction there exists $W \in \mathbb{Y}_*$ so that $d_W(x, \gamma_Y \gamma x) > \mathfrak{T}(\theta)$. Hence, we get $d_{\gamma_Y^{-1} W}(x, \gamma x) > \mathfrak{T}(\theta)$ (we used $\gamma_Y^{-1} x = x$ again).

The case $\gamma x \in \text{Fix}(\Gamma_Y)$ is similar: $\gamma_Y \gamma x \neq x$ since $\gamma_Y \gamma x = \gamma x$, and $c(\gamma_Y \gamma) < c(\gamma)$, so by induction there exists $W \in \mathbb{Y}_*$ so that $d_W(x, \gamma_Y \gamma x) > \mathfrak{T}(\theta)$. Hence, $d_W(x, \gamma x) > \mathfrak{T}(\theta)$, again because $\gamma_Y \gamma x = \gamma x$. \square

Proof of Theorem 7.1.(I) and Proposition 8.13.(V). Let $f, g \in F$ be distinct. We have to prove that $f^{-1}g \notin N_i$. If $f^{-1}g$ has finite order, then $f^{-1}g$ cannot be in H , since H is torsion-free, and hence in particular not in N_i . If not, let $x \in X$ so that $f(x) \neq g(x)$, which exists by induction for $i > -1$, while for $i = -1$ it exists because infinite order elements of $MCG(S)$ act non-trivially on $\mathcal{C}(S)$. Then by Lemma 8.43, if we had $f^{-1}g \in N_i$ we would have some $Y \in \mathbb{Y}_*$ so that $d_Y(f(x), g(x)) > \mathfrak{T}(\theta)$. This contradicts the choice of θ in Notation 8.30. \square

Lemma 8.44. *Let κ be as in Lemma 2.10. Let $x, y \in X$ be so that $d_{[\Delta]}(x, y) \leq \kappa$ for all non-empty non-maximal simplices Δ of X . Let $[x, y]$ be a geodesic from x to y . Then $q|_{[x, y]}$ is an isometric embedding.*

Proof. Consider a lift $[x, y']$ to $\mathcal{C}(S)$ of a geodesic $[\bar{x}, \bar{y}]$ in X/N connecting the images \bar{x}, \bar{y} of x, y . Then $y' = \gamma y$ for some $\gamma \in N$. Since \mathfrak{C} is well-ordered, we can choose γ to have minimal $c(\gamma)$ among:

- the N_i -orbits of the pair (x, y) and of the geodesic $[x, y]$,
- all lifts $[x, y']$ of $[\bar{x}, \bar{y}]$,
- all choices of γ with $y' = \gamma y$.

We claim that $\gamma = 1$, which will show that $d_X(x, y) = d_{X/N}(\bar{x}, \bar{y})$, which readily implies the desired conclusion.

Suppose $\gamma \neq 1$. Consider $Y \in \mathbb{Y}_*$ and $\gamma_Y \in \Gamma_Y$ with $c(\gamma_Y \gamma) < c(\gamma)$ as in Lemma 8.28/Lemma 8.10. There are three cases.

- If $y \in \text{Fix}(\Gamma_Y)$, then we can replace $[x, y]$ with $\gamma_Y[x, y]$, and $[x, y']$ with $\gamma_Y[x, y']$. Then $\gamma_Y \gamma$ maps the second endpoint of the original geodesic (i.e., $\gamma_Y y = y$) to the second endpoint of the lift (i.e., $\gamma_Y y'$), meaning $\gamma_Y y' = \gamma_Y \gamma y$. This contradicts minimality of $c(\gamma)$.
- If $y' \in \text{Fix}(\Gamma_Y)$, then $y' = \gamma_Y y' = \gamma_Y \gamma y$, contradicting minimality of γ .
- Otherwise, $d_Y(y', y) > \mathfrak{T}(\theta)$. If $x \in \text{Fix}(\Gamma_Y)$, we can replace $[x, y]$ with $\gamma_Y[x, y]$, and $[x, y']$ with $\gamma_Y[x, y']$, and we contradict minimality of $c(\gamma)$ since $\gamma_Y y' = \gamma_Y \gamma y$. Otherwise, $d_Y(x, y)$ is well-defined and $\leq \kappa$. By the choice of θ (much larger than κ) in Notation 8.6 or Notation 8.30, we get that $d_Y(x, y')$ is well-defined and large enough that $[x, y']$ intersects $\text{Fix}(\Gamma_Y)$ (by the defining property of C_i , see Proposition 8.13.(VIII)).

We can now replace a terminal subgeodesic $[v, y']$ of the lift $[x, y']$ starting at some $v \in \text{Fix}(\Gamma_Y)$ by its translate $\gamma_Y[v, y']$, thereby obtaining a new lift. This contradicts minimality of $c(\gamma)$ since $\gamma_Y y' = \gamma_Y \gamma y$.

This completes the proof of the lemma. \square

We conclude by proving the remaining statement in Theorem 7.1.

Proof of Theorem 7.1.(III). Recall that we need to prove that the map ϕ restricts to an injective map on Q , and orbit maps from Q to $\mathcal{C}(S)/N_i$ are quasi-isometric embeddings.

We first prove the quasi-isometric embedding statement. By the choice of κ (coming from Lemma 2.10) and Lemma 8.44, for any $x_0 \in X$ and $g \in Q$, we have that any geodesic $[x_0, gx_0]$

in X projects to a geodesic $[\bar{x}_0, g\bar{x}_0]$ in X/N . Since the length of $d_X(x_0, gx_0)$ is comparable up to multiplicative and additive constants with the word length of g , by induction, the same holds for $d_{X/N}(\bar{x}_0, g\bar{x}_0)$. This suffices to show that Q -orbit maps are quasi-isometric embeddings.

To show the injectivity statement, note that Q -orbit maps being quasi-isometric embeddings implies that the kernel of $\phi|_Q$ is finite. But this implies that the kernel of $\phi|_Q$ must be trivial since it is contained in N_i , whence in H , which is torsion-free by construction (see Lemma 8.1). Hence, $\phi|_Q$ is injective. \square

REFERENCES

- [AAS07] James W. Anderson, Javier Aramayona, and Kenneth J. Shackleton. An obstruction to the strong relative hyperbolicity of a group. *J. Group Theory*, 10(6):749–756, 2007.
- [AB18] Carolyn Abbott and Jason Behrstock. Conjugator lengths in hierarchically hyperbolic groups. *arXiv preprint arXiv:1808.09604*, 2018.
- [ABD17] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham. Largest acylindrical actions and stability in hierarchically hyperbolic groups. *To appear in Transactions of the AMS*, 2017.
- [AF19] Javier Aramayona and Louis Funar. Quotients of the mapping class group by power subgroups. *Bull. Lond. Math. Soc.*, 51(3):385–398, 2019.
- [AGM09] Ian Agol, Daniel Groves, and Jason Fox Manning. Residual finiteness, QCERF and fillings of hyperbolic groups. *Geom. Topol.*, 13(2):1043–1073, 2009.
- [Ago13] I. Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.
- [ANS19] Carolyn Abbott, Thomas Ng, and Davide Spriano. Hierarchically hyperbolic groups and uniform exponential growth. *arXiv preprint arXiv:1909.00439*, 2019.
- [BB01] Stephen J. Bigelow and Ryan D. Budney. The mapping class group of a genus two surface is linear. *Algebr. Geom. Topol.*, 1:699–708, 2001.
- [BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math. Inst. Hautes Études Sci.*, 122:1–64, 2015.
- [BDM09] Jason Behrstock, Cornelia Druţu, and Lee Mosher. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. *Math. Ann.*, 344(3):543–595, 2009.
- [BF92] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.
- [BFG⁺07] Uri Bader, Alex Furman, Tsachik Gelander, Nicolas Monod, et al. Property (t) and rigidity for actions on banach spaces. *Acta mathematica*, 198(1):57–105, 2007.
- [BHS] J. Behrstock, M. F. Hagen, and A. Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.* To appear.
- [BHS17a] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. Lond. Math. Soc. (3)*, 114(5):890–926, 2017.
- [BHS17b] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.
- [BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299(2):257–338, 2019.
- [Bir74] J. Birman. *Braids, links, and mapping class groups*, volume 82 of *Annals of Math. Studies*. Princeton University Press, 1974.
- [Bow12] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.
- [Bow13] Brian H. Bowditch. Coarse median spaces and groups. *Pacific J. Math.*, 261(1):53–93, 2013.
- [Bow15] Brian H Bowditch. Large-scale rank and rigidity of the Weil-Petersson metric. *To appear in Groups, Geom. Dyn.*, 2015.
- [BR18] Federico Berlai and Bruno Robbio. A refined combination theorem for hierarchically hyperbolic groups. *Groups, Geometry, and Dynamics*, 2018. To appear.
- [BR20] Daniel Berlyne and Jacob Russell. Hierarchical hyperbolicity of graph products. *arXiv preprint arXiv:2006.03085*, 2020.
- [Bri10] Martin R. Bridson. Semisimple actions of mapping class groups on CAT(0) spaces. In *Geometry of Riemann surfaces*, volume 368 of *London Math. Soc. Lecture Note Ser.*, pages 1–14. Cambridge Univ. Press, Cambridge, 2010.
- [BRW17] Martin R. Bridson, Alan W. Reid, and Henry Wilton. Profinite rigidity and surface bundles over the circle. *Bull. Lond. Math. Soc.*, 49(5):831–841, 2017.

- [CMM20] Matt Clay, Johanna Mangahas, and Dan Margalit. Right-angled artin groups as normal subgroups of mapping class groups. *arXiv preprint arXiv:2001.10587*, 2020.
- [Dah18] François Dahmani. The normal closure of big Dehn twists, and plate spinning with rotating families. *Geom. Topol.*, 22:4113–4144, 2018.
- [DDLS20] Spencer Dowdall, Matthew G Durham, Christopher J Leininger, and Alessandro Sisto. Extensions of veech groups are hierarchically hyperbolic. *arXiv preprint arXiv:2006.16425*, 2020.
- [DG18] François Dahmani and Vincent Guirardel. Recognizing a relatively hyperbolic group by its Dehn fillings. *Duke Math. J.*, 167(12):2189–2241, 2018.
- [DHS17] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.
- [DHS19] Matthew Gentry Durham, Mark F Hagen, and Alessandro Sisto. Corrigendum to: Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geometry and Topology*, 2019. To appear.
- [DHS21] François Dahmani, Mark Hagen, and Alessandro Sisto. Dehn filling Dehn twists. *Proc. Roy. Soc. Edinburgh Sect. A*, 151(1):28–51, 2021.
- [DT15] Matthew Durham and Samuel J Taylor. Convex cocompactness and stability in mapping class groups. *Algebraic & Geometric Topology*, 15(5):2837–2857, 2015.
- [DT19] François Dahmani and Nicholas Touikan. Deciding isomorphism using Dehn fillings: the splitting case. *Invent. Math.*, 2019. to appear.
- [EK95] Rob Kirby (Ed.) and Edited Rob Kirby. Problems in low-dimensional topology. In *Proceedings of Georgia Topology Conference, Part 2*, pages 35–473. Press, 1995.
- [FM02] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152, 2002.
- [FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.
- [Fuj15] Koji Fujiwara. Subgroups generated by two pseudo-Anosov elements in a mapping class group. II. Uniform bound on exponents. *Trans. Amer. Math. Soc.*, 367(6):4377–4405, 2015.
- [Fun99] Louis Funar. On the TQFT representations of the mapping class groups. *Pacific J. Math.*, 188(2):251–274, 1999.
- [GM08] D. Groves and J. F. Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.
- [Gro75] Edna K. Grossman. On the residual finiteness of certain mapping class groups. *J. London Math. Soc. (2)*, 9:160–164, 1974/75.
- [Hae16] Thomas Haettel. Higher rank lattices are not coarse median. *Algebraic & Geometric Topology*, 16(5):2895–2910, 2016.
- [Ham05] Ursula Hamenstädt. Word hyperbolic extensions of surface groups. *arXiv preprint math/0505244*, 2005.
- [HMS22] Mark Hagen, Alexandre Martin, and Alessandro Sisto. Extra-large type artin groups are hierarchically hyperbolic. *Mathematische Annalen*, 2022.
- [HQR20] Camille Horbez, Yulan Qing, and Kasra Rafi. Big mapping class groups with hyperbolic actions: classification and applications. *Preprint, arXiv: math*, 2020.
- [HRSS22] Mark Hagen, Jacob Russell, Alessandro Sisto, and Davide Spriano. Equivariant hierarchically hyperbolic structures for 3-manifold groups via quasimorphisms. *arXiv preprint arXiv:2206.12244*, 2022.
- [Hum92] Stephen P. Humphries. Normal closures of powers of Dehn twists in mapping class groups. *Glasgow Math. J.*, 34(3):313–317, 1992.
- [HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.
- [Iva06] Nikolai V. Ivanov. Fifteen problems about the mapping class groups. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 71–80. Amer. Math. Soc., Providence, RI, 2006.
- [KK13] Sang-hyun Kim and Thomas Koberda. Embedability between right-angled Artin groups. *Geom. Topol.*, 17(1):493–530, 2013.
- [KK14] Sang-Hyun Kim and Thomas Koberda. The geometry of the curve graph of a right-angled Artin group. *Internat. J. Algebra Comput.*, 24(2):121–169, 2014.
- [KL96] Michael Kapovich and Bernhard Leeb. Actions of discrete groups on nonpositively curved spaces. *Math. Ann.*, 306(2):341–352, 1996.
- [KL08] Autumn E Kent and Christopher J Leininger. Shadows of mapping class groups: capturing convex cocompactness. *Geometric and Functional Analysis*, 18(4):1270–1325, 2008.
- [KM12] Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.

- [LM07] Christopher J. Leininger and D. B. McReynolds. Separable subgroups of mapping class groups. *Topology Appl.*, 154(1):1–10, 2007.
- [MM99] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999.
- [MM00] H. A. Masur and Y. N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.
- [MS22] Giorgio Mangioni and Alessandro Sisto. Rigidity of mapping class groups mod powers of twists. *arXiv preprint arXiv:2212.11014*, 2022.
- [Nic13] Bogdan Nica. Proper isometric actions of hyperbolic groups on L^p -spaces. *Compos. Math.*, 149(5):773–792, 2013.
- [Osi07] D. V. Osin. Peripheral fillings of relatively hyperbolic groups. *Invent. Math.*, 167(2):295–326, 2007.
- [Par09] Luis Paris. Residual p properties of mapping class groups and surface groups. *Trans. Amer. Math. Soc.*, 361(5):2487–2507, 2009.
- [Pic81] Émile Picard. Sur une extension aux fonctions de deux variables du problème de riemann relatif aux fonctions hypergéométriques. In *Annales scientifiques de l'École Normale Supérieure*, volume 10, pages 305–322, 1881.
- [Pic85] Émile Picard. Sur les fonctions hyperfuchsienues provenant des séries hypergéométriques de deux variables. In *Annales scientifiques de l'École Normale Supérieure*, volume 2, pages 357–384, 1885.
- [PS20] Harry Petyt and Davide Spriano. Unbounded domains in hierarchically hyperbolic groups. *arXiv preprint arXiv:2007.12535*, 2020.
- [Rei06] Alan W. Reid. Surface subgroups of mapping class groups. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 257–268. Amer. Math. Soc., Providence, RI, 2006.
- [RS20] Bruno Robbio and Davide Spriano. Hierarchical hyperbolicity of hyperbolic-2-decomposable groups. *arXiv preprint arXiv:2007.13383*, 2020.
- [RST18] Jacob Russell, Davide Spriano, and Hung Cong Tran. Convexity in hierarchically hyperbolic spaces. *arXiv preprint arXiv:1809.09303*, 2018.
- [RST19] Jacob Russell, Davide Spriano, and Hung Cong Tran. The local-to-global property for Morse quasi-geodesics. *arXiv preprint arXiv:1908.11292*, 2019.
- [Rus19] Jacob Russell. From hierarchical to relative hyperbolicity. *arXiv preprint arXiv:1905.12489*, 2019.
- [Sco78] Peter Scott. Subgroups of surface groups are almost geometric. *J. London Math. Soc. (2)*, 17(3):555–565, 1978.
- [Thu98] William P. Thurston. Shapes of polyhedra and triangulations of the sphere. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 511–549. Geom. Topol. Publ., Coventry, 1998.
- [Vok22] Kate M. Vokes. Hierarchical hyperbolicity of graphs of multicurves. *Algebr. Geom. Topol.*, 22(1):113–151, 2022.
- [Wis04] D. T. Wise. Cubulating small cancellation groups. *Geom. Funct. Anal.*, 14(1):150–214, 2004.
- [Yu05] Guoliang Yu. Hyperbolic groups admit proper affine isometric actions on L^p -spaces. *Geom. Funct. Anal.*, 15(5):1144–1151, 2005.

LEHMAN COLLEGE AND THE GRADUATE CENTER, CUNY, NEW YORK, NEW YORK, USA
Email address: `jason.behrstock@lehman.cuny.edu`

SCHOOL OF MATHEMATICS, UNIVERSITY OF BRISTOL, BRISTOL, UK
Email address: `markfhagen@posteo.net`

MAXWELL INSTITUTE AND DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY, EDINBURGH, UK
Email address: `alexandre.martin@hw.ac.uk`

MAXWELL INSTITUTE AND DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY, EDINBURGH, UK
Email address: `a.sisto@hw.ac.uk`