Inventiones mathematicae (2024) 236:925–997
https://doi.org/10.1007/s00222-024-01241-9

Special cubulation of strict hyperbolization

Jean-François Lafont[1] · Lorenzo Ruffoni[2]

Received: 28 June 2022 / Accepted: 6 February 2024 / Published online: 2 May 2024

Abstract
We prove that the Gromov hyperbolic groups obtained by the strict hyperbolization procedure of Charney and Davis are virtually compact special, hence linear and residually finite. Our strategy consists in constructing an action of a hyperbolized group on a certain dual CAT(0) cubical complex. As a result, all the common applications of strict hyperbolization are shown to provide manifolds with virtually compact special fundamental group. In particular, we obtain examples of closed negatively curved Riemannian manifolds whose fundamental groups are linear and virtually algebraically fiber.

Mathematics Subject Classification 20F67 · 53C23 · 20E26 · 57Q05

Contents

✉ J.-F. Lafont
jlafont@math.ohio-state.edu

L. Ruffoni
lorenzo.ruffoni2@gmail.com

1 Department of Mathematics, The Ohio State University, 231 West 18th Avenue, Columbus, OH 43210-1174, USA

2 Department of Mathematics, Tufts University, 177 College Avenue, Medford, MA 02155, USA

1 Introduction

Closed aspherical manifolds occupy a central place in manifold topology. For this class of manifolds, the Borel Conjecture predicts that two such manifolds are homeomorphic if and only if they have isomorphic fundamental groups – in other words, that the topology is entirely encoded in the fundamental group. A challenging problem is the question of examples. The fundamental group will always satisfy Poincaré Duality over $\mathbb{Z}$ (i.e. they are PD_n groups), and the Wall Conjecture predicts that conversely, any PD_n group is the fundamental group of an aspherical manifold. Classically, there were two sources of examples of aspherical manifolds: they either arose from Lie theory, as quotients of contractible Lie groups by discrete subgroups, or from differential geometry, as non-positively curved manifolds.

In the late 1970's, Gromov introduced two metric versions of non-positive curvature, CAT(0) spaces and Gromov hyperbolic spaces. Simply connected, complete, locally CAT(0) spaces are automatically contractible. In dimensions ≥ 4, manifolds that support locally CAT(0) metrics form a new source of aspherical manifolds. Moreover, it is easy to produce such manifolds, through a process known as hyperbolization. This was originally outlined by Gromov in [32], and subsequently developed by Davis and Januszkiewicz in [18]. Hyperbolization is a functorial procedure, which inputs a simplicial complex, and outputs a locally CAT(0) space. In a later refinement, Charney and Davis in [14] developed a strict hyperbolization procedure, where the output is locally CAT(-1), i.e. admits a metric of negative curvature (as opposed to just non-positive curvature). The hyperbolization procedures have been used to produce examples of aspherical manifolds with various unexpected properties. In this work we show that one can construct hyperbolizations that have some additional algebraic regularity.

Theorem 1.1 *Given a dimension $n > 0$, there exists a strict hyperbolization procedure $\mathcal{H}$ with the following property. Let K be any n-dimensional simplicial complex, which is compact, homogeneous, and without boundary. Then the resulting hyperbolized space $\mathcal{H}(K)$ has fundamental group $G = \pi_1(\mathcal{H}(K))$ which acts cocompactly on a* CAT(0) *cubical complex by cubical isometries.*

Most of the paper is concerned with the proof of this result (see §3 and §4). The cubical complex in this statement is n-dimensional, but it is not locally compact and the action is not proper. Nevertheless, since the hyperbolization procedure is strict (i.e. $\mathcal{H}(K)$ is locally CAT(-1)), the fundamental group G is Gromov hyperbolic. Therefore the work of Agol, Haglund–Wise, and Groves–Manning about special cube complexes (see [2, 35, 39]) can be used to extract information about G. A cubical complex is *special* if it admits a local isometry to the Salvetti complex of a right-angled Artin group (RAAG) (see [39]). A group is *virtually compact special* if it has a finite index subgroup which is the fundamental group of a compact special cubical complex.

Theorem 1.2 *Given a dimension $n > 0$, there exists a strict hyperbolization procedure $\mathcal{H}$ with the following property. Let K be any n-dimensional compact simplicial complex. Then the resulting hyperbolized space $\mathcal{H}(K)$ has fundamental group $G = \pi_1(\mathcal{H}(K))$ which is Gromov hyperbolic and virtually compact special. In particular, G enjoys the following properties.*

(1) *G virtually embeds in a right–angled Artin group (RAAG) (see [39]).*
(2) *G is linear over $\mathbb{Z}$ (see [19, 39]), hence is residually finite.*
(3) *G has separable quasiconvex subgroups (see [39]).*
(4) *G is virtually residually finite rationally solvable (RFRS) (see [1]).*
(5) *G has the Haagerup property, hence does not have property (T) (see [15, 58]).*
(6) *G satisfies the strong Atiyah conjecture (see [68]).*
(7) *G is virtually bi-orderable (see [25]).*
(8) *G virtually embeds in the mapping class group of a closed surface, in a braid group, and in the group of diffeomorphisms of $\mathbb{R}$ (see [3, 49, 50]).*
(9) *G admits a proper affine action on $\mathbb{R}^n$ for some $n \geq 1$ (see [17]).*
(10) *G admits Anosov representations (see [24]).*

This is achieved in §5 via a study of the cube stabilizers for the action from Theorem 1.1 and using a criterion for improper actions from [35]. The special cubical complex in Theorem 1.2 comes from a geometric action on a CAT(0) cubical complex different from the one in Theorem 1.1; its dimension is in general larger than n and not easy to bound. We note that the fact that hyperbolized groups do not have property (T) was already observed by Belegradek in [9] without using cubical methods.

The use of strict hyperbolization (as opposed to non–strict hyperbolization procedures) is crucial here. Indeed, there are closed aspherical manifolds whose fundamental group is not Gromov hyperbolic and not residually finite (see [8, 55]). A well–known question by Gromov asks whether all Gromov hyperbolic groups are residually finite. Theorem 1.2 implies that the strict hyperbolization procedure introduced by Charney and Davis in [14] does not provide counterexamples to this question.

In [36] we obtained results analogous to Theorems 1.1 and 1.2 for the relative strict hyperbolization procedure that was considered in [9, 20] to construct aspherical manifolds with boundary and relatively hyperbolic groups.

1.1 The main arguments

The hyperbolization procedure in Theorems 1.1 and 1.2 is the composition of two hyperbolization procedures. The first one is Gromov's cylinder construction, which turns the simplicial complex K into a non–positively curved cubical complex $\mathcal{G}(K)$ (see §2.4, and [32, §3.4.A]). The second one is the strict hyperbolization procedure of Charney and Davis, which turns a non–positively curved cubical complex X into a locally CAT(-1) piecewise hyperbolic polyhedron X_Γ (see §3, and [14]). Here Γ is a certain uniform arithmetic lattice of simple type in $\mathrm{SO}_0(n, 1) = \mathrm{Isom}^+(\mathbb{H}^n)$, which needs to be chosen to define the strict hyperbolization procedure. The hyperbolization procedure in Theorem 1.1 is then given by $\mathcal{H}(K) = (\mathcal{G}(K))_\Gamma$, i.e. by the composition

$$K \mapsto X = \mathcal{G}(K) \mapsto X_\Gamma = (\mathcal{G}(K))_\Gamma\,,$$

and in this paper we are mostly concerned with the study of the second part, i.e. the strict hyperbolization of a cubical complex. Sections §3 and §4 lead to the following argument.

Proof of Theorem 1.1 Let K be an n-dimensional simplicial complex, which is compact, homogeneous, and without boundary. Then the cubical complex $X = \mathcal{G}(K)$ is an n-dimensional cubical complex, which is compact, homogeneous, and without boundary. Moreover, up to a barycentric subdivision of K, the cubical complex $X = \mathcal{G}(K)$ can be assumed to be foldable, i.e. to admit a combinatorial map $f : X \to \square^n$ to the standard cube which is injective on each cube (see Proposition 2.7).

Foldability provides a collection of subspaces of X that we call mirrors. A mirror is defined as a connected component of the full preimage in X of a codimension-1 face of the standard cube $\square^n$ under the folding $f : X \to \square^n$ (see §3.4). Mirrors of X give rise to nice locally convex codimension-1 subspaces of the hyperbolized complex X_Γ, which we still call mirrors. Lifting the collection of mirrors of X_Γ to the universal cover $\widetilde{X_\Gamma}$ of X_Γ provides a stratification of $\widetilde{X_\Gamma}$: a point is in the k-stratum if it is contained in $n-k$ mirrors (where $n = \dim K = \dim X = \dim X_\Gamma$).

We construct a dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ in which vertices are given by cells in this stratification, and edges correspond to codimension-1 inclusion of cells (see §4). The complex $\mathcal{C}(\widetilde{X_\Gamma})$ comes with a natural height function on its vertices, recording the dimension of the corresponding cell. In particular, the link of each vertex splits into an ascending sublink and a descending sublink. The former is flag because the cubical complex $X = \mathcal{G}(K)$ is non–positively curved, and the latter is flag because of a Helly property satisfied by collections of pairwise orthogonal hyperplanes in $\mathbb{H}^n$ (see Lemma 4.9). It follows that links of vertices in $\mathcal{C}(\widetilde{X_\Gamma})$ are flag (see Proposition 4.10), hence $\mathcal{C}(\widetilde{X_\Gamma})$ is a non–positively curved cubical complex. Moreover, the separation properties of the collection of mirrors (see §3.7) imply that $\mathcal{C}(\widetilde{X_\Gamma})$ is simply–connected, hence CAT(0) (see Theorem 4.29).

Finally, note that the action of $G = \pi_1(X_\Gamma) = \pi_1(\mathcal{H}(K))$ on $\widetilde{X_\Gamma}$ by deck transformations induces an action of G on $\mathcal{C}(\widetilde{X_\Gamma})$, as desired (see Lemma 4.30). □

The reader should note that the dimension of the dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ in Theorem 1.1 is the same as the dimension of the input simplicial complex K. Moreover, the Charney-Davis strict hyperbolization procedure from [14] relies on the careful choice of a suitable arithmetic lattice Γ. Such a lattice can be chosen with a certain flexibility, and the proof of Theorem 1.1 works for any choice of Γ for which the strict-hyperbolization procedure is defined.

The action from Theorem 1.1 is further studied in §5.

Proof of Theorem 1.2 First of all, let us prove the theorem under the hypothesis and in the setting of Theorem 1.1, i.e. with the additional assumption that K is homogeneous and without boundary. Then the Gromov hyperbolic group $G = \pi_1(X_\Gamma) = \pi_1(\mathcal{H}(K))$ acts on the dual CAT(0) cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ cocompactly and by cubical isometries.

The action is not proper, but the cube stabilizers can be identified with suitable cell stabilizers for the action of G by deck transformations on the universal cover $\widetilde{X_\Gamma}$ of $\mathcal{H}(K)$ (see §5.1). These stabilizers are quasiconvex subgroups both of G and of Γ (see §5.2). Arithmetic lattices like Γ are known to be virtually compact special by [40]. In particular, we obtain that cell stabilizers for the action of G on $\mathcal{C}(\widetilde{X_\Gamma})$ are virtually compact special. It then follows from [35, Theorem D] that G itself is virtually compact special (see Theorem 5.15).

Finally, let us prove the theorem for an arbitrary compact simplicial complex K, without additional assumptions. To this end, let K' be a compact and homogeneous simplicial complex in which K embeds as a subcomplex. (For instance, first embed K in the complete simplex on its vertex set, and then embed this simplex in a triangulation K' of a sphere.) Since Gromov's cylinder construction maps subcomplexes to locally convex subcomplexes, and the Charney-Davis hyperbolization preserves local convexity, we have that $\mathcal{H}(K)$ is a locally convex subspace of $\mathcal{H}(K')$. It follows that $G = \pi_1(\mathcal{H}(K))$ is a quasiconvex subgroup of $G' = \pi_1(\mathcal{H}(K'))$. But G' is hyperbolic and virtually compact special by the first part of this proof, hence G is virtually compact special too by Lemma 5.10. □

For the sake of clarity: the cubical complex that witnesses the specialness of G is **not** the cubical complex from Theorem 1.1. It is obtained via the construction in [35], and its dimension is in general higher than $n = \dim K$. One of the benefits of working with the dual CAT(0) cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ (as opposed to other available CAT(0) cubical complexes, such as $\widetilde{X}$) is that the stabilizers for the action of G on $\mathcal{C}(\widetilde{X_\Gamma})$ can be related to the stabilizers for the action on $\widetilde{X_\Gamma}$, which are more geometric in nature and easier to understand.

1.2 Classical applications of hyperbolization procedures

The interest in hyperbolization procedures is that they can be used to construct closed aspherical manifolds with various interesting properties. As a result of our Theorem 1.2, many applications of the strict hyperbolization procedure introduced by

Charney and Davis in [14] can be obtained with additional algebraic features (e.g. the properties (1)-(10) listed in Theorem 1.2). We now collect some of these applications.

1.2.1 Riemannian hyperbolization

The strict hyperbolization procedure introduced by Charney and Davis in [14] outputs a space with a metric which is locally CAT(-1) and piecewise hyperbolic: the space is obtained by gluing together copies of the hyperbolizing cube $\square^n_\Gamma$ (see §3). When the cell complex X used in the hyperbolization procedure is homeomorphic to a smooth manifold, the hyperbolized complex X_Γ is homeomorphic to a manifold too, but the locally CAT(-1) metric can a priori have singularities where the boundaries of different copies of the hyperbolizing cube $\square^n_\Gamma$ meet. It was recently shown by Ontaneda in [60] that the construction can be tweaked in such a way that the manifold X_Γ supports a smooth Riemannian metric with strictly negative sectional curvatures (possibly with respect to a different smooth structure).

This was used in [60, Corollary 5] to construct examples in any dimension $n \geq 4$ of closed Riemannian n–manifolds of pinched negative curvature which are "new" in the sense that they are not homeomorphic to any of the previously known examples of Riemannian manifold of negative curvature, such as closed real hyperbolic manifolds (or more generally locally symmetric spaces of rank 1), or the Gromov-Thurston branched covers in [33], or the examples of Mostow-Siu in [57] or Deraux in [22]. These manifolds are also distinct from the recent examples constructed by Stover–Toledo in [71, 72], as the latter are Kähler, while the result of strict hyperbolization cannot be Kähler by [9, Theorem 1.8]. Our construction does not require the smoothness provided by Ontaneda's work, but it is compatible with it, so we get the following.

Corollary 1.3 *For any $\varepsilon > 0$ and $n \geq 4$ there are closed Riemannian n-manifolds with the following properties*:

- *they have sectional curvatures in the interval* $[-1-\varepsilon, -1]$;
- *they are not homeomorphic to a locally symmetric space of rank* 1, *or one of the manifolds constructed by Gromov–Thurston*, *Mostow–Siu*, *Deraux*, *or Stover–Toledo*;
- *their fundamental groups are Gromov hyperbolic and virtually compact special* (*in particular, they satisfy properties* (1)-(10) *in Theorem* 1.2).

Remark 1.4 Thanks to the solution of the Borel Conjecture for closed aspherical n-manifolds with Gromov hyperbolic fundamental group in dimension $n \geq 5$ (see Bartels-Lück in [5]), the fundamental groups of these manifolds provide examples of Gromov hyperbolic groups that are not isomorphic to lattices in SO$(n, 1)$ or the other real simple Lie groups of rank 1. While it is not a priori clear from their construction whether these groups are linear, they actually turn out to be virtually compact special, hence linear over $\mathbb{Z}$ and residually finite. We note that Giralt proved in [31] that the fundamental groups of the Gromov-Thurston manifolds are also virtually compact special.

Similarly, other applications obtained by Ontaneda in [60] can be taken to have additional algebraic features. For example, we have the following versions of Corollary 2 in [60].

Corollary 1.5 *Let $\varepsilon > 0$. The cohomology ring of any finite CW–complex embeds in the cohomology ring of a closed Riemannian manifold which has sectional curvatures in $[-1-\varepsilon, -1]$ and whose fundamental group is Gromov hyperbolic and virtually compact special (hence satisfies properties (1)-(10) in Theorem 1.2). In particular, it can be embedded into the cohomology ring of a Poincaré Duality subgroup of $\mathrm{SL}_N(\mathbb{Z})$ (for N large).*

For another application in the spirit of Corollary 1.3 see [64], where a relative version of the Charney-Davis hyperbolization procedure is used to construct closed manifolds with hyperbolic fundamental group that do not admit any real projective or flat conformal structures, in any dimension at least 5. It follows from the present work (or from [36]) that the fundamental groups of these manifolds are virtually compact special too.

1.2.2 Pathological aspherical manifolds

Davis and Januszkiewicz used the hyperbolization procedures to construct aspherical manifolds exhibiting a variety of pathological behavior (see [18]). As a consequence of our Theorem 1.2, these examples can now be constructed to have the added property that their fundamental groups are virtually compact special, hence satisfy properties (1)-(10) from Theorem 1.2. For the convenience of the reader, we collate some of their examples.

Corollary 1.6 *It is possible to construct (topological) manifolds of the following types which are piecewise hyperbolic and locally* $\mathrm{CAT}(-1)$.

- *A closed 4-manifold which is not homotopy equivalent to any PL 4-manifold (see [18, §5a]).*
- *For $n = 4k$, $k \geq 2$, a closed n-manifold which is not homotopy equivalent to any smooth manifold (see [6, Example 5.2]).*
- *For $n \geq 5$, a closed n-manifold whose universal cover is not homeomorphic to $\mathbb{R}^n$ (see [18, §5b]).*
- *For $n \geq 5$, a closed n-manifold whose universal cover is homeomorphic to $\mathbb{R}^n$, but whose boundary at infinity is not homeomorphic to S^{n-1} (see [18, §5c]).*

Moreover, in all these examples, the fundamental groups of the manifolds are Gromov hyperbolic and virtually compact special (in particular, they satisfy properties (1)-(10) from Theorem 1.2).

Remark 1.7 Concerning the first example in Corollary 1.6, taking products with tori yields examples in all dimensions $n \geq 4$ of closed aspherical n-manifolds not homotopy equivalent to any PL n-manifold. These manifolds will have fundamental group which is linear over $\mathbb{Z}$, but when $n \geq 5$ will only support a locally $\mathrm{CAT}(0)$ metric due to the product structure. It would be interesting to produce examples in dimensions $n \geq 5$ which support locally $\mathrm{CAT}(-1)$ metrics.

1.2.3 Representing cobordism classes

As another application, we can obtain representatives for cobordism classes that are both topologically and algebraically nice.

Corollary 1.8 *Let M be an arbitrary closed smooth manifold. Then M is cobordant to an aspherical manifold M', where $\pi_1(M')$ is a Gromov hyperbolic and virtually compact special* (*in particular, it satisfies properties* (1)-(10) *from Theorem* 1.2).

Following an idea of Gromov (see [32, 61]), one lets K be the cone over a smooth triangulation τ of M. Then we apply the strict hyperbolization $\mathcal{H}(K)$, and note that since hyperbolization preserves links, the point $p \in \mathcal{H}(K)$ corresponding to the cone point will have link a copy of τ. Thus, removing a small neighborhood of p leaves us with a cobordism W between M and $M' := \mathcal{H}(\tau)$. Our Theorem 1.2 then applies to M'. Note that $\pi_1(W)$ itself contains $\pi_1(M)$, hence might not be linear (for instance, if $\pi_1(M)$ is a non–linear group). On the other hand, if $\pi_1(M)$ is residually finite, then $\pi_1(M')$ is also residually finite by our Theorem 1.2, and we showed in [36, Corollary 5.9] that in this case there is a cobordism between M and M' that has residually finite fundamental group.

Remark 1.9 Thom's work showed that oriented cobordism classes are rationally represented by products of even dimensional complex projective spaces (see [56, Sect. 17]). So every smooth oriented closed manifold has a multiple which is cobordant to a non–negatively curved Riemannian manifold. In analogy, combining Davis–Januszkiewicz–Weinberger [20], Charney–Davis [14], and Ontaneda [60], one obtains that every smooth oriented closed manifold is cobordant to a strictly negatively curved Riemannian manifold.

In dimensions ≥ 5, the Borel Conjecture is known to hold for aspherical manifolds with Gromov hyperbolic groups (see Bartels–Lück [5]). As such, the topological manifold $M' := \mathcal{H}(\tau)$ is completely determined, up to homeomorphism, by its fundamental group. So the discussion above in principle reduces the study of cobordism classes of manifolds of dimension $n \geq 5$, to the study of the corresponding $\pi_1(M')$. Our corollary further reduces it to the linear case.

Remark 1.10 More generally, Corollary 1.8 works for a PL manifold, or even for a triangulable topological manifold. Note that in all dimension $n \geq 4$ there exist closed topological manifolds that are not triangulable (see [30, 54]). Moreover, in all dimension $n \geq 6$ such manifolds can be chosen to be aspherical by [21], and have virtually compact special (hence residually finite) fundamental group by [36, Theorem 5.12].

1.2.4 Prescribing the Gromov boundary

The groups obtained by strict hyperbolization are Gromov hyperbolic groups, so it is natural to ask what their Gromov boundary looks like. For example, the groups obtained by Riemannian hyperbolization in [60] (see §1.2.1) are fundamental groups of smooth Riemannian manifolds of negative curvature, hence their Gromov boundaries are spheres of the appropriate dimensions.

Corollary 1.11 *Let $n \geq 1$, and let M be a closed connected orientable PL n-manifold that bounds a compact orientable PL $(n+1)$-manifold. Then there exists a Gromov hyperbolic group G such that*

- *the Gromov boundary of G is homeomorphic to the tree of manifolds $\mathcal{X}(M)$;*
- *G is virtually compact special (hence satisfies (1)-(10) in Theorem 1.2).*

The groups in this statement are the ones obtained by Świątkowski in [73] via strict hyperbolization of certain pseudomanifolds in which the link of a point is either a sphere or a copy of the manifold M. The tree of manifolds $\mathcal{X}(M)$ is a compact metrizable space which is obtained, roughly speaking, as a certain limit of connected sums of copies of M.

1.2.5 Manifolds with exotic symmetries

The hyperbolization procedures satisfy a certain functorial property: automorphisms of the simplicial complex K induce isometries of the hyperbolized complex $\mathcal{H}(K)$. This has been used by various authors to produce closed manifolds with interesting symmetries.

For example, if G is a Gromov hyperbolic group which is a Poincaré Duality group over $\mathbb{Z}$, an easy application of Smith theory shows that the fixed subgroup G^σ of an involution $\sigma \in Aut(G)$ is still a Poincaré Duality group, but over $\mathbb{Z}_2$. Farrell–Lafont in [26] used an exotic symmetry produced via strict hyperbolization, to give examples whose fixed subgroups are **not** Poincaré Duality over $\mathbb{Z}$. Our results now show that these examples can also be chosen to satisfy properties (1)-(10) in Theorem 1.2.

For another application, recall that in their seminal paper [7], Baum–Connes defined a trace map $tr : K_0(C_r^* G) \to \mathbb{R}$, where $C_r^* G$ is the reduced C^*-algebra of the discrete group G. They also formulated the *trace conjecture*, which predicted that when G is a group with torsion, the image of the trace map is contained in the additive *subgroup* of $\mathbb{Q}$ generated by $1/n$, where n ranges over the order of finite subgroups of G. A counterexample to this conjecture was constructed by Roy [63], using the Davis–Januszkiewicz (non–strict) hyperbolization procedure. She constructed a group G whose only finite subgroups are isomorphic to $\mathbb{Z}_3$, and an element in $K_0(C_r^* G)$ whose trace equals $-1105/9$. Nevertheless, there is always the possibility that the original Baum–Connes trace conjecture might hold for certain restricted classes of groups. The computations carried out by Roy ([63], pgs. 210-213) apply verbatim if one instead uses the Charney–Davis strict hyperbolization, so our results have the following consequence.

Corollary 1.12 *There exists a Gromov hyperbolic group G whose only finite subgroups are isomorphic to $\mathbb{Z}_3$, but where the image of the trace map contains $-1105/9$. Moreover, this group satisfies properties (1)-(10) in Theorem 1.2. In particular, the original Baum–Connes trace conjecture does not hold for the classes of groups (1)-(10) in Theorem 1.2.*

Remark 1.13 Lück formulated a refinement of the original Baum–Connes trace conjecture: the image of the trace map is contained in the *subring* $\mathbb{Z}[1/|Fin(G)|]$, obtained from $\mathbb{Z}$ by inverting all the orders of finite subgroups of G. Lück showed that this refined Trace Conjecture holds for any group that satisfies the Baum–Connes Conjecture (see [53]). In the subsequent literature, this refined version is what is commonly referred to as the Trace Conjecture. For Gromov hyperbolic groups, the Baum–Connes Conjecture was established by Lafforgue (see [52]). Thus the group appearing in our Corollary 1.12 satisfies the refined trace conjecture.

1.3 Virtual algebraic fibering

In this section we present new applications of a more algebraic flavor. We say a group G *algebraically fibers* if it admits a surjective homomorphism to $\mathbb{Z}$ with finitely generated kernel. We say it *virtually algebraically fibers* if it has a finite index subgroup that algebraically fibers. Agol introduced the notion of *residually finite rationally solvable* (or RFRS) group in [1] as a major ingredient in the solution of the Virtual Haken Conjecture and Virtual Fibering Conjecture. Kielak proved in [48, Thoerem 5.3] that a finitely generated virtually RFRS group virtually algebraically fibers if and only if its first L^2–Betti number vanishes. Fisher has extended this result in [28, Theorem A] to relate the vanishing of higher L^2–Betti numbers of G to higher finiteness properties of the kernel of a virtual algebraic fibration.

All of the groups constructed in this paper via strict hyperbolization are virtually compact special, hence virtually RFRS, see [1, Corollary 2.3]. In some cases, it is possible to prove vanishing of many L^2–Betti numbers (for instance for all the examples obtained by Ontaneda in [60], provided the curvatures are sufficiently pinched; see below for details). Hence, we get several new examples of virtually compact special Gromov hyperbolic groups that admit a virtual algebraic fibration, whose kernel has good algebraic finiteness properties. On the other hand, these groups can often be seen to be incoherent, and in some cases it is possible to see that the kernel of a virtual algebraic fibration is itself a witness to incoherence (i.e. is finitely generated but not finitely presented).

Before providing the details for our case, we note that similar arguments also work for arithmetic hyperbolic manifolds of simple type and for Gromov–Thurston manifolds. These are known to be virtually specially cubulated (hence RFRS) by [40] and [31] respectively.

1.3.1 Kernels with good algebraic finiteness properties

We start by constructing Gromov hyperbolic groups that virtually algebraically fiber, and are not isomorphic to groups that were previously known to have this property.

Corollary 1.14 *For all $n \geq 4$ there is a closed Riemannian n-manifold M with negative sectional curvatures and such that*

- $\pi_1(M)$ *virtually algebraically fibers*;
- $\pi_1(M)$ *is Gromov hyperbolic and virtually compact special* (*hence satisfies* (1)-(10) *in Theorem* 1.2);
- $\pi_1(M)$ *is not isomorphic to a uniform lattice in* SO$(n, 1)$ (*or other real simple Lie group of rank* 1), *or to the fundamental group of a Gromov-Thurston*, *Mostow-Siu*, *Deraux*, *or Stover–Toledo manifold.*

The manifolds in this statement are the ones constructed by Ontaneda in [60] (see Corollary 1.3 above). As a result of our Theorem 1.2 the fundamental group of such a manifold M is virtually compact special, and in particular it is virtually RFRS. Moreover, M can be chosen to have sectional curvatures pinched in the interval $[-1-\varepsilon, -1]$ for an arbitrarily small $\varepsilon > 0$. By a result of Donnelly and Xavier (see [23, §4], and also [45, Theorem 2.3]), if the curvatures of M are sufficiently pinched (i.e. ε is sufficiently small with respect to the dimension n), then M does not have any non–zero L^2–harmonic p-forms, for p in a certain range. In particular, $b_1^{(2)}(\pi_1(M)) = 0$. By [48] we see that if ε is small enough then $\pi_1(M)$ virtually algebraically fibers.

Furthermore, one can pinch the curvatures even more to force the vanishing of the L^2–Betti numbers for $p = 0, 1, \dots, \lfloor \frac{n}{2} \rfloor - 1$. In particular, using results from [28] one can obtain examples in which $\pi_1(M)$ virtually algebraically fibers with kernel of type $FP_{\lfloor \frac{n}{2} \rfloor - 1}(\mathbb{Q})$. Also note that in the even dimensional case M will then satisfy the (weak) Hopf conjecture, i.e. $(-1)^{\frac{n}{2}} \chi(M) \geq 0$.

1.3.2 Kernels witness incoherence in dimension 4

We have discussed how to use strict hyperbolization to obtain Gromov hyperbolic groups that virtually algebraically fiber with kernel of type $FP_{\lfloor \frac{n}{2} \rfloor}(\mathbb{Q})$. On the other hand, these kernels should not be expected to have better finiteness properties. Indeed, in the context of the previous paragraph, we can show that in dimension $n = 4$ these kernels are not finitely presented (i.e. not of type F_2).

To see this, notice that Chern–Weil theory implies that the Euler characteristic of a closed negatively curved 4-manifold is strictly positive (see [16]). This prevents the kernel of an algebraic fibration of $\pi_1(M)$ from being finitely presented, as we now describe. We thank Genevieve Walsh for sharing the following argument with us. (This appears in [51].)

Lemma 1.15 *Let M be a closed aspherical 4-manifold such that $\chi(M) \neq 0$. If $\pi_1(M)$ virtually algebraically fibers, then $\pi_1(M)$ is incoherent (the kernel is not finitely presented).*

Proof Suppose $\pi_1(M)$ virtually algebraically fibers, and let G be the finite index subgroup of $\pi_1(M)$ which surjects to $\mathbb{Z}$ with finitely generated kernel K. Notice that G is a PD_4 group with $\chi(G) \neq 0$ (since Euler characteristic is multiplicative by index), and that $\mathbb{Z}$ is a PD_1 group. Assume by contradiction that K is finitely presented (i.e. type F_2). Then K is in particular of type FP_2, and it follows from [41, Corollary 1.1] that K is a PD_3 group. In particular K has finite homological type (and the same is true for $\mathbb{Z}$). So, by the properties of Euler characteristics on short exact sequences (see [12, Chapter IX, 7.3(d)]) we can conclude that $\chi(G) = \chi(K)\chi(\mathbb{Z}) = 0$. This contradicts the fact that $\chi(G) \neq 0$. □

An alternative argument for this Lemma, under the additional assumption that $\pi_1(M)$ is virtually RFRS, was shared with us by Kevin Schreve.

Proof In the same set up, if by contradiction K is finitely presented, then it is in particular of type $\mathrm{FP}_2(\mathbb{Q})$. So, since G is also virtually RFRS, by [28, 48] we get that $b_1^{(2)}(G) = b_2^{(2)}(G) = 0$. But G is a PD_4 group, so by duality this implies that all L^2-Betti numbers vanish. This gives $\chi(G) = 0$, which is again absurd. □

As a result we obtain the following statement. The manifold in it is once again one of the manifolds obtained by Ontaneda, with curvatures sufficiently pinched.

Corollary 1.16 *There exists a closed* 4*-dimensional Riemannian manifold* M *with negative sectional curvatures and such that*

- $\pi_1(M)$ *is incoherent* (*it virtually algebraically fibers with non finitely presented kernel*);
- $\pi_1(M)$ *is Gromov hyperbolic and virtually compact special* (*hence satisfies* (1)-(10) *in Theorem* 1.2);
- $\pi_1(M)$ *is not isomorphic to a uniform lattice in* $\mathrm{SO}(n, 1)$ (*or other real simple Lie group of rank* 1), *or to the fundamental group of a Gromov-Thurston*, *Mostow-Siu*, *or Stover–Toledo manifold.*

Remark 1.17 The situation in dimension 4 is quite different from that in dimension 5. Indeed, Italiano, Martelli, Migliorini in [44] obtained a 5-dimensional cusped hyperbolic manifold that fibers over the circle. Its fundamental group algebraically fibers, with kernel of finite type (in particular finitely presented). The hyperbolic groups obtained by suitable Dehn filling on these examples were shown to fiber with kernel of finite type. Moreover, recent work of Groves and Manning shows that some of these groups are virtually compact special (see [34]).

Remark 1.18 When $n \geq 5$, the groups obtained by strict hyperbolization done with a sufficiently large piece (as in Ontaneda) contain subgroups isomorphic to uniform arithmetic lattices in SO(4, 1). The incoherence of these subgroups (see [1, 46, 47]) gives incoherence of the hyperbolized groups, but these subgroups are not fibers themselves (for instance because they are quasiconvex). Notice that this approach does not work in dimension $n = 4$, as uniform lattices in SO(3, 1) are coherent.

Structure of the paper This paper is structured as follows. In §1 we presented the motivation, the context, the statements, and the major applications of our results. In §2, we provide combinatorial and metric background about cell complexes and hyperbolization procedures. §3 is devoted to a description of Charney–Davis strict hyperbolization procedure for a cubical complex X. In particular, we study the geometry of the universal cover $\widetilde{X_\Gamma}$ of the hyperbolized complex X_Γ in terms of a certain collection of convex subspaces called mirrors. This provides a graph of spaces decomposition of X_Γ. In §4 we construct and study a dual CAT(0) cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$. Finally in §5 we study the action of the hyperbolized group $\Gamma_X = \pi_1(X_\Gamma)$ on this dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$, and prove that Γ_X is virtually compact special.

Common terminology and notation The numbers in parentheses refer to the section(s) in which each item is introduced or discussed.

- The hyperbolizing lattice Γ (§3.1) and the cubical complex X.
- The hyperbolized complex X_Γ (§3.2), and its universal cover $\widetilde{X_\Gamma}$ (§3.3).
- The hyperbolized cube $\square^n_\Gamma$ (§3.1), and its universal cover $\widetilde{\square^n_\Gamma}$ (§3.3).
- The hyperbolized groups $\Gamma_{\square^n} = \pi_1(\square^n_\Gamma)$ (§3.1) and $\Gamma_X = \pi_1(X_\Gamma)$ (§3.2).
- The folding map $f : X \to \square^n$ of a foldable complex (§2.2), and the induced map $f_\Gamma : X_\Gamma \to \square^n_\Gamma$ on the hyperbolized complex (§3.2,§3.3).
- The Charney-Davis map $g : \square^n_\Gamma \to \square^n$, and the induced map $g_X : X_\Gamma \to X$ on the hyperbolized complex (§3.2).
- The dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ (§4).
- A face $F \subseteq \square^n$.
- A cube $C \subseteq X$ (if X is a cubical complex) or $Q \subseteq \mathcal{C}(\widetilde{X_\Gamma})$.
- A cell $\sigma \subseteq X_\Gamma$, $\widetilde{X_\Gamma}$ (§3.5), $\mathbb{H}^n$, $\square^n_\Gamma$, $\widetilde{\square^n_\Gamma}$ (§3.3).
- A tile τ in $\widetilde{X_\Gamma}$ (§3.3), and the dual tile $\mathcal{C}(\tau)$ in $\mathcal{C}(\widetilde{X_\Gamma})$ (§4.2).
- A mirror M in $\widetilde{X_\Gamma}$ (§3.4), and the dual mirror $\mathcal{C}(M)$ in $\mathcal{C}(\widetilde{X_\Gamma})$ (§4.3).
- An edge–path p in $\mathcal{C}(\widetilde{X_\Gamma})$, its length $\ell(p)$, its height $\mathrm{h}(p)$ (§4.1, §4.2), the number of (p, M)–crossings $\mathfrak{m}(p, M)$ with respect to a mirror M, and its total mirror complexity $\mathfrak{m}(p)$ (§4.3).

2 Cell complexes and hyperbolization procedures

We collect in this section some background material used in our constructions. In §2.1 we review the basics about cell complexes, and in §2.2 we focus on foldable complexes, i.e. complexes that can be folded down to a standard simplex or cube. In §2.3 we introduce a general template for the study of hyperbolization procedures for foldable complexes. In §2.4 we review a specific hyperbolization procedure due to Gromov.

2.1 Combinatorial and metric geometry of cell complexes

In this section we collect background material about cell complexes, mainly to fix notation and terminology; for a detailed treatment see [11, §I.7, §II.5]. Let us denote by $\mathbb{M}^n_k$ the simply connected Riemannian manifold of dimension n and constant sectional curvature k: for instance $\mathbb{M}^n_1 = \mathbb{S}^n$ is the round sphere, $\mathbb{M}^n_0 = \mathbb{E}^n$ is the Euclidean space, and $\mathbb{M}^n_{-1} = \mathbb{H}^n$ is the hyperbolic space. An isometrically embedded copy of $\mathbb{M}^d_k$ inside $\mathbb{M}^n_k$ will be called a *d–plane*, or a *hyperplane* if $d = n - 1$.

2.1.1 Cells

A *cell* in $\mathbb{M}^n_k$ is defined to be the convex hull of a finite set of points; if $k > 0$ we are going to also require that it is contained in an open ball of radius $\frac{\pi}{2\sqrt{k}}$. The *dimension* of a cell C is the smallest d such that C is contained in a d–plane. A cell of dimension d will also be called a *d–cell*. The *interior* of C is its interior inside this d–plane. A *face* F of C is a subspace of the form $F = H \cap C$ where H is a hyperplane such that C lives in one of the two closed half–spaces bounded by H, and $H \cap C \neq \varnothing$. A face is itself a cell, and we call *vertices* and *edges* of C the faces of dimension 0 and 1 respectively.

2.1.2 Cell complexes

An $\mathbb{M}^n_k$-*cell complex* is a topological space X obtained by gluing together cells from $\mathbb{M}^n_k$ by isometries of their faces, in such a way that each cell embeds in X and the intersection of any two cells is either empty or a cell. Notice that this definition is slightly more restrictive than the one in [11, Definition I.7.37] (which allows one to glue two faces of the same cell), and the one in [14] (in which cells are allowed to intersect in a proper union of faces). Both conditions can be satisfied by performing a cellular subdivision. On the other hand, we do not require cell complexes to be locally compact at this stage, i.e. a vertex can be contained in infinitely many cells.

We call an $\mathbb{M}^n_k$-cell complex simply a *cell complex* when we do not need to keep track of $\mathbb{M}^n_k$. For instance we will denote by $\triangle^n$ the standard n-simplex and by $\square^n = [0, 1]^n$ the standard n-cube; these are cells in $\mathbb{M}^n_0$. A *simplicial complex* is a cell complex obtained by gluing simplices, and a *cubical complex* is a cell complex obtained by gluing cubes.

The *dimension* of a cell complex is the maximum dimension of its cells. We say that an n-dimensional cell complex is *homogeneous* if every cell is contained in a cell of dimension n, and that it is *without boundary* if every $(n-1)$-cell is contained in at least two different n-cells. For all $k = 0, \dots, n$, the k-*skeleton* of X is the subspace consisting of all the cells of dimension at most k, and will be denoted by $X^{(k)}$. A *subcomplex* of X is a closed subspace $Y \subseteq X$ which is a union of cells of X. If X and Y are cell complexes, a continuous function $f : X \to Y$ is a *combinatorial map* if for every cell C of X we have that f is a homeomorphism from C to a cell $f(C)$ of Y.

Given a cell complex X, its *barycentric subdivision* $\mathcal{B}(X)$ is the simplicial complex whose k-simplices correspond to strictly ascending sequences of faces $F_0 \subset \cdots \subset F_k$ of X. There exists a natural (non–combinatorial) homeomorphism $X \to \mathcal{B}(X)$. We refer to [11, §I.7.44-48] for more details. Similarly, if X is a cubical complex, then its *cubical subdivision* is the cubical complex obtained by subdividing each n–cube along midcubes into 2^n cubes.

Remark 2.1 By definition, a cell is compact, it has finitely many faces, and it can be realized as the intersection of finitely many closed half-spaces (see [11, §I.7]). We want to warn the reader that one of the main object under investigation in this paper (see §3.5) is obtained by gluing together certain "generalized cells", i.e. subsets of $\mathbb{H}^n$ which are given by the intersection of an infinite but locally finite collection of closed half-spaces. These subsets are convex but not compact, so the resulting space is not strictly speaking a cell complex. However, some of the usual tools for the study of cell complexes can be applied in this context (e.g. links). We will highlight the subtleties in the construction whenever relevant.

2.1.3 Links

Let X be a cell complex. We define the link of points and cells as follows (see [11, §I.7] for more details). Let p be a point of an n–cell $C \subseteq \mathbb{M}^n_k$. Then we define the *link* $\mathrm{lk}\,(p, C)$ to be the space of unit vectors in the tangent space at p inside C. Measuring the angle between vectors endows $\mathrm{lk}\,(p, C)$ with a natural length metric, which makes it isometric to an $(n-1)$–cell in $\mathbb{S}^{n-1}$. The *link* $\mathrm{lk}(p, X)$ of p in X is then defined by gluing together the links $\mathrm{lk}(p, C_i)$, where C_i ranges over the cells of X containing p. This is naturally an $\mathbb{M}_1^{n-1}$–cell complex. If Y is a sufficiently regular subspace of X containing p (e.g. a subcomplex), then the *link* $\mathrm{lk}(p, Y)$ is defined analogously, restricting to vectors along Y.

Let F be a d–face of an n–cell $C \subseteq \mathbb{M}_k^n$. Then we define the *link* $\mathrm{lk}(F, C)$ to be the subspace of unit vectors in the tangent space at an interior point of F, which are pointing into C and are orthogonal to F. As before, this is naturally an $(n-d-1)$–cell in $\mathbb{S}^{n-1}$. The *link* $\mathrm{lk}(C, X)$ of a d–cell $C \subseteq X$ is then defined by gluing together the links $\mathrm{lk}(C, C_i)$, where C_i ranges over the cells of X containing C. It is naturally an $\mathbb{M}_1^{n-d-1}$–cell complex. Finally, if $Y \subseteq X$ is a subcomplex of X containing C, the *link* $\mathrm{lk}(C, Y)$ of C in Y is defined analogously, by restricting to the cells of Y that contain C. Observe that if X is a simplicial or cubical complex, then the link of a d–cell C is a simplicial complex in which vertices are given by the $(d+1)$–cells containing C, and in which $m+1$ vertices span an m–simplex if and only if the corresponding cells are contained in a $(d+m+1)$–cell.

2.1.4 Spaces and complexes of bounded curvature

We will consider the usual notions of curvature for metric spaces, such as being locally CAT(k) or Gromov hyperbolic (see [11, §II.1, §III.H.1] for more details). In particular, we will say a space is *non–positively curved* if it is locally CAT(0), and *negatively curved* if it is locally CAT(k) for some $k < 0$. Note that if $k < 0$ then a CAT(k) space is in particular Gromov hyperbolic (see [11, Proposition III.H.1.2]), and that if $k \leq 0$ then a CAT(k) space is uniquely geodesic (see [11, Proposition II.1.4]). Whenever x, y are points in a uniquely geodesic space, we denote by $[x, y]$ the unique geodesic between them.

Let X be an $\mathbb{M}_k^n$-cell complex. Each cell can be naturally endowed with a metric from $\mathbb{M}_k^n$, and these can be glued together to make X into a complete geodesic metric space, as soon as there are only finitely many isometry classes of cells in X (see [11, Theorem I.7.50]). When equipped with this metric, X is said to be a cell complex of piecewise constant curvature k; we say it is *piecewise spherical*, *Euclidean*, or *hyperbolic* when $k = 1, 0, -1$ respectively. If not otherwise specified, a simplicial or cubical complex is always endowed with its standard piecewise Euclidean metric.

It is natural to ask for conditions under which a complex of piecewise constant curvature is a space of bounded curvature, namely a locally CAT(k) space. For cubical complexes this is completely controlled by the links of vertices. In a cubical complex, cells are isometric to the standard Euclidean cube $\square^n = [0, 1]^n$, so the link of a vertex is a piecewise spherical simplicial complex, in which all edges have length $\frac{\pi}{2}$. The following is known as Gromov's link condition (see [11, Theorems II.5.18, II.5.20]). A simplicial complex is *flag* if any $k+1$ pairwise adjacent vertices span a k–simplex.

Lemma 2.2 *Let X be a cubical complex. Then the following are equivalent.*

(1) *X is non–positively curved* (*i.e. locally* CAT(0)).
(2) *The link of each vertex is a flag simplicial complex.*
(3) *The link of each vertex is a* CAT(1) *simplicial complex.*

2.2 Foldable complexes

Here we consider the notion of foldability for simplicial and cubical complexes that we will require later. The first definition is essentially from [4, §1] (but see also [14, Definition 7.2], and [76] for a more recent discussion).

A simplicial (respectively cubical) n-dimensional complex X is *foldable* if it admits a combinatorial map $f : X \to \triangle^n$ (respectively $f : X \to \square^n$) such that its restriction to each cell of X is injective. Such a map will be called a *folding* for X. Notice that in a foldable complex the cells are necessarily embedded. This is the main reason why we have incorporated this condition in the definition of cell complex in §2.1.

Foldability has some immediate consequences. If X is foldable, and $p : Y \to X$ is a combinatorial map which is injective on each cell, then Y is foldable too. In particular any covering of a foldable complex is foldable. Moreover if X is foldable, then the links of cells of codimension 2 are bipartite graphs. We collect below some examples in the cubical case; analogous ones can be constructed for the simplicial case.

Example 2.3 (Foldable and not foldable cubical complexes)

(1) A graph is foldable if and only if it is bipartite (Fig. 1, left).
(2) The rose R_m consisting of m squares with a vertex in common is foldable if and only if m is even (Fig. 1, right).
(3) Foldability of X implies that links of codimension 2 cells are bipartite. However, foldability is not completely determined by this property. For example, let X be the cubical complex obtained by taking the product $\partial\Delta^2 \times \mathbb{R}$, where $\mathbb{R}$ is endowed with the standard cell structure induced by $\mathbb{Z}$. Then the links of vertices are cycles of length 4 (hence they are bipartite), but X is not foldable; notice that the universal cover of X identifies the square complex defined by $\mathbb{Z}^2$ in $\mathbb{R}^2$, which is foldable (compare [4, Lemma 1.2]).

A main source of foldability comes from barycentric subdivisions; the following is well–known (see [4, Lemma 2.1]), we include a proof for completeness (see left of Fig. 2 for an example).

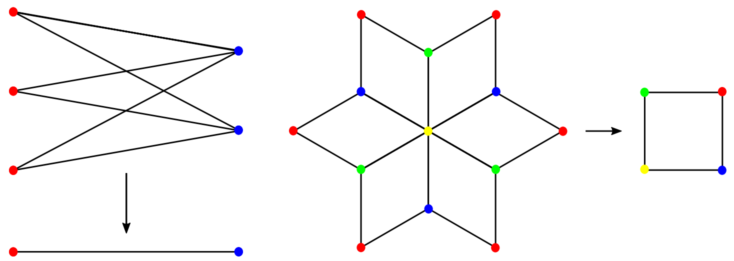

Fig. 1 Foldable cubical complexes

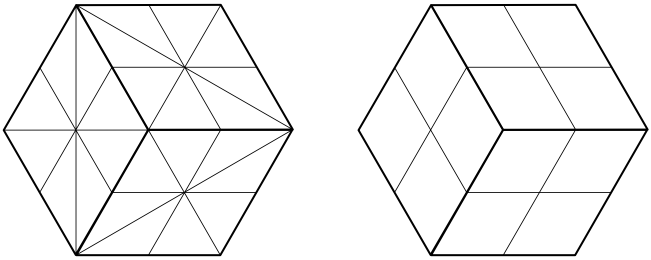

Fig. 2 The barycentric subdivision of the rose of 3 squares is a foldable simplicial complex, but its cubical subdivision is not a foldable cubical complex

Lemma 2.4 *If X is a cell complex, then $\mathcal{B}(X)$ is a foldable simplicial complex.*

Proof Let X have dimension n, and consider the simplex spanned by $\{0, \ldots, n\}$; this is just the standard simplex Δ^n. Then we can define a map $f : \mathcal{B}(X) \to \Delta^n$ by sending a vertex of $\mathcal{B}(X)$ to the number which is equal to the dimension of the corresponding cell in X. □

On the other hand, if X is a non–foldable cubical complex of dimension at least 2, then its cubical subdivision is still non–foldable (see Fig. 2, right). In §2.4 we will review Gromov's construction and show that it can be used to turn any cubical complex into a foldable one (mildly changing the topology).

2.3 Hyperbolization procedures

In this section we set a framework for the study of certain constructions, which take a cell complex as input and return a non-positively curved space as output. The resulting space is in particular always aspherical, so the topology of the original complex is altered. What is interesting is that this can happen in a controlled way that allows to preserve some features of the original complex. Constructions of this type are generally known as hyperbolization procedures (or asphericalization procedures). They were first introduced by Gromov (see [32, §3.4.A]), and then popularized by several authors (see [14, 18, 21, 60, 61]).

All the hyperbolization procedures we will consider in this paper are obtained by different incarnations of the same abstract construction, which we now review briefly, referring the reader to [74] or [18, §1] for more details. The naive idea is to fix some topological space S and then replace every top-dimensional cell of a complex X by a copy of S. For this gluing to be well–defined, it is common to assume that both X and S come equipped with chosen maps to a reference space.

For concreteness let us consider the following set up. Let us denote by σ^n the standard simplex Δ^n or the standard cube $\square^n$, and let us fix a topological space S, equipped with a continuous map $g : S \to \sigma^n$, and a foldable simplicial or cubical complex X, equipped with a fixed folding $f : X \to \sigma^n$. One then considers the fibered product $\mathcal{H}_S(X) = \{(x, s) \in X \times S \mid f(x) = g(s)\}$, i.e. the space obtained via the pullback square in Fig. 3.

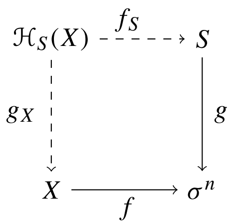

Fig. 3 A template for hyperbolization procedures

Note that the construction endows $\mathcal{H}_S(X)$ with natural continuous maps $g_X : \mathcal{H}_S(X) \to X$ and $f_S : \mathcal{H}_S(X) \to S$, which are just the restrictions of the projections onto the factors of $X \times S$, and which make the diagram commute. Properties of the pair (S, g) will result in properties of the space $\mathcal{H}_S(X)$, and the art of hyperbolization consists in crafting a pair (S, g) which yields some interesting properties on $\mathcal{H}_S(X)$. For a trivial example, consider the case S consists of a single point. Then $\mathcal{H}_S(X)$ is just the discrete set $f^{-1}(g(S))$.

The following lemma identifies a mild condition under which the space $\mathcal{H}_S(X)$ looks like a collection of copies of S (compare the remark on page 321 of [74]). We explicitly remark that we do not assume S to be compact until part (3) of this lemma. This will be relevant in §3.3 for the study of a certain combinatorial decomposition of a space into non–compact pieces.

Lemma 2.5 *Let $g : S \to \sigma^n$ be surjective, and let $C \subseteq X$ be an n–cell. Then the following hold.*

(1) *The map f_S restricts to a homeomorphism $\varphi : g_X^{-1}(C) \to S$.*
(2) *The map $\varphi^{-1} \circ f_S : \mathcal{H}_S(X) \to g_X^{-1}(C)$ is a retraction.*
(3) *If X and S are compact, then $\mathcal{H}_S(X)$ is compact too.*

Proof Let us denote by f_C the restriction of the folding map f to C. Note that $f_C : C \to \sigma^n$ is a homeomorphism. To prove (1), let $\varphi : g_X^{-1}(C) \to S$ be the restriction of f_S to $g_X^{-1}(C)$. Then φ is continuous, because it is just the restriction of the projection $X \times S \to S$. Injectivity and surjectivity of φ follow respectively from those of f_C. To conclude, we construct an explicit continuous inverse. Consider the map $\lambda : S \to C$, $\lambda(s) = f_C^{-1}(g(s))$. Notice it is well–defined (because g is surjective), and continuous. Then the map $\psi : S \to g_X^{-1}(C) \subseteq X \times S$, $\psi(s) = (\lambda(s), s)$ provides a continuous inverse to φ.

Now (2) follows from (1), as every element of $g_X^{-1}(C)$ is of the form $(\lambda(s), s)$. Finally, to prove (3), observe that if X is compact, then it consists of finitely many n–cells. As a result of the previous argument, $\mathcal{H}_S(X)$ is covered by finitely many copies of the compact space S, hence it is compact. □

Depending on the applications in which they are interested, authors differ on what additional geometric conditions they require on the association $X \to \mathcal{H}_S(X)$, hence they start with different spaces (S, g). We refer the reader to [18] for a very general treatment of how properties of (S, g) imply properties of $\mathcal{H}_S(X)$. Some commonly required conditions are the following

Fig. 4 Example for Remark 2.6: a face of the hyperbolizing cell S is not connected and $\mathcal{H}_S(X)$ is not simply connected. The maps f and g here are defined by the vertex coloring

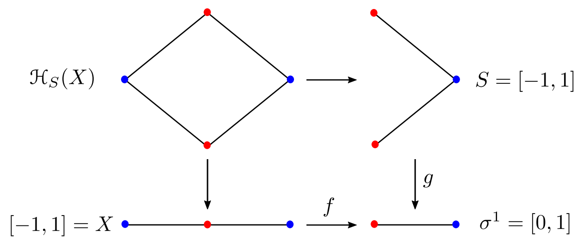

(1) (Hyperbolicity): $\mathcal{H}_S(X)$ admits a non-positively curved metric.
(2) (Functoriality): if $Z \subseteq X$ is a locally convex subcomplex, then $\mathcal{H}_S(Z) \subseteq \mathcal{H}_S(X)$ is a locally convex subspace.
(3) (Local structure): if $C \subseteq X$ is an n–cell, then $\mathcal{H}_S(C)$ is an n-manifold with boundary and corners, and $\operatorname{lk}(\mathcal{H}_S(C), \mathcal{H}_S(X)) \cong \operatorname{lk}(C, X)$. In particular, if X is a manifold, then $\mathcal{H}_S(X)$ is a manifold too.
(4) (Homology surjectivity): the map $g_X : \mathcal{H}_S(X) \to X$ induces a surjection on homology.

The association $X \to \mathcal{H}_S(X)$ is then called the *hyperbolization procedure* defined by (S, g). We call S the *hyperbolizing cell*, and $\mathcal{H}_S(X)$ the *hyperbolized complex*. Despite the name (established in the literature), the output $\mathcal{H}_S(X)$ of a hyperbolization procedure is a metric space which a priori is just non-positively curved. A *strict hyperbolization* is one for which $\mathcal{H}_S(X)$ is negatively curved. In this paper we will consider a (non–strict) hyperbolization for simplicial complexes due to Gromov (see §2.4), and a strict hyperbolization for cubical complexes due to Charney and Davis (see §3).

Remark 2.6 If (S, g) is a given hyperbolizing cell, $g : S \to \sigma^n$ is surjective, and $F \subseteq \sigma^n$ is a closed face of the n–cell σ^n, then the subspace $g^{-1}(F)$ will be called a *face* of S. The dimension of a face of S is defined to be simply the dimension of the corresponding face of σ^n. Note that a face of S does not need to be connected. When this happens, $\mathcal{H}_S(X)$ may fail to be simply connected, even if both X and S are. For some interesting examples, see [18, 1b.1], or consider the elementary one in Fig. 4. Despite their non–trivial role in the construction, most of the times the maps f and g are omitted from the notation.

2.4 Gromov's cylinder construction

In this section we review a construction, due to Gromov, which turns a simplicial complex K into a foldable cubical complex $\mathcal{G}(K)$ having non-positive curvature (see [32, §3.4.A] for the original source, or [18, §4c], [61, §4], and references therein, for expository accounts).

The construction uses induction on dimension and pullback simultaneously, following this scheme. For each dimension $n \geq 1$ we will first define $\mathcal{G}(\Delta^n)$ and a map $g : \mathcal{G}(\Delta^n) \to \Delta^n$, then for any foldable n–dimensional simplicial complex K, with a folding $f : K \to \Delta^n$, we will define $\mathcal{G}(K)$ via the pullback square (compare §2.3)

$$\begin{array}{ccc} \mathcal{G}(K) & \overset{f_{\mathcal{G}(\triangle^n)}}{\dashrightarrow} & \mathcal{G}(\triangle^n) \\ {\scriptstyle g_K}\big\downarrow & & \big\downarrow{\scriptstyle g} \\ K & \underset{f}{\longrightarrow} & \triangle^n \end{array}$$

Finally, for a general K (not necessarily foldable), we will define $\mathcal{G}(K) = \mathcal{G}(\mathcal{B}(K))$ (recall that the barycentric subdivision is always foldable by Lemma 2.4). Note that in any case the construction equips $\mathcal{G}(K)$ with a natural map to $\triangle^n$.

For $n = 1$ we set $\mathcal{G}(\triangle^1) = \triangle^1$, and we define $g : \mathcal{G}(\triangle^1) \to \triangle^1$ to be just the identity. By the pullback construction this defines $\mathcal{G}(K)$ and a map $g : \mathcal{G}(K) \to \triangle^1$ for all simplicial graphs K. Concretely, when K is a simplicial graph, then $\mathcal{G}(K) = K$ if K is bipartite, and $\mathcal{G}(K) = \mathcal{B}(K)$ otherwise; the folding to $\triangle^1$ is induced by the bipartition.

Let us now assume by induction that for any simplicial complex K of dimension at most $n-1$ the space $\mathcal{G}(K)$ is defined, and is endowed with a map to $\triangle^{n-1}$. In order to define $\mathcal{G}(\triangle^n)$, consider a reflection of $\partial\triangle^n$, and induce a reflection on $\mathcal{G}(\partial\triangle^n)$. Let U, V be the two closed half-spaces exchanged by the reflection, and define

$$\mathcal{G}(\triangle^n) = \mathcal{G}(\partial\triangle^n) \times [-1, 1]/\sim$$

where $(u, t) \sim (u', t')$ if and only if $|t| = |t'| = 1$ and $u = u' \in U$. Notice that taking a further quotient which identifies also points on V, one would get a map $\mathcal{G}(\triangle^n) \to \mathcal{G}(\partial\triangle^n) \times S^1$, and we can think of $\mathcal{G}(\triangle^n)$ as being obtained from $\mathcal{G}(\partial\triangle^n) \times S^1$ by cutting a slit in it along a half–fiber (see Fig. 5).

By induction, $\mathcal{G}(\partial\triangle^n)$ is well–defined, and it comes with a map $\mathcal{G}(\partial\triangle^n) \to \triangle^{n-1}$. Notice that the boundary of $\mathcal{G}(\triangle^n)$ consists of two copies of V, glued along a subspace identifiable with $U \cap V$. In other words, $\partial\mathcal{G}(\triangle^n)$ can be naturally identified with $\mathcal{G}(\partial\triangle^n)$, hence $\partial\mathcal{G}(\triangle^n)$ comes with a map to $\partial\triangle^n$. This map can be extended to a map $\mathcal{G}(\triangle^n) \to \triangle^n$ as follows: take a regular neighborhood $N \cong \partial\mathcal{G}(\triangle^n) \times [0, 1]$

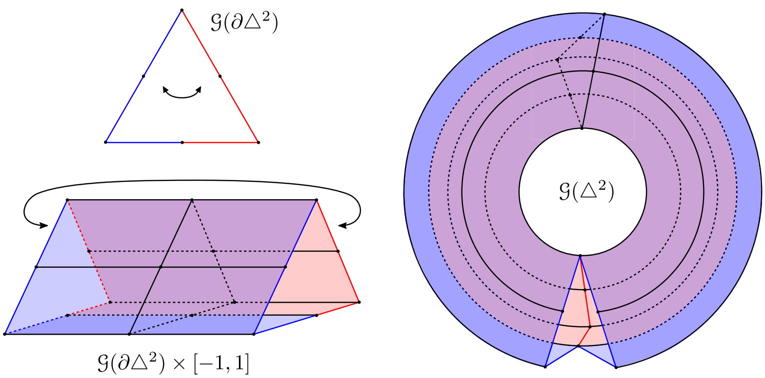

Fig. 5 Gromov's cylinder construction

inside $\mathcal{G}(\triangle^n)$, and identify $\triangle^n$ with the cone over $\partial\triangle^n$. Then extend the map over N along the cone direction, and collapse the complement of N to the cone point. This completes the construction of $\mathcal{G}(\triangle^n)$ and a map $g : \mathcal{G}(\triangle^n) \to \triangle^n$. Arguing as above (i.e. with the template from §2.3), this also defines $\mathcal{G}(K)$ for any simplicial complex K.

Proposition 2.7 *If K is a simplicial complex, then $\mathcal{G}(K)$ is a foldable cubical complex of non–positive curvature. Moreover if K is homogeneous (respectively without boundary, locally compact, or compact), then $\mathcal{G}(K)$ is also homogeneous (respectively without boundary, locally finite, or compact).*

Proof First we show that $\mathcal{G}(K)$ admits the structure of a cubical complex, starting with the case $K = \triangle^n$. This is clear for $\mathcal{G}(\triangle^1) = [0, 1] = \square^1$. Then, arguing by induction, $\mathcal{G}(\triangle^n)$ inherits a cubical structure from the one of $\mathcal{G}(\partial\triangle^n) \times [-1, 1]$. Here we think of $[-1, 1]$ as being given the standard cubical structures as a union of two unit intervals, and we give $\mathcal{G}(\partial\triangle^n) \times [-1, 1]$ the standard cubical structure coming from the fact that $\square^{n-1} \times \square^1 = \square^n$. Since $\mathcal{G}(K)$ is in general defined via the pullback construction (see §2.3), it inherits a natural cubical structure from $\mathcal{G}(\triangle^n)$.

We now prove that the cubical complex $\mathcal{G}(K)$ has the desired properties. Foldability is proven in [14, Lemma 7.5]. Non–positive curvature is proven in [18, Proposition 4c.2(3)]. For the other properties we argue as follows. Note that for each n the hyperbolizing cell $\mathcal{G}(\triangle^n)$ is homogeneous, has a single boundary component, and satisfies $\partial\mathcal{G}(\triangle^n) = g^{-1}(\partial\triangle^n)$. So, if K is homogeneous then $\mathcal{G}(K)$ is homogeneous, and if K is without boundary, the same holds for $\mathcal{G}(K)$. It is proved in [18, Lemma 1e.1 and §4c] that Gromov's construction preserves the local structure (e.g. links). This implies that if K is locally finite, then so is $\mathcal{G}(K)$. In particular, by (3) in Lemma 2.5, if K is compact, then so is $\mathcal{G}(K)$, because $\mathcal{G}(\triangle^n)$ is compact. □

We have defined Gromov's construction for simplicial complexes. Given any cell complex X we can first take its barycentric subdivision $\mathcal{B}(X)$ (which is a simplicial complex), and then apply Gromov's construction to it.

Corollary 2.8 *If X is any cell complex, then $\mathcal{G}(\mathcal{B}(X))$ is a foldable cubical complex of non–positive curvature. Moreover if X is homogeneous (respectively without boundary, locally compact, or compact), then $\mathcal{G}(\mathcal{B}(X))$ is also homogeneous (respectively without boundary, locally compact, or compact).*

Proof We know $K = \mathcal{B}(X)$ is a (foldable) simplicial complex (by Lemma 2.4), homeomorphic to X. Then the statements follow from Proposition 2.7. □

Gromov's construction is known to satisfy even more properties, namely conditions (1)-(6) in [14] and (1), (2'), (3)-(5) in [18]. Some of these are versions of conditions (1)–(4) from §2.3, while others deal with preservation of stable tangent bundles and rational Pontryagin classes, when they are defined. This is needed in the applications of the hyperbolization procedure to construct examples of closed aspherical manifolds with various prescribed features or pathologies (as in [18, 60]).

3 Strict hyperbolization

The hyperbolization procedure introduced by Charney and Davis in [14] is defined for cubical complexes, and fits in the framework outlined in §2.3, in the sense that it is determined by the choice of a hyperbolizing cell. Differently from Gromov's cylinder construction (described in §2.4), this procedure is not defined by induction. Rather, for each dimension $n > 0$ a hyperbolizing cell is defined independently, and defines a hyperbolization procedure for n–dimensional cubical complexes.

While the original construction is a bit more general than the version we use here, we find it convenient to impose some mild restrictions on the cubical complex in order to simplify the presentation. From now on assume X is *admissible*, i.e. it satisfies the following conditions (see §2 for definitions):

(1) cubical;
(2) homogeneous, without boundary;
(3) foldable;
(4) non-positively curved;
(5) locally compact.

This setting, consistent with that of [76], is more general than the one in [4], as we do not require gallery–connectedness. In particular, we allow X to be a pseudomanifold. On the other hand, the first two conditions are a bit more restrictive than the corresponding ones in [14], while the other ones are the same. More precisely, if X is foldable, then necessarily cubes of X are embedded. In [14] they allow two cubes to meet in a proper union of faces; note that such faces have to be disjoint in each cube, because non–positive curvature guarantees that links of vertices are simplicial. In particular, up to performing cubical subdivision, one can always assume that X is cubical. Finally we remark that at this stage we are only assuming local finiteness instead of compactness of X. While in our main theorems (Theorems 1.1 and 1.2) we assume that the complex is compact (in order to get a hyperbolic group), most of the geometric and combinatorial arguments do not need that, and in §5.8 we actually need to consider a certain hyperbolization of $\mathbb{R}^n$.

The main contribution of this section is to define some subspaces of the space that results from strict hyperbolization on an admissible cubical complex X. We call such subspaces *mirrors*, and prove that their lifts to the universal cover are convex and separating (see Proposition 3.14 and Proposition 3.37 respectively). Along the way, we also study a combinatorial decomposition of the universal cover (see §3.3 and §3.5) that will be the starting point for the construction of the dual cubical complex in §4.

3.1 The hyperbolizing cell

The hyperbolizing cell used in this hyperbolization procedure is a certain hyperbolic manifold with boundary and corners, obtained by cutting a closed hyperbolic manifold along a suitable collection of pairwise orthogonal totally geodesic codimension–1 submanifolds. While the existence of such an object is clear in dimension 2 (see Fig. 6), the construction in higher dimension requires some arithmetic methods involving quadratic forms (see §3.1.1 below for more details). Specifically, the construction relies on the choice of a suitable congruence subgroup Γ of an arithmetic lattice in $SO_0(n, 1)$, so we will denote the hyperbolizing cell by $\square^n_\Gamma$. Here and in the following we denote by B_n the group of Euclidean isometries of the standard cube $\square^n$. Also recall from Remark 2.6 that a k–face of $\square^n_\Gamma$ is by definition a subspace of the form $g^{-1}(\square^k)$, where $\square^k$ is a k–face of $\square^n$.

Lemma 3.1 (Corollary 6.2 in [14]) *For every $n \geq 2$ there exists a compact, connected, orientable hyperbolic n–manifold with corners $\square^n_\Gamma$, an isometric action of B_n on $\square^n_\Gamma$, and a B_n–equivariant and face–preserving map $g : \square^n_\Gamma \to \square^n$, such that the following hold.*

(1) *The poset of faces of $\square^n_\Gamma$ is B_n–equivariantly isomorphic to that of $\square^n$.*
(2) *Each face of $\square^n_\Gamma$ is totally geodesic.*
(3) *The faces of $\square^n_\Gamma$ intersect orthogonally.*
(4) *Each 0–dimensional face is a single point.*
(5) *The map $g : \square^n_\Gamma \to \square^n$ and its restriction to each face have degree one.*

We call $\square^n_\Gamma$ the *hyperbolizing cube*, and g the *Charney–Davis map*. We denote by $\Gamma_{\square^n} = \pi_1(\square^n_\Gamma)$ the fundamental group of the hyperbolizing cube.

Remark 3.2 In this hyperbolization procedure, a k–face of $\square^n_\Gamma$ is guaranteed to be connected when $k = 0, n$, but may be disconnected otherwise (see Remark 2.6, and the Remark after Corollary 6.2 in [14]). Nevertheless, by abuse of notation, we will denote by $\square^k_\Gamma = g^{-1}(\square^k)$ the k–face of $\square^n_\Gamma$, even when $0 < k < n$. Notice that $\square^k_\Gamma$ is a priori different from the k–dimensional hyperbolizing cube, i.e. the hyperbolizing cell that one obtains by hyperbolizing a k–dimensional cube with a hyperbolizing lattice $\Lambda \subseteq SO_0(k, 1)$ for $0 < k < n$. Namely, $\square^k_\Lambda$ is always connected by construction. Finally, with respect to (5) in Lemma 3.1, when $\square^k_\Gamma$ is disconnected, there is a preferred component of $\square^k_\Gamma$ on which g has degree one, while it has degree zero on the other components (see [14, Lemma 5.9(b)] and §3.1.1 for details).

3.1.1 The construction of $\square^n_\Gamma$

To construct the hyperbolizing cube $\square^n_\Gamma$, Charney and Davis consider the hyperboloid model for $\mathbb{H}^n$ inside Minkowski space $\mathbb{R}^{n,1}$, i.e. the space $\mathbb{R}^{n+1}$ equipped with a quadratic form of signature $(n, 1)$. The isometry group of $\mathbb{R}^{n,1}$ is naturally identified with the indefinite orthogonal group $O(n, 1)$, and its connected component $SO_0(n, 1)$ is naturally identified with the group of orientation preserving isometries of $\mathbb{H}^n$. Then they show that $SO_0(n, 1)$ contains an arithmetic lattice Γ which enjoys some key properties, from which the properties of $\square^n_\Gamma$ in Lemma 3.1 follow. In particular, Γ is a cocompact torsion–free lattice of $SO_0(n, 1)$, whose normalizer in $O(n, 1)$ contains all the permutations of the coordinates $x_1, \ldots, x_n$, and all the reflections in the coordinate hyperplanes $H_i = \{(x_1, \ldots, x_n, x_{n+1}) \in \mathbb{R}^{n,1} \mid x_i = 0\}$ for $i = 1, \ldots, n$. Note that these generate a group of isometries isomorphic to B_n. We will refer to the lattice constructed in [14, §6] as the *hyperbolizing lattice*.

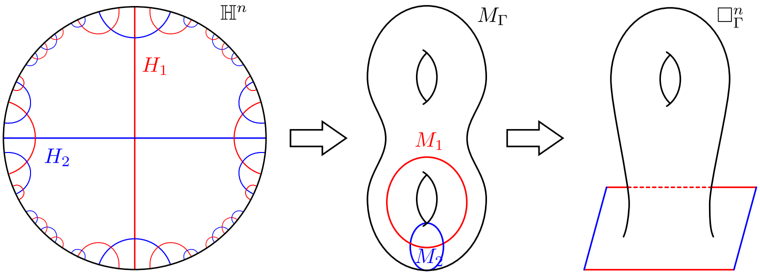

Fig. 6 A hyperbolizing cube

If Γ is such a lattice, then it acts freely, properly discontinuously, and cocompactly by orientation–preserving isometries on $\mathbb{H}^n$. We can consider the closed connected oriented hyperbolic n–manifold $M_\Gamma = \mathbb{H}^n / \Gamma$. The hyperplanes H_i descend to codimension–1 submanifolds $M_i = H_i / \mathrm{Stab}_\Gamma(H_i)$ which are closed, oriented, totally geodesic and pairwise orthogonal (see Fig. 6). Then the hyperbolizing cell $\square^n_\Gamma$ is defined to be the metric completion of the space $M_\Gamma \setminus \cup_{i=1}^n M_i$, with respect to the length metric induced on the complement of $\cup_{i=1}^n M_i$. This is the manifold with boundary and corners obtained by cutting M_Γ open along the submanifolds $M_1, \ldots, M_n$ (see [14, §5]). In particular, the map $g : \square^n_\Gamma \to \square^n$ is induced by the collapse map $g_0 : M_\Gamma \to (S^1)^n$ obtained by applying the Pontryagin-Thom construction to M_Γ with respect to each of the codimension–1 submanifolds $M_1, \ldots, M_n$.

Remark 3.3 It is implicit in [14] that a hyperbolizing lattice Γ contains infinitely many other hyperbolizing lattices as proper subgroups. They still enjoy the properties which are relevant for the construction, and provide corresponding hyperbolizing cubes. As observed by Ontaneda in [59, Lemma 2.1], this can be used to produce hyperbolizing cubes for which the normal injectivity radius of the faces is arbitrarily large.

3.2 The hyperbolized complex

Following the template of §2.3, to define the strict hyperbolization procedure of [14] we proceed as follows. For each dimension $n > 0$, we choose the hyperbolizing cell to be the hyperbolizing cube $(\square^n_\Gamma, g)$ defined in §3.1. Then for any foldable cubical complex X of dimension n, we define the *hyperbolized complex* to be the space X_Γ obtained as the fiber product of the folding map $f : X \to \square^n$ and the Charney-Davis map $g : \square^n_\Gamma \to \square^n$, i.e. by the pullback square in Fig. 7.

Remark 3.4 By (5) in Lemma 3.1 we know that g is surjective. So, Lemma 2.5 allows us to think of X_Γ as being obtained by replacing every n–cube of X by a hyperbolizing cube $\square^n_\Gamma$, in the following sense (see Fig. 8). If C is a top–dimensional cube of X, then its preimage $g_X^{-1}(C)$ in X_Γ is homeomorphic to $\square^n_\Gamma$ (see (1) in Lemma 2.5). Then one can endow X_Γ with a length metric by gluing together these local metrics. In particular, $f_\Gamma : X_\Gamma \to \square^n_\Gamma$ induces an isometry $g_X^{-1}(C) \to \square^n_\Gamma$ for each top–dimensional cube $C \subseteq X$. For a concrete example, if X is (a suitable cubical

Fig. 7 The hyperbolized complex X_Γ as a fibered product

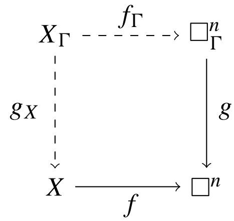

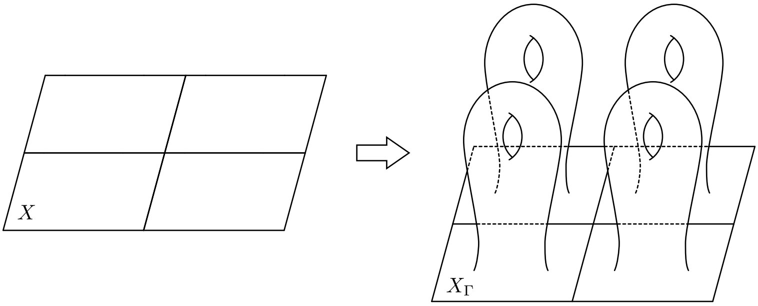

Fig. 8 Strict hyperbolization of a square complex

subdivision of) the standard cubical structure on the n–torus, then X_Γ is a closed hyperbolic manifold (see [9, Lemma 3.2] for details). Indeed, the piecewise hyperbolic metric obtained by gluing the hyperbolizing cubes together has no singularity and is in fact globally smooth and hyperbolic.

We collect here some of the main properties of this construction which are relevant for our work.

Proposition 3.5 (Proposition 7.1 in [14]) *For every $n \geq 2$ and every n–dimensional foldable cubical complex X, the space X_Γ carries the structure of an n–dimensional piecewise hyperbolic cell complex, and is endowed with a map $g_X : X_\Gamma \to X$, such that the following hold.*

(1) *If $C \subseteq X$ is a k–cube, then $g_X^{-1}(C) \subseteq X_\Gamma$ is isometric to a k–face of $\square^n_\Gamma$, and $\mathrm{lk}\left(g_X^{-1}(C), X_\Gamma\right)$ is a piecewise spherical cell complex, isomorphic to $\mathrm{lk}\,(C, X)$.*
(2) *If $Z \subseteq X$ is locally convex subcomplex of X, then $g_X^{-1}(Z)$ is a locally convex subspace of X_Γ.*
(3) *If X is locally* $\mathrm{CAT}(0)$, *then X_Γ is locally* $\mathrm{CAT}(-1)$.
(4) *If X is compact and locally* $\mathrm{CAT}(0)$, *then $\Gamma_X = \pi_1(X_\Gamma)$ is a Gromov hyperbolic group.*

Remark 3.6 The statement says in particular that if C is a top–dimensional cube of X then $g_X^{-1}(C)$ is isometric to $\square^n_\Gamma$ (compare Remark 3.4). On the other hand, if C is a k–cube with $k < n$, then $g_X^{-1}(C)$ is isometric to $\square^k_\Gamma = g^{-1}(\square^k)$, i.e. the hyperbolization of a lower dimensional cell, as introduced in Remark 3.2.

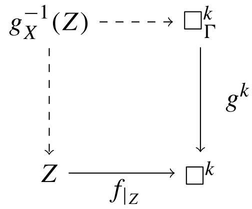

Fig. 9 Hyperbolization of lower dimensional subcomplexes

If $Z \subseteq X$ is a k–dimensional subcomplex, the subspace $g_X^{-1}(Z)$ can be identified with the fibered product of the maps $f_{|Z} : Z \to \square^k$ and $g^k : \square^k_\Gamma \to \square^k$, respectively obtained by restricting the folding map $f : X \to \square^n$ to Z and the Charney–Davis map $g : \square^n_\Gamma \to \square^n$ to $\square^k_\Gamma$ (see Fig. 9). Loosely speaking, hyperbolization trickles down to the lower dimensional skeletons of the complex X.

Remark 3.7 In this construction the choice of X and Γ are essentially independent. In particular for any fixed cubical complex X one can consider deeper hyperbolizations by taking deeper hyperbolizing lattices (see Remark 3.3). While the combinatorial geometry of the hyperbolized complex, controlled by X, remains unchanged under different choices of the hyperbolizing lattice, its hyperbolic geometry can be quantitatively improved by an appropriate choice of the hyperbolizing lattice, as observed by Ontaneda in [59, Lemma 2.1].

Remark 3.8 (Finding codimension-1 subspaces) The original approaches to cubulating a group G relied on producing sufficiently many codimension one subgroups inside G (see [10, 42, 65, 66]).

Since the copies of $\square^n_\Gamma$ in the hyperbolized complex X_Γ are obtained from an arithmetic hyperbolic manifold, they contain a large supply of compact totally geodesic codimension one submanifolds. It is tempting to try and use these to produce codimension one subgroups in the hyperbolized group $\Gamma_X = \pi_1(X_\Gamma)$. The difficulty with this approach is due to lack of control on the angles at which these totally geodesic codimension one hypersurfaces intersect the boundary of $\square^n_\Gamma$. This makes it unclear how to extend the proposed subspace past the boundary. One could take a geodesic extension, but it would not be clear what the global behaviour of the subspace would be (see left of Fig. 10). Or one could take a geodesic reflection, but that would give rise to a kink angle (see right of Fig. 10). Given that $\square^n_\Gamma$ has fixed finite diameter, kink angles too far from right angles might prevent the subspace from even being quasiconvex.

You can try to control the kink angle, for instance by requiring the codimension one submanifold to be orthogonal to all faces of $\square^n_\Gamma$. In this case, the extension would be a locally convex subspace of X_Γ. Examples of orthogonal subspaces can be obtained by noting that the symmetry group of the cube B_n acts on $\square^n_\Gamma$ (see Lemma 3.1). Each reflection of B_n has some fixed point set, which meets the boundary orthogonally and is totally geodesic.

However, one can only find finitely many such subspaces, both in the orthogonal case and in the case of kink angles bounded away from 0 (see [70] and [29, §5]). This would make it quite delicate to ensure that one can find enough such subspaces to

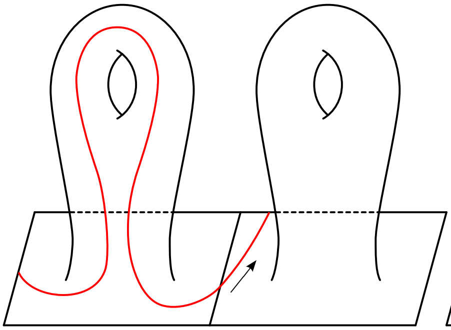

Fig. 10 Failure of the attempt to create hyperplane–like subspaces. Left: geodesic extension. Right: geodesic reflection

apply the standard criteria for properness of the induced cubulation (such as those in [10, 42]). To address these issues, we turn to a different type of subspaces, which we call mirrors. These are defined in §3.4 using the foldability of X, and enjoy properties reminiscent of those of hyperplanes in a CAT(0) cube complex. For the sake of clarity, the collection of mirrors also fails to provide a proper action of $\Gamma_X = \pi_1(X_\Gamma)$ on a CAT(0) cubical complex in the usual way. Nevertheless, in §4 we will be able to use mirrors to construct an action of Γ_X on a CAT(0) cubical complex for which the cube stabilizers are manageable, and are in a certain sense already detected by the action of Γ_X by deck transformations on the universal cover $\widetilde{X_\Gamma}$ (see §5.1). The reader interested in these remarks should also compare this discussion with that in Remark 4.2 below.

3.3 Tiling, folding, and developing the universal cover

Recall that we are assuming X is an admissible complex, as defined at the beginning of §3. It follows from Proposition 3.5 (see also Lemma 2.5) that the hyperbolized complex X_Γ admits a decomposition into hyperbolized cubes, analogous to the decomposition of X into cubes. In this section we show how to obtain an analogous decomposition of the universal cover $\widetilde{X_\Gamma}$ of X_Γ into pieces which are isometric to the universal cover $\widetilde{\square^n_\Gamma}$ of the hyperbolizing cube. Let us denote by $\pi : \widetilde{X_\Gamma} \to X_\Gamma$ and $\pi_\square : \widetilde{\square^n_\Gamma} \to \square^n_\Gamma$ the universal covering projections.

We start by realizing the space $\widetilde{\square^n_\Gamma}$ as a convex subset of $\mathbb{H}^n$. Let us consider once again the coordinate hyperplanes $H_i = \{(x_1, \dots, x_n, x_{n+1}) \in \mathbb{R}^{n,1} \mid x_i = 0\}$ for $i = 1, \dots, n$ (introduced in §3.1.1). An *open* Γ*–cell* is a connected component of the complement in $\mathbb{H}^n$ of the collection of Γ–orbits of the hyperplanes H_i. A Γ*–cell* is the closure of an open Γ–cell. Notice that all Γ–cells are convex, isometric to each other, and that Γ permutes them transitively. It follows from the construction of $\square^n_\Gamma$ in §3.1.1 that the universal cover $\widetilde{\square^n_\Gamma}$ of $\square^n_\Gamma$ can be isometrically identified with any Γ–cell (see Fig. 11).

While it might be tempting to think that $\widetilde{X_\Gamma}$ is obtained via some simple fibered product construction involving $\widetilde{X}$ and $\widetilde{\square^n_\Gamma}$, that is not the case.

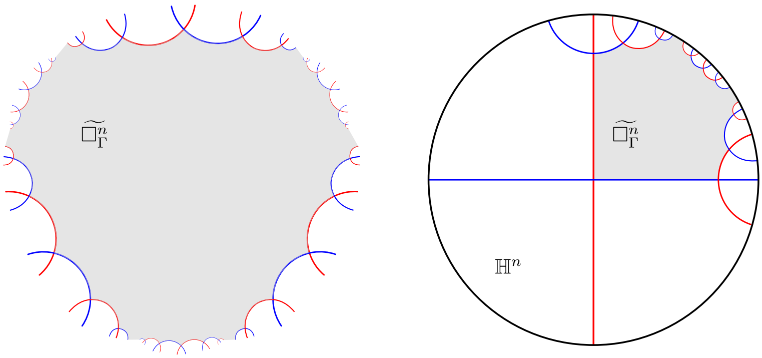

Fig. 11 The universal cover $\widetilde{\square}^n_\Gamma$ of $\square^n_\Gamma$, and its isometric embedding in $\mathbb{H}^n$ as a Γ–cell

Fig. 12 The hyperbolized complex X_Γ, its covering spaces, and the folding map

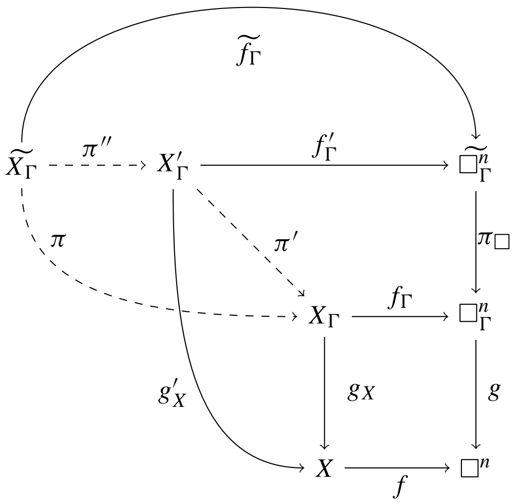

Remark 3.9 (What $\widetilde{X_\Gamma}$ is not) Note that $\widetilde{X_\Gamma} \neq (\widetilde{X})_\Gamma$, i.e. the universal cover of the hyperbolization of X is not the hyperbolization of the universal cover of X. Indeed, $(\widetilde{X})_\Gamma$ is not simply connected, because it retracts to $\square^n_\Gamma$ by (2) in Lemma 2.5. Analogously, $\widetilde{X_\Gamma}$ is not the fiber product of $\widetilde{X}$ and $\widetilde{\square}^n_\Gamma$ either. Indeed, note that the faces of $\widetilde{\square}^n_\Gamma$ (i.e. the preimages of faces of $\square^n$ via the map $g \circ \pi_\square$) are disconnected (see Fig. 11). This prevents the fiber product of $\widetilde{X}$ and $\widetilde{\square}^n_\Gamma$ from being simply connected, as observed in Remark 2.6.

In order to address this, and get a working understanding of $\widetilde{X_\Gamma}$, we consider the intermediate space X'_Γ obtained as a fibered product of X and $\widetilde{\square}^n_\Gamma$ along the maps $f : X \to \square^n$ and $g \circ \pi_\square : \widetilde{\square}^n_\Gamma \to \square^n_\Gamma \to \square^n$ (see Fig. 12). By the universal property of pullbacks we have an induced map $\pi' : X'_\Gamma \to X_\Gamma$.

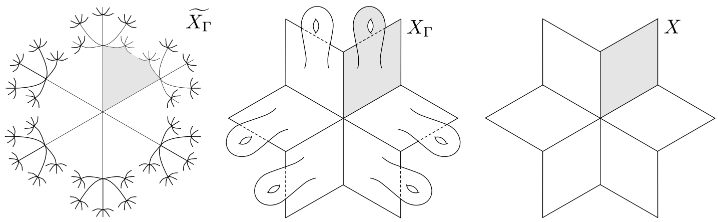

Fig. 13 Tiles in $\widetilde{X_\Gamma}$, X_Γ and X

Lemma 3.10 *The map $\pi' : X'_\Gamma \to X_\Gamma$ is a covering map.*

Proof By the composition law for pullbacks the space X'_Γ is actually the same as the pullback of $f_\Gamma : X_\Gamma \to \square^n_\Gamma$ and $\pi_\square : \widetilde{\square^n_\Gamma} \to \square^n_\Gamma$. In particular, the map π' is the pullback of the universal covering projection $\pi_\square$ along the map f_Γ, hence is itself a covering map. $\square$

In particular, X'_Γ can be endowed with a length metric that makes π' a local isometry (see [11, Proposition I.3.25]), and the universal cover $\widetilde{X_\Gamma}$ can be realized as the universal cover of this space X'_Γ, even in a metric sense. Let $\pi'' : \widetilde{X_\Gamma} \to X'_\Gamma$ denote the universal covering projection.

We define a *tile* of X_Γ to be a subspace of the form $g_X^{-1}(C)$, for C a top–dimensional cube of X. Recall from Remark 3.4 that each tile of X_Γ is isometric to $\square^n_\Gamma$. In complete analogy, we define a *tile* in X'_Γ and in $\widetilde{X_\Gamma}$ to be a connected component of the lift of a tile from X_Γ via the covering maps π' and $\pi = \pi' \circ \pi''$ respectively. We refer to this decomposition into tiles as the *tiling* of each of these spaces (see Fig. 13). Note that, since the complex X is assumed to be admissible, each point of X is either contained in the interior of a tile, or in the intersection of at least two tiles. Moreover the folding map f of X induces an analogous map on X_Γ and its covering spaces, as established in the next lemma.

Lemma 3.11 *The map $\widetilde{f_\Gamma} = f'_\Gamma \circ \pi'' : \widetilde{X_\Gamma} \to X'_\Gamma \to \widetilde{\square^n_\Gamma}$ restricts to an isometry between each tile of $\widetilde{X_\Gamma}$ and $\widetilde{\square^n_\Gamma}$.*

Proof Recall that X'_Γ is defined via a pullback construction, in the sense of §2.3. Therefore, by (1) in Lemma 2.5, the map $f'_\Gamma : X'_\Gamma \to \widetilde{\square^n_\Gamma}$ restricts to a homeomorphism between each tile of X'_Γ and $\widetilde{\square^n_\Gamma}$. Since the metric on X'_Γ is lifted from X_Γ via π', and f_Γ restricts to an isometry between each tile of X_Γ and $\square^n_\Gamma$ (see Remark 3.4), the map f'_Γ actually gives an isometry between a tile of X'_Γ and $\widetilde{\square^n_\Gamma}$. Since the tiles of X'_Γ are simply connected, they lift isometrically to tiles of $\widetilde{X_\Gamma}$ via π''. In particular, π'' maps a tile of $\widetilde{X_\Gamma}$ isometrically onto a tile of X'_Γ. Therefore the map $\widetilde{f_\Gamma} = f'_\Gamma \circ \pi''$ maps a tile of $\widetilde{X_\Gamma}$ isometrically onto $\widetilde{\square^n_\Gamma}$, just by composition. $\square$

The map $\widetilde{f_\Gamma}$ from Lemma 3.11 will be called the *folding map* of $\widetilde{X_\Gamma}$. The composition of the folding map $\widetilde{f_\Gamma}$ with any isometric embedding $\varphi : \widetilde{\square^n_\Gamma} \to C$ onto a Γ–cell $C \subseteq \mathbb{H}^n$ will be called a *developing map* for $\widetilde{X_\Gamma}$.

Remark 3.12 The restriction of a developing map to a tile is an isometric embedding of a tile into $\mathbb{H}^n$ as a Γ–cell. Moreover if T_1, T_2 are two tiles of $\widetilde{X_\Gamma}$ meeting along a codimension–1 subspace Z, and $\varphi_1 : T_1 \to C_1 \subseteq \mathbb{H}^n$ is an isometric embedding onto a Γ–cell that maps Z into some hyperplane H, then post–composing φ_1 with the reflection across H provides an isometric embedding φ_2 of T_2 as a Γ–cell C_2 adjacent to C_1 along H. The two embeddings can be glued together to give an isometric embedding of $T_1 \cup T_2$ onto the union of two Γ–cells $C_1 \cup C_2$ adjacent along H. This can be "analytically continued" by sequentially extending over adjacent tiles, to obtain a globally defined map $\widetilde{X_\Gamma} \to \mathbb{H}^n$. However, this does not result in a global isometric embedding $\widetilde{X_\Gamma} \to \mathbb{H}^n$ in general. This is due to the fact that links in X can be very large, which gives rise to overlaps and singularities.

3.4 Mirrors: convexity

In this section we exploit foldability to define a collection of convex subcomplexes of X, and induce corresponding subspaces in X_Γ and $\widetilde{X_\Gamma}$. Let Y be a foldable cubical complex of dimension n (in the following we will consider $Y = X$ and $Y = \widetilde{X}$ depending on the situation). If $f : Y \to \square^n$ is a fixed folding and $F \subseteq \square^n$ is a codimension-1 face, then we define a *mirror* in Y to be a connected component of $f^{-1}(F)$.

Proposition 3.13 *Let Y be an admissible cubical complex. Then each mirror is a locally convex and geodesically complete subcomplex of Y. In particular, if Y is* CAT(0), *then each mirror is convex.*

Proof For the first statement see [76, Proposition 2.3] (and references therein such as [4, Lemma 3.2(4)]). In the CAT(0) case, local convexity implies global convexity (see for instance [13, Theorem 1.6,1.10], or [62, Theorem 1.1]). $\square$

We now define a *mirror* in $\widetilde{X_\Gamma}$ to be a connected components of $\widetilde{f}^{-1}(F)$, where F is a codimension-1 face of $\square^n$ and $\widetilde{f}$ is the map given by the composition $\widetilde{f} = f \circ g_X \circ \pi : \widetilde{X_\Gamma} \to X_\Gamma \to X \to \square^n$ (see Fig. 14 and Fig. 15). Equivalently, we could define it as a connected component of the full preimage of a mirror of X via the map $\widetilde{g_X} = g_X \circ \pi$, but we find it convenient to use this definition. We will say that *M folds* to F, and we will denote by $\mathcal{M}$ the collection of all mirrors in $\widetilde{X_\Gamma}$. Mirrors in X_Γ are defined in the analogous way using the map $f \circ g_X$.

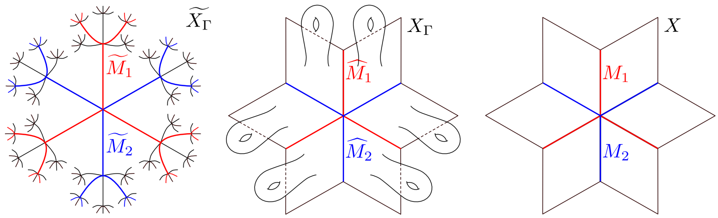

Fig. 14 Right to left: mirrors M_1, M_2 in X, their preimages $\widehat{M}_1 = g_X^{-1}(M_1)$, $\widehat{M}_2 = g_X^{-1}(M_2)$ in X_Γ, and the lifts $\widetilde{M}_1$, $\widetilde{M}_2$ to $\widetilde{X_\Gamma}$

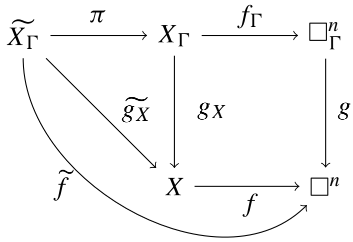

Fig. 15 The hyperbolized complex X_Γ and the maps used to define mirrors

Proposition 3.14 *Let X be an admissible cubical complex. Then each mirror of $\widetilde{X_\Gamma}$ is a closed connected convex subspace of $\widetilde{X_\Gamma}$.*

Proof Let M be a mirror of $\widetilde{X_\Gamma}$, and let $F \subseteq \square^n$ be the codimension-1 face to which it folds. By definition M is connected and closed. To prove convexity we argue as follows. Let $Z = g_X(\pi(M)) \subseteq X$, and notice that Z is a mirror of X that folds to F. By Proposition 3.13 we know that Z is locally convex in X. By (2) in Proposition 3.5, we also know that $g_X^{-1}(Z)$ is locally convex in X_Γ, and therefore $M \subseteq \widetilde{X_\Gamma}$ is locally convex too. By (3) in Proposition 3.5 we also know that X_Γ is locally CAT(−1). In particular M is a closed and locally convex subspace in the CAT(0) space $\widetilde{X_\Gamma}$. Therefore it is convex (again by [13, Theorem 1.6,1.10], or [62, Theorem 1.1]). □

3.5 Stratification of $\widetilde{X_\Gamma}$

In this section we use the collection $\mathcal{M}$ of mirrors, introduced in §3.4, to define a stratification of $\widetilde{X_\Gamma}$. The *open k–stratum* Σ^k of $\widetilde{X_\Gamma}$ is the subspace consisting of points that fold into the interior of a k-face of $\square^n$ via the map $\widetilde{f} = f \circ g_X \circ \pi : \widetilde{X_\Gamma} \to X_\Gamma \to X \to \square^n$, or equivalently to the interior of a k–cube of X via the map $\widetilde{g_X} = \pi \circ g_X : \widetilde{X_\Gamma} \to X_\Gamma \to X$ (see Fig. 15). Notice that Σ^k is a locally closed subspace. An *open k–cell* is a connected component of Σ^k. A *k–cell* is the closure of an open k–cell. We say that a cell σ *folds* to the face $F = \widetilde{f}(\sigma) \subseteq \square^n$ and to the cube $C = \widetilde{g_X}(\sigma) \subseteq X$. The integer k will be referred to as the *dimension* of a k–cell. An $(n-k)$–cell is a proper subset of the intersection of k mirrors. In particular 0–cells are points, and n-cells are tiles (as defined in §3.3). We call 0–cells *vertices*, and 1–cells *edges* of the stratification.

Remark 3.15 (Cellular structure) We explicitly observe that this choice of strata does not define a stratified space structure on $\widetilde{X_\Gamma}$ in the sense of [11, Definition II.12.1]. Moreover, the decomposition of $\widetilde{X_\Gamma}$ into cells does not turn it into a genuine cell complex, as defined in §2.1. Indeed, while an open k-cell is homeomorphic to an open disk of dimension k, a k-cell is not homeomorphic to a closed disk of dimension k as soon as $k \geq 2$. Its boundary in $\widetilde{X_\Gamma}$ consists of an infinite union of lower–dimensional cells, so it is neither connected nor compact. For instance, an n–cell (i.e. a tile) is isometric to a Γ–cell (see Fig. 11).

Nevertheless, we can still recover a lot of the classical behavior and tools, by observing that cells are convex and that the link of cells and points can be defined in analogy to the classical case (see §2.1.3). We gather here preliminary results about this, that will be useful in the following. For the sake of clarity, we emphasize that in our terminology cells are closed.

Lemma 3.16 *Let $\sigma \subseteq \widetilde{X_\Gamma}$ be a cell. Then σ is convex.*

Proof Let $C = \widetilde{g_X}(\sigma) \subseteq X$ be the cube of X to which σ folds. By (2) in Proposition 3.5, we know $g_X^{-1}(C)$ is locally convex in X_Γ. Since σ is by definition a connected component of $\pi^{-1}(g_X^{-1}(C))$ and π is a local isometry, we can conclude that it is a locally convex subspace of $\widetilde{X_\Gamma}$. Arguing similarly to previous proofs of convexity, we can conclude that σ is convex, because it is closed and locally convex in the CAT(0) space $\widetilde{X_\Gamma}$ (see (3) in Proposition 3.5 and [13, Theorem 1.6,1.10], or [62, Theorem 1.1]). $\square$

We now proceed to the study of links. Consider the universal covering map $\pi : \widetilde{X_\Gamma} \to X_\Gamma$. By Proposition 3.5, X_Γ is a piecewise hyperbolic cell complex, so the link of points and cells in X_Γ is well–defined (see §2.1.3). Since π is a local isometry, we can just identify the *link* of points and cells in $\widetilde{X_\Gamma}$ with the links of the corresponding points and cells in X_Γ.

Lemma 3.17 *Let $\sigma \subseteq \widetilde{X_\Gamma}$ be a cell.*

(1) *Let $C = \widetilde{g_X}(\sigma) \subseteq X$ be the cube to which it folds. Then $\widetilde{g_X}$ induces an isomorphism between* $\mathrm{lk}\left(\sigma, \widetilde{X_\Gamma}\right)$ *and* $\mathrm{lk}\,(C, X)$.
(2) *Let σ be contained in another cell τ. Let $F = \widetilde{f}(\sigma)$, $E = \widetilde{f}(\tau) \subseteq \square^n$ be the faces to which they fold. Then $\widetilde{f}$ induces an isomorphism between* $\mathrm{lk}\,(\sigma, \tau)$ *and* $\mathrm{lk}\,(F, E)$.
(3) *Let σ be a k–cell. Then* $\mathrm{lk}\left(\sigma, \widetilde{X_\Gamma}\right)$ *is a piecewise spherical simplicial complex with vertices given by the $(k+1)$–cells containing σ, and in which $m+1$ vertices span an m–simplex if and only if the corresponding $(k+1)$–cells are contained in a $(k+m+1)$–cell.*

Proof The map $\widetilde{g_X} : \widetilde{X_\Gamma} \to X$ is the composition of the map $\pi : \widetilde{X_\Gamma} \to X_\Gamma$, which preserves links because it is a covering map, and the map $g_X : X_\Gamma \to X$, which preserves links thanks to (1) in Proposition 3.5. This proves (1).

To prove (2) we argue similarly. The map $\widetilde{f} : \widetilde{X_\Gamma} \to \square^n$ is the composition of the map $\widetilde{g_X} : \widetilde{X_\Gamma} \to X$, which preserves links by (1), and the folding map $f : X \to \square^n$. By definition of folding, f is a combinatorial isomorphism on each cube of X. If B is the cube to which τ folds, the folding induces an isomorphism between $\mathrm{lk}\,(C, B)$ and $\mathrm{lk}\,(F, E)$.

Finally, (3) follows from (1), the fact that $\widetilde{g_X} : \widetilde{X_\Gamma} \to X$ maps cells of $\widetilde{X_\Gamma}$ to cubes of X preserving inclusion relations, and the fact that the link of a cell in a cubical complex carries a piecewise spherical simplicial structure as described in the statement. $\square$

Lemma 3.18 *Let $\sigma_1, \sigma_2 \subseteq \widetilde{X_\Gamma}$ be two cells. Then either $\sigma_1 \cap \sigma_2$ is empty or it is a cell.*

Proof Let $\sigma_1 \cap \sigma_2$ be non empty. If it contains either a single vertex, or a single edge, then we are done. So let us assume that it contains at least two edges. Also note that, since cells are convex by Lemma 3.16, the intersection $\sigma_1 \cap \sigma_2$ is convex. Therefore, if there are several edges then they cannot all be disjoint.

Let $v \in \sigma_1 \cap \sigma_2$ be a vertex, and let e, e' be two edges of $\sigma_1 \cap \sigma_2$ meeting at v. Note that by Lemma 3.17 links in $\widetilde{X_\Gamma}$ are isomorphic to the corresponding links in X. In particular, e and e' are edges of the cell σ_1 meeting at a vertex, so there is a 2–face $\tau_1 \subseteq \sigma_1$ containing both e and e'. Analogously, we get a 2–face $\tau_2 \subseteq \sigma_2$ with the same property. Since these links are simplicial, necessarily we have $\tau_1 = \tau_2$, otherwise we would see a bigon in the link of v. In particular, $\tau_1 = \tau_2 \subseteq \sigma_1 \cap \sigma_2$. This shows that any two edges of $\sigma_1 \cap \sigma_2$ meeting at v are adjacent in $\mathrm{lk}\,(v, \sigma_1 \cap \sigma_2)$. By Lemma 3.17 we know that $\mathrm{lk}\left(v, \widetilde{X_\Gamma}\right) \cong \mathrm{lk}\,(\widetilde{g_X}(v), X)$, and this is a flag simplicial complex because X is non–positively curved (see Lemma 2.2). The same holds for $\mathrm{lk}\,(v, \sigma_1 \cap \sigma_2)$ because $\sigma_1 \cap \sigma_2$ is convex in $\widetilde{X_\Gamma}$. In particular, all the edges of $\sigma_1 \cap \sigma_2$ that contain v are actually contained in a unique cell of minimal dimension in $\sigma_1 \cap \sigma_2$; we denote this cell by σ_v. Now, if v, w are adjacent vertices of $\sigma_1 \cap \sigma_2$, then by uniqueness we have $\sigma_v = \sigma_w$. Finally, by connectedness of $\sigma_1 \cap \sigma_2$, it follows that all vertices of $\sigma_1 \cap \sigma_2$ are contained in a single cell. □

Lemma 3.19 *Let $\{\sigma_j \mid j \in J\}$ be a collection of cells of $\widetilde{X_\Gamma}$. Then the following statements hold.*

(1) *If $\sigma = \bigcap_{j \in J} \sigma_j$ is not empty, then σ is the unique cell of maximal dimension contained in σ_j for all $j \in J$.*
(2) *If $\bigcup_{j \in J} \sigma_j$ is contained in a single cell, then there exists a unique cell σ of minimal dimension containing σ_j for all $j \in J$.*

We refer to the cell in (1) (respectively (2)) of Lemma 3.19 as the *lower cell* (respectively *upper cell*) of the collection $\{\sigma_j \mid j \in J\}$.

Proof Since X is finite–dimensional and locally compact, if J is infinite, then $\sigma = \bigcap_{j \in J} \sigma_j$ is empty. So let us assume that J is finite. By Lemma 3.18 the intersection of finitely many cells is either empty or made of a single cell. This proves (1). To prove (2), assume by contradiction that there are two different cells of minimal dimension σ, σ' containing each σ_j. Then $\sigma \cap \sigma'$ is a proper union of cells, against Lemma 3.18. □

Lemma 3.20 *Let $\tau \subseteq \widetilde{X_\Gamma}$ be a cell. Let $\sigma_1, \sigma_2 \subseteq \tau$ be cells of lower dimension, and let $F_1, F_2 \subseteq \square^n$ be the faces to which they fold. If $F_1 = F_2$, then σ_1, σ_2 are either disjoint or equal.*

Proof Assume that σ_1, σ_2 are not disjoint, and let $v \in \sigma_1 \cap \sigma_2$ be a vertex. Let $E \subseteq \square^n$ be the face to which τ folds. We have that $\widetilde{f}(\sigma_1) = F_1 = F_2 = \widetilde{f}(\sigma_2)$, and by (2) in

Lemma 3.17 (with $\sigma = v$) the map $\widetilde{f}$ induces an isomorphism between $\mathrm{lk}\,(v, \tau)$ and $\mathrm{lk}\left(\widetilde{f}(v), E\right)$. Therefore, $\sigma_1 = \sigma_2$. □

Lemma 3.21 *Let M be a mirror and let τ be a k–cell of $\widetilde{X_\Gamma}$ not entirely contained in M. If $M \cap \tau \neq \varnothing$, then $M \cap \tau$ is a $(k-1)$–cell.*

Proof First we show that $M \cap \tau$ is a union of $(k-1)$–cells. Then we show that the union actually consists of a single cell.

Let $\widetilde{f} = f \circ g_X \circ \pi : \widetilde{X_\Gamma} \to X_\Gamma \to X \to \square^n$ be the map that folds $\widetilde{X_\Gamma}$ to $\square^n$, and let $F \subseteq \square^n$ be the codimension–1 face to which M folds (i.e. $F = \widetilde{f}(M)$). Similarly, let $E \subseteq \square^n$ be the k–face to which τ folds (i.e. $E = \widetilde{f}(\tau)$). Since $\tau \nsubseteq M$ we have $E \nsubseteq F$, and therefore $E \cap F$ is a $(k-1)$–face of $\square^n$. Let $p \in M \cap \tau$ be an arbitrary point. By Lemma 3.20, among the $(k-1)$–cells of τ that contain p, there is exactly one that folds to $E \cap F$; denote it by σ_p. Clearly $\sigma_p \subseteq \tau$. Moreover, since σ_p and M are non–disjoint, both fold into F, and M is a mirror, we also have $\sigma_p \subseteq M$. Therefore we have $M \cap \tau = \bigcup_{p \in M \cap \tau} \sigma_p$, i.e. $M \cap \tau$ is a union of $(k-1)$–cells that fold to $E \cap F$.

To see that $M \cap \tau$ actually consist of only one cell, assume by contradiction that $M \cap \tau$ contains two distinct $(k-1)$–cells σ_1, σ_2. Let $p_i \in \sigma_i$ and let $\gamma = [p_1, p_2]$ be the unique geodesic between them. Since M and τ are both convex (by Proposition 3.14 and Lemma 3.16 respectively), we have that $\gamma \subset \tau \cap M$. Hence we find a path of cells in the boundary of τ that all fold to $E \cap F$. But this is absurd because different boundary cells of τ folding to the same face of $\square^n$ are necessarily disjoint, again by Lemma 3.20. □

3.6 Graph of spaces decomposition for $\widetilde{X_\Gamma}$

Our goal in §3.7 will be to prove that mirrors in $\widetilde{X_\Gamma}$ enjoy a strong separation property. Our strategy will be to exploit a certain graph of spaces decomposition for $\widetilde{X_\Gamma}$ (in the sense of [69]), which we introduce in this section, using the foldability of X (see [4, 76] for analogous constructions).

Recall from §3.4 that $\widetilde{X_\Gamma}$ is equipped with a collection $\mathcal{M}$ of closed convex subspaces called mirrors. For each $i = 1, \dots, n$, let $\mathcal{M}_i$ be the collection of mirrors of $\widetilde{X_\Gamma}$ that fold to one of the two parallel ith faces of $\square^n = [0, 1]^n$, i.e. $\{x_i = 0\}$ and $\{x_i = 1\}$. Notice that by construction any two elements of $\mathcal{M}_i$ are disjoint, and even have disjoint ε–neighborhoods for ε sufficiently small (because Γ is cocompact).

Let $\mathcal{C}_i$ be the collection of connected components of $\widetilde{X_\Gamma} \setminus \cup_{M \in \mathcal{M}_i} M$. For each mirror $M \in \mathcal{M}_i$ and for each component $C \in \mathcal{C}_i$, consider the following *equidistant space*, obtained by pushing the mirror M into the component C (see Fig. 16).

$$E^{\varepsilon}_{M,C} = \{x \in C \mid d(x, M) = \varepsilon\}.$$

Notice that while we know M is convex by Proposition 3.14, it is not clear whether C is convex. A priori, C could meet M on more than one side, i.e. the closure of C could contain a piece of M in its interior. We will see this is not the case by considering a suitable graph of spaces decomposition of $\widetilde{X_\Gamma}$. Our first step is to show that $E^{\varepsilon}_{M,C}$ is simply connected; in the process, we actually show it is a CAT(k)–space

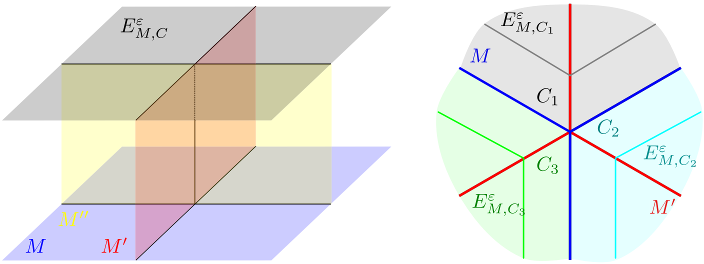

Fig. 16 Some examples of equidistant spaces in dimension 2 and 3. Left: an equidistant space relative to a mirror M in the vicinity of the intersection with two other mirrors. Here $\dim X = 3$, and all mirrors are locally Euclidean. Right: three equidistant spaces relative to the same mirror M but three different complementary components, in the vicinity of the intersection with another mirror. Here $\dim X = 2$, and the mirrors branch, i.e. are not locally Euclidean

Fig. 17 Equidistant surface from a hyperplane

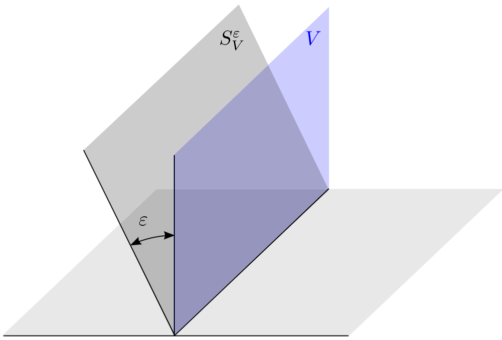

for some $k \in (-1, 0)$. The idea for this can be summarized as follows: inside each tile, $E^{\varepsilon}_{M,C}$ looks like an equidistant hypersurface from a hyperplane in $\mathbb{H}^n$, and this is a non–positively curved hypersurface in $\mathbb{H}^n$ (see Fig. 17). Then contribution from different tiles come together in a way that does not introduce any positive curvature along mirrors. We start from a preliminary lemma from classical hyperbolic geometry. Recall that a hyperplane in $\mathbb{H}^n$ is a totally geodesic copy of $\mathbb{H}^{n-1}$.

Lemma 3.22 *Let $V \subseteq \mathbb{H}^n$ be a hyperplane, and let $\pi_V : \mathbb{H}^n \to V$ be the nearest point projection to V. Let $\varepsilon > 0$ and $S^{\varepsilon}_V = \{x \in \mathbb{H}^n \mid d(x, V) = \varepsilon\}$. Then the following hold.*

(1) *S^{ε}_V is a smooth $(n-1)$–dimensional submanifold of $\mathbb{H}^n$.*
(2) *For each $p \in S^{\varepsilon}_V$, the geodesic $[p, \pi_V(p)]$ is orthogonal to V and S^{ε}_V.*
(3) *For every other hyperplane W, if $V \cap W \neq \varnothing$, then $S^{\varepsilon}_V \cap W \neq \varnothing$.*
(4) *For every hyperplane W, W is orthogonal to S^{ε}_V if and only if W is orthogonal to V.*
(5) *$\pi_V : S^{\varepsilon}_V \to V$ is a $\cosh^2(\varepsilon)$–conformal diffeomorphism.*
(6) *The induced metric on S^{ε}_V has constant sectional curvature $\frac{-1}{\cosh^2(\varepsilon)}$.*

Proof The first five statements can be proved by explicit computations in the upper half–space model of $\mathbb{H}^n$, normalizing so that V is a vertical hyperplane (see Fig. 17).

The computation for dimension $n = 3$ is carried out in detail in [27, IV.5, page 58], and readily generalizes to higher dimensions. Finally, (6) follows from (5) and the general formula for the behavior of the sectional curvatures under rescaling. □

For the next lemma, recall from §3.3 that tiles are closed by definition, and that a developing map is an isometric embedding of a tile into $\mathbb{H}^n$ as a Γ–cell.

Lemma 3.23 *Let $M \in \mathcal{M}_i$ and $C \in \mathcal{C}_i$. Then for $\varepsilon > 0$ small enough the following hold.*

(1) *For every mirror $N \in \mathcal{M}$, if $E^{\varepsilon}_{M,C} \cap N \neq \varnothing$ then $M \cap N \neq \varnothing$ and $C \cap N \neq \varnothing$.*
(2) *For every tile τ, if $E^{\varepsilon}_{M,C} \cap \tau \neq \varnothing$ then $M \cap \tau \neq \varnothing$ and $C \cap \tau \neq \varnothing$.*
(3) *For every tile τ such that $E^{\varepsilon}_{M,C} \cap \tau \neq \varnothing$, and any developing map $\varphi : \tau \to \mathbb{H}^n$, φ induces an isometry between $E^{\varepsilon}_{M,C} \cap \tau$ and $S^{\varepsilon}_V \cap \varphi(\tau)$, where V is the hyperplane containing $\varphi(M \cap \tau)$.*
(4) *For every mirror $N \in \mathcal{M}$, if $E^{\varepsilon}_{M,C} \cap N \neq \varnothing$ then $E^{\varepsilon}_{M,C}$ is orthogonal to N.*

Proof To prove (1) note that if $E^{\varepsilon}_{M,C} \cap N \neq \varnothing$, then in particular $C \cap N \neq \varnothing$. Since Γ is cocompact, there is a uniform lower bound $D > 0$ on the distance between disjoint mirrors. But $E^{\varepsilon}_{M,C} \cap N \neq \varnothing$ means that N comes ε close to M. By choosing $\varepsilon < D$ we can force N to actually intersect M.

The proof of (2) is analogous to that of (1). Suppose $E^{\varepsilon}_{M,C} \cap \tau \neq \varnothing$. Then clearly $C \cap \tau \neq \varnothing$. Moreover, a point in $E^{\varepsilon}_{M,C} \cap \tau$ witnesses that $d(M, \tau) < \varepsilon$, and by choosing ε small enough we can ensure that this forces an intersection, again by cocompactness of Γ.

Now we consider (3). Suppose that $E^{\varepsilon}_{M,C} \cap \tau \neq \varnothing$. Then by (2) we know that $M \cap \tau \neq \varnothing$ and $C \cap \tau \neq \varnothing$. In particular M appears as an $(n-1)$–cell in the boundary of τ thanks to Lemma 3.21. If we pick a developing map φ for τ, then $\varphi(\tau)$ is a Γ–cell, and $\varphi(M)$ is some hyperplane V on its boundary (see Lemma 3.11 and Remark 3.12). Then the statement follows from the fact that φ is an isometric embedding of τ into $\mathbb{H}^n$.

Finally, to prove (4), suppose that $E^{\varepsilon}_{M,C} \cap N \neq \varnothing$. Then by (1) we know that $N \cap M \neq \varnothing$. In particular by construction N is orthogonal to M. Then the statement follows from (3), together with (4) in Lemma 3.22. □

Next, our goal is to prove that equidistant spaces are negatively curved. In order to do this, we will study the geometry of links of points in $\widetilde{X_\Gamma}$, along various subspaces (we refer the reader to §3.5 for definitions). Recall that the link of a point in $\widetilde{X_\Gamma}$ is identified to the link of its projection to X_Γ.

Remark 3.24 All the subspaces of $\widetilde{X_\Gamma}$ considered here (such as a mirror M, and the induced space $E^{\varepsilon}_{M,C}$) carry a natural locally finite cellular structure induced by that of $\widetilde{X_\Gamma}$. Even if they are not genuine cell complexes (as in Remark 3.15), their projections to X_Γ are, and links can be defined in analogy to the classical case.

For a mirror M we denote by $\pi_M : \widetilde{X_\Gamma} \to M$ the nearest point projection. This is well–defined because $\widetilde{X_\Gamma}$ is CAT(0) and M is convex by Proposition 3.14.

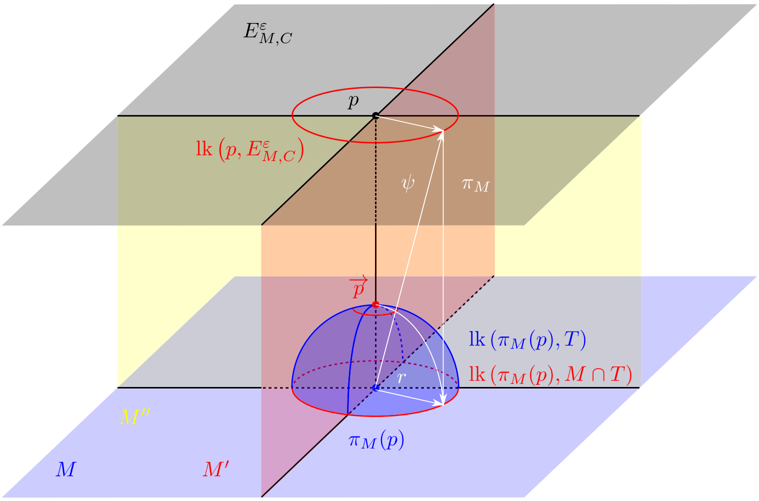

Fig. 18 Links of points along various subspaces in the proof of Lemma 3.25. Here p is contained in four tiles and sits on the intersection of two mirrors M', M''. The vertical projection is the nearest point projection $\pi_M : E^{\varepsilon}_{M,C} \to M$

Lemma 3.25 *Let $M \in \mathcal{M}_i$, $C \in \mathcal{C}_i$, $p \in E^{\varepsilon}_{M,C}$. Then for $\varepsilon > 0$ small enough the following holds. Let $\tau_1, \ldots, \tau_m$ be the collection of tiles containing p, and let $T = \tau_1 \cup \cdots \cup \tau_m$. Then the following hold.*

(1) $\mathrm{lk}\,(\pi_M(p), M \cap T)$ *is* CAT(1).
(2) $\pi_M : E^{\varepsilon}_{M,C} \to M$ *induces an isometry* $\lambda_p : \mathrm{lk}\left(p, E^{\varepsilon}_{M,C}\right) \to \mathrm{lk}\,(\pi_M(p), M \cap T)$.
(3) $\mathrm{lk}\left(p, E^{\varepsilon}_{M,C}\right)$ *is* CAT(1).

Proof Of course, (3) follows from (1) and (2). For convenience, let us denote $L = \mathrm{lk}\left(\pi_M(p), \widetilde{X_\Gamma}\right)$, $L_T = \mathrm{lk}\,(\pi_M(p), T)$, and $L_{M\cap T} = \mathrm{lk}\,(\pi_M(p), M \cap T)$. We have $L_{M\cap T} \subseteq L_T \subseteq L$. Equip $L_{M\cap T}$ and L_T with the induced length metric. Let $\overrightarrow{p} \in L_T$ be the direction at $\pi_M(p)$ pointing to p (see Fig. 18).

We start by proving (1). Since $\widetilde{X_\Gamma}$ is negatively curved, L is CAT(1). In particular, balls of radius at most $\pi/2$ are π-convex and CAT(1). Since $\widetilde{X_\Gamma}$ is piecewise hyperbolic, L is piecewise spherical. Moreover, all the mirrors containing p intersect M orthogonally by construction. Therefore, L has a natural structure of all-right spherical complex in which $\overrightarrow{p}$ is a vertex (possibly up to subdivision if $\pi_M(p)$ is not a vertex). In particular, we have natural identifications $L_T = \overline{B\left(\overrightarrow{p}, \frac{\pi}{2}\right)}$ and $L_{M\cap T} = \partial B\left(\overrightarrow{p}, \frac{\pi}{2}\right)$.

Let $C_1(Y)$ denote the spherical cone over a space Y, and denote the cone point by 0. Since L is an all-right spherical complex, we have a natural isometry

$$\varphi : C_1\left(\partial B\left(\overrightarrow{p}, \frac{\pi}{2}\right)\right) \to \overline{B\left(\overrightarrow{p}, \frac{\pi}{2}\right)}$$

defined as follows: $\varphi(0) = \overrightarrow{p}$, and for each $\overrightarrow{q} \in \partial B\left(\overrightarrow{p}, \frac{\pi}{2}\right)$ and $0 < t \leq \frac{\pi}{2}$ let $\varphi(t, \overrightarrow{q})$ be the point at distance t from $\overrightarrow{p}$ along the geodesic $[\overrightarrow{p}, \overrightarrow{q}]$. As a result, $C_1(L_{M\cap T}) = C_1\left(\partial B\left(\overrightarrow{p}, \frac{\pi}{2}\right)\right)$ is CAT(1). By Berestovskii's Theorem (see [11, II.3.14]) we conclude that $L_{M\cap T}$ is CAT(1) as desired.

To prove (2) we argue as follows. By (3) in Lemma 3.23 and (5) in Lemma 3.22 we know that within each tile τ_k the projection π_M is a conformal diffeomorphism, so it induces an isometry $\lambda_p^{\tau_k} : \mathrm{lk}\left(p, E^{\varepsilon}_{M,C} \cap \tau_k\right) \to L_k = \mathrm{lk}\left(\pi_M(p), M \cap \tau_k\right)$. This is enough in the case $m = 1$, i.e. when p is contained in a single tile. When $m \geq 2$, by gluing together the maps $\lambda_p^{\tau_k}$, we obtain a map $\lambda_p : \mathrm{lk}\left(p, E^{\varepsilon}_{M,C}\right) \to L_1 \cup \cdots \cup L_m = L_{M\cap T}$. Notice that shooting geodesic rays from $\pi_M(p)$ into T along directions in L_T provides an isometry

$$\psi : \mathrm{lk}\left(\overrightarrow{p}, L_T\right) \to \mathrm{lk}\left(p, E^{\varepsilon}_{M,C}\right).$$

Combining this with the natural isometry

$$r : \mathrm{lk}\left(\overrightarrow{p}, B\left(\overrightarrow{p}, \frac{\pi}{2}\right)\right) \to \partial B\left(\overrightarrow{p}, \frac{\pi}{2}\right)$$

and using the aforementioned identifications, we obtain the desired isometry

$$\mathrm{lk}\left(p, E^{\varepsilon}_{M,C}\right) \xrightarrow{\psi^{-1}} \mathrm{lk}\left(\overrightarrow{p}, L_T\right) = \mathrm{lk}\left(\overrightarrow{p}, \overline{B\left(\overrightarrow{p}, \frac{\pi}{2}\right)}\right) \xrightarrow{r} \partial B\left(\overrightarrow{p}, \frac{\pi}{2}\right) = L_{M\cap T}. \qquad \square$$

Remark 3.26 Note that, in the notation of Lemma 3.25, $L_{M\cap T} = \mathrm{lk}\left(\pi_M(p), M \cap T\right)$ is a closed subspace of $\mathrm{lk}\left(\pi_M(p), M\right)$ which is possibly proper. Indeed, $\pi_M(p)$ might live on a lower dimensional cell, where M might branch off away from T, as in Fig. 19. However, all the branches make an angle of at least π with each other, because M is convex.

Lemma 3.27 *Let $M \in \mathcal{M}_i$ and $C \in \mathcal{C}_i$. Then for $\varepsilon > 0$ small enough there is $k \in (-1, 0)$ such that the following hold.*

(1) *The metric induced on $E^{\varepsilon}_{M,C}$ is locally* CAT(k).
(2) *The nearest point projection $\pi_M : E^{\varepsilon}_{M,C} \to M$ maps non–constant local geodesics to non–constant local geodesics.*
(3) *The metric induced on $E^{\varepsilon}_{M,C}$ is* CAT(k).

Proof To prove (1) we argue as follows. By (3) in Lemma 3.23, we know that, away from the intersection with mirrors, $E^{\varepsilon}_{M,C}$ is locally isometric (via a developing map) to an equidistant hypersurface in $\mathbb{H}^n$. Such a hypersurface is a manifold of negative curvature $k \in (-1, 0)$ by (5) in Lemma 3.22. By Remark 3.24, $E^{\varepsilon}_{M,C}$ is essentially a cell complex, so by [11, Theorem II.5.2] $E^{\varepsilon}_{M,C}$ is locally CAT(k) if and only if the link of every vertex is a CAT(1) space. This condition is verified by (3) in Lemma 3.25.

Now we consider (2). By (3) in Lemma 3.23 and (5) in Lemma 3.22, we know that in the interior of each tile π_M is a conformal diffeomorphism with constant

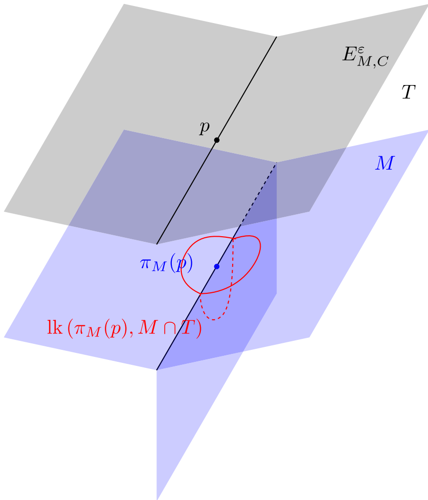

Fig. 19 A mirror M branching away from T, the union of tiles containing p (other mirrors not displayed)

conformal factor. Therefore it sends a local geodesic on $E^{\varepsilon}_{M,C}$ to a piecewise local geodesic on M, possibly broken at points where two or more tiles meet. To take care of those possibly singular points, we invoke (2) in Lemma 3.25, which guarantees that π_M induces an isometric embedding of links also at those points. Indeed, if $p \in E^{\varepsilon}_{M,C}$ is such a break point, and c is a geodesic on $E^{\varepsilon}_{M,C}$ through p, then the incoming and outgoing directions are at distance $D \geq \pi$ in $\mathrm{lk}\left(p, E^{\varepsilon}_{M,C}\right)$. Let $c' = \pi_M(c)$. Then c' is a piecewise geodesic in M through $\pi_M(p)$. With the notations of Lemma 3.25, the distance in $\mathrm{lk}\,(\pi_M(p), M \cap T)$ between the incoming and outgoing directions is the same $D \geq \pi$. The distance in the full $\mathrm{lk}\,(\pi_M(p), M)$ is not smaller, as $\mathrm{lk}\,(\pi_M(p), M)$ does not contain geodesic loops shorter than 2π by convexity. So, c' is a local geodesic in M at $\pi_M(p)$. Moreover if c is non-constant then c' is non–constant because π_M is locally injective.

To conclude, we prove (3). By (1) we know that $E^{\varepsilon}_{M,C}$ is locally CAT(k), so we only need to prove that it is also simply connected. By contradiction, let $\gamma \in \pi_1(E^{\varepsilon}_{M,C})$ be a non–trivial homotopy class. Since $E^{\varepsilon}_{M,C}$ is complete and non–positively curved, γ is represented by a unique non–constant local geodesic c_γ. By (2) $\pi_M(c_\gamma)$ is a non–constant local geodesic on M. Since M is complete and non–positively curved, $\pi_M(c_\gamma)$ is not nullhomotopic, which contradicts the fact that M is contractible. □

Remark 3.28 Note that if for a mirror M and a tile τ the intersection $M \cap \tau$ was lower–dimensional, then the equidistant space would develop to an equidistant hypersurface from a lower–dimensional totally geodesic subspace of $\mathbb{H}^n$, which has some positive curvature. So, Lemma 3.21 (establishing that if a mirror intersects a tile then the intersection is a codimension-1 cell) is a key tool to prove that edge spaces are non–positively curved.

Proposition 3.29 $\widetilde{X_\Gamma}$ *admits the structure of a graph of spaces, with underlying graph a connected tree.*

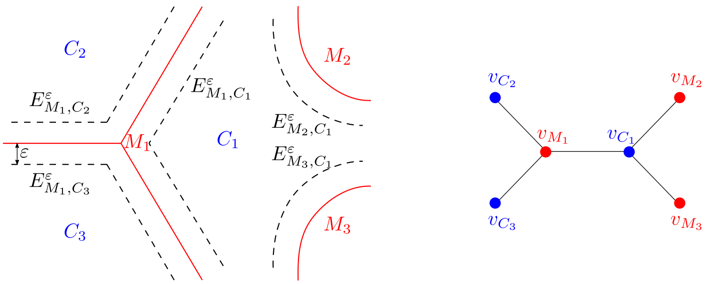

Fig. 20 The graph of spaces decomposition of $\widetilde{X_\Gamma}$

Proof We define a graph T_i as follows (see Fig. 20). Vertices come in two different families, namely a vertex v_M for each mirror $M \in \mathcal{M}_i$ and a vertex v_C for each component $C \in \mathcal{C}_i$. Then we place one edge $e_{M,C}$ between v_M and v_C whenever M intersects the closure of C. Vertex and edge spaces are defined as follows: we associate M to v_M, C to v_c, and $E^{\varepsilon}_{M,C}$ to the edge $e_{M,C}$ between them.

The edge maps to the two types of vertices are respectively given by the nearest point projection $\pi_M : E^{\varepsilon}_{M,C} \to M$ and the natural inclusion $i : E^{\varepsilon}_{M,C} \hookrightarrow C$. Note that i is an embedding, and that by (5) in Lemma 3.22 we know that the restriction of the projection π_M to each tile is an embedding too. Moreover, edge spaces are contractible by (3) in Lemma 3.27, so the gluing maps are automatically injective on fundamental groups.

Recall that by construction any two mirrors from $\mathcal{M}_i$ are disjoint, and even have disjoint ε–neighborhoods for ε sufficiently small (because Γ is cocompact). Similarly, any two components from $\mathcal{C}_i$ are disjoint. Moreover, $\widetilde{X_\Gamma} \setminus \bigsqcup_{M\in\mathcal{M}_i} M = \bigsqcup_{C\in\mathcal{C}_i} C$ and the boundary of a component consists of a disjoint union of closed subspaces, each of which sits inside a different mirror from $\mathcal{M}_i$. The resulting graph of spaces is homeomorphic to $\widetilde{X_\Gamma}$.

We are left to show that T_i is a connected tree. Connectedness of T_i follows directly from that of X. There is a natural map $r_i : \widetilde{X_\Gamma} \to T_i$ obtained by collapsing all the vertex spaces to points and all the cylinders over edge spaces to edges. Notice that r_i is a retraction and $\widetilde{X_\Gamma}$ is contractible, which forces T_i to be simply connected. □

Remark 3.30 In this graph of spaces decomposition all the spaces involved are non–positively curved, but the edge maps are not local isometries. Moreover, further pathological behavior can arise depending on the structure of the mirrors, as we now discuss. Note that the following phenomena already arise in the setting of cubical complexes, i.e. are not introduced by the hyperbolization procedure.

On one hand, if the mirror M branches (i.e. has non–locally Euclidean points) in such a way that different branches meet the closure of different complementary components, then the nearest point projections $\pi_M : E^{\varepsilon}_{M,C} \to M$ from the individual edge spaces fail to be surjective.

On the other hand, if the mirror M is such that a complementary component C wraps around M and meets it on different sides, then the map $\pi_M : E^{\varepsilon}_{M,C} \to M$ fails to be injective. This would be the case for a mirror that separates locally but not globally, e.g. one that is contained in the closure of a single complementary component.

In this case the corresponding vertex would be a boundary vertex for the tree T_i. We will see in §3.7 that this failure of injectivity does not occur in our setting.

Remark 3.31 (A graph of groups decomposition for Γ_X) Note that $\Gamma_X = \pi_1(X_\Gamma)$ acts on $\widetilde{X}_\Gamma$ sending mirrors to mirrors and preserving the coloring, i.e. each family $\mathcal{M}_i$. In particular it preserves this graph of spaces decomposition, hence it acts on the underlying graph, which has been seen to be a tree. The action is without global fix points and without inversions. This realizes $\Gamma_X = \pi_1(X_\Gamma)$ as a graph of groups. It is worth noticing that combination theorems are available in the literature, which provide a way to construct a cubulation of a group expressed as a graph of cubulated groups, when certain conditions are met (see for instance [40, 43, 75]). In our context, the vertex groups are given by the fundamental groups of the mirrors from $\mathcal{M}_i$ and the components from $\mathcal{C}_i$. While it is reasonable to expect that the former are cubulated (e.g. arguing by induction on dimension), it is not at all clear that the latter should be. The guiding idea for the rest of the paper is that nevertheless those components can be further decomposed into tiles. The fundamental group of a tile can be shown to be cubulated (see Lemma 5.12), and the results of Groves and Manning from [35] then provide a way to combine the cubulation from each tile into a global cubulation.

3.7 Mirrors: separation

In this section we will prove a strong separation property for mirrors in $\widetilde{X}_\Gamma$. In order to obtain convexity of the mirrors, in the proof of Proposition 3.14 we have used the fundamental fact that in a CAT(0) space local convexity implies global convexity. The same local–to–global property fails for separation, as shown by the following example.

Example 3.32 Consider the square complex Y in the center of Fig. 21. Consider the subcomplex Z consisting of the central thick (red) edge. The subspace Z is locally separating in Y, in the sense that for any $z \in Z$ and any arbitrarily small neighborhood U_z of z in Y, $U_z \setminus Z$ is disconnected. However, Z is not separating, i.e. $Y \setminus Z$ is connected. Notice that Y is a CAT(0) and foldable cubical complex, but Z is not a full connected component of the preimage of a codimension-1 face, i.e. not a mirror.

In this example both Y and Z have boundary, but it can be modified to obtain an example without boundary. We start by attaching eight more squares following the pattern in Fig. 21, and extending Z with two more edges. In the resulting complex, no edge meeting Z is a boundary edge, so we can keep adding squares (and extending Z) to get an admissible complex which displays the same pathology as the original one.

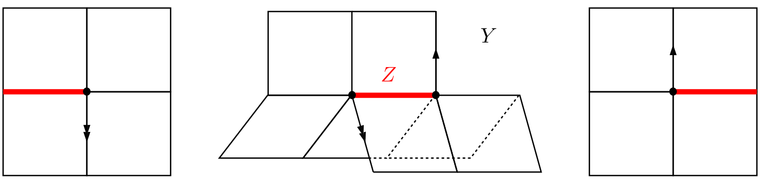

Fig. 21 A locally separating but not separating subcomplex in a CAT(0) square complex

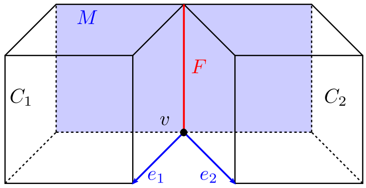

Fig. 22 A framing for a cube F on a mirror M

When Y is a homogeneous cubical complex of dimension n, every k–cube F of Y is contained in some n-cell. When Y has no boundary, F is contained in at least two distinct n-cubes. This motivates the following definition. Let M be a mirror of Y and let F be a k–cube of M. A *framing* for F is a choice of two n–cubes $\{C_1, C_2\}$ of Y such that $F \subseteq C_1 \cap C_2 \subseteq M$. We note explicitly that this definition is relative to the fixed mirror M. For the next proof, we will make use of some properties of hyperplanes in CAT(0) cubical complexes. We refer the reader to [65, Theorem 4.10] or [39, Example 3.3.(3), Lemma 13.3] for details and proofs.

Lemma 3.33 *Let Y be a* CAT(0) *admissible cubical complex. Then each mirror separates Y. More precisely, let $M \subseteq Y$ be a mirror, let $F \subseteq M$ be a k–cube, and let $\{C_1, C_2\}$ be a framing for F. Then C_1, C_2 are contained in the closure of two distinct connected components of $Y \setminus M$.*

Proof Let v be a vertex on F, let e_i be the edge of C_i with starting point v and endpoint in $C_i \setminus M$ (see Figure 22). Note that this edge exists because C_i is n–dimensional, while M is $(n-1)$–dimensional and convex, so that $M \cap C_i$ is some $(n-1)$–dimensional face E_i of C_i, by an argument similar to that of Lemma 3.21. Also note that by definition of framing, $C_1 \cap C_2 \subseteq M$, and therefore $e_1 \neq e_2$. Let H_i be the hyperplane of Y dual to e_i. In particular this means that H_i meets C_i in the midcube orthogonal to e_i. Since Y is CAT(0), Y is special in the sense of [39]. Since foldability of Y prevents e_1, e_2 from being contained in the same square, we then get that $H_1 \neq H_2$ (hyperplanes do not self-osculate) and $H_1 \cap H_2 = \varnothing$ (hyperplanes do not inter-osculate). Moreover, $H_k \cap M = \varnothing$ and $Y \setminus H_k$ consists of exactly two components, one containing M and one not containing M.

The carrier of a hyperplane H in a CAT(0) is isomorphic to $H \times [0, 1]$. By definition of mirror, if M contains an $(n-1)$-cube of $H \times \{0\}$ then actually $H \times \{0\} \subseteq M$. Since M contains the $(n-1)$–cell $E_i = C_i \cap M$ of C_i, and $E_i \subseteq H_i \times \{0\}$ by construction, we can conclude that M contains $H_i \times \{0\}$ for $i = 1, 2$. It follows that any path from H_1 to H_2 must intersect M. In particular, M separates Y in at least two components, one containing H_1 and one containing H_2. The closures of such components contain C_1 and C_2 respectively. □

We want to extend this result to mirrors in $\widetilde{X_\Gamma}$. To do this, we introduce the following terminology, in analogy with the cubical case. Let M be a mirror of $\widetilde{X_\Gamma}$, and let σ be a k–cell of M. A *framing* for σ is the choice of two distinct n–cells τ_1, τ_2 such that $\sigma \subseteq \tau_1 \cap \tau_2 \subseteq M$. We begin by obtaining a weak separation property.

Lemma 3.34 *Let $M \in \mathcal{M}_i$, let $\sigma \subseteq M$ be a k–cell, and let $\{\tau_1, \tau_2\}$ be a framing for σ. Then there exist two different components $C_1, C_2 \in \mathcal{C}_i$ whose closure contain τ_1, τ_2 respectively.*

Proof The map $g_X : X_\Gamma \to X$ lifts to a map $\alpha : \widetilde{X_\Gamma} \to \widetilde{X}$ between the universal covers. Note that it sends mirrors to mirrors. In particular we obtain a mirror $\alpha(M)$ and a k–cube $\alpha(\sigma) \subseteq \alpha(M)$ with a framing $\{\alpha(\tau_1), \alpha(\tau_2)\}$. By Lemma 3.33 we can conclude that $\alpha(\tau_1)$ and $\alpha(\tau_2)$ are separated by $\alpha(M)$ in $\widetilde{X}$. This implies that $\alpha^{-1}(\alpha(\tau_1))$ and $\alpha^{-1}(\alpha(\tau_2))$ are separated in $\widetilde{X_\Gamma}$ by $\alpha^{-1}(\alpha(M))$, i.e. the full preimage of the mirror $\alpha(M)$ in $\widetilde{X_\Gamma}$. Note that $\alpha^{-1}(\alpha(M))$ consist of infinitely many mirrors from $\mathcal{M}_i$: indeed, recall from Lemma 3.21 that disjoint $(n-1)$–cells of a tile belong to different mirrors. A fortiori, τ_1 and τ_2 are separated by the entire collection $\mathcal{M}_i$. In particular, there exists two different components $C_1, C_2 \in \mathcal{C}_i$ whose closure contain τ_1, τ_2 respectively, as desired. □

Remark 3.35 Observe that in the proof of Lemma 3.34, it is not clear whether the framing is separated by M itself. While the entire collection of mirrors $\mathcal{M}_i$ disconnects $\widetilde{X_\Gamma}$ into a collection of complementary components, it is not a priori clear that any single mirror separates $\widetilde{X_\Gamma}$.

Recall from Proposition 3.29 that $\widetilde{X_\Gamma}$ admits the structure of a graph of spaces over a connected tree T_i, and that there is a natural retraction $r_i : \widetilde{X_\Gamma} \to T_i$ obtained by collapsing all the vertex spaces to points and all the cylinders over edge spaces to edges.

Lemma 3.36 *The tree T_i has no boundary.*

Proof It is enough to show that each vertex has at least two neighboring vertices. Vertices of T_i are either associated to mirrors from $\mathcal{M}_i$ or to components from $\mathcal{C}_i$. We analyze the two different cases separately. Let v_C be the vertex associated to a component $C \in \mathcal{C}_i$. Then v_C has infinitely many edges coming into it, because C has infinitely many mirrors from $\mathcal{M}_i$ on its boundary (this is already true for a single tile: by Lemma 3.21, disjoint $(n-1)$–cells in the boundary of a tile belong to different mirrors).

Now let v_M be the vertex associated to a mirror $M \in \mathcal{M}_i$. Let $\sigma \subseteq M$ be an $(n-1)$-cell on it, and pick a framing $\{\tau_1, \tau_2\}$. By Lemma 3.34, there exist two different components $C_1, C_2 \in \mathcal{C}_i$ whose closure contain τ_1, τ_2 respectively. The corresponding vertices v_{C_1}, v_{C_2} in T_i are both adjacent to the vertex v_M corresponding to M, as desired. □

The next result is the analogue of Lemma 3.33 from the cubical case.

Proposition 3.37 *Each $M \in \mathcal{M}_i$ separates $\widetilde{X_\Gamma}$. More precisely, let $M \in \mathcal{M}_i$ be a mirror, let $\sigma \subseteq M$ be a k–cell, and let $\{\tau_1, \tau_2\}$ be a framing for σ. Then τ_1, τ_2 are contained in the closure of two distinct connected components of $\widetilde{X_\Gamma} \setminus M$.*

Proof For the first statement, consider the natural retraction $r_i : \widetilde{X}_\Gamma \to T_i$. Note that for each mirror $M \in \mathcal{M}_i$ there is a corresponding vertex $v_M \in T_i$, and $M = r_i^{-1}(v_M)$. By Lemma 3.36 we know that T_i is a tree with no boundary, hence any of its vertices disconnects it. Therefore $M = r_i^{-1}(v_M)$ disconnects $\widetilde{X}_\Gamma$.

For the second statement, we fix a k–cell $\sigma \subseteq M$ and a framing $\{\tau_1, \tau_2\}$. By Lemma 3.34 we get two components $C_1, C_2 \in \mathcal{C}_i$ containing τ_1, τ_2 in their closures. Note that these are complementary components of the entire collection of mirrors $\mathcal{M}_i$, not complementary components of the mirror M. The corresponding vertices v_{C_1}, v_{C_2} in T_i are both adjacent to the vertex v_M corresponding to M, and are separated by v_M in T_i, since T_i is a tree (see Proposition 3.29). Arguing as above via the natural retraction $r_i : \widetilde{X}_\Gamma \to T_i$, we can conclude that τ_1, τ_2 are separated by M in $\widetilde{X}_\Gamma$. □

We conclude this section with some remarks about the construction that we have described.

Remark 3.38 (Foldability is key) Foldability of X has played the role of some sort of *combinatorial completeness*, as it guarantees that if a mirror M intersects a tile T, then M goes across T along a top dimensional subcomplex of the boundary. This has provided both features of non–positive curvature (see Remark 3.28) and separation properties (as in the proof of Lemma 3.33). Example 3.32 shows that neither is available if foldability is not taken into account in the definition of mirrors (even on a foldable complex).

Remark 3.39 (Complexes with boundary) The construction from §3.6 can be generalized to cubical complexes that have enough *good* mirrors (i.e. mirrors that admit a cell which locally separates a framing), and keeping track only of such mirrors in the construction of the tree of spaces. For instance, one could drop the assumption that X is without boundary, and ignore the mirrors that are entirely contained in the boundary. One still gets a decomposition as a graph of spaces over a tree without boundary. Indeed, vertices associated to good mirrors still have degree at least 2. One may worry about vertices associated to components. Even if there is a cube of X with only one face F contained in a good mirror, each of the components $C \in \mathcal{C}_i$ of $\widetilde{X}_\Gamma$ arising from it still has infinitely many boundary cells corresponding to F. This guarantees that the vertices of the tree which are associated to components in $\mathcal{C}_i$ still have infinite degree.

Alternatively, one can work with a relative version of the Charney-Davis hyperbolization procedure that is designed to hyperbolize complexes with boundary without altering the boundary components (see [9, 14]). We consider the problem of cubulating the resulting relatively hyperbolic groups in [36].

4 The dual cubical complex

We define a cubical complex associated to the stratification of $\widetilde{X}_\Gamma$ introduced in §3.5, and prove that it is a CAT(0) cubical complex (see Theorem 4.29). Recall that X is

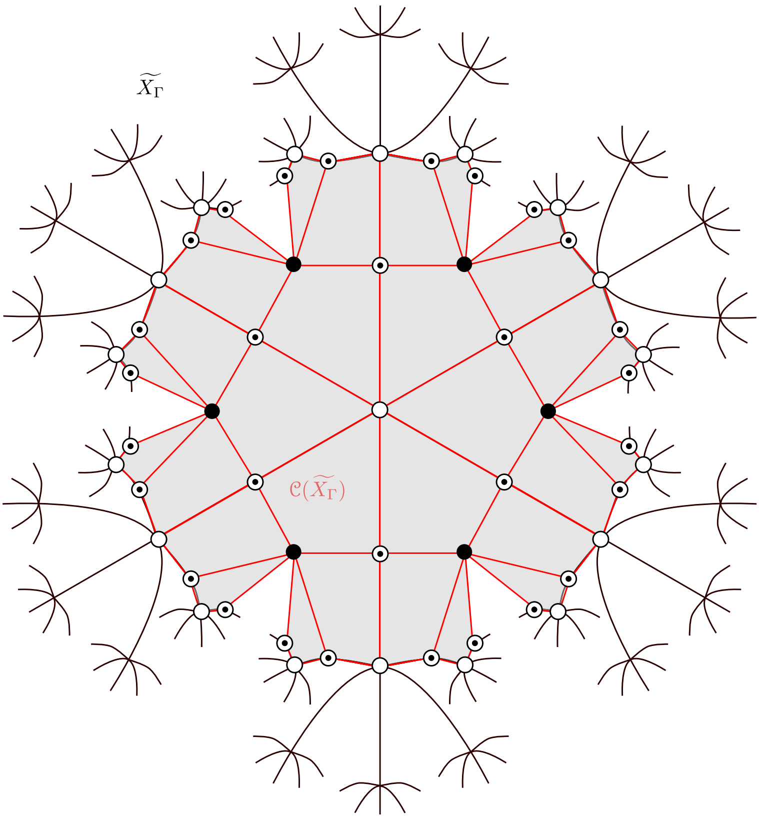

Fig. 23 The dual cubical complex $\mathcal{C}(\widetilde{X}_\Gamma)$ superimposed on the stratification of $\widetilde{X}_\Gamma$; compare Fig. 13. (In this picture the dimension is $n = 2$. Key: ○, ⊙, and ● denote a vertex of height 0, 1, 2 respectively)

assumed to be an admissible cubical complex (as defined at the beginning of §3). Let $n = \dim(X)$ be its dimension. The *dual cubical complex* is denoted $\mathcal{C}(\widetilde{X}_\Gamma)$ and defined as follows.

- Vertices are given by all the k-cells in $\widetilde{X}_\Gamma$ for $k = 0, \dots, n$.
- Two vertices corresponding to cells σ and τ are connected by an edge if and only if $|\dim(\sigma) - \dim(\tau)| = 1$, and either $\sigma \subseteq \tau$ or $\tau \subseteq \sigma$.
- For $k > 1$, we attach one k-dimensional cube whenever we see its 1-skeleton.

The resulting cell complex $\mathcal{C}(\widetilde{X}_\Gamma)$ is a cubical complex (see Fig. 23). Moreover, we can label its 0–skeleton by integers $0 \le k \le n$: if v is a vertex dual to a k–cell σ, then we define the *height* of v to be $\mathrm{h}(v) = \dim(\sigma) = k$.

Note that if one applies this construction to the standard n-simplex, one obtains a natural cubulation of the n-simplex by $n+1$ cubes of the same dimension. However, it is not clear when this construction preserves asphericity. For instance, applying this construction to a solid octahedron results in a cube complex with non–trivial π_2.

In this section we study the combinatorial geometry of $\mathcal{C}(\widetilde{X}_\Gamma)$, by analyzing cubes and links in §4.1, some notions of complexity for edge–paths in §4.2 and §4.3, and how to use them to prove that $\mathcal{C}(\widetilde{X_\Gamma})$ is simply connected in §4.4. Before starting, the following two remarks address the relation between $\mathcal{C}(\widetilde{X_\Gamma})$ and other natural combinatorial structures associated to $\widetilde{X_\Gamma}$ and its collection of mirrors $\mathcal{M}$.

Remark 4.1 (The associated graded poset) The set of cells in $\widetilde{X_\Gamma}$ can be partially ordered by inclusion. The result is a graded poset, whose rank function is given by the dimension of the corresponding cell. The height we just defined is induced by this rank function. One could construct the order complex of such a poset, by taking a simplex for every chain. This would result in a simplicial complex, and is not what we are considering here.

Remark 4.2 (The associated wallspace) Since mirrors are separating subspaces (see Proposition 3.37), the collection of mirrors can be used to define a wallspace structure $(\widetilde{X_\Gamma}, \mathcal{M})$ on $\widetilde{X_\Gamma}$, and one could consider the dual CAT(0) cubical complex $\mathcal{C}(\widetilde{X_\Gamma}, \mathcal{M})$ associated to this wallspace by Sageev's construction. We refer the reader to [38, 42, 65] for details about this construction, and we only review the main ingredients here. Given a mirror M, any partition of the complementary components into two classes is called a wall associated to M. An orientation of a wall is a choice of one of the two classes. A vertex of $\mathcal{C}(\widetilde{X_\Gamma}, \mathcal{M})$ can then be described as a consistent choice of orientation for each mirror.

When X and all mirrors are homeomorphic to manifolds, each mirror of $\widetilde{X_\Gamma}$ has exactly two complementary components. In this quite restrictive case, an orientation of a wall is just a choice of one of the two complementary components. Therefore vertices of $\mathcal{C}(\widetilde{X_\Gamma}, \mathcal{M})$ correspond to tiles (i.e. n–cells) in the stratification of $\widetilde{X_\Gamma}$, and two vertices are connected by an edge when the corresponding tiles are adjacent along a mirror. In particular, $\mathcal{C}(\widetilde{X_\Gamma}, \mathcal{M})$ is an n–dimensional cubical complex that can be subdivided to recover $\mathcal{C}(\widetilde{X_\Gamma})$. However, if there are mirrors which have more than two complementary components (such as in Figs. 14 and 23), then we find vertices in $\mathcal{C}(\widetilde{X_\Gamma}, \mathcal{M})$ which do not correspond to tiles from the stratification of $\widetilde{X_\Gamma}$ (they are not canonical vertices, in the terminology of [42]). As a result, the dimension of $\mathcal{C}(\widetilde{X_\Gamma}, \mathcal{M})$ is usually higher than that of X, and it is more challenging to relate the actions of Γ_X on $\widetilde{X_\Gamma}$ and on $\mathcal{C}(\widetilde{X_\Gamma}, \mathcal{M})$.

4.1 Cubes and links

In this section we explore basic facts about the cubical geometry of $\mathcal{C}(\widetilde{X_\Gamma})$. While this complex is not locally compact (see Remark 4.8), its dimension is the same as that of X (see Lemma 4.5), and the links of vertices are flag complexes (see Proposition 4.10).

The first two lemmas show that squares and cubes in $\mathcal{C}(\widetilde{X_\Gamma})$ admit unique vertices of minimum and maximum height. Recall that the height of a vertex is the dimension of its dual cell, and notice that, by definition of $\mathcal{C}(\widetilde{X_\Gamma})$, if u, v are adjacent vertices, then $|\,\mathrm{h}(u) - \mathrm{h}(v)| = 1$.

Lemma 4.3 *Let S be a square of $\mathcal{C}(\widetilde{X_\Gamma})$. Let v_1, v_2, v_3, v_4 be its vertices, with v_2 and v_4 non-adjacent in S. If $\mathrm{h}(v_2) = \mathrm{h}(v_4)$, then $|\,\mathrm{h}(v_1) - \mathrm{h}(v_3)| = 2$. In particular there is a unique vertex of maximal (respectively minimal) height, and the cell dual to it contains (respectively is contained in) each of the cells dual to the other vertices.*

Fig. 24 A square in $\mathcal{C}(\widetilde{X_\Gamma})$

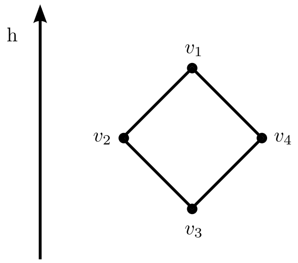

Proof Let $h = \mathrm{h}(v_2) = \mathrm{h}(v_4)$ be the common value of the height of v_2 and v_4. Since v_1 is adjacent to v_2 and v_4, we have $\mathrm{h}(v_1) = h \pm 1$, and similarly for v_3 (see Fig. 24). In particular $|\,\mathrm{h}(v_1) - \mathrm{h}(v_3)|$ is either 0 or 2. By contradiction let us assume that $|\,\mathrm{h}(v_1) - \mathrm{h}(v_3)| = 0$, i.e. $\mathrm{h}(v_1) = \mathrm{h}(v_3) = h \pm 1$. Without loss of generality we can assume that $\mathrm{h}(v_1) = \mathrm{h}(v_3) = h + 1$. (The case $\mathrm{h}(v_1) = \mathrm{h}(v_3) = h - 1$ is completely analogous, via a dual argument). For $j = 1, 2, 3, 4$, let σ_j be the cell of $\widetilde{X_\Gamma}$ dual to the vertex v_j. Since v_1 is adjacent to v_2 and v_4, and has higher height, σ_1 contains σ_2 and σ_4; the same holds for σ_3. So $\sigma_1 \cap \sigma_3$ contains $\sigma_2 \cup \sigma_4$, contradicting Lemma 3.18.

To prove the final statement, let us assume without loss of generality that v_1 is the vertex of maximal height and v_3 is the one of minimal height, i.e. $\mathrm{h}(v_1) - 1 = h = \mathrm{h}(v_3) + 1$. Then we have that $\sigma_3 \subseteq \sigma_2$, $\sigma_4 \subseteq \sigma_1$. $\square$

In the next lemma we extend this result to higher dimensional cubes of $\mathcal{C}(\widetilde{X_\Gamma})$. By an *edge–path* in $\mathcal{C}(\widetilde{X_\Gamma})$ we will mean a continuous path which is entirely contained in the 1–skeleton (i.e. is a sequence of edges). If an edge–path p goes through vertices $v_0, \ldots, v_s$ of $\mathcal{C}(\widetilde{X_\Gamma})$, we will write $p = (v_0, \ldots, v_s)$; note that the sequence of vertices completely determines the sequence of edges, hence the path. We call p an *edge–loop* if it is a closed loop, i.e. $v_0 = v_s$. For an edge–path $p = (v_0, \ldots, v_s)$ we define $\ell(p) = s$ to be the *length* of p, i.e. the number of edges in it. We also define the *height* of p to be $\mathrm{h}(p) = \max\{\mathrm{h}(v_0), \ldots, \mathrm{h}(v_s)\}$. Notice that along each edge of p the height must increase or decrease exactly by 1.

Lemma 4.4 *Let Q be a cube of $\mathcal{C}(\widetilde{X_\Gamma})$. Then the following hold.*

(1) *There is a unique vertex $v \in Q$ of minimal height. The cell dual to it is contained in each of the cells dual to the vertices of Q.*
(2) *There is a unique vertex $w \in Q$ of maximal height. The cell dual to it contains each of the cells dual to the vertices of Q.*

Proof We prove the first statement; the second is obtained by an analogous argument. Let k be the minimal height of vertices of Q, and assume by contradiction that there is at least a pair of vertices of Q of height k. Consider an edge–path $p = (v_0, \ldots, v_s)$ in Q such that $\mathrm{h}(v_0) = \mathrm{h}(v_s) = k$, $v_0 \neq v_s$, and such that p is an edge–path of minimal height among all edge–paths in Q joining a pair of vertices of height k. This is well–defined since the height of such a path can only be an integer between 0 and n. Let $\mathrm{h}(p) = h$ be the height of p.

Let v_j be a vertex of p of maximal height $\mathrm{h}(v_j) = h = \mathrm{h}(p)$. Then $\mathrm{h}(v_{j\pm1}) = h-1$ (notice that $k \geq 0$ and $h \geq k + 1 \geq 1$). Since (v_{j-1}, v_j, v_{j+1}) is part of the cube Q,

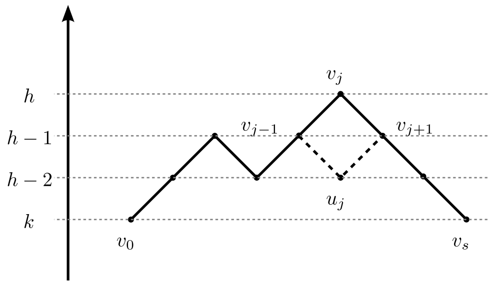

Fig. 25 Lowering a vertex of maximal height on a path in $\mathcal{C}(\widetilde{X_\Gamma})$

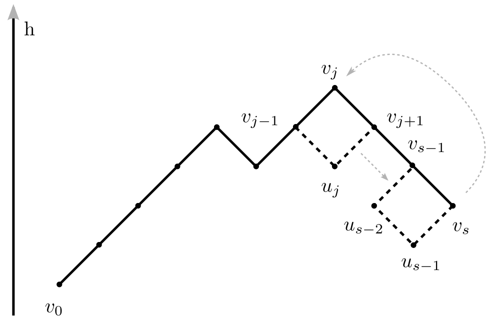

Fig. 26 Walking backwards along a path in $\mathcal{C}(\widetilde{X_\Gamma})$ to find a vertex v_j that can be lowered, and then walking forward to lower vertices until the exceptional endpoint v_s

it must be contained in a square, i.e. there exists a vertex u_j of Q (not necessarily on p), such that $\{v_{j-1}, v_j, v_{j+1}, u_j\}$ span a square in Q. Lemma 4.3 implies that $\mathrm{h}(u_j) = h - 2$. We can construct a new path in Q starting from the path p by lowering the vertex of maximal height, i.e. by replacing v_j with u_j (see Fig. 25). We repeat the same operation on all vertices of height h along the path, and let p' be the resulting path in Q. We have that $\mathrm{h}(p') = h - 1 < h = \mathrm{h}(p)$, contradicting the minimality of the height of p. This concludes the proof by contradiction, and proves the uniqueness of a vertex v of minimal height k in Q.

We are left to show that the cell σ dual to v is contained in all the cells dual to the other vertices of Q. By contradiction suppose there are vertices in Q whose dual cells do not contain σ; call such vertices exceptional. Let $p = (v_0, \dots, v_s)$ be an edge–path in Q with $v_0 = v$, v_s an exceptional vertex, and having minimal length among all edge–paths of Q between v and an exceptional vertex. We have $\mathrm{h}(v_{s-1}) = \mathrm{h}(v_s) \pm 1$. If $\mathrm{h}(v_{s-1}) = \mathrm{h}(v_s) - 1$, then the cell dual to v_{s-1} is contained in the cell dual to v_s. By minimality of p, we have that v_{s-1} is not exceptional, so the cell dual to v_{s-1} contains σ, and hence v_s cannot be exceptional. Therefore $\mathrm{h}(v_{s-1}) = \mathrm{h}(v_s) + 1$ (as in Fig. 26).

We keep walking backwards along p until we find a triple of vertices $\{v_{j-1}, v_j, v_{j+1}\}$ such that $\mathrm{h}(v_j) = \mathrm{h}(v_{j\pm1}) + 1$ (notice j is well defined and positive, since v is the unique vertex of minimal height in the whole Q). Arguing as before we complete to a square in Q with vertices $\{v_{j-1}, v_j, v_{j+1}, u_j\}$; again by Lemma 4.3 we have $\mathrm{h}(u_j) = \mathrm{h}(v_j) - 2$ (see Fig. 26). By minimality of p, u_j must be non–exceptional, and so we can change p by replacing v_j with u_j, without changing its length. Walking forward along p, we can keep changing the path without changing its length, until we are able to complete $\{v_{s-1}, v_s\}$ to a square $\{u_{s-2}, v_{s-1}, v_s, u_{s-1}\}$ in Q with u_{s-1}

non-exceptional and with height $\mathrm{h}(u_{s-1}) = \mathrm{h}(v_{s-1}) - 2 = \mathrm{h}(v_s) - 1$ (once again, see Fig. 26). In particular, the cell dual to u_{s-1} contains σ and is contained in the cell dual to v_s, which contradicts the fact that the last vertex v_s was chosen to be exceptional. □

As a consequence, we obtain the following statement.

Lemma 4.5 *The complex* $\mathcal{C}(\widetilde{X_\Gamma})$ *has dimension* $\dim \mathcal{C}(\widetilde{X_\Gamma}) = \dim X = n$.

Proof If τ is a tile of $\widetilde{X_\Gamma}$ and x one of its vertices, then the collection of cells containing x and contained in τ provides a cube of dimension exactly n, so $\dim \mathcal{C}(\widetilde{X_\Gamma}) \geq n$, so we focus on the other inequality.

Let Q be a cube of $\mathcal{C}(\widetilde{X_\Gamma})$, and let $v_{\min}$ be the vertex of minimal height in Q (see Lemma 4.4). We claim that for each vertex $v \in Q$ we have

$$\mathrm{h}(v) = \mathrm{h}(v_{\min}) + d_Q(v_{\min}, v)$$

where $d_Q(v_{\min}, v)$ is the distance in Q of v from $v_{\min}$. Since the height can take values only between 0 and $n = \dim X$, this directly implies that

$$\dim Q = \max\{d_Q(v_{\min}, v)\} = \max\{\mathrm{h}(v) - \mathrm{h}(v_{\min})\} \leq n.$$

In order to prove the claim, pick a vertex $v \in Q$, and let $p = (v_0, \ldots, v_s)$ be an edge–path of minimal length $s = d_Q(v_{\min}, v)$ in Q from $v_0 = v_{\min}$ to $v_s = v$. Since the height can at most increase by 1 along each edge of p, we have the inequality $\mathrm{h}(v) \leq \mathrm{h}(v_{\min}) + d_Q(v_{\min}, v)$. Assume by contradiction that the inequality is strict. Then the height is not monotonically increasing along p. Let v_k be the first vertex of p which is a local maximum for the height function. Arguing as above via Lemma 4.3, we look at the triple v_{k-1}, v_k, v_{k+1}, and complete it to a square with a fourth vertex u_k such that $\mathrm{h}(u_k) = \mathrm{h}(v_k) - 2$. We can even assume without loss of generality that $k = 2$ (otherwise we proceed as in the proof of Lemma 4.4 and change p along squares walking backwards towards $v_{\min}$). But then $\mathrm{h}(u_2) = \mathrm{h}(v_2) - 2 = \mathrm{h}(v_{\min})$. Minimality of $v_{\min}$ implies $u_2 = v_{\min}$, and therefore we get that v_3 was already adjacent to $v_{\min}$. This provides a path from $v_{\min}$ to v of length at most $d_Q(v_{\min}, v) - 2$, which is a contradiction. □

We now turn to the study of links of vertices of $\mathcal{C}(\widetilde{X_\Gamma})$. Recall that $\mathcal{C}(\widetilde{X_\Gamma})$ is a cubical complex, hence its links are simplicial complexes (see §2.1.3). In particular, if $v \in \mathcal{C}(\widetilde{X_\Gamma})$ is a vertex, then vertices in $\mathrm{lk}\big(v, \mathcal{C}(\widetilde{X_\Gamma})\big)$ correspond to vertices in $\mathcal{C}(\widetilde{X_\Gamma})$ which are adjacent to v. We begin with the following combinatorial characterization of simplices in the link of a vertex.

Lemma 4.6 *Let* σ *be a* k*-cell of* $\widetilde{X_\Gamma}$, *and let* v *be the dual vertex in* $\mathcal{C}(\widetilde{X_\Gamma})$. *Let* $v_0, \ldots, v_m$ *be a collection of vertices of* $\mathcal{C}(\widetilde{X_\Gamma})$ *adjacent to* v, *and let* $\tau_0, \ldots, \tau_m$ *be the dual cells in* $\widetilde{X_\Gamma}$. *Then* $v_0, \ldots, v_m$ *induce a simplex in* $\mathrm{lk}\big(v, \mathcal{C}(\widetilde{X_\Gamma})\big)$ *if and only if the following two conditions are satisfied*

($\downarrow$) *there exists a cell* λ *of* $\widetilde{X_\Gamma}$ *such that* $\lambda \subseteq \tau_j$, $j = 0, \ldots, m$,
($\uparrow$) *there exists a cell* μ *of* $\widetilde{X_\Gamma}$ *such that* $\tau_j \subseteq \mu$, $j = 0, \ldots, m$.

Proof First of all, note that since $\mathcal{C}(\widetilde{X_\Gamma})$ is a cubical complex, the vertices $v_0, \dots, v_m$ induce a simplex in $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ if and only if there exists a cube Q of $\mathcal{C}(\widetilde{X_\Gamma})$ containing $v, v_0, \dots, v_m$.

Assume that they induce a simplex, and let Q be the corresponding cube. From Lemma 4.4 we know that Q has a unique vertex of minimal height, and a unique vertex of maximal height. Let λ, μ be the dual cells. Lemma 4.4 then implies that λ, μ satisfy the conditions ($\downarrow$) and ($\uparrow$) in the statement.

Vice versa suppose that the conditions ($\downarrow$) and ($\uparrow$) are satisfied. Notice that we have $\lambda \subseteq \tau_j \subseteq \mu$ for all $j = 0, \dots, m$. Let $C_\lambda = \widetilde{g_X}(\lambda)$ and $C_\mu = \widetilde{g_X}(\mu)$ be the corresponding cubes of X, under the map $\widetilde{g_X} = g_X \circ \pi : \widetilde{X_\Gamma} \to X_\Gamma \to X$. Notice that $\mathrm{lk}\,(\lambda, \mu) \cong \mathrm{lk}\left(C_\lambda, C_\mu\right)$ by Lemma 3.17. In particular, we see that in $\widetilde{X_\Gamma}$ there is a collection of cells containing λ and contained in μ (among which we find the cells τ_j) that gives rise to a cube Q in $\mathcal{C}(\widetilde{X_\Gamma})$ containing the vertices $v, v_0, \dots, v_m$. Therefore $v_0, \dots, v_m$ induce a simplex in $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$, as desired. □

Remark 4.7 When the conditions ($\downarrow$) and ($\uparrow$) from Lemma 4.6 are satisfied, the cells λ, μ can be chosen to be the lower and upper cell provided by Lemma 3.19.

Remark 4.8 A cell of dimension at least 2 in $\widetilde{X_\Gamma}$ always admits infinitely many codimension-1 cells (see Fig. 11). Lemma 4.6 implies that the link of the dual vertex is neither compact nor connected. In particular the cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ is not locally compact. As a result, even though $\mathcal{C}(\widetilde{X_\Gamma})$ is constructed as a sort of dual cubical barycentric subdivision with respect to the combinatorial decomposition of $\widetilde{X_\Gamma}$ into cells, $\mathcal{C}(\widetilde{X_\Gamma})$ is not homeomorphic to $\widetilde{X_\Gamma}$. Namely, $\widetilde{X_\Gamma}$ is locally compact, while $\mathcal{C}(\widetilde{X_\Gamma})$ is not locally compact.

As recalled above, if $v \in \mathcal{C}(\widetilde{X_\Gamma})$ is a vertex, then the vertices appearing in $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ correspond to vertices of $\mathcal{C}(\widetilde{X_\Gamma})$ that are adjacent to v, and these vertices have height equal to $\mathrm{h}(v) \pm 1$. We find it useful to decompose $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ into two subcomplexes: we denote by $\mathrm{lk}_\downarrow\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ the full subcomplex of $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ generated by vertices of height $\mathrm{h}(v) - 1$, and by $\mathrm{lk}^\uparrow\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ the full subcomplex of $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ generated by vertices of height $\mathrm{h}(v) + 1$. As we will see, their geometry is controlled respectively by a certain Helly property for orthogonal hyperplanes in $\mathbb{H}^n$, and by the non–positive curvature of X. The following statement provides the Helly property. Notice that orthogonality is a key feature here: without the orthogonality requirement, the statement already fails for three geodesics in $\mathbb{H}^2$. On the other hand, the interested reader will notice that the argument generalizes to a collection of pairwise orthogonal and totally geodesic hypersurfaces in a simply connected complete manifold of non–positive curvature. We will not need this generality in this paper.

Lemma 4.9 (Helly property for orthogonal hyperplanes in $\mathbb{H}^n$) *Let $\mathcal{V}$ be a collection of pairwise orthogonal hyperplanes in $\mathbb{H}^n$. Then $|\mathcal{V}| \leq n$, and for all $k \in \{2, \dots, n\}$ all the k–fold intersections are non-empty.*

Proof We begin with some preliminary observation about orthogonal subspaces. Let $V_1, \dots, V_k \in \mathcal{V}$ be a collection of hyperplanes from $\mathcal{V}$, and let $N = \cap_{j=1}^k V_j$ be their

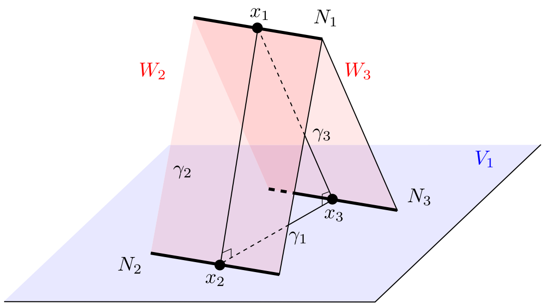

Fig. 27 The Helly property in Lemma 4.9

intersection. For $x \in N$, let $\mathrm{T}_x(\mathbb{H}^n)$ denote the tangent space of $\mathbb{H}^n$ at x, and let $v_j \in \mathrm{T}_x(\mathbb{H}^n)$ be a unit vector orthogonal to V_j (i.e. to all vectors in the tangent space $\mathrm{T}_x(V_j)$). The fact that V_i and V_j are orthogonal hyperplanes means that v_i and v_j are orthogonal vectors for all $i \neq j$. Then a direct computation shows that if $\{n_1, \dots, n_m\}$ is an orthonormal basis for the tangent space of $\mathrm{T}_x(N)$, then $\{n_1, \dots, n_m, v_1, \dots, v_k\}$ is an orthonormal basis for $\mathrm{T}_x(\mathbb{H}^n)$. This shows in particular that $k \leq n$.

To prove the statement about non–emptiness of intersections, we notice that the case $k = 2$ is exactly the hypothesis that any pair of hyperplanes from $\mathcal{V}$ intersect. For $k \geq 3$, we argue that if all the h–fold intersections of elements from $\{V_1, \dots, V_k\}$ are non–empty for all $h < k$, then the k–fold intersection is non–empty too.

Let $N_j = \cap_{i \neq j} V_i$. By assumption we have $N_j \neq \varnothing$. Assume by contradiction that $V_1 \cap \dots \cap V_k = \varnothing$. Then for any choice of indices $j_1 \neq j_2$ we have that $N_{j_1} \cap N_{j_2} = \varnothing$. In particular, N_2 and N_3 are non–empty disjoint subspaces of V_1 (see Fig. 27). Let γ_1 be the common perpendicular between them in V_1, and let $x_k \in N_k$ be its endpoint for $k = 2, 3$. Now in the tangent space $\mathrm{T}_{x_2}(\mathbb{H}^n)$ we consider an orthonormal basis $\{n_1, \dots, n_m, v_1, v_3, \dots, v_k\}$ constructed as above by adding to an orthonormal basis $\{n_1, \dots, n_m\}$ for $\mathrm{T}_{x_2}(N_2)$ unit vectors $v_1, v_3, \dots, v_k$ orthogonal to $V_1, V_3, \dots, V_k$. If w denotes a tangent vector at x_2 along γ_1, then a direct computation shows that w is orthogonal to $\{n_1, \dots, n_m\}$, because γ_1 is orthogonal to N_2, and it is also orthogonal to v_1, because $\gamma_1 \subseteq V_1$. Therefore w is in the subspace of $\mathrm{T}_{x_2}(\mathbb{H}^n)$ generated by $v_3, \dots, v_k$. If we define $W_2 = \cap_{j=3}^k V_j$, then this means that γ_1 is orthogonal to W_2 at x_2. Arguing in the same way at the point x_3, we find that γ_1 is orthogonal at x_3 to the subspace $W_3 = \cap_{j=2, j \neq 3}^k V_j$. Note that $W_2 \cap W_3 = \cap_{j=2}^k V_j = N_1$ is non–empty. Moreover, as observed above, it is disjoint from N_2 and from N_3. Therefore we can connect x_2 (respectively x_3) to a point x_1 in N_1 with a geodesic arc γ_2 contained in W_2 (respectively γ_3 contained in W_3). Since all the spaces involved are totally geodesic, the arcs $\gamma_1, \gamma_2, \gamma_3$ are geodesic arcs in $\mathbb{H}^n$, so we have obtained a geodesic triangle with two right angles, which leads to the desired contradiction. □

The next statement completes our investigation of the combinatorial geometry of $\mathcal{C}(\widetilde{X_\Gamma})$. Thanks to Gromov's link condition (see Lemma 2.2), it already implies that $\mathcal{C}(\widetilde{X_\Gamma})$ is locally CAT(0). We will show in Theorem 4.29 that it is actually CAT(0).

Proposition 4.10 *Let σ be a k-cell of $\widetilde{X_\Gamma}$, and let v be the dual vertex in $\mathcal{C}(\widetilde{X_\Gamma})$. Then the following hold*:

(1) $\mathrm{lk}_{\downarrow}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ *is a flag simplicial complex.*
(2) $\mathrm{lk}^{\uparrow}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ *is a flag simplicial complex.*
(3) $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ *is a flag simplicial complex.*

Proof Throughout this proof, w_j will denote a vertex in $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$, v_j the corresponding vertex of $\mathcal{C}(\widetilde{X_\Gamma})$ adjacent to v, and τ_j the cell of $\widetilde{X_\Gamma}$ dual to v_j. Notice that two vertices w_i, w_j are adjacent in $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ precisely when v, v_i, v_j are contained in a square of $\mathcal{C}(\widetilde{X_\Gamma})$.

We first prove (1). Let $w_0, \dots, w_p$ be pairwise adjacent vertices in $\mathrm{lk}_{\downarrow}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$. Notice that $\tau_0, \dots, \tau_p$ are all cells of codimension 1 in the boundary of σ. For each $i \neq j$, v, v_i, v_j are contained in a square of $\mathcal{C}(\widetilde{X_\Gamma})$. By Lemma 4.3, the fourth vertex of the square is dual to a cell contained in $\tau_i \cap \tau_j$. This shows that the cells τ_j intersect pairwise. We claim that actually $\tau_0 \cap \dots \cap \tau_p \neq \varnothing$. To see this, embed σ into a hyperbolic space of dimension $\dim \sigma$ (as in §3.3). The family of hyperplanes $V_0, \dots, V_p$ supporting the cells $\tau_0, \dots, \tau_p$ is a collection of pairwise orthogonal hyperplanes, and the boundary of the cell τ_j in V_j is given by subspaces that are orthogonal to the other V_i's. Note that by Lemma 4.9 we know that $V_0 \cap \dots \cap V_p \neq \varnothing$. Arguing by induction on p, we can assume that $\tau_0 \cap \dots \cap \tau_{p-1} \neq \varnothing$, and then we can leverage the orthogonality structure as in the proof of Lemma 4.9 to obtain that $\varnothing \neq \tau_0 \cap \dots \cap \tau_{p-1} \cap V_p \subseteq \tau_0 \cap \dots \cap \tau_p$, which proves the claim. Since the latter intersection is non-empty, by Lemma 3.19, it consists of a single cell $\lambda \subseteq \tau_j$. We use Lemma 4.6 with this cell λ and $\mu = \sigma$ to conclude that $w_0, \dots, w_k$ span a simplex.

We argue via a dual argument to prove (2). Let $w_0, \dots, w_p$ be pairwise adjacent vertices in $\mathrm{lk}^{\uparrow}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$. Notice that $\tau_0, \dots, \tau_p$ are cells containing σ as a cell of codimension 1 in their boundary. For each $i \neq j$, v, v_i, v_j are contained in a square of $\mathcal{C}(\widetilde{X_\Gamma})$. By Lemma 4.3, the fourth vertex of the square is dual to a cell containing $\tau_i \cup \tau_j$. So τ_i, τ_j are adjacent in $\mathrm{lk}\left(\sigma, \widetilde{X_\Gamma}\right)$. By (1) in Lemma 3.17 this link is isomorphic to the link of the corresponding cube in X. Since X is non-positively curved, this link is a flag simplicial complex. Therefore there is a cell μ containing all the cells τ_j; this can actually be taken to be the upper cell provided by Lemma 3.19. We use Lemma 4.6 with this cell μ and $\lambda = \sigma$ to conclude that $w_0, \dots, w_k$ span a simplex.

Finally, in order to prove (3), let $w_0, \dots, w_p$ be pairwise adjacent vertices in $\mathrm{lk}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$, ordered so that for some m we have $w_0 \dots, w_m \in \mathrm{lk}_{\downarrow}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$ and $w_{m+1} \dots, w_p \in \mathrm{lk}^{\uparrow}\left(v, \mathcal{C}(\widetilde{X_\Gamma})\right)$. By (1) we know that $w_0 \dots, w_m$ span a simplex, hence by Lemma 4.6 there exists a cell λ in $\cap_{j=0}^{m} \tau_j$. Similarly, by (2) we know that $w_{m+1} \dots, w_p$ span a simplex, hence by Lemma 4.6 there exists a cell μ containing $\tau_{m+1}, \dots, \tau_p$. Notice that $\lambda \subseteq \tau_i \subseteq \sigma \subseteq \tau_j \subseteq \mu$ for all $i = 0, \dots, m$ and $j = m+1, \dots, p$. In particular we have $\lambda \subseteq \tau_j \subseteq \mu$ for all $j = 0, \dots, p$. Using Lemma 4.6 again we obtain that $w_0, \dots, w_p$ spans a simplex. □

4.2 Efficiency

In this section we study a notion of complexity for edge–paths in $\mathcal{C}(\widetilde{X_\Gamma})$, which is based on the height function, and use it to find suitable representatives of homotopy classes of edge–paths and edge–loops. Recall that if p is an edge–path in the 1–skeleton of a cubical complex, an *elementary homotopy* of p is a homotopy which

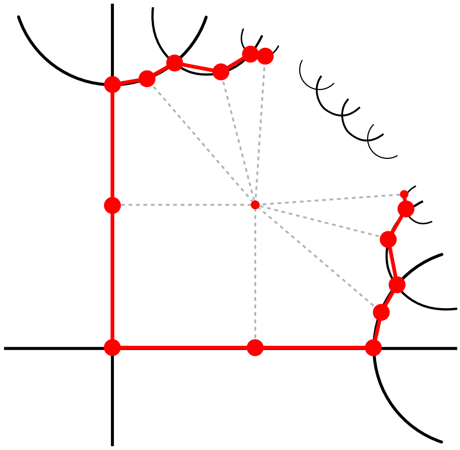

Fig. 28 A dual tile in $\mathcal{C}(\widetilde{X_\Gamma})$, and a long edge–path that stays in a tile

is contained in the 2–skeleton and is obtained by a finite sequence of the following two moves:

- remove a backtracking subpath, i.e. replace (v_1, v_2, v_1) with v_1;
- slide across a square, i.e. replace (v_1, v_2, v_3) with (v_1, v_4, v_3) if v_1, v_2, v_3, v_4 appear in this order on the boundary of a square (as in Fig. 24).

An edge–path $p = (v_0, \dots, v_s)$ is said to be *efficient* if $\exists\, k \in \{0, \dots, s\}$ such that the height strictly increases from v_0 to v_k and strictly decreases from v_k to v_s, i.e. v_k is the unique point of maximum for the height along p. We allow $k = 0$ or $k = s$, i.e. that the height is strictly monotone along p. In any case, $\mathrm{h}(p) = \mathrm{h}(v_k)$, and the cell dual to v_k contains the cells dual to all the other vertices of p. This implies that an efficient edge–path is contained in the union of at most two cubes which share at least a vertex. In particular, an efficient edge–loop is entirely contained in a single cube. These observations motivate the following definitions and constructions.

If τ is a tile of $\widetilde{X_\Gamma}$, we define the *dual tile* $\mathcal{C}(\tau)$ to be the full subcomplex of $\mathcal{C}(\widetilde{X_\Gamma})$ whose vertices are dual to the cells of τ. If v is the vertex of $\mathcal{C}(\widetilde{X_\Gamma})$ which is dual to τ, then $\mathcal{C}(\tau)$ consists of all the cubes of $\mathcal{C}(\widetilde{X_\Gamma})$ that contain v, i.e. $\mathcal{C}(\tau)$ is the combinatorial 1–neighborhood of v. Notice that v is the only vertex of height n in $\mathcal{C}(\tau)$ (see Fig. 28). We say that an edge–path p in $\mathcal{C}(\widetilde{X_\Gamma})$ *stays in a tile* if there exists a tile τ of $\widetilde{X_\Gamma}$ such that $p \subseteq \mathcal{C}(\tau)$.

Lemma 4.11 *Let p be an edge–path in $\mathcal{C}(\widetilde{X_\Gamma})$. If p stays in a tile, then there is an elementary homotopy relative to endpoints between p and an efficient path.*

Proof Let $p = (v_0, \dots, v_s)$. First of all, notice that if $s = 0, 1$ then p is already efficient. Moreover, by an elementary homotopy relative to endpoints, we can assume that p has no backtracking subpath. Since p stays in a tile, p goes through at most one vertex of height n (possibly several times, possibly at the endpoints v_0, v_s).

For $0 < j < s$, we say v_j is a local minimum (with respect to the height function along p) if $\mathrm{h}(v_{j\pm 1}) = \mathrm{h}(v_j) + 1$, and we consider the following quantity

$$\mathfrak{h}(p) = \min\{\mathrm{h}(v_j) \mid v_j \text{ is a local minimum}\}.$$

If there is no local minimum, set $\mathfrak{h}(p) = \infty$; in this case p is already efficient. So let us assume that there are some local minima, i.e. $\mathfrak{h}(p) < \infty$. Notice that $\mathfrak{h}(p)$ is in general larger than the minimum of the height along p. If $\mathfrak{h}(p) = n$ then p is constant, hence efficient. If $\mathfrak{h}(p) = n-1$, then p has a backtracking subpath, because p goes through at most one vertex of height n. By an elementary homotopy relative to endpoints we can remove this local minimum. Repeating this process, we obtain a path p' with $\mathfrak{h}(p') = n$, and we reduce to the previous case. So let us assume in the following that $\mathfrak{h}(p) \leq n-2$.

We now claim that, by deforming p locally at local minima, we can produce an elementary homotopy relative to endpoints to a path p' such that $\mathfrak{h}(p') \geq \mathfrak{h}(p)+1$. To prove the claim, let v_j be a local minimum, and let its height be $\mathrm{h}(p_j) = h_j$ for some $0 < j < s$. Consider the subpath (v_{j-1}, v_j, v_{j+1}), and note that the cells dual to v_{j-1}, v_{j+1} meet along the cell dual to v_j. Since p stays in a tile, there is a cell containing all these cells, namely the tile itself. By Lemma 4.6 we get that (v_{j-1}, v_j, v_{j+1}) is part of a square in $\mathcal{C}(\widetilde{X_\Gamma})$, whose fourth vertex is some v_j', of height $\mathrm{h}(v_j') = h_j+2$. Then we can homotope (v_{j-1}, v_j, v_{j+1}) to the other side (v_{j-1}, v_j', v_{j+1}) of the square via an elementary homotopy relative to endpoints (see Lemma 4.3). This process can be applied to all local minima at the same time, since no two local minima can be adjacent along p. Then we remove all backtracking subpaths, if needed, keeping endpoints fixed. The result is an elementary homotopy relative to endpoints between p and an edge–path p' with $\mathfrak{h}(p') \geq \mathfrak{h}(p) + 1$. It is even possible that $\mathfrak{h}(p') = \infty$ but in any case this proves the claim.

We repeat this process of elevating local minima, and after a finite number of steps we obtain a path p'' with $\mathfrak{h}(p'') \geq n-1$ (again, possibly $\mathfrak{h}(p'') = \infty$). Hence, we reduce to the previously discussed cases to conclude that p'' (hence p) admits an elementary homotopy relative to endpoints to an efficient path. □

In the previous lemma we allow p to be an edge–loop, i.e. $v_0 = v_s$. All the homotopies in it are relative to the base point $v_0 = v_s$. In the following statement we consider free homotopies, i.e. homotopies that are not required to fix any point.

Corollary 4.12 *Let p be an edge–loop in $\mathcal{C}(\widetilde{X_\Gamma})$. If p stays in a tile, then there is an elementary homotopy between p and a constant path.*

Proof Pick a basepoint v_0 on p to be a vertex of maximal height on p, and write $p = (v_0, \ldots, v_s)$, for $v_0 = v_s$. Apply the previous argument (from Lemma 4.11) to p. At every iteration we allow ourselves to change the basepoint on p to always be a vertex of maximal height. At the end there can be no local minimum, hence the path is constant. □

A simple way for an edge-loop to satisfy the condition of Corollary 4.12 is to be short. Recall from §4.1 that the length $\ell(p)$ of an edge–path p is defined to be the number of edges of p.

Corollary 4.13 *Let p be an edge–loop in $\mathcal{C}(\widetilde{X_\Gamma})$. Then $\ell(p)$ is even. Moreover, if $\ell(p) \leq 4$ then p stays in a tile, and there is an elementary homotopy between p and a constant path.*

Proof The first statement follows from the fact that if an edge e has endpoints v, w then $|h(v) - h(w)| = 1$, so if an edge–path has odd length then the endpoints have different height.

Suppose now $\ell(p) \leq 4$. If $\ell(p) = 2$ then $p = (v, w, v)$ for two adjacent vertices v, w. In particular the cell dual to v contains the one dual to w, or vice versa. If $\ell(p) = 4$ then p is the boundary path of a square. It follows from Lemma 4.3 that p contains a unique point of maximal height, and that the cell dual to it contains every other cell. In either case, there is a cell containing all the cells dual to the vertices of p. If τ is a tile of $\widetilde{X_\Gamma}$ containing that cell, then p is entirely contained in $\mathcal{C}(\tau)$ by construction. In particular, p stays in a tile, so the statement follows from Corollary 4.12. □

Remark 4.14 From now on, our main goal in this section will be to show that every edge–loop in $\mathcal{C}(\widetilde{X_\Gamma})$ can be written as a product of nullhomotopic edge–loops, i.e. $\mathcal{C}(\widetilde{X_\Gamma})$ is simply connected. A naive approach would consist in splitting an edge–loop along mirrors into shorter edge–paths, until they are short enough to be contracted (in the sense of Corollary 4.12). However, there are arbitrarily long edge–paths that stay in a tile (see Fig. 28). Therefore, an inductive argument based on length alone would not suffice, and this idea requires some additional tools which we develop in §4.3, before returning to the problem of simple connectedness of $\mathcal{C}(\widetilde{X_\Gamma})$ in §4.4.

Remark 4.15 Given two cells σ, σ' contained in the same tile τ, let $\mu = \mu(\sigma, \sigma')$ be their upper cell (i.e. the smallest cell that contains both of them, as defined in Lemma 3.19). If v, v' and w are the vertices dual to σ, σ' and μ respectively, then an edge–path of minimal length in $\mathcal{C}(\widetilde{X_\Gamma})$ from v to v' can be obtained as an efficient path p in $\mathcal{C}(\tau)$ going through w. Such an efficient edge–path is not unique, but the length of any such path is given by

$$\ell(p) = 2\,\mathrm{h}(w) - \mathrm{h}(v) - \mathrm{h}(v') = 2\dim\mu - \dim\sigma - \dim\sigma'.$$

It should be noted that if $\mu \subsetneq \tau$ then there are edge–paths from v to v' which are strictly longer than p but still efficient.

4.3 Mirror complexity

Here we define an additional notion of complexity for an edge–path, based on the relative position in $\widetilde{X_\Gamma}$ between mirrors and the cells dual to the vertices of the edge–path. We start with the following definition, in analogy to that of a dual tile. If M is a mirror of $\widetilde{X_\Gamma}$, we define the *dual mirror* $\mathcal{C}(M)$ to be the full subcomplex of $\mathcal{C}(\widetilde{X_\Gamma})$ whose vertices are dual to the cells of M. Since we have not proved yet that $\mathcal{C}(\widetilde{X_\Gamma})$ is simply connected, a priori it is not clear that a dual mirror enjoys properties reminiscent of those of a mirror of $\widetilde{X_\Gamma}$; for instance, it is not clear yet whether it is convex. Nevertheless, we can obtain the following statement about separation (analogous to Proposition 3.37).

Lemma 4.16 *Let M be a mirror of $\widetilde{X_\Gamma}$ and let $\mathcal{C}(M)$ be the dual mirror in $\mathcal{C}(\widetilde{X_\Gamma})$. Let z_1, z_2 be two points in $\widetilde{X_\Gamma} \setminus M$, let σ_1, σ_2 be cells in $\widetilde{X_\Gamma}$ such that $z_k \in \sigma_k$, and* *let v_k be the vertex of $\mathcal{C}(\widetilde{X_\Gamma})$ dual to σ_k. Then M separates z_1 and z_2 if and only if $\mathcal{C}(M)$ separates v_1 and v_2. In particular, $\mathcal{C}(M)$ separates $\mathcal{C}(\widetilde{X_\Gamma})$.*

Proof Suppose M separates z_1 and z_2, and assume by contradiction that there is an edge–path p in $\mathcal{C}(\widetilde{X_\Gamma})$ from v_1 to v_2 avoiding $\mathcal{C}(M)$. Then the union of the cells dual to the vertices of p contains a path–connected subspace of $\widetilde{X_\Gamma} \setminus M$ that contains both z_1 and z_2. This is in contradiction with the fact that M separates z_1 from z_2.

Vice versa, suppose $\mathcal{C}(M)$ separates v_1 and v_2, and assume by contradiction that there is a path γ in $\widetilde{X_\Gamma}$ from z_1 to z_2 avoiding M. By a small perturbation, we can assume that γ intersects the strata of $\widetilde{X_\Gamma}$ in such a way that the sequence of the minimal cells that it visits gives rise to an edge–path in $\mathcal{C}(\widetilde{X_\Gamma})$ (i.e. their dimension jumps by 1 at a time along γ). By construction, such an edge–path connects v_1 to v_2 in the complement of $\mathcal{C}(M)$, which is not possible.

In particular, it follows that $\mathcal{C}(M)$ separates $\mathcal{C}(\widetilde{X_\Gamma})$, because M separates $\widetilde{X_\Gamma}$ by Proposition 3.37. □

This provides a correspondence between complementary components of a mirror M in $\widetilde{X_\Gamma}$ and complementary components of the dual mirror $\mathcal{C}(M)$ in $\mathcal{C}(\widetilde{X_\Gamma})$.

4.3.1 Crossings

Let $p = (v_0, \dots, v_s)$ be an edge–path (possibly an edge–loop) in $\mathcal{C}(\widetilde{X_\Gamma})$, and let $\sigma_0, \dots, \sigma_s$ be the cells of $\mathcal{C}(\widetilde{X_\Gamma})$ dual to its vertices. Let M be a mirror in $\widetilde{X_\Gamma}$, and let $\mathcal{C}(M)$ be the dual mirror in $\mathcal{C}(\widetilde{X_\Gamma})$. Recall from Lemma 4.16 that $\mathcal{C}(\widetilde{X_\Gamma}) \setminus \mathcal{C}(M)$ is disconnected. We say that p *crosses* M if $p \cap \mathcal{C}(M) \neq \varnothing$ and there are at least two connected components C_1, C_2 of $\mathcal{C}(\widetilde{X_\Gamma}) \setminus \mathcal{C}(M)$ such that $p \cap C_k \neq \varnothing$. This means that among the cells $\sigma_0, \dots, \sigma_s$, some are contained in M, but at least two of them are such that their interiors are contained in different complementary components of M. (Recall that in our setting cells are closed and complementary components of mirrors are open.) Let $q = (v_j, \dots, v_{j+m})$ be a subpath of p. We say that q is a (p, M)*–crossing* if $v_j, \dots, v_{j+m} \in \mathcal{C}(M)$, but v_{j-1} and v_{j+m+1} lie in different connected components of $\mathcal{C}(\widetilde{X_\Gamma}) \setminus \mathcal{C}(M)$. (See Fig. 29 for some examples.) We denote by $\mathfrak{m}(p, M)$ the number of (p, M)–crossings. The *mirror complexity* of p is defined by taking into account the family $\mathcal{M}$ of all mirrors of $\widetilde{X_\Gamma}$, i.e. by the following formula:

$$\mathfrak{m}(p) = \sum_{M \in \mathcal{M}} \mathfrak{m}(p, M).$$

The relevance of this notion of complexity with respect to Remark 4.14 is showcased by the following two lemmas.

Lemma 4.17 *Let p be an edge–path in $\mathcal{C}(\widetilde{X_\Gamma})$. Then p stays in a tile if and only if p does not cross any mirror.*

Proof Suppose that p stays in a tile, i.e. there exists a tile τ such that $p \subseteq \mathcal{C}(\tau)$. Assume by contradiction that p crosses a mirror M. So there are two vertices v_1, v_2 of p which are separated by $\mathcal{C}(M)$. Let σ_k be the cell dual to v_k. By Lemma 4.16 M

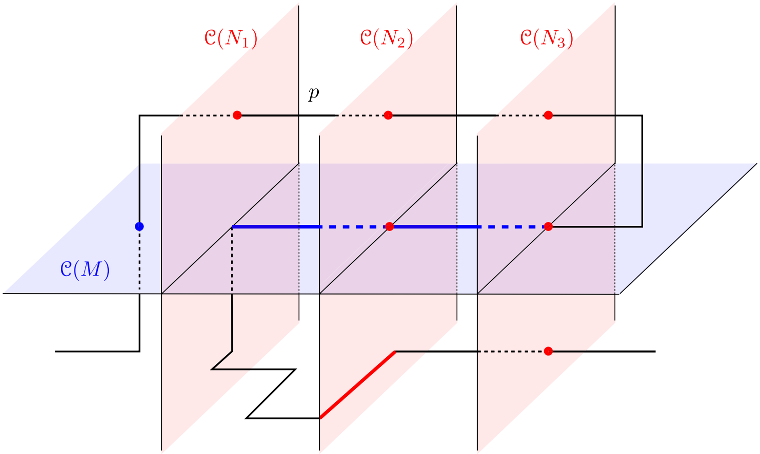

Fig. 29 An edge–path p in $\mathcal{C}(\widetilde{X_\Gamma})$ crossing some mirrors. Mirror crossings are highlighted. We have $\mathfrak{m}(p, M) = 2$, $\mathfrak{m}(p, N_1) = 1$, $\mathfrak{m}(p, N_2) = 3$, and $\mathfrak{m}(p, N_3) = 3$. In particular, notice that even if p intersects $\mathcal{C}(N_1)$ twice, there is only one (p, N_1)-crossing

separates the interior of σ_1 from the interior of σ_2. In particular, there is no tile of $\widetilde{X_\Gamma}$ that contains both of them, which contradicts the hypothesis that p stays in a tile.

Vice versa, suppose p does not cross any mirror, and assume by contradiction that there are two vertices v_1, v_2 on p such that the dual cells are not contained in the same tile. Let τ_1, τ_2 be different tiles containing them. Up to choosing v_1, v_2 closer to each other along p, we can assume that the tiles are adjacent, i.e. $\tau_1 \cap \tau_2 \neq \varnothing$. In particular, $\sigma = \tau_1 \cap \tau_2$ is a cell and it is contained in some mirror M. Then p intersects M between v_1 and v_2. Moreover the tiles τ_1, τ_2 provide a framing in the sense of §3.7. Proposition 3.37 implies that the interiors of τ_1, τ_2 are separated by M. The same holds for the interiors of the cells dual to v_1, v_2. So by Lemma 4.16 we have that v_1, v_2 are separated by $\mathcal{C}(M)$, i.e. p crosses M, a contradiction. ☐

Lemma 4.18 *Let p be an edge–path in $\mathcal{C}(\widetilde{X_\Gamma})$, and let M be a mirror in $\widetilde{X_\Gamma}$. Then the following hold.*

(1) $\mathfrak{m}(p, M) = 0$ *if and only if p does not cross M.*
(2) $\mathfrak{m}(p) = 0$ *if and only if p stays in a tile.*
(3) *If p is a loop and* $\mathfrak{m}(p, M) \geq 1$, *then* $\mathfrak{m}(p, M) \geq 2$.
(4) *If p has finite length, then* $\mathfrak{m}(p, M)$ *and* $\mathfrak{m}(p)$ *are finite.*

Proof For (1), note that $\mathfrak{m}(p, M)$ is by definition the number of (p, M)-crossings. For (2), note that $\mathfrak{m}(p)$ is a sum of non-negative numbers, so it is zero if and only if $\mathfrak{m}(p, M) = 0$ for every mirror M. By (1) this is equivalent to saying that p does not cross any mirror. Then the statement follows from Lemma 4.17. To prove (3), note that if p is a loop that crosses M at least once, then it must cross it at least twice, because $\mathcal{C}(M)$ separates $\mathcal{C}(\widetilde{X_\Gamma})$ by Lemma 4.16.

Finally, to prove (4) notice that each (p, M)–crossing contributes to at least one vertex of p, dual to a cell of M. Since p has finite length, there can be only finitely many (p, M)–crossings. Then the finiteness of $\mathfrak{m}(p)$ follows from the fact that X (hence the collection of mirrors $\mathcal{M}$) is locally finite. □

Remark 4.19 In this framework, Corollary 4.13 can be restated by saying that $\ell(p) \leq 4$ implies $\mathfrak{m}(p) = 0$.

4.3.2 Bridges

Let $p = (v_0, \dots, v_s)$ be an edge–path in $\mathcal{C}(\widetilde{X_\Gamma})$, and let $\sigma_0, \dots, \sigma_s$ be the cells of $\widetilde{X_\Gamma}$ dual to its vertices. Let M be a mirror in $\widetilde{X_\Gamma}$, and let $\mathcal{C}(M)$ be the dual mirror in $\mathcal{C}(\widetilde{X_\Gamma})$. We say that p is a *bridge* if there exists a mirror M of $\widetilde{X_\Gamma}$ such that $v_0, v_s \in \mathcal{C}(M)$, but $p \nsubseteq \mathcal{C}(M)$. In other words, $\sigma_0, \sigma_s \subseteq M$ but some of the other cells $\sigma_1, \dots, \sigma_{s-1}$ are not contained in M. In this case, we say that p is *supported* by M. We say p is a *minimal bridge* if none of its subpaths is a bridge (see Fig. 30).

Lemma 4.20 *Let p be an edge–path in $\mathcal{C}(\widetilde{X_\Gamma})$. If p is a bridge, then there exists a subpath $q \subseteq p$ that is a minimal bridge.*

Proof Let us consider the collection of subpaths of p which are bridges. Notice that this collection contains p itself, it is partially ordered by inclusion, and it is finite. Therefore it contains a minimal element. □

Lemma 4.21 *Let p be a minimal bridge supported by a mirror M, and let N be a mirror such that $\mathfrak{m}(p, N) > 0$. Then the following hold.*

(1) $\mathfrak{m}(p, N) = 1$.
(2) $\mathcal{C}(M) \cap \mathcal{C}(N) \neq \varnothing$ *and* $M \cap N \neq \varnothing$.

Proof The assumption that $\mathfrak{m}(p, N) > 0$ means that p crosses N at least once. If p crossed N twice, then any subpath between two consecutive (p, N)-crossings would be a bridge supported by N. But this would contradict minimality, hence $\mathfrak{m}(p, N) = 1$, which proves (1). In particular the endpoints of p lie in different connected components of $\mathcal{C}(\widetilde{X_\Gamma}) \setminus \mathcal{C}(N)$. Since they also live on the same dual mirror

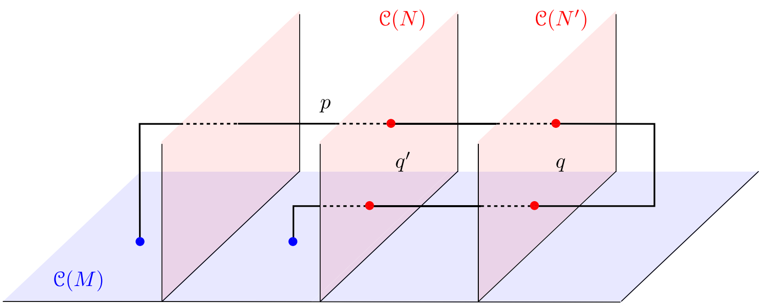

Fig. 30 A bridge p supported by a mirror M, with subpaths q and q' which are bridges supported by the mirrors N and N' respectively. Notice that only q is a minimal bridge

$\mathcal{C}(M)$, which is connected, and dual mirrors separate by Lemma 4.16, we can conclude that $\mathcal{C}(M) \cap \mathcal{C}(N) \neq \varnothing$. Finally, the cell dual to a vertex in their intersection is contained in $M \cap N$, hence we obtain (2). □

Recall that for a mirror M in $\widetilde{X_\Gamma}$ we have a nearest point projection $\pi_M : \widetilde{X_\Gamma} \to M$, as discussed in §3.6. If p is a minimal bridge supported on M, then we can use π_M to induce a projection of p to $\mathcal{C}(M)$, as established by the next results.

Lemma 4.22 *Let M, N be mirrors in $\widetilde{X_\Gamma}$ and let τ be a tile in $\widetilde{X_\Gamma}$.*

(1) *If $M \cap N \neq \varnothing$, then $\pi_M(N) = M \cap N$.*
(2) *If $M \cap \tau \neq \varnothing$, then $\pi_M(\tau) = M \cap \tau$.*

Proof We start by proving $\pi_M(N) \subseteq M \cap N$. By contradiction, let $x \in N$ be such that $\pi_M(x) \notin N$. Let $y = \pi_{M\cap N}(x) \in M \cap N$, where $\pi_{M\cap N}$ denotes the nearest point projection to the closed convex subspace $M \cap N$. Since $\pi_M(x) \notin N$, we have $y \neq \pi_M(x)$, so we can consider the geodesic triangle with vertices x, $\pi_M(x)$ and y. By convexity of M, the geodesic $[\pi_M(x), y]$ is contained in M. By convexity of N, the geodesic $[x, y]$ is contained in N. Moreover, since π_M is the nearest point projection to M, the angle between $[x, \pi_M(x)]$ and $[\pi_M(x), y]$ at $\pi_M(x)$ is at least $\frac{\pi}{2}$. Analogously, the angle between $[x, y]$ and any geodesic in $M \cap N$ at y is at least $\frac{\pi}{2}$ too, and since N meets M orthogonally, this is enough to ensure that the angle between $[x, y]$ and $[\pi_M(x), y]$ at y is at least $\frac{\pi}{2}$ too. We obtained a geodesic triangle with two angles at least $\frac{\pi}{2}$, which is impossible in the CAT(0) space $\widetilde{X_\Gamma}$. Vice versa, if $x \in M \cap N$, then $x = \pi_M(x)$, so $x \in \pi_M(N)$ already.

The second statement can be obtained via an analogous argument. Indeed, recall that τ is isometric to a convex subset of $\mathbb{H}^n$ bounded by orthogonal hyperplanes (see Lemma 3.11). In particular, the nearest point projection to a boundary face of τ is entirely contained in τ. □

The next lemma is a combinatorial statement about the stratification of $\widetilde{X_\Gamma}$ introduced in §3.5, and will be needed in the following lemma.

Lemma 4.23 *Let τ, τ' be non-disjoint tiles of $\widetilde{X_\Gamma}$. Let $W_1, \dots, W_q$ be the collection of mirrors of $\widetilde{X_\Gamma}$ that separate τ and τ'. Then we have that*

(1) *$W_1, \dots, W_q$ coincides with the collection of mirrors of $\widetilde{X_\Gamma}$ that contain $\tau \cap \tau'$.*
(2) *$\tau \cap W_1 \cap \cdots \cap W_q = \tau \cap \tau' = \tau' \cap W_1 \cap \cdots \cap W_q$.*

Proof First of all, notice that the collection of mirrors is not empty since τ and τ' are different tiles. We start by proving (1). Let W be a mirror containing $\tau \cap \tau'$. Then the two tiles provide a framing for the cell $\tau \cap \tau'$. In particular we get from Proposition 3.37 that W separates the two tiles, hence W is in the collection $\{W_1, \dots, W_q\}$. Conversely, if $\tau \cap \tau'$ was not inside one W_i, then we could connect the two tiles with a path that goes through the intersection but avoids W_i, contradicting the fact that W_i separates them.

To prove (2) we argue as follows. By (1) we know that $\tau \cap \tau' \subseteq W_1 \cap \cdots \cap W_q$, so we have that $\tau \cap \tau' \subseteq \tau \cap W_1 \cap \cdots \cap W_q$. Now note that, by definition of the stratification, if $\tau \cap \tau'$ is a k-cell, then it must be contained in $n-k$ mirrors, so $q = n-k$. But then the two sides of the inclusion are cells of the same dimension k, so they must be equal. Switching the roles of τ and τ' proves the second equality in (2). □

Lemma 4.24 *Let τ be a tile and M be a mirror in $\widetilde{X_\Gamma}$, such that $M \cap \tau \neq \varnothing$. Let σ be a cell of τ, and let $N_1, \dots, N_r \neq M$ be all the mirrors containing σ and such that $M \cap N_j \neq \varnothing$ for $j = 1, \dots, r$. (Possibly $r = 0$ if there are no such mirrors.) Then the following hold.*

(1) *$\tau \cap M \cap N_1 \cap \cdots \cap N_r$ is an $(n-1-r)$–cell that contains $\pi_M(\sigma)$.*
(2) *The cell $\tau \cap M \cap N_1 \cap \cdots \cap N_r$ only depends on σ and M.*

Proof We start by proving (1). It follows from Lemma 4.22 that $\pi_M(\sigma) \subseteq \tau \cap M \cap N_j$ for each $j = 1, \dots, r$. So, we obtain that $\pi_M(\sigma) \subseteq \tau \cap M \cap N_1 \cap \cdots \cap N_r$. To show that this intersection is a cell, note that τ is an n–cell. So, by Lemma 3.21 we have that $M \cap \tau$ is an $(n-1)$–cell and then for each $j = 1, \dots, r$ we have that $\tau \cap M \cap N_1 \cap \cdots \cap N_j$ is an $(n-1-j)$–cell.

To prove (2) we argue as follows. Suppose τ' is another tile as in the statement, i.e. $\sigma \subseteq \tau'$ and $M \cap \tau' \neq \varnothing$. Let $W_1, \dots, W_q$ be the collection of mirrors of $\widetilde{X_\Gamma}$ that separate τ and τ'. (Note that this collection depends on τ and τ', while the collection $N_1, \dots, N_r$ only depends on σ and M.) Since $\sigma \subseteq \tau \cap \tau'$, we also have that σ is contained in each W_i thanks to (1) in Lemma 4.23. We now claim that each W_i meets M. This is clear if $\sigma \subseteq M$. On the other hand, if σ is not inside M, then we can take an efficient edge–path p in $\mathcal{C}(\widetilde{X_\Gamma})$ from the vertex dual to σ to the vertex dual to $M \cap \tau$ which is contained in $\mathcal{C}(\tau)$ and meets $\mathcal{C}(M)$ only at the endpoint $M \cap \tau$. Take an analogous path p' in $\mathcal{C}(\tau')$, and concatenate p and p' to obtain a minimal bridge $\widehat{p}$ supported on M. Since W_i separates τ and τ', we see that $\widehat{p}$ crosses W_i. So by (2) in Lemma 4.21 we conclude that $M \cap W_i \neq \varnothing$, which proves the claim.

As a result, we have that the collection $\{W_1, \dots, W_q\}$ is a subcollection of $\{M, N_1, \dots, N_r\}$. (Note that M could be one of the mirrors separating τ and τ', but $N_i \neq M$ by definition.) In particular, using (2) from 4.23, we obtain that

$$\tau \cap M \cap N_1 \cap \cdots \cap N_r \subseteq \tau \cap W_1 \cap \cdots \cap W_r \overset{4.23}{=} \tau \cap \tau' \subseteq \tau'.$$

Therefore it follows that $\tau \cap M \cap N_1 \cap \cdots \cap N_r \subseteq \tau' \cap M \cap N_1 \cap \cdots \cap N_r$. Reversing the roles of τ and τ' provides the other inclusion, and shows that the cell defined in (1) does not depend on the choice of the tile. □

In the notation and setting of Lemma 4.24, if $v \in \mathcal{C}(\widetilde{X_\Gamma})$ is the vertex dual to σ, then we denote by $\pi_M(v)$ the vertex dual to the cell constructed in (1) of the lemma, and call it the *projection of v to $\mathcal{C}(M)$*. This is well defined by (2) in the same lemma. Note that in general $0 \leq r \leq n - \dim \sigma$, as σ could be contained in some mirrors that are disjoint from M. Nevertheless, this provides the desired projection to $\mathcal{C}(M)$ for vertices of $\mathcal{C}(\widetilde{X_\Gamma})$ which are contained in the cubical 2–neighborhood of $\mathcal{C}(M)$, i.e. the union of all the dual tiles corresponding to all the tiles that intersect M in $\widetilde{X_\Gamma}$.

The content of the next two lemmas is that a minimal bridge supported by a mirror M is completely contained in such a neighborhood of $\mathcal{C}(M)$ (see Lemma 4.25), so we can define a projection of a minimal bridge to $\mathcal{C}(M)$ (see Lemma 4.26). We note that the minimality assumption is necessary, see the difference between q and q' in Fig. 30.

Lemma 4.25 *Let p be a minimal bridge supported on a mirror M. Then for each vertex v of p there exists a tile τ such that $v \in \mathcal{C}(\tau)$ and $\tau \cap M \neq \varnothing$.*

Proof Let $p = (v_0, \dots, v_s)$, let σ_k be the cell of $\widetilde{X_\Gamma}$ dual to v_k, and assume by contradiction that some vertices do not satisfy the statement. Let v_k be the first one. Since p is a bridge, its endpoints are on $\mathcal{C}(M)$, so $k \neq 0, s$. Let $\tau_\pm$ be tiles such that $v_{k-1} \in \mathcal{C}(\tau_-)$ and $v_k \in \mathcal{C}(\tau_+)$. In particular this means that $\sigma_{k-1} \subseteq \tau_-$ and $\sigma_k \subseteq \tau_+$. By construction, we can choose τ_- so that $\tau_- \cap M \neq \varnothing$, while necessarily τ_+ is disjoint from M. Moreover, if we had $\sigma_k \subseteq \sigma_{k-1}$ then we would have $v_k \in \mathcal{C}(\tau_-)$, against the choice of v_k. But since v_{k-1} and v_k are adjacent in $\mathcal{C}(\widetilde{X_\Gamma})$, this forces $\sigma_{k-1} \subseteq \sigma_k$. In particular, the intersection $\tau_- \cap \tau_+$ is non empty: it contains at least the cell σ_{k-1}.

Consider the cell $\sigma = \tau_- \cap \tau_+$. For any mirror N containing σ, we claim that N must intersect M. Indeed, the tiles $\tau_\pm$ form a framing for σ in the sense of §3.7. By Proposition 3.37 we know that $\tau_\pm$ belong to the closure of distinct complementary components of N. In particular, a maximal subpath of $p \cap \mathcal{C}(N)$ whose vertices are dual to cells contained in σ gives rise to a (p, N)-crossing, hence $\mathfrak{m}(p, N) > 0$. By (2) in Lemma 4.21 we know $M \cap N \neq \varnothing$, as claimed.

Let $N_1, \dots, N_r$ be the collection of all mirrors containing σ. Since σ is a cell of τ_-, we can write $\sigma = \tau_- \cap N_1 \cap \dots \cap N_r$. Using (1) in Lemma 4.24, we have that

$$\pi_M(\sigma) \subseteq \tau_- \cap M \cap N_1 \cap \dots \cap N_r =$$

$$= (M \cap \tau_-) \cap (\tau_- \cap N_1 \cap \dots \cap N_r) \subseteq M \cap \sigma \subseteq M \cap \tau_+.$$

This contradicts the fact that τ_+ is disjoint from M. □

In the next lemma we finally obtain a projection of a minimal bridge to a supporting mirror. As it might be expected, such a projection is length–decreasing (see Fig. 31).

Lemma 4.26 *Let p be a minimal bridge supported on a mirror M. Then there exists an edge–path $p^M \subseteq \mathcal{C}(M)$, such that p^M has the same endpoints as p and $\ell(p^M) \leq \ell(p) - 2$.*

Proof Let $p = (v_0, \dots, v_s)$, and let $\sigma_0, \dots, \sigma_s$ be the cells dual to its vertices. Since p is a minimal bridge supported on M, by Lemma 4.25 we know that for each vertex v_k there exists a tile τ_k of $\widetilde{X_\Gamma}$ such that $v_k \in \mathcal{C}(\tau_k)$ and $\tau_k \cap M \neq \varnothing$. Let $w_k = \pi_M(v_k)$ be the projection of v_k to $\mathcal{C}(M)$, constructed in Lemma 4.24. We claim that for each k the vertices w_{k-1} and w_k are either the same vertex or adjacent vertices.

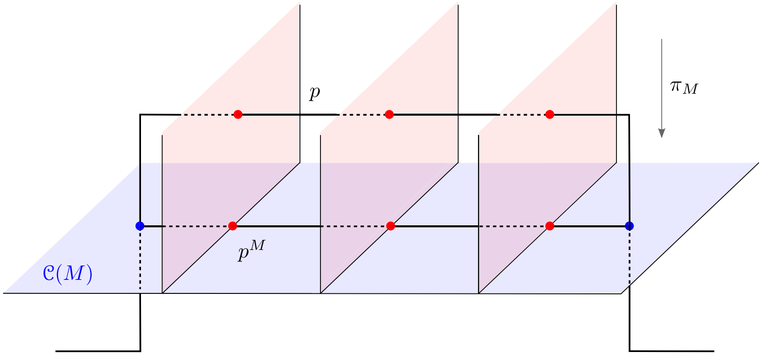

Fig. 31 Projection of the minimal bridge p to $\mathcal{C}(M)$, for a supporting mirror M

To see this, consider two vertices v_{k-1} and v_k adjacent along p. Without loss of generality (i.e., possibly reversing the orientation of p) we can assume that σ_{k-1} is a cell of codimension 1 in σ_k. In particular, we can take $\tau_{k-1} = \tau_k$, and there is exactly one mirror $\widehat{N}_k$ that contains σ_{k-1} but not σ_k. Let $\{N_1, \dots, N_r\}$ be the collection of all the mirrors that contain σ_k and intersect M, but are different from M. Then the analogous collection for σ_{k-1} consists of the same mirrors, possibly with the addition of $\widehat{N}_k$. (Note that since p is a minimal bridge supported on M, any mirror containing $\sigma_1, \dots, \sigma_{s-1}$ is guaranteed to be different from M, while $\widehat{N}_k = M$ for $k = 1$.) By (1) in Lemma 4.24 we have that $\pi_M(\sigma_k) \subseteq \tau_k \cap M \cap N_1 \cap \dots \cap N_r$ and that either $\pi_M(\sigma_{k-1}) \subseteq \tau_k \cap M \cap N_1 \cap \dots \cap N_r$ or $\pi_M(\sigma_{k-1}) \subseteq \tau_k \cap M \cap N_1 \cap \dots \cap N_r \cap \widehat{N}_k$. In the first case we have that $\pi_M(\sigma_{k-1})$ and $\pi_M(\sigma_k)$ are contained in the intersection of the same mirrors, hence $w_{k-1} = w_k$; in the second case $\tau_k \cap M \cap N_1 \cap \dots \cap N_r \cap \widehat{N}_k$ is a codimension-1 cell of $\tau_k \cap M \cap N_1 \cap \dots \cap N_r$, hence w_{k-1} is adjacent to w_k. This proves the claim.

Notice in particular that in the case $k = 1$ we have $\widehat{N}_k = M$, so we have proved that $w_0 = w_1$. Analogously, we also have $w_s = w_{s-1}$. As a result, $(w_0, \dots, w_s)$ is an edge–path in $\mathcal{C}(M)$. Let p^M be the edge–path obtained from $(w_0, \dots, w_s)$ by removing all backtracking subpaths and repeated vertices. In particular, since $w_0 = w_1$ and $w_s = w_{s-1}$, we have that $\ell(p^M) \leq s - 2 = \ell(p) - 2$. Moreover, since p is a bridge supported on M, we have that $\sigma_0, \sigma_s \subseteq M$, so that $v_0 = w_0$, $v_s = w_s$, i.e. p and p^M have the same endpoints. □

4.4 Surgeries on edge–loops

We are now ready to apply the above technology to the study of edge–loops in $\mathcal{C}(\widetilde{X_\Gamma})$. The goal is to show that $\mathcal{C}(\widetilde{X_\Gamma})$ is simply connected. The strategy will be to reduce the length and mirror complexity of an edge–loop enough to ensure that it stays in a tile, so that Corollary 4.12 can be applied. The following statement is the key surgery step. Roughly speaking, we chop an edge–loop p along a mirror M that it crosses, use the projection p^M of p to M to introduce a shortcut along M and obtain two edge–loops p_1, p_2 such that p and $p_1 p_2$ are elementary homotopic, and finally then check that the lengths have dropped.

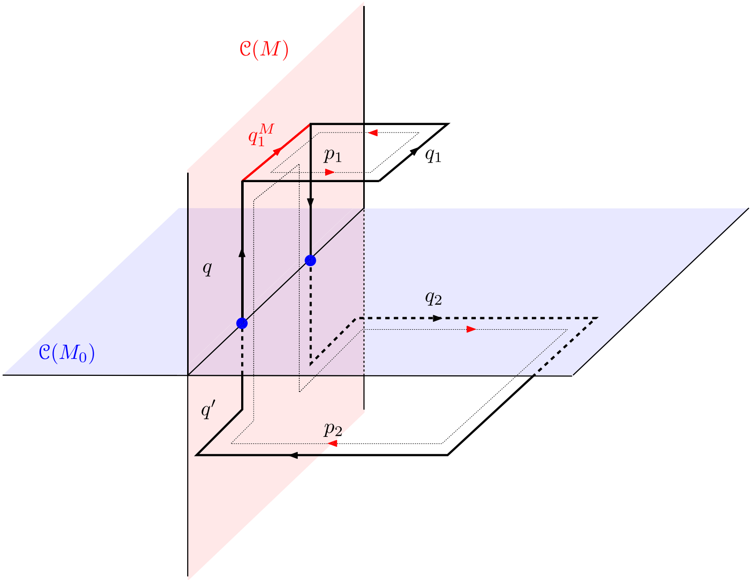

Fig. 32 The surgery in Proposition 4.27. Note that all the minimal bridges in p are supported on M but p does not cross M, hence the need to first split p into q and q', and then into q_1 and q_2

Proposition 4.27 *Let p be an edge–loop in $\mathcal{C}(\widetilde{X_\Gamma})$. If $\mathfrak{m}(p) > 0$, then there exist two edge–loops p_1, p_2 in $\mathcal{C}(\widetilde{X_\Gamma})$ such that $\ell(p_1), \ell(p_2) < \ell(p)$, and there is an elementary homotopy between p and $p_1 p_2$*

Proof By assumption, there is a mirror M_0 that is crossed by p, so $\mathfrak{m}(p, M_0) \geq 1$, hence by (3) in Lemma 4.18 we have that $\mathfrak{m}(p, M_0) \geq 2$, i.e. p crosses M_0 at least twice. It follows from the definitions that any subarc of p between any two (p, M_0)-crossings is a bridge supported by M_0.

Choose a (p, M_0)–crossing and a bridge q supported by M_0 (in general this cannot be chosen to be minimal, see Fig. 32). Let q' be the complement of q in p, i.e. the edge–path such that $p = qq'$. Of course we have

$$\ell(q) + \ell(q') = \ell(p). \tag{4.1}$$

Moreover, without loss of generality we can assume that

$$\ell(q) \leq \ell(q'). \tag{4.2}$$

By Lemma 4.20 we can find a minimal bridge $q_1 \subseteq q \subseteq p$. In particular we have

$$\ell(q_1) \leq \ell(q). \tag{4.3}$$

Let q_2 be the complement of q_1 in p, i.e. the edge–path such that $p = q_1 q_2$. We can compute that

$$\ell(q_2) = \ell(p) - \ell(q_1) \overset{(4.3)}{\geq} \ell(p) - \ell(q) \overset{(4.1)}{=} \ell(q') \overset{(4.2)}{\geq} \ell(q) \overset{(4.3)}{\geq} \ell(q_1). \tag{4.4}$$

Let M be a mirror supporting the minimal bridge q_1, and let q_1^M be the projection of q_1 to $\mathcal{C}(M)$, i.e. the edge–path obtained in Lemma 4.26. In particular we have

$$\ell(q_1^M) \leq \ell(q_1) - 2 < \ell(q_1). \tag{4.5}$$

Define the edge–loops $p_1 = q_1\overline{q_1^M}$ and $p_2 = q_1^M q_2$, where $\overline{q_1^M}$ denotes the edge–path q_1^M with the opposite orientation. There is an elementary homotopy between the edge–loops $p = q_1 q_2$ and $p_1 p_2 = q_1\overline{q_1^M} q_1^M q_2$, obtained by removing the backtracking subpath $\overline{q_1^M} q_1^M$. We can compute the desired inequality on the length of p_1 and p_2 as follows:

$$\ell(p_1) = \ell(q_1) + \ell(q_1^M) \overset{(4.5)}{<} \ell(q_1) + \ell(q_1) \overset{(4.4)}{\leq} \ell(q_1) + \ell(q_2) = \ell(p).$$

$$\ell(p_2) = \ell(q_1^M) + \ell(q_2) \overset{(4.5)}{<} \ell(q_1) + \ell(q_2) = \ell(p).$$ □

Remark 4.28 Note that it is possible to have an edge–loop p for which the surgery from Proposition 4.27 strictly reduces the mirror complexity for only one subloop. For an example see Fig. 32, where the mirror complexity of the loop p_2 is the same as that of the original loop p.

We are now ready to prove the main result of this section.

Theorem 4.29 *The complex* $\mathcal{C}(\widetilde{X_\Gamma})$ *is a connected* CAT(0) *cubical complex.*

Proof By construction, the complex $\mathcal{C}(\widetilde{X_\Gamma})$ is a cubical complex. Moreover, Gromov's link condition from Lemma 2.2 implies that $\mathcal{C}(\widetilde{X_\Gamma})$ is non–positively curved, since the link of any vertex is a flag simplicial complex by Proposition 4.10.

Next, $\mathcal{C}(\widetilde{X_\Gamma})$ is path–connected because $\widetilde{X_\Gamma}$ is path–connected. Indeed, let v, w be vertices in $\mathcal{C}(\widetilde{X_\Gamma})$, and let σ_v, σ_w be the dual cells. Pick any continuous path η connecting the two cells in $\widetilde{X_\Gamma}$, and keep track of the list of cells that are intersected by η. By isotoping η into lower-dimensional cells, we can ensure that the difference between the dimension of two consecutive cells in this list is exactly 1. The dual vertices in $\mathcal{C}(\widetilde{X_\Gamma})$ give rise to an edge–path from v to w.

To conclude, we need to show that $\mathcal{C}(\widetilde{X_\Gamma})$ is simply connected. To do this, we argue that edge–loops are nullhomotopic by induction on their length. Let p be an edge–loop in $\mathcal{C}(\widetilde{X_\Gamma})$, homotopically non–trivial and of minimal length. If p does not cross any mirror, then by Lemma 4.17 p stays in a tile. Hence by Corollary 4.12 there is an elementary homotopy between p and a constant path. So let us assume that $\mathfrak{m}(p) > 0$. Then by Proposition 4.27 there exist two edge–loops p_1, p_2 in $\mathcal{C}(\widetilde{X_\Gamma})$ such that p is homotopic to $p_1 p_2$ and $\ell(p_1), \ell(p_2) < \ell(p)$. By minimality, both p_1 and p_2 are homotopically trivial, and so is p. □

We conclude this section by noting that the action of the hyperbolized group $\Gamma_X = \pi_1(X_\Gamma)$ on $\widetilde{X_\Gamma}$ induces an action on the dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$.

Lemma 4.30 *The group $\Gamma_X = \pi_1(X_\Gamma)$ acts on $\mathcal{C}(\widetilde{X_\Gamma})$ by cubical isometries. Moreover, if X is compact, then the action is cocompact.*

Proof The group Γ_X acts on $\widetilde{X_\Gamma}$ preserving the family of mirrors, hence the stratification defined in §3.5. The action permutes the cells, so Γ_X acts on vertices of the dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ described in §4. Moreover, the action of Γ_X on $\widetilde{X_\Gamma}$ preserves the incidence relations between cells, hence we can extend the action of Γ_X on vertices to a combinatorial action of Γ_X on the entire $\mathcal{C}(\widetilde{X_\Gamma})$. Since Γ_X acts on $\mathcal{C}(\widetilde{X_\Gamma})$ combinatorially, it preserves the standard cubical metric.

When X is compact, X_Γ is compact as well, by (3) in Lemma 2.5. The action of Γ_X on $\widetilde{X_\Gamma}$ has finitely many orbits of cells, so its action on $\mathcal{C}(\widetilde{X_\Gamma})$ has finitely many orbits of vertices. Since $\mathcal{C}(\widetilde{X_\Gamma})$ is finite-dimensional (see Lemma 4.5), the quotient has finitely many cubes, so it is compact. □

5 Special cubulation

The purpose of this section is to study the action of the hyperbolized group $\Gamma_X = \pi_1(X_\Gamma)$ on the dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ (see Lemma 4.30), and prove that the group Γ_X is virtually compact special in the sense of [39]. When X is compact and admissible, Γ_X is a Gromov hyperbolic group and $\mathcal{C}(\widetilde{X_\Gamma})$ is a CAT(0) cubical complex (see (4) in Proposition 3.5, and Theorem 4.29). Therefore, one could hope to obtain virtual specialness directly from Agol's result from [2]. However, the action of Γ_X on $\mathcal{C}(\widetilde{X_\Gamma})$ is not proper (see Remark 5.3). To remedy this, we will use a result of Groves and Manning from [35, Theorem D] that deals with improper actions. This requires a study of stabilizers of cubes.

In §5.1 we show that cube stabilizers for the action of Γ_X on $\mathcal{C}(\widetilde{X_\Gamma})$ coincide with cell stabilizers for the action of Γ_X on $\widetilde{X_\Gamma}$. Then in §5.2 we show that such stabilizers are quasiconvex and virtually compact special. The complex X is always assumed to be admissible in the sense of §3. In some statements (such as Theorem 5.15) we also assume that it is compact.

Remark 5.1 (Why we consider the action on $\mathcal{C}(\widetilde{X_\Gamma})$ instead of $\widetilde{X}$) It is worth noting that when X is admissible, $\widetilde{X}$ is already a CAT(0) cube complex. Moreover the map $g_X : X_\Gamma \to X$ from Proposition 3.5 induces a surjection $\Gamma_X \to \pi_1(X)$ that can be used to obtain an action of Γ_X on $\widetilde{X}$. However, this action has a very large kernel. For example, in the case in which X is already simply connected the map $\Gamma_X \to \pi_1(X)$ is trivial, but Γ_X is an infinite group; indeed, it retracts to $\Gamma_{\square^n}$, as discussed in Remark 3.9.

5.1 Cube stabilizers for the action of Γ_X on $\mathcal{C}(\widetilde{X_\Gamma})$

In this section we relate the cube stabilizers for the action of Γ_X on $\mathcal{C}(\widetilde{X_\Gamma})$ to the cell stabilizers for the action of Γ_X on $\widetilde{X_\Gamma}$.

Lemma 5.2 *The stabilizer of a vertex in $\mathcal{C}(\widetilde{X_\Gamma})$ coincides with the stabilizer of its dual cell in $\widetilde{X_\Gamma}$.*

Remark 5.3 It follows from Lemma 5.2 that the action of Γ_X on $\mathcal{C}(\widetilde{X_\Gamma})$ is in general not proper. Namely, the stabilizer of a vertex dual to a cell of dimension at least 2 is infinite (compare Remark 4.8 and Fig. 11).

We now proceed to the study of stabilizers of higher–dimensional cubes for the action of Γ_X on $\mathcal{C}(\widetilde{X_\Gamma})$. Recall that the dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ is equipped with a Γ_X-invariant height function: the vertex dual to a k-cell has height k. We proved in Lemma 4.4 that every cube in $\mathcal{C}(\widetilde{X_\Gamma})$ has a unique vertex of minimal height.

Lemma 5.4 *The stabilizer of a cube in $\mathcal{C}(\widetilde{X_\Gamma})$ coincides with the stabilizer of its vertex of minimal height in $\widetilde{X_\Gamma}$.*

Proof Let C be a cube in $\mathcal{C}(\widetilde{X_\Gamma})$, let v be its vertex of minimal height, and let σ be the dual cell in $\widetilde{X_\Gamma}$. Let $g \in \Gamma_X$ be an element that stabilizes C. Since the height function is invariant, g must fix v, by uniqueness of the vertex of minimal height.

Conversely let g fix v. By Lemma 5.2 we get that g stabilizes σ, i.e. $g.\sigma = \sigma$. Let w be another vertex of C and let τ be the dual cell. By Lemma 4.4, we have that $\sigma \subseteq \tau$. It follows that $\sigma = g.\sigma \subseteq g.\tau$, so that both τ and $g.\tau$ appear in the link of σ in the combinatorial structure of $\widetilde{X_\Gamma}$ (see (3) in Lemma 3.17). Since the covering projection $\pi : \widetilde{X_\Gamma} \to X_\Gamma$ induces isomorphisms on links, if τ and $g.\tau$ were distinct, then in X_Γ we would see a tile $\pi(\tau) = \pi(g.\tau)$ isometric to a copy of $\square^n_\Gamma$ with some identification along the boundary (namely along the subspace corresponding to $\pi(\sigma)$). However, by (1) in Lemma 2.5 we know that tiles of X_Γ are embedded copies of $\square^n_\Gamma$, so we must have $g.\tau = \tau$. By Lemma 5.2, this means $g.w = w$. Therefore g fixes C pointwise. □

Remark 5.5 In the proof of Lemma 5.4 we established that the stabilizer of a cell in $\mathcal{C}(\widetilde{X_\Gamma})$ actually fixes the cell pointwise.

5.2 Cell stabilizers are quasiconvex and virtually compact special

The goal of this section is to study the stabilizers of cells for the action of Γ_X on $\widetilde{X_\Gamma}$ by covering transformations. In particular, note that by Lemma 3.11 stabilizers of tiles (i.e. top–dimensional cells) are isomorphic to the fundamental group of the hyperbolizing cube $\Gamma_{\square^n} = \pi_1(\square^n_\Gamma)$. More precisely, our goal is to show that cell stabilizers for the action of Γ_X on $\widetilde{X_\Gamma}$ are quasiconvex in Γ_X, and virtually compact special.

5.2.1 Quasiconvexity

In the following we say that an action of a group on a metric space is *geometric* if it is proper, cocompact and isometric. We will make use of the following standard fact.

Lemma 5.6 *Let Z be a proper Gromov hyperbolic geodesic space, and let $Y \subseteq Z$ be a quasiconvex subset. Let G be a finitely generated group acting geometrically on Z, and let H be the stabilizer of Y in G. If H acts cocompactly on Y, then H is quasiconvex in G.*

We apply this lemma to the cases $G = \Gamma_X$, $H = \Gamma_{\square^n}$ and $G = \Gamma$, $H = \Gamma_{\square^n}$. As noted, Γ_X is a Gromov hyperbolic group when X is compact. In both cases, before using the lemma we need to ensure that H is a subgroup of G. This is not obvious, because a priori H is just defined as the fundamental group of the hyperbolizing cube $\square^n_\Gamma$.

Lemma 5.7 *Let X be compact. Then $\Gamma_{\square^n}$ is a quasiconvex subgroup of Γ_X.*

Proof By Lemma 2.5, we know that a hyperbolized complex retracts to each of its tiles, each of which is homeomorphic to the hyperbolizing cell. In our setting this means that X_Γ retracts to $\square^n_\Gamma$, so in particular the inclusion $\square^n_\Gamma \hookrightarrow X_\Gamma$ as a tile induces an injection $\Gamma_{\square^n} \hookrightarrow \Gamma_X$. Since X is compact, the group Γ_X acts geometrically on $\widetilde{X_\Gamma}$. Moreover, the subgroup $\Gamma_{\square^n}$ acts geometrically on a tile, which is a closed convex subspace by Lemma 3.16. Therefore $\Gamma_{\square^n}$ is quasiconvex by Lemma 5.6. $\square$

Lemma 5.8 *The group $\Gamma_{\square^n}$ is a quasiconvex subgroup of Γ.*

Proof First of all we will prove that $\Gamma_{\square^n}$ naturally injects in Γ, by showing that there exists a (normal) cover Y_Γ of $M_\Gamma = \mathbb{H}^n/\Gamma$ which retracts to $\square^n_\Gamma$ (see Fig. 33). This would provide the desired injection

$$\Gamma_{\square^n} = \pi_1(\square^n_\Gamma) \hookrightarrow \pi_1(Y_\Gamma) \hookrightarrow \pi_1(M_\Gamma) = \Gamma.$$

To construct this cover, consider the cubical complex Y given by the standard cubulation of $\mathbb{R}^n$ with vertices on $\mathbb{Z}^n$. Notice that Y admits a standard folding $f : Y \to \square^n$, and that Y is an admissible cubical complex. Therefore we can consider the hyperbolized complex Y_Γ. As in the proof of Lemma 5.7, Lemma 2.5 implies that Y_Γ retracts onto any of its tiles, hence $\Gamma_{\square^n}$ injects in $\pi_1(Y_\Gamma)$.

We now claim that Y_Γ is a (normal) covering space of M_Γ. For each $i = 1, \ldots, n$ consider the mirror of Y given by $M_i = \{y_i = 0\}$, and the hyperplane of Y given by $H_i = \{y_i = \frac{1}{2}\}$. Let m_i and h_i be the reflections in M_i and H_i respectively, i.e.

$$m_i : Y \to Y, m_i(y_1, \ldots, y_i, \ldots, y_n) = (y_1, \ldots, -y_i, \ldots, y_n),$$

$$h_i : Y \to Y, h_i(y_1, \ldots, y_i, \ldots, y_n) = (y_1, \ldots, 1 - y_i, \ldots, y_n).$$

For each $i = 1, \ldots, n$, the group $D_i = \langle m_i, h_i \rangle$ is an infinite dihedral group of cubical isometries of Y. The group $D = \langle m_1, h_1, \ldots m_n, h_n \rangle$ is isomorphic to the direct

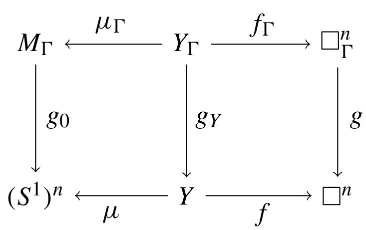

Fig. 33 Y_Γ as a covering space of M_Γ that retracts to $\square^n_\Gamma$

product $D_1 \times \cdots \times D_n$, and admits a representation into the group B_n of Euclidean isometries of the standard cube $\square^n$, in which m_i acts trivially and h_i acts as the standard reflection of $\square^n$ in the i–th coordinate. By Lemma 3.1 we have an action of B_n on $\square^n_\Gamma$ by isometries, hence we can induce an action of D on $\square^n_\Gamma$ by isometries such that the Charney–Davis map $g : \square^n_\Gamma \to \square^n$ is equivariant. Moreover, the standard folding $f : Y \to \square^n$ is clearly D–equivariant too, because it can be obtained by reflecting in the mirrors of Y. Since the two maps in the pullback square defining Y_Γ are D–equivariant (see Fig. 33), we obtain an action of D on Y_Γ by isometries.

Note that $t_i = h_i m_i$ is the unit integer translation of Y in the ith direction. As a result, D contains a (normal) subgroup T isomorphic to the group of integer translations $\mathbb{Z}^n$. The action of D on Y_Γ restricts to a free and properly discontinuous action of T on Y_Γ. A fundamental domain for this action is given by a single tile. Each tile is isometric to a hyperbolizing cube $\square^n_\Gamma$, and the induced action identifies corresponding cells on opposite mirrors, recovering M_Γ (see §3.1.1 for more details about the construction of $\square^n_\Gamma$.) In particular $\mu_\Gamma : Y_\Gamma \to Y_\Gamma / T \cong M_\Gamma$ realizes the desired covering map, which covers the standard universal covering map $\mu : Y = \mathbb{R}^n \to (S^1)^n$ (see Fig. 33).

Finally, let us prove that $\Gamma_{\square^n}$ is quasiconvex in Γ. We know Γ acts geometrically on $\mathbb{H}^n$, permuting the stratification induced by the coordinate mirrors and their translates. The subgroup $\Gamma_{\square^n}$ stabilizes a Γ–cell, i.e. the closure of a connected component of the complement of such collection. This is a closed convex subspace, on which $\Gamma_{\square^n}$ acts geometrically (see §3.3 for details). In particular $\Gamma_{\square^n}$ is quasiconvex in Γ by Lemma 5.6. $\square$

Remark 5.9 In Lemma 5.8 we have constructed a normal covering space of M_Γ by producing an action of $T = \mathbb{Z}^n$ by deck transformations on the hyperbolization Y_Γ of the standard integral cubulation of $\mathbb{R}^n$. This covering space can also be defined as the covering space of M_Γ corresponding to the kernel of the homomorphism $\Gamma = \pi_1(M_\Gamma) \to \mathbb{Z}^n$ induced by the collapse map $g_0 : M_\Gamma \to (S^1)^n$ obtained by applying the Pontryagin–Thom construction to M_Γ with respect to suitable codimension–1 submanifolds (see §3.1.1 for details, and Fig. 33). Compact quotients of Y_Γ, provide examples of closed hyperbolized manifolds which are finite covers of M_Γ. These are all genuine arithmetic hyperbolic manifolds.

5.2.2 Virtual specialness

A cubical complex is *special* if it admits a local isometry into the Salvetti complex of a right-angled Artin group (see [39]). A group G is *virtually compact special* if there exist a finite index subgroup $G' \subseteq G$ and a compact special cubical complex B such that $G' = \pi_1(B)$. This property passes from a Gromov hyperbolic group to its quasiconvex subgroups, as established in the following statement. This kind of arguments has appeared in the literature (see for instance [39, Corollary 7.8]). We include a proof for the reader's convenience.

Lemma 5.10 *Let G be a Gromov hyperbolic group, and let H be a quasiconvex subgroup. If G is virtually compact special, then so is H.*

Proof Let G' be a finite index subgroup of G and B a compact special cubical complex such that $G' = \pi_1(B)$. By [39, Remark 3.4, Lemma 3.13] we can assume without loss of generality that B is also non–positively curved. The universal cover $\widetilde{B}$ is a CAT(0) cubical complex. It is finite dimensional, uniformly locally finite, and Gromov hyperbolic, because G' acts geometrically on it by covering transformations.

Let $H' = H \cap G'$. This is a finite–index subgroup of H, and a quasiconvex subgroup of G'. Since G' is Gromov hyperbolic and acts geometrically on $\widetilde{B}$, it follows that H'–orbits are quasiconvex. By [37, Theorem H, or Corollary 2.29] or [67, Theorem 1.2], there exists a cocompact convex core for H', i.e. a convex subcomplex $\widetilde{Y} \subseteq \widetilde{B}$ on which H' acts cocompactly. Moreover, H' acts by deck transformations, and the quotient $Y = \widetilde{Y}/H'$ is a compact non–positively curved cubical complex with $\pi_1(Y) = H'$. The convex embedding $\widetilde{Y} \hookrightarrow \widetilde{B}$ descends to a local isometry $Y \to B$. Since B is special, by [39, Corollary 3.9] we obtain that Y is special too. Therefore H is virtually compact special, as desired. □

Remark 5.11 In the previous proof we have the Gromov hyperbolic group H' acting geometrically on the CAT(0) cubical complex $\widetilde{Y}$, so the fact that H' is virtually compact special also follows from the celebrated theorem of Agol in [2]. However here everything happens inside the special group G', so one does not need Agol's result.

We now apply the previous lemma to the cell stabilizers for the action of $\square^n_\Gamma$ on $\widetilde{X_\Gamma}$, starting with the stabilizer of a tile.

Lemma 5.12 *The group $\Gamma_{\square^n}$ is virtually compact special.*

Proof $\Gamma_{\square^n}$ is a quasiconvex subgroup of Γ by Lemma 5.8 and Γ is virtually compact special by [40, Theorem 1.6]. Indeed, it is an arithmetic lattice in $\mathrm{SO}_0(n, 1)$ by construction (see §3.1 or [14] for details). The statement then follows from Lemma 5.10. □

Finally we prove the same result for all cell stabilizers.

Lemma 5.13 *Let X be compact. Then the cell stabilizers for the action of Γ_X on $\widetilde{X_\Gamma}$ are quasiconvex and virtually compact special.*

Proof Let σ be a cell in $\widetilde{X_\Gamma}$ and let H be the stabilizer of σ for the action of Γ_X on $\widetilde{X_\Gamma}$. Since σ is a convex subset of $\widetilde{X_\Gamma}$ and H acts geometrically on it, we conclude by Lemma 5.6 that H is quasiconvex in Γ_X.

Arguing as in the proof of Lemma 5.4, if τ is a tile containing σ, and K is its stabilizer, then $H \subseteq K$. Note that the folding map $X_\Gamma \to \square^n_\Gamma$ provides an isomorphism of $K \cong \Gamma_{\square^n}$, under which H is isomorphic to a quasiconvex subgroup of $\Gamma_{\square^n}$ (again by Lemma 5.6). We know that $\Gamma_{\square^n}$ is Gromov hyperbolic (by Lemma 5.7 or Lemma 5.8) and virtually compact special (by Lemma 5.12). So it follows from Lemma 5.10 that H is virtually compact special too. □

5.3 Specialization

We are now ready to prove that the fundamental group Γ_X of the hyperbolized complex X_Γ is virtually compact special, when the original cubical complex X is admissible and compact. If the action of Γ_X on $\mathcal{C}(\widetilde{X_\Gamma})$ was proper, this would follow from Theorem 4.29, Lemma 4.30 and Agol's main result from [2]. However, as observed in Remark 5.3, the action on $\mathcal{C}(\widetilde{X_\Gamma})$ is **not** proper. We will use a result by Groves and Manning (see Theorem D in [35]), which is designed to deal with this situation. We report here their statement for the reader's convenience.

Theorem 5.14 [35, *Theorem D*] *Suppose that G is a Gromov hyperbolic group acting cocompactly on a* CAT(0) *cubical complex so that cell stabilizers are quasiconvex and virtually compact special. Then G is virtually compact special.*

Note that when authors of [35] say "virtually special" they imply that the quotient is compact (see page 3 in [35]). Also notice that they explicitly do not assume their complexes to be locally compact (see page 2).

Theorem 5.15 *If X is a compact admissible cubical complex and Γ is a hyperbolizing lattice, then Γ_X is virtually compact special Gromov hyperbolic group.*

Proof First of all, since X is admissible, by Theorem 4.29 the dual cubical complex $\mathcal{C}(\widetilde{X_\Gamma})$ is a CAT(0) cubical complex. Moreover, since X is compact, Γ_X is a Gromov hyperbolic group by (4) in Proposition 3.5. By Lemma 4.30 Γ_X acts on $\mathcal{C}(\widetilde{X_\Gamma})$ cocompactly by isometries.

Let C be a cube in $\mathcal{C}(\widetilde{X_\Gamma})$ and let H be its stabilizer. By Lemma 5.4 H coincides with the stabilizer of the vertex of minimal height in C. By Lemma 5.2 this in turn coincides with the stabilizer of the corresponding dual cell in $\widetilde{X_\Gamma}$. Therefore by Lemma 5.13 H is a quasiconvex subgroup of Γ_X and it is also virtually compact special. Finally, by [35, Theorem D] the group Γ_X is virtually compact special. □

Acknowledgements We thank Igor Belegradek, Dick Canary, Mike Davis, Jingyin Huang, Ben McReynolds, Nick Miller for useful conversations, and in particular Genevieve Walsh and Kevin Schreve for pointing out the arguments for Lemma 1.15. J.-F. Lafont was partly supported by the NSF Grant number DMS-1812028 and DMS-2109683. L. Ruffoni acknowledges support by the AMS, the Simons Foundation, the INDAM group GNSAGA, and thanks the Department of Mathematics at The Ohio State University for the hospitality during some visits in which this work was conducted. We thank the referee for their very careful reading and their suggestions.

References

1. Agol, I.: Criteria for virtual fibering. J. Topol. **1**(2), 269–284 (2008). https://doi.org/10.1112/jtopol/jtn003
2. Agol, I.: The virtual Haken conjecture. Doc. Math. **18**, 1045–1087 (2013). With an appendix by Agol, Daniel Groves, and Jason Manning. https://doi.org/10.4171/dm/421
3. Baik, H., Kim, S.-H., Koberda, T.: Right-angled Artin groups in the C^∞ diffeomorphism group of the real line. Isr. J. Math. **213**(1), 175–182 (2016). https://doi.org/10.1007/s11856-016-1307-8
4. Ballmann, W., Świątkowski, J.: On groups acting on nonpositively curved cubical complexes. Enseign. Math. (2) **45**(1–2), 51–81 (1999)
5. Bartels, A., Lück, W.: The Borel conjecture for hyperbolic and CAT(0)-groups. Ann. Math. (2) **175**(2), 631–689 (2012). https://doi.org/10.4007/annals.2012.175.2.5
6. Bartels, A., Lück, W., Weinberger, S.: On hyperbolic groups with spheres as boundary. J. Differ. Geom. **86**(1), 1–16 (2010). https://doi.org/10.4310/jdg/1299766682
7. Baum, P., Connes, A.: Geometric K-theory for Lie groups and foliations. Enseign. Math. (2) **46**(1–2), 3–42 (2000). https://doi.org/10.4171/LEM
8. Belegradek, I.: Aspherical manifolds, relative hyperbolicity, simplicial volume and assembly maps. Algebraic Geom. Topol. **6**, 1341–1354 (2006). https://doi.org/10.2140/agt.2006.6.1341
9. Belegradek, I.: Aspherical manifolds with relatively hyperbolic fundamental groups. Geom. Dedic. **129**, 119–144 (2007). https://doi.org/10.1007/s10711-007-9199-8
10. Bergeron, N., Wise, D.T.: A boundary criterion for cubulation. Am. J. Math. **134**(3), 843–859 (2012). https://doi.org/10.1353/ajm.2012.0020
11. Bridson, M.R., Haefliger, A.: Metric Spaces of Non-positive Curvature. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, pp. xxii+643. Springer, Berlin (1999). https://doi.org/10.1007/978-3-662-12494-9
12. Brown, K.S.: Cohomology of Groups. Graduate Texts in Mathematics, vol. 87, pp. x+306. Springer, New York (1982). https://doi.org/10.1007/978-1-4684-9327-6
13. Bux, K., Witzel, S.: Local convexity in CAT(k)-spaces. arXiv e-prints (2012). arXiv:1211.1871 [math.MG]
14. Charney, R.M., Davis, M.W.: Strict hyperbolization. Topology **34**(2), 329–350 (1995). https://doi.org/10.1016/0040-9383(94)00027-I
15. Cherix, P.-A., Martin, F., Valette, A.: Spaces with measured walls, the Haagerup property and property (T). Ergod. Theory Dyn. Syst. **24**(6), 1895–1908 (2004). https://doi.org/10.1017/S0143385704000185
16. Chern, S.-S.: On curvature and characteristic classes of a Riemann manifold. Abh. Math. Semin. Univ. Hamb. **20**, 117–126 (1955). https://doi.org/10.1007/BF02960745
17. Danciger, J., Guéritaud, F., Kassel, F.: Proper affine actions for right-angled Coxeter groups. Duke Math. J. **169**(12), 2231–2280 (2020). https://doi.org/10.1215/00127094-2019-0084
18. Davis, M.W., Januszkiewicz, T.: Hyperbolization of polyhedra. J. Differ. Geom. **34**(2), 347–388 (1991). https://doi.org/10.4310/jdg/1214447212
19. Davis, M.W., Januszkiewicz, T.: Right-angled Artin groups are commensurable with right-angled Coxeter groups. J. Pure Appl. Algebra **153**(3), 229–235 (2000). https://doi.org/10.1016/S0022-4049(99)00175-9
20. Davis, M.W., Januszkiewicz, T., Weinberger, S.: Relative hyperbolization and aspherical bordisms: an addendum to “Hyperbolization of polyhedra” [J. Differential Geom. **34** (1991), no. 2, 347–388; MR 1131435 (92h:57036)] by Davis and Januszkiewicz. J. Differ. Geom. **58**(3), 535–541 (2001). https://doi.org/10.4310/jdg/1090348358
21. Davis, M.W., Fowler, J., Lafont, J.-F.: Aspherical manifolds that cannot be triangulated. Algebraic Geom. Topol. **14**(2), 795–803 (2014). https://doi.org/10.2140/agt.2014.14.795
22. Deraux, M.: A negatively curved Kähler threefold not covered by the ball. Invent. Math. **160**(3), 501–525 (2005). https://doi.org/10.1007/s00222-004-0414-z
23. Donnelly, H., Xavier, F.: On the differential form spectrum of negatively curved Riemannian manifolds. Am. J. Math. **106**(1), 169–185 (1984). https://doi.org/10.2307/2374434
24. Douba, S., Fléchelles, B., Weisman, T., Zhu, F.: Cubulated hyperbolic groups admit Anosov representations. arXiv e-prints (2023). arXiv:2309.03695 [math.GR]
25. Duchamp, G., Krob, D.: The lower central series of the free partially commutative group. Semigroup Forum **45**(3), 385–394 (1992). https://doi.org/10.1007/BF03025778
26. Farrell, F.T., Lafont, J.-F.: Finite automorphisms of negatively curved Poincaré duality groups. Geom. Funct. Anal. **14**(2), 283–294 (2004). https://doi.org/10.1007/s00039-004-0457-8

27. Fenchel, W.: Elementary Geometry in Hyperbolic Space. De Gruyter Studies in Mathematics, vol. 11, pp. xii+225. de Gruyter, Berlin (1989). With an editorial by Heinz Bauer. https://doi.org/10.1515/9783110849455
28. Fisher, S.P.: Improved algebraic fibrings. arXiv e-prints (2021). arXiv:2112.00397 [math.GR]
29. Fisher, D., Lafont, J.-F., Miller, N., Stover, M.: Finiteness of maximal geodesic submanifolds in hyperbolic hybrids. J. Eur. Math. Soc. **23**(11), 3591–3623 (2021). https://doi.org/10.4171/jems/1077
30. Freedman, M.H.: The topology of four-dimensional manifolds. J. Differ. Geom. **17**(3), 357–453 (1982). https://doi.org/10.4310/jdg/1214437136
31. Giralt, A.: Cubulation of Gromov-Thurston manifolds. Groups Geom. Dyn. **11**(2), 393–414 (2017). https://doi.org/10.4171/GGD/401
32. Gromov, M.: Hyperbolic groups. In: Essays in Group Theory. Math. Sci. Res. Inst. Publ., vol. 8, pp. 75–263. Springer, New York (1987). https://doi.org/10.1007/978-1-4613-9586-7_3
33. Gromov, M., Thurston, W.: Pinching constants for hyperbolic manifolds. Invent. Math. **89**(1), 1–12 (1987). https://doi.org/10.1007/BF01404671
34. Groves, D., Manning, J.F.: Special IMM groups. Bull. Lond. Math. Soc. **55**(2), 717–722 (2023). https://doi.org/10.1112/blms.12751
35. Groves, D., Manning, J.F.: Hyperbolic groups acting improperly. Geom. Topol. **27**(9), 3387–3460 (2023). https://doi.org/10.2140/gt.2023.27.3387
36. Groves, D., Lafont, J.-F., Fox Manning, J., Ruffoni, L.: Relative cubulation of relative strict hyperbolization. arXiv e-prints (2023). arXiv:2304.14946 [math.GR]
37. Haglund, F.: Finite index subgroups of graph products. Geom. Dedic. **135**, 167–209 (2008). https://doi.org/10.1007/s10711-008-9270-0
38. Haglund, F., Paulin, F.: Simplicité de groupes d'automorphismes d'espaces à courbure négative. In: The Epstein Birthday Schrift. Geom. Topol. Monogr., vol. 1, pp. 181–248. Geom. Topol. Publ., Coventry (1998). https://doi.org/10.2140/gtm.1998.1.181
39. Haglund, F., Wise, D.T.: Special cube complexes. Geom. Funct. Anal. **17**(5), 1551–1620 (2008). https://doi.org/10.1007/s00039-007-0629-4
40. Haglund, F., Wise, D.T.: A combination theorem for special cube complexes. Ann. Math. (2) **176**(3), 1427–1482 (2012). https://doi.org/10.4007/annals.2012.176.3.2
41. Hillman, J.A., Kochloukova, D.H.: Finiteness conditions and PD_r-group covers of PD_n-complexes. Math. Z. **256**(1), 45–56 (2007). https://doi.org/10.1007/s00209-006-0058-3
42. Hruska, G.C., Wise, D.T.: Finiteness properties of cubulated groups. Compos. Math. **150**(3), 453–506 (2014). https://doi.org/10.1112/S0010437X13007112
43. Huang, J., Wise, D.T.: Virtual specialness of certain graphs of special cube complexes. Math. Ann. **388**(1), 329–357 (2024). https://doi.org/10.1007/s00208-022-02527-0
44. Italiano, G., Martelli, B., Migliorini, M.: Hyperbolic 5-manifolds that fiber over S^1. Invent. Math. **231**(1), 1–38 (2023). https://doi.org/10.1007/s00222-022-01141-w
45. Jost, J., Xin, Y.L.: Vanishing theorems for L^2-cohomology groups. J. Reine Angew. Math. **525**, 95–112 (2000). https://doi.org/10.1515/crll.2000.071
46. Kapovich, M.: Noncoherence of arithmetic hyperbolic lattices. Geom. Topol. **17**(1), 39–71 (2013). https://doi.org/10.2140/gt.2013.17.39
47. Kapovich, M., Potyagailo, L., Vinberg, E.: Noncoherence of some lattices in $\mathrm{Isom}(\mathbb{H}^n)$. In: The Zieschang Gedenkschrift. Geom. Topol. Monogr., vol. 14, pp. 335–351. Geom. Topol. Publ., Coventry (2008). https://doi.org/10.2140/gtm.2008.14.335
48. Kielak, D.: Residually finite rationally solvable groups and virtual fibring. J. Am. Math. Soc. **33**(2), 451–486 (2020). https://doi.org/10.1090/jams/936
49. Kim, S.-H., Koberda, T.: Anti-trees and right-angled Artin subgroups of braid groups. Geom. Topol. **19**(6), 3289–3306 (2015). https://doi.org/10.2140/gt.2015.19.3289
50. Koberda, T.: Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups. Geom. Funct. Anal. **22**(6), 1541–1590 (2012). https://doi.org/10.1007/s00039-012-0198-z
51. Kropholler, R., Vidussi, S., Walsh, G.: Virtual algebraic fibrations of surface-by-surface groups and orbits of the mapping class group. arXiv e-prints (2021). arXiv:2103.06930 [math.GT]
52. Lafforgue, V.: K-théorie bivariante pour les algèbres de Banach et conjecture de Baum-Connes. Invent. Math. **149**(1), 1–95 (2002). https://doi.org/10.1007/s002220200213
53. Lück, W.: The relation between the Baum-Connes conjecture and the trace conjecture. Invent. Math. **149**(1), 123–152 (2002). https://doi.org/10.1007/s002220200215
54. Manolescu, C.: Pin(2)-equivariant Seiberg-Witten Floer homology and the triangulation conjecture. J. Am. Math. Soc. **29**(1), 147–176 (2016). https://doi.org/10.1090/jams829

55. Mess, G.: Examples of Poincaré duality groups. Proc. Am. Math. Soc. **110**(4), 1145–1146 (1990). https://doi.org/10.2307/2047770
56. Milnor, J.W., Stasheff, J.D.: Characteristic Classes. Annals of Mathematics Studies, vol. 76, pp. vii+331. Princeton University Press, Princeton (1974)
57. Mostow, G.D., Siu, Y.T.: A compact Kähler surface of negative curvature not covered by the ball. Ann. Math. (2) **112**(2), 321–360 (1980). https://doi.org/10.2307/1971149
58. Niblo, G., Reeves, L.: Groups acting on CAT(0) cube complexes. Geom. Topol. **1**, 1–7 (1997). https://doi.org/10.2140/gt.1997.1.1
59. Ontaneda, P.: Normal smoothings for Charney-Davis strict hyperbolizations. J. Topol. Anal. **9**(1), 127–165 (2017). https://doi.org/10.1142/S1793525317500054
60. Ontaneda, P.: Riemannian hyperbolization. Publ. Math. Inst. Hautes Études Sci. **131**, 1–72 (2020). https://doi.org/10.1007/s10240-020-00113-1
61. Paulin, F.: Constructions of hyperbolic groups via hyperbolizations of polyhedra. In: Group Theory from a Geometrical Viewpoint (Trieste, 1990), pp. 313–372. World Sci. Publ., River Edge (1991). https://doi.org/10.1142/1235
62. Ramos-Cuevas, C.: Convexity is a local property in CAT(κ) spaces. In: Mexican Mathematicians Abroad: Recent Contributions. Contemp. Math., vol. 657, pp. 189–196. Am. Math. Soc., Providence (2016). https://doi.org/10.1090/conm/657/13097
63. Roy, R.: The trace conjecture—a counterexample. K-Theory **17**(3), 209–213 (1999). https://doi.org/10.1023/A:1007787132573
64. Ruffoni, L.: Manifolds without real projective or flat conformal structures. Proc. Am. Math. Soc. **151**(8), 3611–3620 (2023). https://doi.org/10.1090/proc/16293
65. Sageev, M.: Ends of group pairs and non-positively curved cube complexes. Proc. Lond. Math. Soc. (3) **71**(3), 585–617 (1995). https://doi.org/10.1112/plms/s3-71.3.585
66. Sageev, M.: Codimension-1 subgroups and splittings of groups. J. Algebra **189**(2), 377–389 (1997). https://doi.org/10.1006/jabr.1996.6884
67. Sageev, M., Wise, D.T.: Cores for quasiconvex actions. Proc. Am. Math. Soc. **143**(7), 2731–2741 (2015). https://doi.org/10.1090/S0002-9939-2015-12297-6
68. Schreve, K.: The strong Atiyah conjecture for virtually cocompact special groups. Math. Ann. **359**(3–4), 629–636 (2014). https://doi.org/10.1007/s00208-014-1007-9
69. Scott, P., Wall, T.: Topological methods in group theory. In: Homological Group Theory (Proc. Sympos., Durham, 1977). London Math. Soc. Lecture Note Ser., vol. 36, pp. 137–203. Cambridge University Press, Cambridge (1979)
70. Shah, N.A.: Closures of totally geodesic immersions in manifolds of constant negative curvature. In: Group Theory from a Geometrical Viewpoint (Trieste, 1990), pp. 718–732. World Sci. Publ., River Edge (1991). https://doi.org/10.1142/1235
71. Stover, M., Toledo, D.: Residual finiteness for central extensions of lattices in PU(n, 1) and negatively curved projective varieties. Pure Appl. Math. Q. **18**(4), 1771–1797 (2022). https://doi.org/10.4310/PAMQ.2022.v18.n4.a15
72. Stover, M., Toledo, D.: Residually finite lattices in $\widetilde{\mathrm{PU}(2, 1)}$ and fundamental groups of smooth projective surfaces. Mich. Math. J. **72**, 559–597 (2022). https://doi.org/10.1307/mmj/20217215
73. Świątkowski, J.: Trees of manifolds as boundaries of spaces and groups. Geom. Topol. **24**(2), 593–622 (2020). https://doi.org/10.2140/gt.2020.24.593
74. Williams, R.F.: A useful functor and three famous examples in topology. Trans. Am. Math. Soc. **106**, 319–329 (1963). https://doi.org/10.2307/1993773
75. Wise, D.T.: The Structure of Groups with a Quasiconvex Hierarchy. Annals of Mathematics Studies, vol. 209, pp. x+357. Princeton University Press, Princeton (2021). https://doi.org/10.2307/j.ctv1574pr6
76. Xie, X.: Foldable cubical complexes of nonpositive curvature. Algebraic Geom. Topol. **4**, 603–622 (2004). https://doi.org/10.2140/agt.2004.4.603